FINAL TRANSCRIPT

Thomson StreetEvents℠

VZ - Verizon at Goldman Sachs Communacopia XIX Conference

Event Date/Time: Sep. 23. 2010 / 12:00PM GMT

THOMSON REUTERS

CORPORATE PARTICIPANTS

Ivan Seidenberg
Verizon Communications, Inc. - Chairman and CEO

CONFERENCE CALL PARTICIPANTS

Jason Armstrong
Goldman Sachs - Analyst

PRESENTATION

Jason Armstrong - *Goldman Sachs - Analyst*

Okay, let's go ahead and kick things off this morning. We have obviously had a great start to the last day of Communacopia. Joining us today to start the session is Ivan Seidenberg, the Chairman and Chief Executive Officer of Verizon. Welcome, Ivan.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Good morning, Jason. Glad to be here.

Jason Armstrong - *Goldman Sachs - Analyst*

Thanks for joining us today. Maybe start off with what is topical recently. We have obviously had some releases in the last couple of weeks around personnel changes. Love for you to sort of set the stage for us in terms of your perspective and sort of tie this all together.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Well, I think it is pretty easy. We have a wonderful selection as a Chief Operating Officer who someday will be CEO and Chairman of the Company, Lowell McAdam. He is an extraordinary guy. He has had a 10-year run at various jobs in wireless. I couldn't think of a better executive in the wireless industry that has produced better results and has a feel for the business and also wonderful character, good long-standing industry guy. So we are very fortunate.

We have a great choice as the CFO. A guy who has been Controller. He has had operating jobs in both the wireline and wireless business. He integrated MCI. So I couldn't think of a better choice for CFO.

As far as John Killian is concerned, I have worked with him for 26 years. He is a close colleague. He is a really good guy. He has done a great job. I have asked him to move four or five times including to the UK, took his whole family and moved to the UK for us. I think I asked him to move one too many times.

So I think I guess he told me the beginning of September, he is out of gas. He is not going to live in New Jersey and Massachusetts any longer. And what I appreciated a great deal is he told me early in the cycle as we were getting ready to sort of plan the final stages of transition. He said, let's make this easy, Ivan, I am ready to go. I sent him home, told him to think about it for two weeks. He came back in and said his mind is made up.

We announced it that afternoon and I thought that once John clarified his plans that we should move quickly to solidify the extraordinary bench we have because the last thing I needed was to put anybody in limbo when we had clarity about what we wanted to do.

So, Jason, I am delighted with the outcome here. I think one of my key jobs here is to make sure that we have a deep, energized, confident bench and I guess based on shareholders and the reaction of the Street to our investors, we have gotten a lot of calls from people who seem pleased with it.

And so the key thing is whether or not their final comment would be is everybody happy with Lowell coming or me going? I am not sure of that but either way, it works for our shareholders.

Jason Armstrong - *Goldman Sachs - Analyst*

This has to take some weight off your shoulders obviously with clarity on the path now. Does this change the way investors should think about the succession timeline?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

No, actually there is no real change in my mind. I think our Board has been working on this issue for the last six years. We have built flexibility and depth in our team. It should come as no surprise that John leaves, we have a plan. If Lowell happened to have left, we would have had a plan but we end up with a great lineup, Lowell is a great choice. He is going to be a great CEO.

And so I think what we will do is my plan is to work for the next whatever it takes, probably no later than the end of next year and work with Lowell and the team and make sure that that the transition is complete and there shouldn't be any change. I think the best thing I can say is I don't think you're going to see much shift in approach or strategy from our Company.

Jason Armstrong - *Goldman Sachs - Analyst*

Great, thanks. Maybe we can shift gears and just talk about the macroeconomic climate. I think some of the things that we have heard over the course of this conference is just the polarization of what is going on in the high end of the market whether it is consumer or enterprise versus what is going on maybe in lower tiers of the market both in cable and to some extent telco. But more focused on cable actually in the video market. Can you help us think through segmentation in your base and how you think macro is playing in?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Yes. So what we would say if we look at the industry in general I think the industry is going into a new period of really exciting growth and for the past couple of years, we have been talking about, for example, penetration that will go well beyond 100%. Investors haven't been quite comfortable with that notion but I was pleased to see some of the comments Randall made when he talked about Smartphone penetration going up, we agree with that. He talked about tablets. We think tablets are an extraordinary new opportunity, data growth all going to be fueled by expansion of data.

So when we see the industry, we look at the US 30% data growth or thereabouts, should go to 55% to 60% by 2015, Smartphone penetration should go to 70% of the market by 2015. Tablets not really quite sure what the penetration of tablets would be but certainly tablets plus machine to machine, we are sure everybody in this room certainly will have two devices on them going forward.

So then you get down to the high-end, low-end. I think even at the lower end you are watching the cost of devices will come down sufficiently that the data phenomenon will affect the mid market and the low market. So I think the segmentation you are talking about now is occurring because of the technological state of the industry, this is where we are. But I think as data and video penetrate all segments of the market, I think you will see [strata] that we'll do that.

So we are very excited that in a macro sense, the wireless industry will have enormous opportunity for step function growth here in the next two to three years both in revenues and in volumes. And so I think that is a good spot.

I think, if I look at the macro picture, I said this just a minute ago, I think that again, if you look at what T has done, even a little bit what Sprint has done and to some extent T-Mobile, I think the US carriers have made investments and improvements in the operations so that the industry has the capacity to grow a lot better.

I think if you look at Europe, they are struggling. They will grow but they are strapped into lack of spectrum, a lot of capital in front of them. So we feel that to answer your question, I am very bullish about the state of the industry.

Now if you take it down to our Company, we obviously believe that we are well positioned within that community to do well on that regard and if you like we can chat more about that.

Jason Armstrong - *Goldman Sachs - Analyst*

Lots of topics for discussion. If I could just ask one about higher-level strategy. You have obviously spent several years repositioning the Company more towards growth segments and segments that probably carry a heck of a lot less secular risk than some of the legacy segments you were exposed to. Should investors in a baseball analogy think of this as the ninth inning of the strategy? You acquired Alltel, you spun off of Frontier and FairPoint properties and now the business is kind of set or is there more to go, more repositioning?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Well, I don't think -- the game is ever over so I don't think there is a ninth inning here. My view is what I have been trying to do is to reduce the exposure of the Company to the secular shifts on the landline side. I have been clear about needing to do that, increase the Company's exposure to growth markets which -- and on the landline side by the way would include enterprise and the very high bandwidth services like FiOS.

So I think we are at a point where we have done as much as we can do with the existing asset base we have. We have I think the least exposure of probably any global carrier to its landline business. 60% of our footprint is covered by FiOS. Based on all of the end term franchise agreements we have, we will probably bring that number closer to 70%. So we are feeling good about that so we have less exposure there and so we are feeling good.

I think there is no other asset sale that comes to mind that would be a high priority. But I do think when you look at the broader landscape, you look at telecommunications wireless, you look at content, you look at cable, you look at all of that very interesting phenomenon you are beginning to see multiples starting to settle in the same space which means to me that you have too many companies, you have got multiples starting to contract and to consolidate. And so I think that probably in the next three to five years, Jason, I think you will see another wave of consolidation occur.

Our view is to get our Company in a position where we are dealing from a position of strength so our multiple is better than others and we have -- actually it is kind of an anomaly but we have done fairly well on that basis. So a dollar of our earnings is pretty competitive with what other companies (inaudible).

So we are not in the ninth inning but we are starting a new game and that new game will probably will play out over at the next three or four years.

Jason Armstrong - *Goldman Sachs - Analyst*

To ask the sort of the obligatory Vodafone question. We had Vodafone's CEO here yesterday. He sort of laid out options around the Verizon wireless stake and really pegged it around dividends or a split and really seemed to rule out a merger talking about valuation levels being different. But is your view the same? Is it that the order of priority would probably be dividends or split and mergers a distant likelihood?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

So I want to help Vittorio because I don't know how you go from dividends to split and I don't know what the logic of that is. But I think what our investors should expect is this, that as our business generates cash and as we reach a point where net debt gets closer to that of balance, then a distribution of cash to the owners is probably the right thing to do. I don't have any problem with that.

The issue of the dividend to me in the past has always been asymmetrical benefit that a distribution of the dividend would have created for the two sets of shareholders. In many of the Vodafone properties, debt to EBITDA levels were 2, 2.5 times so we were not going to take a capital structure, stick it inside of Verizon wireless that was different than the capital structure we want to put in our company.

So I think as we get our capital structure where we think it needs to be which is for the total Company should be in the 1.2, 1.3 range that kind of a thing, which is what we've said we think that makes sense for us. And you drive your wireless debt down, a distribution of cash would probably be okay. So that will occur.

We don't think there is a clock to make that occur. Along the way something like Alltel came along and we bought it. But our current course of speed, the business is so successful that it is generating cash faster than you could find uses for that cash so therefore, distribution is probably going to be something that will occur in the future.

Jason Armstrong - *Goldman Sachs - Analyst*

If I could just shift to sort of the policy environment. In the last couple of months, Verizon and Google obviously took a leadership position in proposing a constructive step for the industry as it relates to net neutrality. This audience and the investor base has sort of debated the threat around Title II and what that means around net neutrality but other extensions of that and we have had a lot of back and forth. Is this something investors should still be focused on or is this off the near-term agenda in your mind?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

I don't think it is ever off our agenda. I think we have to view it but I think Randall said this well, if Glenn was here, he would probably say the same thing. We didn't think Title II was a particularly good idea. Obviously everybody understood that because it was an extension of prescriptive powers for the FCC that was unneeded in this environment.

Google who had led the charge on net neutrality I think has done a really wonderful job at studying and learning it and it is pretty simple. I think just to remind everybody, we all agree that an open Internet in which you have access to any content, you can attach any device to it and you can reach any website, any application could run on it is something we all agree with, the public Internet should be that way.

The other two things are, we wanted a separate treatment for wireless and we wanted to make sure there was an innovation opportunity around managed services. So we didn't get stuck with making big investments and then not having the ability to drive specialty or one-off type of events.

Because think about it, even in the Internet space, we are moving into -- you could argue that Facebook is a managed service that people might pay something different for that experience than they would over the public Internet. So I think Google has come to the conclusion that their interest is served by us pumping a lot of iron into the Internet but having separate market based solutions to it.

So now getting back to this, it always takes the government a little extra time to sort of work through their issue on it. To me this business about restoring the FCC's powers to where it was prior to the BitTorrent case is a little bit of a red herring. But I think we will end up with helping the FCC figure out what it should enforce, what it has prescriptive powers over and what should be left to the market. And I think there are a lot of discussions going on.

I don't think Title II sits in the middle of those discussions at this point. I don't think investors should feel that that is a particular threat to any of the industries, cable or us. But I would suggest that the industries do need to come to a conclusion of these discussions so that we can settle the need to have firm rules and a fair set of responsibilities cut across these agencies in a way that doesn't have this bubble up every year.

I think the FCC is in a better place than it was. I think Congress is in a slightly better place than it was and so I think the industry frankly has moved a little bit to make sure that we help out. So I hope this helps. I don't want to think the issue goes away but I don't think Title II the way we have thought about it is a particularly big threat because it is not a well thought out idea anyway.

Jason Armstrong - *Goldman Sachs - Analyst*

Shifting gears to the wireless segment, maybe a couple of questions related to iPhone. Obviously a lot of people like to speculate when Verizon is getting the iPhone for investing and probably for personal reasons. You have had obviously significant improvements in your device line if you think about devices you have launched over the course of the summer. It seems to us that the gap is actually narrowing pretty quickly. SO I am wondering how this plays into the desire/the need to get the iPhone on your network?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

This is like the Knicks getting Carmelo Anthony, you know. Like it would be really good if the Knicks got Carmelo Anthony but they have to play the game to get him. Right? So here is kind of like the way I would look at that. We have worked hard to make our network a destination for all the suppliers to develop devices and equipment. Because we think we have the best value proposition to attract customers. So that is our view.

So the iPhone, AT&T has done a wonderful job with it, Apple obviously has done a wonderful job with it. We were not in the game because it was GSM-based, it was global standard, really not there. We have worked hard at building a franchise out of Droid which proves that if there is an alternative, an acceptable alternative to the iPhone, and a great network and a great distribution channel, the market will go that direction.

So I'm hoping that the success we have had with the Droid will lead us to the place where Apple, Samsung, Motorola, Huawei even, will want to build 4G devices and accelerate the ecosystem development so that the number of devices that will be available on 4G will be extraordinary. And our view would be we would want to carry them all.

I can't speak for Apple but my view is there is a lot of momentum in the industry for people to move on the 4G issues quicker. Now there was a fear that we build a network and that the devices wouldn't be there. Well, we are going to have a tablet. We have already announced one. There will be other tablets coming out. We have got a device lineup that would include some 4G capability starting in the first quarter.

So we are feeling very good that the answer to your question is we don't feel like we have an iPhone deficit. We would love to carry it when we get there but we have to earn it. And I think what we need to do is show all the suppliers that we have equal treatment and we have the best network for people to put all their equipment on.

So the answer to your question is I am believing that 4G will accelerate the process and any other decisions that Apple makes would be fine with us and I think we are running our business to make ourselves the destination for all of the key tablets and devices and hopefully at some point, Apple will get with the program.

Jason Armstrong - *Goldman Sachs - Analyst*

And assuming they do get with the program, I think one of the things that your investors debate as it relates to Verizon's earnings power, there is just so much pent up demand but obviously pent up demand in AT&T's base potentially shift over. As we think about tolerance for dilution around [AT&T] device and how you would think of that, I guess the fear from people is we go into a year where we take $0.30 in dilution and do a device and then a year after that, it looks great but there is a lot of variability in earnings because of this. How would you think about that?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Well, I can't help you with this one because I think that it depends how you think about it. If you want to pay for growth, you pay for growth. If you don't want to pay for growth, you don't pay for growth. I think a lot of times I think we have made it pretty clear over the years that we have run a business in wireless that has got mid 40s margins and with high growth and the growth slows down a little bit, margins pick up a little bit. If it turns out we had to make some investments to accelerate growth, I think you should assume we would do that. Okay? But it doesn't change our view that our base case is to run a business that is producing mid 40 margins and producing a high-growth business.

I think it is really important to make sure we focus on the fact that our scale gives us a different entry point to seek growth than it would have been five or six years ago because we are so much bigger. So I think the recovery period is going to be quicker, whatever it might be. I think the level of dilution might be less if you go through it.

But it all depends on volumes and we don't know what those volumes might look like but certainly in our case, we look at what is in front of us, we have 20% Smartphones. We think the market is going to 70%. When we get to 4G, we are on the same footing as everybody. And that is a huge opportunity for that. I think tablets are going to explode. I mean they are terrific devices.

So we are feeling very bullish about it. So the answer to your question is is Lowell going to be scared of a little dilution for growth? Lowell is going to manage that very well. I don't think he is going to be nervous about suggesting that we go to our Board and we say, if we can double our growth rate, would we accept some minimal dilution to do that? I think the answer is that we would.

Jason Armstrong - *Goldman Sachs - Analyst*

A lot of your comments have tied into the 4G launch which a lot of us are excited about. You talked about it opening up opportunities on the handset side. I am wondering about on the pricing side, Verizon, we have seen AT&T move into tiered data earlier this year. And a lot of people expected others to potentially replicate that fairly quickly. We haven't seen it yet.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Our view on that is we didn't need to be first. I think what we have watched others do is good. We are not sure that we agree yet with how they have valued the data. So the tiers they have created is not necessarily the way we think about it. Our view is that perhaps the data experience we are going to offer is going to be better. That would be our view. So therefore the tiers and the bundles might be different.

Just to give you a sense of that point, T for example, they have done great with the iPhone. They have a lot more Smartphones then we have. Coincidentally the amount of data traffic we both carry is not that far off from each other. You get third party companies that measure the stuff and they find that the volume of data that we both carry is not that different. Remember, we have kind of owned the air card market for a long time so we have the a lot more of that.

And so I think that when we layer in 4G, the experience is going to be very good. When you think about 4G, if I may make a comment about it, we think about it this way. So we have actually had the 4G network in the 30 markets or so we are going to introduce. It has been up and running since around the middle of June. So some of you may be running around on your devices actually have been experiencing the 4G speed and you will probably say what happened? But it is probably not that noticeable because you don't think about it.

So we have been tweaking, tuning, trying to get everything all lined up. So when we cut over on our network, we will have cut over about 30 markets and about 50 airports and all the airports are in the same market as the 30. So the issue is you will get a pretty good nationwide experience right from the get-go.

But think of it this way. So the day we cut this over at sometime between now and the end of the year and we are now up and running at 100 million pops. Starting that day, 24 months later every month there will be markets cut over so we will be at 90% two years later. So the way I think about it is from the day we start, 24 months later we've got an enormously powerful market. Nobody is going to be even close to us in that.

So the issue is I think and coupled with the idea that that 700 MHz spectrum at the low band is great for the suppliers and manufacturers. They all want to design to a nationwide standard. We are kind of thinking like, okay, the other guy had their time with the iPhone. I think our time is now so I think we have all of this in front of us with a chance to really penetrate devices, smart phones, data.

So getting back to your question on pricing, I think we are going to feather in our pricing over the next couple of months as we introduce new services, new products and try to get that value proposition to fit the added excitement and added capabilities that we are going to offer into the marketplace. So I don't think we want to match our pricing to what we think is going on today because we think we are not sure it is the right experience nor is it the right value proposition.

But we do agree with tier pricing and we do think we have to monetize the investments we make and you will see us do that over the next four to six months.

Jason Armstrong - *Goldman Sachs - Analyst*

Maybe switching gears to wireline, one of the things that has been topical at the conference is just what the video opportunity, actually what the sort of end market looks like. You guys are sort of more of in a nascent stage, you have plenty of room to grow into with obviously, it is a top-tier product out there.

But as you sort of step back and see what the cable companies are talking about now, the low end of the video base sort of disappearing a little bit more quickly than they expected and tagging it to cyclical, what is your perspective on this and what is going on?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Let's start from our side. When we look at video and broadband, we think we are doing okay with video. I think John Killian last week was at another conference in which he indicated that we should be exceeding or tracking anyway, we are exceeding our second-quarter number on video net adds on FiOS. I think we had 170,000 something last quarter. Is that right? So we are tracking better than that and we are feeling good so we think we are fine.

Then when you look at the net broadband which is what we are getting out of FiOS plus what we are losing out of DSL, we are kind of net, we are kind of positive. So to us, that is a good accomplishment. So in this market for our net broadband to be positive, we are feeling good.

So what is going on I think is pretty clear. DSL is losing out to cable. I think cable is probably starting to experience what we experienced five, six years ago which is the low end is disappearing into other alternatives. And the first thing when that happens is you deny it. I know the drill, I have been there.

So the issue is you just need to increase the value proposition which they are doing a good job at. So when you think about cable they increase the speeds, they increase HD, they are doing a lot of things with their product but they are still finding that there is lots of substitution going on. Some of it is the economy I wouldn't doubt that, housing starts and I get that. But housing starts will never pick up enough to offset what we now see is the drop.

So I think that the video experience is going to go the direction that we kind of feel we are prepared for which is very high speed, high experience of video. So 3-D is an example. Very high data loads that we use on FiOS. So I think in our case, we think we have a smaller footprint, we are positioned to do that. I do think the economy has a big impact but I also think the secular shifts that are occurring will hit video the way they hit Voice. And I think we are feeling like with our 4G rollout, with our FiOS position, we will be okay.

Jason Armstrong - *Goldman Sachs - Analyst*

Different statement then what we heard yesterday, an interesting perspective. Thank you.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Yesterday from?

Jason Armstrong - *Goldman Sachs - Analyst*

Cable companies and a lot of media companies.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Listen, they are doing well and they have to play their cards. My only comment is they don't know what they don't know. And I think I have all of the bruises of sort of playing out the same strength. We never thought anybody would cut the cord on telco, right? We have got 30% of our customers cutting the cord.

Young people are pretty smart. They're not going to pay for something they don't have to pay for. So you've got to watch the market, over the top there is going to be a pretty big issue for cable.

Jason Armstrong - *Goldman Sachs - Analyst*

On the wireline business, the margins seem to have leveled out. Cost-cutting is just sort of ahead of the revenue trajectory at this point. Positive surprise in 2Q. Should we extrapolate that for a point where margins are stable and potentially move up?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

It is another one. I think we caught the bus last quarter and so I think we have created a little operating momentum and so the goal is to do better this quarter than we did last quarter and continue to do that going forward. So I am pretty confident that we are at a good spot with that and I'm not going to predict where margins need to go to but I think the issue for us is to make sure incrementally we beat or at least match what we have done in the last quarter.

So I think certainly we have the benefit this quarter of significant force reductions because most of the force reductions that were announced hadn't hit the Company's books and they won't even be 100% hit in the third quarter. There will be a good piece of it. By the time we get to the fourth quarter though, we will see a lot of [that].

We have done a lot of other things on cost so the issue on margins is we have every incentive to reduce the leakage out of the wireline business to make sure we book as much of the game we think we are going to get out of wireless. So to me, we are on fire to make sure that we minimize the leakage on wireline and actually for the first time in several years, I feel better about long-term with the wireline. Because we have been in sort of a fire drill mode here. But when I think of FiOS keeps getting stronger and it is doing well and the contribution from that is extremely good.

We are seeing some stabilization in enterprise. It is not great yet but that is clearly cyclically impacted and I think we are doing that. And while this is going on, is we have built a really good portfolio of better products, better services and we have a lot of interest in customers and I think we have accelerated our cost activity.

So I think when I look at the wireline business we still have this problem of access line loss and all of that kind of thing and voice revenues shrinking. But we have more tools in motion today than we did a year or so ago to offset that. Remember, the most important thing to me at this point is to minimize that leakage so that as the wireless team sort of ramps it up here with 4G and tablets and all of these other things we are going to do, we see more of that earnings power go to the bottom line.

Jason Armstrong - *Goldman Sachs - Analyst*

Maybe one more from me and then we will open it up for some audience questions. Think about returning cash to shareholders, your Board made the decision to hike the dividend again recently. And it is interesting because the Frontier spin sort of was a dividend hike in itself and that gave you plenty of room to not have to hike the dividend this year but still have plenty to say listen, to shareholders, we gave you an effective hike over the summer. So in my mind, sort of a positive signal it showed strength. Can you step us through the decision?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

It is a math issue. It is both a policy and a math issue. I think if you took the financial people in the Company, they were banging on me to say that we are paying out too much in earnings and if you are only earning $2 and change, and we are paying out $1.95 in dividend that is kind of not the thing to do. So I got that speech from Killian and everybody else in the Company.

On the other side of the coin is Doherty is laughing because he was part of that speech. The other side of the coin is we feel we have really positive prospects of growth going forward. And if I look back on my own career, one of the mistakes I made on this point was every time I thought I was chasing growth and we ratcheted down the growth and the dividend, some investors thought that was okay, most of them didn't.

So we thought through this cycle, we would make sure we kept that dividend at a place where people would be attracted to the stock. And so by the way for those of you who feel like you want to buy, I think if you drove our yield down to 5%, our stock would be 38 or 39. I am not sure what the number -- it would be 38 or 39. So I think the issue for us is to be really clear about maintaining the sanctity of some level of dividend growth. We are too big for people to look at us as a pure growth play.

So if we could get growth back into that 5% to 7% range and have a yield that is the above 5%, I think we become a pretty attractive looking stock. So that is kind of where I am at.

My own view is you pay down debt and you fix your balance sheet. We have great cash so no one should worry that our dividend is at risk. Then you get to the issue of share buybacks which I will deal with since you raised it. It is kind of like the last issue on my hierarchy of good things to do. Reinvest in the business, you want to do that. You want to make sure you protect your dividend, you want to pay down your debt and if you have anything left over, you buy your shares back.

We have so many shares outstanding. We have to have some massive share buyback program which is probably not in the cards right now so therefore, the dividend becomes more of a central factor how we think about it. I hope that helps a little bit.

Our Board is very good. Our Board worries about the math to make sure that we can afford to pay it and that we are not taking our metrics totally out of line. On the other side of the coin is our Board is excited about the fact we have restructured our Company, we have reduced exposure to the slower growth assets and we are thinking that over the next several years organically, we have a lot in front of us to achieve.

Jason Armstrong - *Goldman Sachs - Analyst*

We will take some audience questions. (Operator Instructions).

Unidentified Audience Member

Ivan, your opinion on what is going on in cable is the same that I have reached, but I get it just looking at the attrition of subscribers. It reminds me so much of the consumer wireline. You have a better look at it than I do. Do you have information sort of in a more granular form that leads you to your opinion?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Well, I think what I am going to say is obvious, yes, we have a lot of information, a lot more granular information of which I'm not going to share. But I think the point I am making is we take the over the top issue with video very seriously so we have a smaller 3.5 million video customers. We think we are not going to keep them all the way the old cable model kept them.

So we are looking at lots of different alternatives to make sure that we have, we participate in how to monetize over the top across a range of platforms which we are working on. But we also feel that -- we do agree with cable that you can keep a lot of the customers if you have very high rich experiences like either HD or maybe 3-D.

The only point I would make about this part of the business which is a little bit of a rub to us but we agree with cable, we all face the problem of 5% and 10% content cost increases every year. And so as long as you are stuck in the sort of bundled model and you have bandwidth limited capabilities, you are going to find that the over the top is going to be somewhat limited. So just put that in perspective, you want to drive all of your channels onto the Internet, the Internet doesn't have the capability to offer 20, 30 channels. Maybe it could offer six or seven.

And so I think the issue is I think cable has some life left in its model because with all of the proliferation of content, the bundle offering of 300, 400 channels is probably the only way a customer is going to be able to buy that. But that is going to get this intermediated over the next several years. It is going to take a little time.

We don't know what tablets will do. We don't know what 4G will do. We don't know what 3-D will do. So in my judgment, all of that is going to further segment the business and what we are going to do in our Company is make sure that our FiOS-based offering and our 4G-based offering trying to take advantage of multiple platforms. And we don't find ourselves building ourselves a cable-like business and getting ourselves stuck in the problem of cord cutters and all these kinds of things (inaudible).

So I hope that helps a little bit. They have a good business. I don't want anybody to think I am suggesting they don't. But I have seen the movie. And I think if you remain static too long, the technology is going to nibble at you at the edges and you need to be prepared for it.

Unidentified Audience Member

(inaudible question - microphone inaccessible)

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

I think I understand the question. My preliminary thoughts, it is a GSM-based -- I mean they wanted to put that on GSM-based devices. So I think we were taken out of the game. I don't think there is any strategic issue there other than they went to Blackberry. They put it on a GSM device and that is what they did. So the answer is it will come around the corner soon I think.

Jason Armstrong - *Goldman Sachs - Analyst*

We have time for one more. Back here.

Unidentified Audience Member

Regarding the introduction of the 4G product looking out both a year from now and two years from now, what percent of your phone sales and what percent of your network usage do you expect to be 4G?

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Well, as I said by the end of 2012 and as we head into 2013, we will have 90% -- 93%, 94% coverage of 4G across the country so that is what we view with that. You want specific model questions so I won't do that but I will give you some broader issues.

I think we will get our 20% of Smartphones way up in the next three years to 40%, 50% heading to 70%. I think everybody will have -- I think penetration of products and services on 4G will be 200% to 300% as opposed to right now at 95%. So I think that there is lots of opportunity.

Revenues, ARPUs should increase because I think we started to see that already as we go up the smartphone ramp, we are watching our postpaid ARPUs be accretive. So I think all of this is heading in the right direction. I think we need a few quarters under our belt so you can model some trends but we are feeling that the volumes are being very strong.

Jason Armstrong - *Goldman Sachs - Analyst*

Great, we will have to leave it there. Ivan, thanks so much for joining us.

Ivan Seidenberg - *Verizon Communications, Inc. - Chairman and CEO*

Thank you.

TOWARDS AN ECONOMIC FRAMEWORK FOR NETWORK NEUTRALITY REGULATION

DR.-ING. BARBARA VAN SCHEWICK, ASS. IUR.*

Network neutrality rules forbid network operators from excluding or discriminating against third-party applications. This analysis shows that calls for network neutrality regulation are justified: absent network neutrality regulation, network providers will likely discriminate against or exclude independent producers of applications, content, or portals from their networks. This threat reduces the amount of innovation in applications, content and portals at significant costs to society. While network neutrality rules remove this threat, they are not without costs. Due to the potentially enormous benefits of application-level innovation for economic growth, however, increasing the amount of application-level innovation through network neutrality regulation is more important than the costs associated with it. This paper also highlights important limitations of the "one monopoly rent" argument, demonstrating previously unidentified exceptions that may be quite common in the Internet context, showing how exclusion may be a profitable strategy even if the excluding actor does not manage to drive its competitors from the complementary market, and proving that competition in the primary market may be insufficient to remove the ability and incentive to engage in exclusionary conduct.

---

* Non-Residential Fellow, Center for Internet and Society, Stanford Law School, CA, USA and Senior Researcher, Telecommunication Networks Group, Department of Electrical Engineering and Computer Science, Technical University Berlin, Germany. E-Mail: schewick@stanford.edu / schewick@tkn.tu-berlin.de.

Thanks to Pio Baake, Yochai Benkler, Marjory Blumenthal, David Clark, Joseph Farrell, Gerald Faulhaber, David Isenberg, Bill Lehr, Lawrence Lessig, Doug Lichtman, Robert Pepper, Arnold Picot, David Reed, Tim Wu, Christopher Yoo and participants of the 33rd Research Conference on Communication, Information and Internet Policy (TPRC 2005), the Berkman Center for Internet and Society Luncheon Series, Harvard Law School, and the MIT Computer Science and Artificial Intelligence Laboratory Speaker Series, Massachusetts Institute of Technology, for comments on an earlier version of this paper and for discussions.

Parts of this paper are based on Barbara van Schewick, Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet, Chapter 9 and Chapter 11 (Ph.D. dissertation, Technical University Berlin 2005, MIT Press forthcoming 2008), for which financial support of the German National Academic Foundation ("Studienstiftung des Deutschen Volkes") and the Gottlieb Daimler- and Karl Benz-Foundation is gratefully acknowledged.

INTRODUCTION

Over the past years, the merits of network neutrality regulation have become a hot topic in telecommunications policy debates. Repeatedly, proponents of network neutrality regulation have asked the Federal Communications Commission to impose rules on the operators of broadband access networks that forbid network operators to discriminate against third-party applications, content or portals ("independent applications") and to exclude them from their network.[1] Congress is currently considering proposals to introduce network neutrality legislation;[2] the House of Representatives and the Senate held hearings on the subject.[3]

---

1. *See, e.g.*, *Ex parte* Submission of Tim Wu and Lawrence Lessig to the *Declaratory Ruling & Notice of Proposed Rulemaking* in Inquiry Concerning High-Speed Access to the Internet, CS Dkt. No. 02-52 (Aug. 22, 2003), *available at* http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6514683885 [hereinafter Wu & Lessig, *Ex parte*]; Comments of the High Tech Broadband Coalition, to the *Declaratory Ruling & Notice of Proposed Rulemaking* in Inquiry Concerning High-Speed Access to the Internet, CS Dkt. No. 02-52 (June 17, 2002), *available at* http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6513198353; *Ex parte* Submission of the Coalition of Broadband Users and Innovators to the *Declaratory Ruling & Notice of Proposed Rulemaking* in Inquiry Concerning High-Speed Access to the Internet, CS Dkt. No. 02-52 (July 17, 2003), *available at* http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6514286197. For proponents of nondiscrimination rules in the scientific arena, see, for example, LAWRENCE LESSIG, THE FUTURE OF IDEAS 248-49 (2001) [hereinafter LESSIG, FUTURE OF IDEAS]; Philip J. Weiser, *Toward a Next Generation Regulatory Strategy*, 35 LOY. U. CHI. L.J. 41 (2003); Wu & Lessig, *Ex parte, supra*; Tim Wu, *Network Neutrality and Broadband Discrimination*, 2 J. ON TELECOMM. & HIGH TECH. L. 141 (2003) [hereinafter Wu, *Network Neutrality*]; Tim Wu, *The Broadband Debate: A User's Guide*, 3 J. ON TELECOMM. & HIGH TECH. L. 69 (2004) [hereinafter Wu, *Broadband Debate*]; Brett M. Frischmann & Barbara van Schewick, *Yoo's Frame and What it Ignores: Network Neutrality and the Economics of an Information Superhighway*, 47 JURIMETRICS (forthcoming 2007); Bill D. Herman, *Opening Bottlenecks: On Behalf of Mandated Network Neutrality*, 59 FED. COMM. L. J. (forthcoming 2007), *available at* http://ssrn.com/abstract=902071; Robert D. Atkinson & Philip J. Weiser, *A Third Way on Network Neutrality*, THE NEW ATLANTIS, Summer 2006, at 47, *available at* http://www.thenewatlantis.com/archive/13/atkinsonweiser.htm (last visited November 23, 2006); Susan P. Crawford, *Network Rules*, LAW & CONTEMP. PROBS. (forthcoming 2007), *available at* http://ssrn.com/abstract=885583.

2. For an overview of the different proposals and their history, see Declan McCullagh, *Republicans Defeat Net Neutrality Proposal*, CNET NEWS.COM, Apr. 5, 2006, http://news.com.com/2100-1028_3-6058223.html. For an overview of government actions and statements of officials concerning network neutrality, see John Windhausen, Public Knowledge, *Good Fences Make Bad Broadband. A Public Knowledge White Paper*, http://static.publicknowledge.org/pdf/pk-net-neutrality-whitep-20060206.pdf (Feb. 6, 2006) [hereinafter *Public Knowledge White Paper*], at 13-16.

3. *Net Neutrality: Hearing Before the S. Comm. on Commerce, Science, and Transportation*, 109th Cong. (2006), *available at* http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=109_senate_hearings&docid=f:30115.pdf [hereinafter *Senate Hearing*]; *Internet Protocol and Broadband Services Legislation: Hearing Before the Subcomm. on Telecommunications and the Internet of the H. Comm. on Energy and Commerce*, 109th Cong. (2005), http://frwebgate.access.gpo.gov/cgi-bin/getdoc.cgi?dbname=109_house_hearings&docid=f:26998.pdf.

Network neutrality proposals are based on the concern that in the absence of such regulation, network operators may discriminate against independent applications and that this behavior may reduce innovation by providers of these products to the detriment of society.

Opponents of regulation deny the need for network neutrality regulation.[4] They argue that regulation is not necessary because network operators do not have an incentive to discriminate against independent applications anyway,[5] or, alternatively,[6] that regulation is harmful because it would reduce network operators' incentive to upgrade their networks in the future.[7]

This paper aims at assessing the economic merits of network neutrality regulation. To this aim, the paper applies insights from game theory, industrial organization, antitrust, evolutionary economics and management strategy to analyze network operators' incentives to discriminate, the impact of potential discriminatory behavior on innovation and social welfare, and the costs of regulation. By focusing on the economic merits of network neutrality, the paper complements theoretical approaches that base calls for network neutrality regulation on non-

4. *See, e.g.*, Ex Parte Submission of the National Cable & Telecommunications Association to the *Declaratory Ruling & Notice of Proposed Rulemaking* in Inquiry Concerning High-Speed Access to the Internet, CS Dkt. No. 02-52 (Sept. 8, 2003), http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6514882243. For opponents of nondiscrimination rules in the scientific arena, see, for example, Bruce M. Owen & Gregory L. Rosston, *Local Broadband Access: Non Nocere or Primum Processi? A Property Rights Approach* (Stanford Law Sch., John M. Olin Program in Law and Econ., Working Paper No. 263, 2003); Christopher S. Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1 (2005) [hereinafter Yoo, *Beyond Network Neutrality*]; Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847 (2006) [hereinafter Yoo, *Economics of Congestion*]; Christopher S. Yoo, *Would Mandating Broadband Network Neutrality Help or Hurt Competition? A Comment on the End-to-End Debate*, 3 J. ON TELECOMM. & HIGH TECH. L. 23 (2004) [hereinafter Yoo, *Mandating Network Neutrality*]; J. Gregory Sidak, *A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet*, 2 J. OF COMPETITION L. & ECON. 349 (2006), *available at* http://jcle.oxfordjournals.org/cgi/reprint/2/3/349.pdf.

5. There are two representative examples of this view in the context of the debate over open access to broadband networks. *See* James B. Speta, *Handicapping the Race for the Last Mile?: A Critique of Open Access Rules for Broadband Platforms*, 17 YALE J. ON REG. 39 (2000) [hereinafter Speta, *Handicapping*]; James B. Speta, *The Vertical Dimension of Cable Open Access*, 71 U. COLO. L. REV. 975 (2000) [hereinafter Speta, *Vertical Dimension*]. On the open access debate, see *infra* note 31.

6. Both arguments are mutually exclusive. If network owners do not have an incentive to discriminate against independent applications anyway, the imposition of a network neutrality regime that prevents such discrimination will not reduce their profits. If it does not reduce their profits, however, it cannot reduce their incentives to invest in upgrades of their network infrastructure in the future.

7. For a representative example of this view, see Adam D. Thierer, *"Net Neutrality" Digital Discrimination or Regulatory Gamesmanship in Cyberspace?* (Cato Inst., Policy Analysis No. 507, 2004).

economic rationales.[8]

Throughout this paper, the term "network neutrality rules" refers to non-discrimination rules that forbid operators of broadband networks to discriminate against third-party applications, content or portals ("independent applications") and to exclude them from their network. This terminology captures the common rationale behind the various network neutrality proposals before Congress and the FCC – to design rules that prevent network operators and ISPs from using their power over the transmission technology to negatively affect competition in complementary markets for applications, content and portals.[9] By contrast, network neutrality opponents sometimes use a much broader definition of network neutrality that includes mandating interconnection, non-discrimination, rate regulation and the adoption of standardized protocol interfaces such as TCP/IP.[10] While providing a convenient straw man for

---

8. *E.g.*, Brett M. Frischmann, *An Economic Theory of Infrastructure and Commons Management*, 89 MINN. L. REV. 917 (2005); Crawford, *supra* note 1. For a critical evaluation of such approaches, see Yoo, *Beyond Network Neutrality*, *supra* note 4, at 53-57.

9. For an overview of the various proposals, see McCullagh, *supra* note 2, and *Public Knowledge White Paper, supra* note 2, at 3-7, 26-27. In addition to the non-discrimination rules discussed in the text, network neutrality proposals often include the right of consumers to attach communication equipment of their choice to the network. *See Public Knowledge White Paper, supra* note 2, at 3-7, 26-27. Network neutrality regulation is not intended to prevent vertical integration between network providers and application providers (i.e., network providers are allowed to offer applications as well). *See* Wu, *Broadband Debate*, *supra* note 1, at 89.

While calls for network neutrality rules share a common rationale, they differ with respect to how these rules should be implemented. For example, in some proposals, the non-discrimination rules take the form of user rights (to access and use the content and applications of their choice), in others the respective rights are vested in the providers of complimentary products (to offer the application and content of their choice). *See, e.g.*, *Public Knowledge White Paper, supra* note 2, at 3-7, 26-27. Proposals differ with respect to the exceptions to the non-discrimination rule they include, (i.e., with respect to the cases in which a deviation from the principle of network neutrality is justified). *Id.* at 27. For example, whether and, if yes, what form of price discrimination should be forbidden under network neutrality regulation, is still an open question. *Compare* Wu, *Network Neutrality*, *supra* note 1, at 151-54 (arguing against price discrimination, if it is based on discrimination between applications), *and* JONATHAN E. NUECHTERLEIN & PHILIP J. WEISER, DIGITAL CROSSROADS 177 (2005); *see also Senate Hearing*, *supra* note 3 (testimony of Prof. Lawrence Lessig), *available at* http://commerce.senate.gov/pdf/lessig-020706.pdf (arguing against "access tiering," i.e. "any policy by network owners to condition content or service providers' right to provide content or service to the network upon the payment of some fee [. . . which is] independent of basic Internet access fee," *id.* at 2 note 2, but supporting "customer tiering," i.e. price discrimination, as long as it is not based on discrimination among content or application providers); *but see* Sidak, *supra* note 4, at 83-99 (arguing in favor of allowing access tiering). For specific implementation proposals from the scientific literature, refer to Wu, *Broadband Debate*, *supra* note 1, appendix A; Weiser, *supra* note 1, at 74-84; Atkinson & Weiser, *supra* note 1, at 55.

10. *E.g.* Yoo, *Beyond Network Neutrality*, *supra* note 4, at 3, 8, 27, 32. As most of Yoo's arguments about the negative impact of network neutrality are based on the negative impact of measures such as the adoption of standardized interfaces that are not part of the network neutrality regime discussed in this paper, his analysis does not carry over to the case of "pure" non-discrimination rules discussed here. *See also* discussion *infra* notes 192, 198. For a

attack, this definition goes far beyond what network neutrality proponents want to achieve: the measures included in the broad definition constitute heavy forms of regulation; by contrast, the non-discrimination rules in network neutrality proposals have been explicitly designed to provide a light form of behavioral regulation that narrowly targets the behavior identified as problematic and is far less intrusive than other forms of regulation such as structural separation or open access regulation.[11]

The analysis proceeds in three steps. Part II explores whether network providers have an incentive to discriminate against applications.[12] This question has not been examined in detail in the existing literature.[13] If, however, network providers do not have such an incentive, there is no need for regulation.[14]

Whether exclusionary conduct in complementary markets is a profitable strategy has been hotly debated over the years. Today, most scholars agree that a monopolist in a primary market does not generally have an incentive to exclude its competitors from a secondary, complementary

---

critical appraisal of Yoo's work on network neutrality, see Frischmann & van Schewick, *supra* note 1; Herman, *supra* note 1.

11. *See, e.g.*, Weiser, *supra* note 1, at 48, 74, 78-80; Wu, *Network Neutrality*, *supra* note 1, at 145-49. By contrast, Yoo derives its definition from statements of network neutrality proponents with respect to Internet policy in general, not to network neutrality in particular, Yoo, *Beyond Network Neutrality*, *supra* note 4, at 3.

12. See *infra* Part II.

13. For a similar assessment with respect to network neutrality proponents, see Weiser, *supra* note 1, at 74-75. For an example of the treatment of the question by a network neutrality proponent, see Wu, *Broadband Debate, supra* note 1, Part II.B. For an example of the treatment of the question by a network neutrality opponent, see Yoo, *Beyond Network Neutrality*, *supra* note 4, Part II. (arguing that network neutrality proponents' focus on safeguarding competition in the markets for application and content is misplaced without examining whether there is indeed a threat of discrimination) and at 60-61 (arguing that competition in the broadband market "should remain sufficiently robust to ameliorate concerns of anticompetitive effects" without covering specific motivations for discrimination); Yoo, *Economics of Congestion, supra* note 4 (manuscript at 49-50). While Farrell and discuss exceptions to the "one monopoly rent" argument in detail, their analysis is not specifically targeted at the economic relationships relevant in the network neutrality context. Joseph Farrell & Philip J. Weiser, *Modularity, Vertical Integration, and Open Access Policies: Towards a Convergence of Antitrust and Regulation in the Internet Age*, 17 HARV. J.L. & TECH. 85, 114 (2003). Similarly, while participants in the open access debate have explored the incentives of network providers to exclude independent ISPs from their network, their analysis focuses on the competitive relationships between the operators of physical networks on the one hand and unaffiliated ISPs that may also offer applications, content or portals on the other hand. *E.g.*, Speta, *Handicapping, supra* note 5; Speta, *Vertical Dimension, supra* note 5 (both denying an incentive to exclude); Daniel L. Rubinfeld & Hal J. Singer, *Vertical Foreclosure in Broadband Access?*, 49 J. INDUS. ECON. 299 (identifying an incentive to exclude). By contrast, network neutrality rules focus on the competitive relationships between the operators of physical networks and/or ISPs on the one hand and unaffiliated providers of complementary applications, content and portals on the other hand. Thus, the open access literature is not directly applicable to the network neutrality debate.

14. For a qualification of this assessment, see *infra* notes 27-28 and accompanying text.

market – the well known "one monopoly rent" argument.[15] There are known exceptions to this rule, but these rarely apply. As a result, when analyzing allegations of exclusionary conduct in a complementary market, most scholars intuitively assume there will not be a problem, in particular if the excluding actor faces competition in the primary market.

The results of the analysis challenge this intuition in several ways:

First, Part II identifies exceptions to the "one monopoly rent" argument that have not been previously thought of, but are quite common in the Internet context.[16]

Second, the paper shows that some of the known exceptions do indeed apply in the Internet context.[17]

Third, researchers commonly assume that discrimination against a complementary product will only be profitable, if the primary good monopolist manages to monopolize the market for the application in question. The paper shows that this assumption is not necessarily correct. A network operator may have an incentive to discriminate against an application even if the operator does not manage to drive all independent applications from the corresponding market.[18] As a result, researchers commonly underestimate the potential for discriminatory behavior by network providers.

Finally, in line with conventional thinking on the profitability of exclusionary conduct, participants in the debate usually share the view that competition in the market for Internet services may be able to mitigate the problem. Two policy proposals, the proposals for facilities-based competition and for open access, are based on this view. The results of Part II contradict this view. The analysis highlights a variety of circumstances under which a network operator may have the ability and incentive to discriminate against independent applications in spite of competition in the market for Internet services.[19]

Thus, Part II highlights important limitations of the "one monopoly rent" argument in the Internet context that may be relevant beyond the network neutrality debate. In the network neutrality context, it shows that in the absence of network neutrality regulation, there is a real threat of discriminatory behavior that is more severe than is commonly assumed.

Part III analyzes the impact of this threat on innovation in the markets for applications, content and portals ("application-level innovation").[20] It shows that the threat of discrimination reduces the amount of

---

15. *See infra* Part II.B.1.
16. *See infra* Part II.B.2.
17. *See infra* Part II.B.3.
18. *See infra* Part II.B.4.
19. *See infra* Part II.C.
20. *See infra* Part III.

application-level innovation by independent producers of complementary products.[21] While discrimination increases network providers' incentives to engage in application-level innovation, this increase cannot offset the reduction in innovation by independent producers.[22] Thus, the threat of discrimination reduces the amount of application-level innovation.

Part IV explores the social benefits and costs of network neutrality regulation.[23] It shows that the increase in application-level innovation resulting from network neutrality rules is socially beneficial.[24] On the cost side, network neutrality rules reduce network providers' incentives to innovate at the network level and to deploy network infrastructure.[25] While regulatory intervention has its own costs, these are not covered in detail. When deciding whether to introduce network neutrality regulation, regulators must trade-off the benefits against the costs. The analysis shows that in the context of the Internet, the benefits of network neutrality regulation are more important than the costs.[26]

I. THREAT OF DISCRIMINATION

Calls for network neutrality regulation are based on the assumption that network providers have an incentive to discriminate against unaffiliated providers of complementary products. If network providers do not have such an incentive, there is no need for regulation. In this case, regulation may still serve an educational function and protect customers and providers of independent content, portals and applications from discriminatory or exclusionary conduct by "incompetent incumbents"[27] that fail to recognize that discrimination is not in their best economic interest.[28] Compared to a threat of discrimination due to a real incentive to discriminate, this constitutes a considerably weaker basis for regulatory intervention.

Network technology gives network providers the ability to discriminate against applications running over their networks or to exclude them from the network. The following part explores, whether network providers have an incentive to actually use this discriminatory power.[29] The analysis is based on a stylized model (Section A). As the answer may dif-

21. *See infra* Part III.A.
22. *See infra* Part III.B.
23. *See infra* Part IV.
24. *See infra* Part IV.A.
25. *See infra* Part IV.B.
26. *See infra* Part IV.C.
27. Farrell & Weiser, *supra* note 13, at 114.
28. *Id.* at 114-17; Wu, *Network Neutrality*, *supra* note 1, at 154-56.
29. There have been various instances of discrimination by network providers in practice, both in the United States and internationally. *See generally Public Knowledge White Paper*, *supra* note 2, at 16-23.

fer depending on the market structure in the market for Internet services, the analysis proceeds in two steps: In the first step, the network provider is a local monopolist (Section B). In the second step, the network provider competes with at least one other network provider (Section C).

The analysis shows that discrimination is much more likely than is commonly assumed.

A. Stylized Model

Network neutrality rules seek to protect competition in complementary products such as Internet applications, content and portals from anti-competitive behavior by network operators or ISPs. To reflect this goal, the analysis focuses on the competitive interactions between "the network" and "applications."[30] Economically, "the network" comprises two distinct layers of economic activity: the operation of physical networks and the provision of Internet access and transport services over these networks. In real life, these activities may or may not be provided by different economic actors with differing economic interests. The resulting competitive interactions between network operators and Internet service providers have featured prominently in the debate over "open access" for independent Internet service providers to broadband cable networks in the United States.[31] To focus on the specific impact of network neutrality

---

30. In the context of the four layer model of the Internet Architecture used by the Internet Engineering Task Force, "the network" consists of the network layer and the Internet layer, while the application domain consists of the transport layer and the application layer. *See, e.g.*, LARRY L. PETERSON & BRUCE S. DAVIE, COMPUTER NETWORKS: A SYSTEMS APPROACH 27-30 (3d ed. 2003).

31. The open access debate focuses on the question whether the owners of cable networks should be required to allow independent Internet service providers to provide Internet access services over their cable networks. Several scholars advocate open access regulation. *See Ex parte* Submission of Mark A. Lemley & Lawrence Lessig, to the *Public Notice*, in Application for Consent to the Transfer of Control of License Licenses from MediaOne Group, Inc. to AT&T Corp., at 1, CS Dkt. No. 99-251 (November 10, 1999), http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6010050443 [hereinafter Lemley & Lessig, *Ex parte*]; Mark A. Lemley & Lawrence Lessig, *The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era*, 48 UCLA L. REV. 925 (2001); LESSIG, FUTURE OF IDEAS, *supra* note 1, at 147-67, 246-49; Francois Bar et al., *Access and Innovation Policy for the Third-Generation Internet*, 24 TELECOMM. POL'Y 489 (2000); Jim Chen, *The Authority to Regulate Broadband Internet Access over Cable*, 16 BERKELEY TECH. L.J. 677 (2001); Mark N. Cooper, *Open Access to the Broadband Internet: Technical and Economic Discrimination in Closed, Proprietary Networks*, 71 U. COLO. L. REV. 1011 (2000); Jerry A. Hausman et al., *Cable Modems and DSL: Broadband Internet Access for Residential Customers*, 91 AM. ECON. REV. 302 (2001) [hereinafter Hausman et al., *Cable Modems*]; Jerry A. Hausman et al., *Residential Demand for Broadband Telecommunications and Consumer Access to Unaffiliated Internet Content Providers*, 18 YALE J. ON REG. 129 (2001) [hereinafter Hausman et al., *Residential Demand*]; William P. Rogerson, *The Regulation of Broadband Telecommunications, the Principle of Regulating Narrowly Defined Input Bottlenecks, and Incentives for Investment and Innovation*, 2000 U. CHI. LEGAL F. 119; Daniel L. Rubinfeld & Hal J. Singer, *Open Access to Broadband Networks: A Case Study of the*

rules, the following analysis abstracts from these issues and treats these players as a single economic entity called the "network provider."

The analysis will be based on the following stylized model: for a given physical network, Internet access and transport services and the operation of the network infrastructure are provided by the same economic entity, the "network provider." The corresponding service will be called "Internet service." The network is assumed to provide the same general functionality as the Internet in that it enables computers attached to distinct physical, but interconnected networks to communicate. Contrary to the original Internet,[32] but similar to networks today, the network is application-aware and can control the execution of applications running over its network. Today, technology is available that enables network operators and ISPs to distinguish between the different applications using the network and to control their execution.[33] For example, network

---

AOL/Time Warner Merger, 16 BERKELEY TECH. L.J. 631 (2001); Rubinfeld & Singer, *supra* note 13. Several experts also opose open access regulation. *See* John E. Lopatka & William H. Page, *Internet Regulation and Consumer Welfare: Innovation, Speculation, and Cable Bundling*, 52 HASTINGS L.J. 891 (2001); Glen O. Robinson, *On Refusing to Deal With Rivals*, 87 CORNELL L. REV. 1177 (2002); Speta, *Handicapping*, *supra* note 5; Speta, *Vertical Dimension*, *supra* note 5; Glenn A. Woroch, *Open Access Rules and the Broadband Race*, 2002 L. REV. M.S.U.-D.C.L. 719 (2002); Christopher S. Yoo, *Vertical Integration and Media Regulation in the New Economy*, 19 YALE J. ON REG. 171 (2002).

32. In the original Internet, the network was application-blind, (i.e., it was unable to distinguish between the applications running over the network). Consequently, network operators were unable to affect the execution of specific applications, shielding independent application developers from strategic behavior by network operators.

The application-blindness was the result of following the broad version of the end-to-end arguments during the design of the Internet, Barbara van Schewick, Architecture and Innovation: The Role of the End-to-End Arguments in the Original Internet 101-03 (Ph.D. dissertation, Technical University Berlin 2005, MIT Press forthcoming 2008). This design principle requires that the lower layers of the network be as general as possible, while all application-specific functionality is concentrated at higher layers at end hosts. (There are two versions of the end-to-end arguments: a narrow version, which was first identified, named and described in a seminal paper by Saltzer, Clark and Reed in 1981. Jerome H. Saltzer et al., *End-to-End Arguments in System Design*, 1981 2ND INT'L CONF. ON DISTRIBUTED COMPUTING SYS. 509 (a revised version of paper was later published as Jerome H. Saltzer et al., *End-to-End Arguments in System Design*, 2 ACM TRANSACTIONS ON COMPUTER SYS.S 277 (1984)). A broad version was the focus of later papers by other authors. *See, e.g.*, David P. Reed et al., *Commentaries on "Active Networking and End-to-End Arguments"*, 12 IEEE NETWORK 69, 69 (1998); Marjory S. Blumenthal & David D. Clark, *Rethinking the Design of the Internet: The End-to-End Arguments vs. the Brave New World*, 1 ACM TRANSACTIONS ON INTERNET TECH. 70, 71 (2001). While both versions have shaped the original architecture of the Internet, only the broad version is responsible for the application-blindness of the network.) For a detailed analysis of the two versions and their relationship to the architecture of the Internet, see van Schewick, *supra*, at 87-129.

33. *See, e.g.*, Cisco Systems, Inc., *Network-Based Application Recognition and Distributed Network-Based Application Recognition*, http://www.cisco.com/en/US/products/ps6350/products_configuration_guide_chapter09186a0080455985.html (last visited Sept. 30, 2006). This technology violates the broad version of the end-to-end arguments, but as the end-to-end arguments are just a design principle, there is

providers can slow down selected applications or content, speed them up or exclude them from the network completely.

In the analysis of Section B, the network provider is a local monopolist.[34] The size of its footprint relative to the size of the nationwide network may differ. In the extreme case, the network provider owns the nationwide network and has a nationwide monopoly in the provision of Internet services.

In Section C, the network provider competes with at least one other network provider.

The network provider also offers products in the market for applications, content or portals.[35] These products may be offered in two different ways:

In the first case, the complementary product is offered to consumers nationwide. Thus, if the size of the provider's footprint is smaller than the nationwide network, the product in question is not only offered to customers of its Internet services, but also to consumers living outside its footprint. A product that is offered this way will be called an affiliated product.

Alternatively, the network provider may only offer the product to customers of its Internet service. If the size of the provider's footprint is smaller than the nationwide network, consumers outside its footprint will not be able to use or buy the product. This kind of product will be referred to as proprietary product.

For a particular product, the two ways of offering the product are mutually exclusive.

This division reflects the way in which network providers' comple-

---

nothing that forces technology to comply with it. *See* van Schewick, *supra* note 32, at 101-03.

34. *See* Applications for Consent to the Transfer of Control of Licenses and Section 214 Authorizations by Time Warner Inc. and America Online, Inc., Transferors, to AOL Time Warner Inc., Transferee, *Memorandum Opinion & Order,* 16 FCC Rcd. 6,547, ¶ 74 (2001) [hereinafter *AOL Memorandum Opinion & Order*] ("The relevant geographic markets for residential high-speed Internet access services are local. That is, a consumer's choices are limited to those companies that offer high-speed Internet access services in his or her area, and the only way to obtain different choices is to move. While high-speed ISPs other than cable operators may offer service over different local areas (e.g., DSL or wireless), or may offer service over much wider areas, even nationally (e.g., satellite), a consumer's choices are dictated by what is offered in his or her locality.") *See also* Hausman et al., *Residential Demand, supra* note 31, at 135 ("From a consumer's perspective, the relevant geographic market is local because one can purchase broadband Internet access only from a local residence. Stated another way, a hypothetical monopoly supplier of broadband Internet access in a given geographic market could exercise market power without controlling the provision of broadband access in neighboring geographic markets").

35. Thus, the analysis assumes that the network provider is vertically integrated into the provision of at least some applications. Vertical integration, however, is not the only case to which the analysis applies. A similar analysis applies to other forms of close vertical relationships between the network provider and a provider of complementary products such as partial integration, partial equity investments, long-term contracts, or other forms of close affiliation.

mentary products are offered in today's Internet market. For example, AOL offers MapQuest, AOL Moviefone or its instant messenger to anybody using the Internet.[36] Similarly, AOL's portal is available both bundled with Internet service and separately.[37] By contrast, T-Online, the dominant German Internet provider, offers its portal only bundled with its Internet service.

The subsequent analysis does not further examine the choice of product provisioning, but takes the result as given.

B. Network Provider is Monopolist in the Market for Internet Services

Economic theory predicts that a network operator that has a monopoly in the market for Internet services does not generally have an incentive to discriminate against independent applications (Section 1). There are known exceptions to this rule, but there is considerable debate over whether these apply in the Internet context. The following analysis shows that the threat of discrimination is more severe than is commonly assumed. First, there are more exceptions than have been previously identified (Section 2). Second, some of the known exceptions may be more relevant in the Internet context than is commonly assumed (Section 3). Third, discrimination may be a profitable strategy, even if the network provider does not manage to drive independent applications from the market (Section 4).

1. No General Incentive to Discriminate

According to the "one monopoly rent" theory, a monopolist has no incentive to monopolize a complementary product market, if the complementary product is used in fixed proportions[38] with the monopoly good and is competitively supplied.[39]

---

36. Time Warner, Inc., *Time Warner Businesses: AOL*, Aug. 2, 2006, http://www.timewarner.com/corp/businesses/detail/aol/index.html.

37. Alan Breznick, *AOL Shifts Broadband Strategy*, CABLE DATACOM NEWS, Jan. 1, 2003, http://www.cabledatacomnews.com/jan03/jan03-3.html.

38. If the two goods are used in variable proportions, the monopolist may have an incentive to monopolize the complementary market, as this creates greater flexibility in its relative pricing of both components. Through appropriate pricing, the monopolist may be able to extract more surplus from consumers. If it needs a monopoly over both products to price discriminate in this fashion, monopolizing the second market will increase its profits. *See, e.g.*, Janusz A. Ordover et al., *Nonprice Anticompetitive Behavior by Dominant Firms toward the Producers of Complementary Products*, *in* ANTITRUST AND REGULATION: ESSAYS IN MEMORY OF JOHN J. MCGOWAN 119 (Franklin M. Fisher ed., 1985).

39. *See, e.g.*, ROBERT H. BORK, THE ANTITRUST PARADOX; A POLICY AT WAR WITH ITSELF 372-75 (Free Press 1993) (1978); RICHARD A. POSNER, ANTITRUST LAW 198-99 (2d ed. 2001).

In this case, there is only one final product, and, therefore, only one monopoly profit available in the market for the final product. The monopolist can extract the complete monopoly profit through its pricing of the monopoly good, and does not gain additional profits by monopolizing the complementary good.

This line of reasoning suggests that the monopolist need not monopolize the secondary market to extract the entire monopoly rent and therefore has no incentive to drive rivals from that market.

Moreover, economists note that the monopolist may benefit from the presence of independent producers in the complementary product market, implying that the monopolist will welcome, not exclude independent producers of complementary products. This argument has been labeled "internalizing complementary efficiencies (ICE)."[40]

If the presence of independent producers of complementary products generates additional surplus, the monopolist may be able to capture some of that surplus through its pricing of the primary good. In this case, the monopolist will earn greater profits when its rivals are in the market than when they are not. In this case, the monopolist does not wish to steal sales in the secondary market, but takes its profits by charging a higher price for the primary good.[41]

Whether the presence of independent producers generates additional surplus, depends on the structure of consumer preferences and on factors such as the intensity of competition in the complementary market or the degree of differentiation in the complementary market.[42]

---

40. Farrell & Weiser, *supra* note 13, at 89.

41. *See, e.g.*, Michael D. Whinston, *Tying, Foreclosure, and Exclusion*, 80 AM. ECON. REV. 837, 840, 850-52 (1990); Joseph Farrell & Michael L. Katz, *Innovation, Rent Extraction, and Integration in Systems Markets*, 48 J. INDUS. ECON. 413 (2000); Farrell & Weiser, *supra* note 13, at 103.

42. As the intensity of competition increases, prices are driven down to marginal costs. Due to the complementarity between both products, the monopolist benefits from lower prices in the complementary market. The lower prices in the complementary market, the higher demand (if demand is responsive to price) or consumer surplus (if demand is inelastic), and, consequently, the higher the profits that can be extracted in the primary market. *Id.*

Given the complementarity between both markets and appropriate consumer preferences, an increase in the quality or variety of complementary goods will increase consumers' valuation of the primary good. For example, consumer surplus rises, if a rival enters with a differentiated complementary product and some consumers prefer that product, *e.g.*, Whinston, *supra* note 41, at 850-52; Dennis W. Carlton & Michael Waldman, *The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries* 11 (George J. Stigler Ctr. for the Study of the Econ. and the State, Graduate Sch. of Bus., Univ. of Chicago, Working Paper No. 145, 2000), *available at* http://gsbwww.uchicago.edu/research/cses/WorkingPapersPDF's/145.pdf. The value consumers derive from greater variety may well differ depending on the type of complementary product. For example, consumers may value the fifth teleconferencing application less than the fifth multiplayer online game.

In general, two goods are complements if a decrease in the price of one increases the de-

While the "one monopoly rent" theory argues that exclusionary conduct in the complementary market will not increase the monopolist's profits, the "internalizing complementary efficiencies" theory suggests that such conduct may even reduce its profits.

Recent research shows that this line of reasoning is incomplete: Contrary to the assumptions of the "one monopoly rent" argument, there are cases in which the monopolist profits from monopolizing the complementary market. In these cases, the monopolist may profit from the presence of independent producers in the complementary market, but the loss of these profits may be more than offset by the gains associated with discriminating in the complementary market. In other words, although the monopolist may profit from the presence of independent producers in the complementary market, it may profit even more by excluding them from the market. In this case, the monopolist will engage in exclusionary conduct, if the associated profits are larger than the associated costs.[43]

2. New Exceptions

The following section highlights three exceptions that have not been previously considered. In the first exception, the complementary product is a source of outside revenues that the monopolist cannot extract in the primary market. In the second exception, which is a variant of the first, only the monopolist's complementary product is a source of outside revenue which is lost when rival producers of the product make the sales. This exception is particularly relevant in the Voice over IP (VoIP) context. In the third exception, the exclusionary conduct in the complementary market preserves a legally acquired monopoly in the complementary market.

The following analysis sets out the theories underlying these exceptions, highlights the conditions under which they apply and shows that these conditions may well be met in the Internet context.

2.1. Complementary Product Source of Outside Revenue

a) Theory

A monopolist in the primary market may be unable to extract the maximum possible profit through its sales of the primary good, if some of the revenue in the complementary market comes from outside

---

mand for the other. HAL R. VARIAN, INTERMEDIATE MICROECONOMICS; A MODERN APPROACH 112 (5th ed. 1999).

43. *See, e.g.*, Whinston, *supra* note 41, at 850-52, 855.

sources.[44]

In conventional markets, firms typically derive their revenue from sales of products or from fees for the provision of services. Firms also have the option of following the example of the media: they offer value to their customers, but at least partly charge third parties such as advertisers. In other words, a part of their revenue stems from selling access to their customers to interested third parties. In the extreme case, consumers get a firm's product or service for free, while all of the firm's revenue comes from outside sources.

If firms in the complementary market derive some of their revenue from outside sources, a monopolist in the primary market may be unable to earn the maximum possible profit unless it monopolizes the complementary market as well. To see this, consider the following example: suppose that firms in the complementary market offer their product or service for free and make all their revenue from selling access to their customers to third parties.

Usually, the monopolist can use a variety of tactics to extract or "squeeze" revenue from its rivals: A common set of tactics forces rival producers of the complementary good to lower the quality-adjusted price of their product.[45] This increases the consumer surplus available for extraction in the primary market. In the example, the price of complementary products already equals zero; thus, these tactics are not feasible.

In another tactic, the monopolist threatens to exclude a rival from the complementary market, unless the rival pays an access charge.[46] To be able to apply this tactic, the monopolist must have the power to exclude its rivals, for example due to intellectual property rights or because rivals' access to the primary good requires the monopolist's cooperation. While this mechanism enables the monopolist to extract its rivals' outside revenue, the monopolist may still earn less than if it excludes its rivals, monopolizes the complementary product market and captures all outside revenue directly: first, by monopolizing the complementary market, the monopolist gains a monopoly in the market for access to the users of its primary good. As a result, it will be able to charge higher prices (per customer) for access to its customers than competing producers of complementary products.[47] Second, due to its relationship with consum-

---

44. This theory is new and has not been covered by the existing literature.

45. For an overview of such tactics, see, e.g., Farrell & Katz, *supra* note 41, at 414-15.

46. *See, e.g., id.* at 422.

47. Ultimately, this will harm consumers, as firms will pass on at least some of the increased costs to their customers. For example, higher advertising fees will ultimately lead to higher prices for the goods that are advertised. *See* Rubinfeld & Singer, *supra* note 13, at 316; Jeffrey K. MacKie-Mason, An AOL/Time Warner Merger Will Harm Competition in Internet Online Services 23 (October 17, 2000) (Report submitted to the U.S. Federal Trade Commission), http://www-personal.umich.edu/~jmm/papers/aol-tw00-public.pdf.

ers in the primary market, the monopolist may have information about its consumers that enables it to charge higher prices to third parties.[48] Third, even if the per customer prices charged to third parties stay the same, the monopolist's profits will be lower in the presence of rivals due to the costs of negotiating and administering the access fees.

Thus, the monopolist will have an incentive to exclude its rivals from the complementary market, if the gains from directly capturing the outside revenue more than offset the reduction in profits that results from the reduction in complementary goods variety.

b) Application to the Internet

In the market for Internet content, portals and applications, firms often derive at least some of their revenue from outside sources by selling access to their customers to advertisers or online merchants.[49]

In the hypothetical network that is the focus of this analysis, the monopolist can extract at least some of its rivals' outside revenue: the network enables the monopolist to exclude applications from the network. Thus, the monopolist can condition the "access" of rivals' products and services on the payment of an access fee that captures some or all of its rivals' outside revenue. That this is not a mere theoretical possibility shows the practice of cable network owners in the United States. Unaffiliated Internet service providers who want to offer their service over a cable network have to pay a fixed fee per customer. In addition, the cable network owner receives a portion of the outside revenue that the Internet service provider earns on that customer.[50]

While the monopolist is able to capture some or all of its rivals' outside revenue by threatening exclusion, its outside revenue will be higher if it excludes its rivals and collects the outside revenue directly.

First, selling access to one large group of customers as a whole may yield substantially more revenue than selling access to subgroups of that group separately. This is obvious, if the monopolist network provider manages to monopolize the market in which access to its Internet service customers is sold.[51]

Second, through its billing relationship with customers of its Inter-

48. *See, e.g.*, CARL SHAPIRO & HAL R. VARIAN, INFORMATION RULES: A STRATEGIC GUIDE TO THE NETWORK ECONOMY 34-35 (1999).

49. ALLAN AFUAH & CHRISTOPHER L. TUCCI, INTERNET BUSINESS MODELS AND STRATEGIES; TEXT AND CASES 56 (2001); SHAPIRO & VARIAN, *supra* note 48, at 162-63.

50. *See* Seth Schiesel, *New Economy: A New Model for AOL May Influence Cable's Future*, N.Y. TIMES, Aug. 26, 2002, at C1 (discussing a contract between AOL and AT&T Comcast).

51. *See, e.g.*, Rubinfeld & Singer, *supra* note 13, at 316; MacKie-Mason, *supra* note 47, at 23. This remains true even if the monopolist does not manage to drive its rivals from the market completely. *See* the analysis *infra* Part II.B.4.2.

net service, the network owner has data on customer demographics that enables it to charge higher advertising fees or commissions for online sales than many of its rivals in the market for Internet content, portals and applications.[52]

Finally, due to the potentially large number of complementary products, negotiating and administering the access charges for unaffiliated content, applications and portals may be prohibitively expensive. In any event, these transaction costs will further decrease the monopolist's profits in the presence of rivals.

Thus, if firms in the market for a particular type of application, content or portal derive some of their revenue from outside sources, a monopolist in Internet services may have an incentive to monopolize that market in order to capture all outside revenue available in that market directly.

2.2. Monopolist's Complementary Product Source of Outside Revenue

a) Theory

In the scenario outlined above, only the network provider, not its rivals in the complementary market can realize higher outside revenues. As a result, letting rivals make the sales and extracting the outside revenue from them is less profitable than making the sales directly.

The following exception is a variant of this line of reasoning. The network provider's offering is a source of outside revenue; the rivals' offering does not provide this revenue. Thus, this revenue is lost if rivals make the sales. As a result, the network provider has an incentive to make as many sales as possible directly.

b) Application to the Internet[53]

Consider a local phone company that offers broadband Internet services over its network. Independent companies such as Vonage or Skype offer Voice over IP (VoIP) services to customers of this network provider. As the costs of long-distance calls using VoIP are usually considerably lower than the costs of long-distance calls using the conventional telephone service, those of the network provider's customers using VoIP will place less long-distance calls using the network provider's legacy

---

52. Even if those rivals require consumers to register before using their product or service, they have no way to verify the information, unless they require payment; in this case, they can verify the information as part of the billing process. *See* SHAPIRO & VARIAN, *supra* note 48, at 34-35; MacKie-Mason, *supra* note 47, at 11.

53. Thanks to Robert Pepper for highlighting this example.

telephone service.

To the network provider, conventional long-distance services are a source of outside revenue that is not similarly available to the providers of VoIP services. In the US, local phone companies are paid so-called access charges by long-distance providers for every long-distance call they originate or terminate. As access charges were traditionally intended to implicitly cross-subsidize local telephone service, regulators have mostly set these access charges significantly above the costs of originating or terminating long-distance calls. Thus, for many local phone companies, access charges are an important source of revenue.[54]

Independent VoIP providers threaten the source of this revenue: The more of the network provider's telephone customers place their long-distance calls using VoIP, the less access charges the network provider will receive. If independent VoIP providers are excluded from the network and the network provider does not offer VoIP itself,[55] customers are forced to make their long-distance calls using the conventional telephone service. Thus, exclusion in the VoIP market serves to preserve the network provider's current profits.[56]

It is not surprising that the first publicly documented incident of VoIP blocking involved a rural telephone company.[57] For rural phone companies, access charges constitute a substantial portion of their revenue. Thus, they have a particularly high incentive to protect this revenue.

---

54. *See* NUECHTERLEIN & WEISER, *supra* note 9, at 195, 204, 294.

55. The access charge is lost if the call is placed using VoIP, regardless of whether VoIP is provided by the network provider or by an independent provider. Thus, the network provider has an incentive not to have VoIP used on its network at all.

56. In the example discussed in the text, the existence of the outside revenue is the result of regulation that requires long-distance providers to pay above-cost access charges to local phone companies. Whether local phone companies that are local monopolists in the market for Internet services (this assumption holds throughout Section II.B) would also have an incentive to block VoIP in the absence of such regulation, is more difficult to determine.

57. In February 2005, Vonage, a US VoIP provider, complained to the Federal Communications Commission that its Internet telephony application was being blocked by Madison River Communications, a rural, local telephone company based in North Carolina. After a short investigation, Madison River and the FCC entered into a consent decree in March 2005. Madison River agreed to voluntarily pay $15,000 as well as to stop blocking VoIP applications; the FCC terminated the investigation. *See* Madison River Communications, LLC and Affiliated Companies, *Order*, 20 FCC Rcd. 4,295 (2005); Ben Charny, *Vonage Says Broadband Provider Blocks Its Calls*, CNET News.com, http://news.com.com/2100-7352_3-5576234.html (last modified Feb 14, 2005); Declan McCullagh, *Telco Agrees to Stop Blocking VoIP Calls*, CNET News.com, http://news.com.com/2100-7352_3-5598633.html (last modified Mar 3, 2005); Madison River Communications, *Who We Are*, *at* http://www.madisonriver.net/about_us/who_we_are.php (last visited Nov 21, 2006).

2.3. Monopoly Preservation in the Complementary Market

a) Theory

The monopolist may also use its monopoly over the primary good to protect a monopoly in the complementary market against dynamic competition. In this case, the exclusionary conduct in the complementary market preserves the monopoly in that market.[58]

For this theory to apply, the following conditions must be met:[59]

First, the monopolized product is not essential for all uses of the complementary good (i.e., there are uses of the complementary good that do not require the primary good). Second, the monopolist can prevent its rivals from selling their version of the complementary good to users of the primary good. Third, the complementary market is subject to economies of scale or network effects. Fourth, the monopolist also has a monopoly in the complementary market.

While the first condition explains why the monopolist will want to maintain its monopoly in the complementary market in spite of its monopoly in the primary market, the second and third condition provide the mechanism that enables the monopolist to protect its monopoly in the complementary market.

The first condition provides the motivation for preserving a monopoly in the complementary market in spite of the monopoly in the primary market: The existence of uses of the complementary good that do not require the primary good deprives the monopolist of its ability to extract all profits through sales of the primary good.

To see this, consider the following example: suppose there is some use of the complementary good that does not require the primary good. As a result, the complementary market consists of two parts: a "systems market" for uses in which the primary good is essential, and a "standalone market" for uses that do not require the primary good; consumers in the systems market desire the primary and the complementary good,

---

58. This theory has not been used as an exception to the "one monopoly rent" argument before. It generalizes from an argument that was used by the Federal Communications Commission in the AOL/ Time Warner merger proceeding with regard to instant messaging. *AOL Memorandum Opinion & Order*, *supra* note 34, at 6603-29, ¶¶ 128-200; Gerald Faulhaber, *Network Effects and Merger Analysis: Instant Messaging and the AOL-Time Warner Case*, 26 TELECOMM. POL'Y 311 (2002). *See infra* Part II.B.2.3.b).

59. The structure of the model and the underlying reasoning are parallel to the "primary good not essential" case outlined *infra* Part II.B.3.1. Whinston, *supra* note 41, at 854-55. However, in the "primary good not essential" case, the monopolist takes advantage of economies of scale and network effects in the complementary market to extend its monopoly to the complementary market by excluding its rivals from the systems part of the market. In the case under consideration here, the monopolist uses the same mechanism to protect a legally acquired monopoly in the complementary market against emerging competition.

whereas consumers in the stand-alone market desire only the complementary good.

Suppose there are rival producers of the complementary good. The monopolist can extract all monopoly profits in the systems market through its pricing of the primary good. As consumers in the stand-alone market do not buy the primary good, however, the monopolist does not derive any profit from its rivals' sales in that market. Moreover, the presence of rivals constrains its ability to price its version of the complementary good in the stand-alone market.

Thus, the monopolist cannot earn monopoly profits in the stand-alone market, unless it has a monopoly in that market. Consequently, keeping competitors out of the complementary market is a prerequisite for preserving current profits.

The second and third condition provide the mechanism that enables the monopolist to preserve the monopoly in the complementary market: In the presence of economies of scale or network effects, the monopolist may be able to drive potential rivals from the complementary market by excluding them from the systems part of the market.

When the second condition is met, the monopolist can deprive rival producers of complementary products of any sales in the systems part of the market.

This behavior does not exclude rivals from the stand-alone market. Given economies of scale[60] in the complementary market, the remaining sales to customers in the stand-alone market may not suffice to reach an economically efficient scale. Thus, being excluded from the systems part of the market, rivals may be forced to exit the stand-alone market as well.

Similarly, in the presence of network effects[61] in the complementary

---

60. Economies of scale exist, if an increase in output causes long run average total costs to decrease. In other words, the more output is produced, the lower the cost per unit. *E.g.*, ROBERT E. HALL & MARC LIEBERMAN, ECONOMICS; PRINCIPLES AND APPLICATIONS 177-78 (2d ed. 2001). For example, economies of scale exist, if fixed costs are large relative to marginal costs. In this case, an increase in output allows the firm to spread the fixed costs of production over greater amounts of output, lowering the costs of unit per output.

61. Network effects exist if the utility an individual customer derives from the consumption of a good depends upon, and increases with, the number of other customers who consume products that are compatible with that good. *See, e.g.*, the definition by Michael L. Katz & Carl Shapiro, *Network Externalities, Competition, and Compatibility*, 75 AM. ECON. REV. 424, 424 (1985) [hereinafter Katz & Shapiro, *Network Externalities*]. Network effects are covered by a large body of literature. *See, e.g.*, Jeffrey Rohlfs, *A Theory of Interdependent Demand for a Communications Service*, 5 BELL J. ECON. & MGMT. SCI. 16 (1974); Paul A. David, *Clio and the Economics of QWERTY*, 75 AM. ECON. REV. 332 (1985); Joseph Farrell & Garth Saloner, *Standardization, Compatibility, and Innovation*, 16 RAND J. ECON. 70 (1985); Katz & Shapiro, *Network Externalities, supra*; Michael L. Katz & Carl Shapiro, *Technology Adoption in the Presence of Network Externalities*, 94 J. POL. ECON. 822 (1986); Carmen Matutes & Pierre Regibeau, *"Mix and Match": Product Compatibility without Network External-*

market, exclusion from the systems part of the market may suffice to drive competitors from the market or into a niche existence. In markets with network effects, the incumbent's large installed base makes it difficult for new entrants to dislodge the incumbent. Exclusion from the customers in the systems part of the market makes it even more difficult for new entrants to reach the critical mass of customers necessary to start the positive feedback required to succeed with their product.

Thus, the exclusion of rivals from the systems part of the market enables the monopolist to protect a legally acquired monopoly in the complementary market against emerging competition.

Such a scenario may be particularly relevant, if the complementary market belongs to an R&D intensive industry subject to dynamic or "Schumpeterian" competition.[62] Due to the presence of intellectual property rights, economies of scale or network effects, R&D intensive industries are prone to short run exercise of market power. In other words, competition in these markets often results in a single firm dominating the market. Thus, firms in these industries typically compete "for the market," not "within the market." While firms with market power (the winners of the competition) are an inherent feature of such industries, their dominance may be temporary, as rapid technological change and drastic

---

ities, 19 RAND J. ECON. 221 (1988); Brian W. Arthur, *Competing Technologies, Increasing Returns, and Lock-In by Historical Events*, 99 ECON. J. 116 (1989); Jeffrey Church & Neil Gandal, *Network Effects, Software Provision, and Standardization*, 40 J. INDUS. ECON. 85 (1992); Nicholas Economides & Steven C. Salop, *Competition and Integration among Complements, and Network Market Structure*, 40 J. INDUS. ECON. 105 (1992); Joseph Farrell & Garth Saloner, *Converters, Compatibility, and the Control of Interfaces*, 40 J. INDUS. ECON. 9 (1992); Michael L. Katz & Carl Shapiro, *Product Introduction with Network Externalities*, 40 J. INDUS. ECON. 55 (1992); Stanley M. Besen & Joseph Farrell, *Choosing How to Compete: Strategies and Tactics in Standardization*, 8 J. ECON. PERSP. 117 (1994); Michael L. Katz & Carl Shapiro, *Systems Competition and Network Effects*, 8 J. ECON. PERSP. 93 (1994) [hereinafter Katz & Shapiro, *Systems Competition*]; Nicholas Economides, *The Economics of Networks*, 14 INT'L J. INDUS. ORG. 673 (1996); *see also* SHAPIRO & VARIAN, *supra* note 48, chapters 7-9 (analyzing network effects in the context of information goods); Joseph Farrell & Paul Klemperer, *Coordination and Lock-In: Competition with Switching Costs and Network Effects*, *in* 3 HANDBOOK OF INDUSTRIAL ORGANIZATION (forthcoming) (providing a recent survey), *available at* http://ssrn.com/abstract=917785; Mark A. Lemley & David McGowan, *Legal Implications of Network Economic Effects*, 86 CAL. L. REV. 479 (1998) (analyzing the legal implications of network economic effects). For some critical voices, see STAN J. LIEBOWITZ & STEPHEN E. MARGOLIS, WINNERS, LOSERS & MICROSOFT. COMPETITION AND ANTITRUST IN HIGH TECHNOLOGY (rev. ed. 2001); William J. Kolasky, *Network Effects: A Contrarian View*, 7 GEO. MASON L. REV. 577 (1999).

62. On dynamic or "Schumpeterian" competition, *see* JOSEPH A. SCHUMPETER, CAPITALISM, SOCIALISM AND DEMOCRACY, 81-86 (Harper Perennial 1975); *see also* Dennis W. Carlton & Robert H. Gertner, *Intellectual Property, Antitrust and Strategic Behavior* 19-22 (Nat'l Bureau of Econ. Research, Working Paper No. 8976, 2002); David S. Evans & Richard L. Schmalensee, *Some Economic Aspects of Antitrust Analysis in Dynamically Competitive Industries* (Nat'l Bureau of Econ. Research, Working Paper No. 8268, 2001); Howard A. Shelanski & Gregory J. Sidak, *Antitrust Divestiture in Network Industries*, 68 U. CHI. L. REV. 1, 10-15 (2001).

innovations may cause demand for their product to collapse: for example, rivals may come up with a vastly superior product or develop a new product that makes the incumbent's product obsolete. Thus, incumbents in these industries are primarily constrained by dynamic competition - by the innovation of other firms seeking to replace the existing firm with market power. To avoid being dislodged by rivals, incumbents are forced to innovate themselves.

In the scenario described above, a monopolist could use its market power in the primary market to preserve the legally obtained market power in the complementary market, distorting the dynamic competition for future market power. Instead of innovating to prevent being dislodged by competitors, the monopolist could simply exclude its rivals from the systems part of the complementary market, preventing them from reaching the scale or network size necessary to displace the incumbent.

b) Application to the Internet

The conditions underlying this model may well be present in the Internet context.

First, a specific provider's Internet service may be non-essential for using applications or accessing content. Consider the market for residential broadband Internet access in the United States.[63] Depending on local conditions, the owner of a cable network that provides broadband Internet access through its affiliated broadband Internet access provider may well be a local monopolist.[64] While this monopolist offers broadband Internet access only in the area covered by its network, it may offer content or applications to Internet users nationwide. In this case, the area covered by its network constitutes the "systems market," while customers outside its footprint make up the "stand-alone market."

Such a situation is not uncommon. For example, where it has been able to strike a deal with cable network owners, AOL offers its portal bundled with broadband Internet access. In addition, consumers nationwide can buy the portal without access, known as the "bring your own access" option.[65] Other AOL services such as MapQuest or AOL Moviefone are also offered to all consumers on the Internet.[66] Similarly, if a narrowband access provider has a monopoly with respect to narrowband access, but offers its portal both to its narrowband access customers and

63. The market for broadband Internet access is considered a distinct market from the narrowband access market, *see, e.g., AOL Memorandum Opinion & Order, supra* note 34, at 6574-77, ¶¶ 68-73; Hausman et al., *Residential Demand, supra* note 31, at 135-57.

64. *See supra* note 34.

65. Breznick, *supra* note 37.

66. *See* Time Warner, Inc., *supra* note 36.

to anybody on the Internet, the narrowband access service will be non-essential for customers accessing the portal via broadband access services.[67]

Second, in the hypothetical network that forms the basis of the analysis, the monopolist can technically exclude rivals' applications, content or portals from running over its network. As a result, the monopolist's Internet service customers (the consumers in the systems market) are unable to access or use these products. Thus, rivals are deprived of any sales in the systems part of the market.

Third, the markets for software applications, Internet content and portals are all subject to significant economies of scale. The development of these products and services is characterized by large fixed costs, while the marginal costs of production and distribution over the Internet are very small. Thus, the marginal cost of production[68] is very low relative to the average cost of production,[69] resulting in significant economies of scale.[70]

In addition, many software applications are subject to direct or indirect network effects.[71] For example, a communication service like instant

---

67. Scott Beardsley et al., *Making Sense of Broadband*, MCKINSEY Q., Issue 2, 2003 at 78-87 (showing that "so far, [. . .] faster and better access to the Internet is the sole killer application of broadband") Thus, the scenario described in the text may be quite common. *See also* Farrell & Weiser, *supra* note 13, at 119.

68. The marginal cost of production is the incremental cost of producing an additional unit of the good. Thus, the marginal cost of production does not include the costs of product development, *e.g.*, HALL & LIEBERMAN, *supra* note 60, at 168-69. In the case of software applications, Internet content and portals, the marginal cost of production is the cost of making an additional digital copy of the product, which is typically very low.

69. The average cost of production indicates a firm's total cost per unit of output. In other words, it denotes the total cost associated with a particular product divided by the quantity of output produced. Thus, contrary to the marginal cost of production which does not include the cost of developing the first unit of the product, the average cost of production includes the cost of development divided by the total number of copies. *E.g.*, *id.* at 168.

70. This cost structure (low marginal costs relative to average costs), which results in significant economies of scale, is generally viewed as a key economic characteristic of the markets for these products. *See, e.g.*, SHAPIRO & VARIAN, *supra* note 48, at 3-4 (discussing information goods in general); Michael L. Katz & Carl Shapiro, *Antitrust in Software Markets*, *in* COMPETITION, INNOVATION, AND THE MICROSOFT MONOPOLY: ANTITRUST IN THE DIGITAL MARKETPLACE 29 (Jeffrey A. Eisenach & Thomas M. Lenard eds., 1999) (discussing software markets), *manuscript available at* http://faculty.haas.berkeley.edu/shapiro/software.pdf; POSNER, *supra* note 39, at 245-46 (discussing Internet content, portals and software); MacKie-Mason, *supra* note 47, at 14 (discussing broadband portals); Rubinfeld & Singer, *supra* note 13, at 307 (discussing broadband content).

71. Network effects are called "direct network effects," if the consumption benefits directly result from the size of the network. *E.g.*, Katz & Shapiro, *Network Externalities*, *supra* note 61, at 424. "Indirect network effects" exist, if consumer demand for the primary good increases with the variety of complementary goods and services. In this case, network effects arise from supply-side economies of scale in the complementary market: a larger installed base for the primary product allows application developers to spread sunk development costs over a

messenger or Internet telephony is more valuable the more people can be contacted using the service.[72] Viewers for multimedia content are subject to indirect network effects:[73] The larger the catalogue of content available in a particular format, the more users value owning viewers compatible with that format. At the same time, content providers are more likely to incur the costs of coding their content in a particular format, the larger the installed base of viewers compatible with that format.

Finally, at least some of these markets are subject to rapid technological change. Not surprisingly, markets for software applications are the canonical example of R&D intensive industries subject to dynamic competition.[74]

Now consider a network provider that is a local monopolist in Internet services and has acquired a dominant position in the nationwide market for a particular application. Such a provider has an incentive to exclude rivals from that market to protect itself from dynamic competition and preserve its monopoly in that market. Whether the monopolist will manage to prevent new entrants from entering the complementary market by excluding them from access to its Internet service customers, depends on the exact size of economies of scale with respect to the product in question, on the strength of any potential network effects and on the size of both the monopolist's network and the remaining network.

This theory played an important role in the FCC's evaluation of the merger between AOL and Time Warner. Time Warner owned a number of broadband cable networks; AOL held a dominant position in the market for instant messaging services and offered its instant messaging program to consumers nationwide. The FCC was concerned that the merged firm could use its control over broadband cable networks to disadvantage competitors seeking to overturn AOL's legally acquired monopoly in instant messaging services. To alleviate this problem, the FCC approved the merger subject to a condition (among others) that required AOL Time Warner to interoperate with instant messenger competitors prior to offering "advanced" instant messaging services.[75]

---

larger potential sales base. Thus, in the presence of economies of scale and free entry into the complementary product market, a larger customer base leads to lower costs and greater variety of complementary products. *See, e.g.*, *id.* at 424; Katz & Shapiro, *Systems Competition*, *supra* note 61, at 99. The existence of direct or indirect network effects is a fundamental economic characteristic of many software markets. *See, e.g.*, EVANS & SCHMALENSEE, *supra* note 62, at 9-11; Katz & Shapiro, *supra* note 70.

72. *E.g.*, Faulhaber, *supra* note 58.

73. *E.g.*, MacKie-Mason, *supra* note 47, at 16.

74. *E.g.*, EVANS & SCHMALENSEE, *supra* note 62, at 4-15.

75. *AOL Memorandum Opinion & Order*, *supra* note 34, at 6603-29, ¶¶ 128-200. For an in-depth analysis of the economic rationale underlying this condition, see Faulhaber, *supra* note 58.

3. Relevance of Known Exceptions

There are a number of known exceptions to the "one monopoly rent" argument and to the "internalizing complementary efficiencies" argument outlined above. The following section describes two exceptions that may be relevant in the network neutrality context, but whose relevance in the network neutrality context has not been discussed in detail yet.[76]

In the first exception, the primary good is not essential for all uses of the complementary good, making it impossible for the monopolist to extract all monopoly profits through its pricing of the primary good.

In the second exception, the monopolist excludes competitors from the complementary market in order to protect its monopoly in the primary market.

3.1. Primary Good not Essential

a) Theory

The structure of models in this category,[77] and the underlying reasoning, is similar to the "monopoly preservation in the complementary market" case described above:

First, the monopolist has a monopoly in the primary market and the primary good is not essential (i.e., there are uses of the complementary good that do not require the primary good). Thus, the complementary market consists of a systems market and a stand-alone market. As a result, the monopolist cannot extract all profits through its pricing of the primary good and profits from extending its monopoly to the complementary market.

Second, there is a mechanism that enables the monopolist to exclude rival producers of the complementary good from the systems part of the market. Third, the complementary market is subject to economies of scale or network effects.

---

76. For a more complete overview of known exceptions to the "one monopoly rent" argument, see Farrell & Weiser, *supra* note 13, at 105-19; van Schewick, *supra* note 32, at 245-67.

77. The following theory was developed by Whinston, *supra* note 41, at 854-55, and is widely accepted as an exception to the "one monopoly rent" argument. *See, e.g.*, Dennis W. Carlton, *A General Analysis of Exclusionary Conduct and Refusal to Deal: Why Aspen and Kodak Are Misguided*, 68 ANTITRUST L.J. 659, 667-68 (2001); Dennis W. Carlton & Michael Waldman, *The Strategic Use of Tying to Preserve and Create Market Power in Evolving Industries*, 33 RAND J. ECON. 194, 195 (2002); Jay Pil Choi & Chris Stefanadis, *Tying, Investment, and the Dynamic Leverage Theory*, 32 RAND J. ECON. 52, 55 (2001); Whinston, *supra* note 41, at 71. For a detailed application of this theory in the context of the open access debate, *see* Rubinfeld & Singer, *supra* note 13. *See also* Farrell & Weiser, *supra* note 13, at 119.

Given economies of scale in the complementary market, the monopolist can force its rivals to exit the stand-alone market by excluding them from the systems part of the market, extending its monopoly to the complementary market.[78]

Similarly, in the presence of network effects[79] in the complementary market, exclusion from the systems part of the market may suffice to drive competitors from the market or into a niche existence:

If the benefits derived from a larger network are large relative to the benefits of product differentiation in the network good, competition between two incompatible technologies will usually result in a single technology dominating the market.[80] The reason is that network effects give rise to strong positive feedback in technology adoption: other things being equal, consumers derive larger benefits from a larger network. As the larger network is more attractive, more consumers will join that network, making it even more valuable, leading to even more consumers joining the network. Once this positive feedback loop sets in, the affected technology will quickly pull away from its rivals in market share, ultimately dominating the market. This phenomenon is referred to as "tipping."[81]

As small initial advantages may quickly get magnified, small differences, in either perception[82] or reality, may determine the outcome of the competition. Therefore, establishing an early lead in installed base[83] that is large enough to start the positive feedback loop is an important strategy in network markets.[84]

---

78. In the "monopoly preservation in the complementary market" case described *supra* Part II.B.2.3, the monopolist uses this mechanism to protect a legally acquired monopoly in the complementary market against emerging competition.

79. On network effects in general, see *supra* note 61. On direct and indirect network effects, see *supra* note 71.

80. Often, competitors will not be driven completely from the market. In particular, some customers with high switching costs or a unique preference for a competitor's product will prefer to stay with that competitor in spite of the strong network effects associated with the winning technology. *See, e.g.*, Faulhaber, *supra* note 58, at 329 n.37.

81. "'Tipping' occurs when a single provider reaches a critical mass of customers that are so attractive to others that competitors must inevitably shrink, in the absence of interoperation." *Id.* at 316.

82. In network markets, consumer expectations about the future size of the network play a crucial role in determining the outcome of the competition. This is due to the costs of belonging to the losing network: A consumer who has chosen the losing network can either switch to the winner, which may be costly, or suffer from the lower value of a small network. To avoid this situation, the consumer will choose the network that it expects to be the winner. *See, e.g.*, Besen & Farrell, *supra* note 61, at 118.

83. The installed base is the total number of consumers who have already bought the network good.

84. A substantial lead in installed base is not the only factor that influences the outcome of the competition. Due to the huge benefits of belonging to the winning network, users have a strong desire to choose the technology that will ultimately prevail. Therefore, consumers expectations of who the winner will be are at least as important. Other factors that may influence customers' expectations and that may therefore result in a competitive advantage are an estab-

Thus, if the monopolist excludes its rivals from the complementary market, it can capture all customers in the systems market. If the systems market is large enough, the monopolist's advantage in that market may enable it to reach a critical mass of customers that are so attractive to others that positive feedback sets in, making it impossible for a rival to catch up.

If the presence of rivals increases consumer surplus, the exclusion of rivals may reduce the monopolist's profits in the systems market.[85] In such a case, monopolizing the complementary market increases the monopolist's profits, if the gain from monopolizing the stand-alone market is larger than the loss resulting from the exclusion of the rival in the systems market.[86]

If the complementary market is subject to network effects, two effects make it even more likely that exclusion is a profitable strategy:

First, the potential profits from winning the competition between incompatible technologies are huge, increasing the benefits of exclusion. Imagine a competition between incompatible technologies that are subject to indirect network effects. If the winning standard is protected by intellectual property, the winner can make money on any primary and complementary product that uses the standard. Given the potentially large number of complementary products in markets with indirect network effects, licensing fees can lead to substantial profits.[87] For example, the winner in the standard competition between competing media player technologies who wins with a proprietary standard protected by intellectual property will not only dominate the market for media players, but will also be able to charge licensing fees for every piece of music or video that is encoded for use with the player.

Second, if the complementary product is subject to network effects, the presence of an independent rival in the complementary market does not necessarily increase the monopolist's profits in the systems market, a fact that reduces the costs of exclusion. If the monopolist's and the rival's complementary product are incompatible, sales to the rival decrease the size of the network of users of the monopolist's complementary

---

lished reputation, a well-known brand name, or ready visible access to capital. Thus, an unknown firm with an early lead may be overtaken by a market leader that enters second, but has a well-known brand name and good reputation. *See, e.g.*, Katz & Shapiro, *Systems Competition*, *supra* note 61, at 107.

85. For example, if the rival produces a differentiated product, the rival's presence creates additional surplus, some of which the monopolist can extract through its sales of the primary good. Thus, the monopolist's profits in the systems market are increased if its rival is in the market.

86. *See* Whinston, *supra* note 41, at 850-52, 855.

87. Due to the cost structure of information products, profits are not even dependent on charging a monopoly price. *See* the analysis *infra* Part II.B.4.1.

product. As a result, the value users can derive from the monopolist's complementary product (and the profit the monopolist can extract from them) is lower than the corresponding value if the rival does not make any sales.[88]

b) Application to the Internet

As has been set out above,[89] the conditions underlying this theory are quite common in the Internet context:

Network providers may be local monopolists in the market for Internet services, but offer applications, content or portals to consumers nationwide. Network technology enables network providers to exclude providers of complementary products from access to its Internet service customers. At the same time, the markets for applications, content or portals are usually subject to significant economies of scale and, potentially, network effects.

As a result, an Internet service provider may be able to force its rivals from the nationwide market (the stand-alone market) by excluding rival portal, content or application providers from access to its Internet service customers (the systems part of the market). Whether exclusion from its Internet service customers suffices to drive its competitors from the nation-wide market[90] depends on the exact size of economies of scale with respect to the product in question, on the strength of any potential network effects and on the size of both the monopolist's network and the remaining network.[91]

Such a provider will have an incentive to monopolize the market for a particular type of application, content or portal, if the increased profit from additional application, content or portal sales nationwide more than offsets the reduction in broadband access revenues due to the reduction in variety resulting from the exclusion of its rivals with respect to its Internet service customers.[92]

---

88. Carlton & Waldman, *supra* note 77, at 206-07.

89. *See supra* Part II.B.2.3.b.

90. *See* Rubinfeld & Singer, *supra* note 13, at 310-13 (providing a numerical example). Their paper assesses the likelihood of content discrimination (i.e., blocking or degrading the quality of outside content) by a broadband network provider that is vertically integrated into the market for broadband content and portals in the context of the merger between AOL and Time Warner.

91. Even if the monopolist's footprint is not large enough to force its rivals to exit the market completely, exclusion from a part of the market may put them at a severe competitive disadvantage by forcing them to operate at a less efficient scale or with a smaller network. *See* the analysis *infra* Part II.B.4.1.

92. *See, e.g.*, Rubinfeld & Singer, *supra* note 13, at 310-13.

3.2. Monopoly Preservation in the Primary Market

a) Theory

In the following class of models, exclusionary behavior in the complementary market maintains the monopoly in the primary market.[93]

In models belonging to this category, the monopolist faces potential competition in the primary market. The monopolist can deter entry to the primary market by engaging in exclusionary conduct in the complementary market. Thus, by deterring entry to the primary market, the exclusionary behavior in the complementary market preserves the monopoly in the primary market.

Economists have come up with a number of explanations of why exclusionary conduct in the complementary market may be able to deter entry to the primary market. The following analysis will focus on an explanation that is particularly relevant in the Internet context: the exclusionary behavior in the complementary market harms future competitors in the primary market by depriving them of a source of complementary products.[94] As a result, in order to make any sales in the primary market, an entrant to the primary market needs to enter the complementary market as well (or otherwise secure a sufficient supply of complementary products). If this is significantly more difficult or costly than entering the primary market alone, potential entrants to the primary market may refrain from entering.

For such a strategy to succeed, two conditions must be met:

First, the exclusionary behavior in the complementary market must deprive a potential entrant to the primary market of a source of complementary products. As a result, the entrant cannot enter the primary market alone, but must enter both markets at once.

Second, simultaneously entering both markets must be more difficult or costly than the sum of the costs of entering both markets on their own.[95] Otherwise, the exclusionary behavior in the complementary mar-

---

93. On this type of monopoly maintenance in general, see, e.g., Carlton, *supra* note 77, at 668-71; Farrell & Weiser, *supra* note 13, at 109-12; Steven C. Salop & R. Craig Romaine, *Preserving Monopoly: Economic Analysis, Legal Standards, and Microsoft*, 7 GEO. MASON L. REV. 617, 623-24 (1999). For specific models, see, for example, Carlton & Waldman, *supra* note 77; Choi & Stefanadis, *supra* note 77.

94. *See, e.g.*, Carlton, *supra* note 77, at 669-70.

95. *E.g.*, U.S. Dep't of Justice & FTC, Non-Horizontal Merger Guidelines, § 4.212 (promulgated in 1984 and reaffirmed in 1992 and 1997) ("The relevant question is whether the need for simultaneous entry to the secondary market gives rise to a substantial incremental difficulty as compared to entry into the primary market alone. If the entry at the secondary level is easy in absolute terms, the requirement of simultaneous entry to that market is unlikely adversely to affect entry to the primary market."), *available at* http://www.usdoj.gov/atr/public/guidelines/2614.htm [hereinafter Non-Horizontal Merger

ket is unlikely to adversely affect entry to the primary market.

Economists have identified four alternative reasons why simultaneous entry to both markets may be significantly more difficult or costly than the sum of the costs of entering each market on its own:

-increased cost of capital,
-differing economies of scale in both markets,
-the uncertainty of innovation, or
-the existence of indirect network effects.

Increased Cost of Capital

An entrant that is forced to enter both markets may face an increased cost of capital, if it only has experience relevant for operating in one of the markets. If the skills and knowledge necessary to succeed in both markets differ considerably, the increased probability of failure due to his inexperience in one of them may lead lenders to charge a higher rate for the necessary capital. The risk premium will be even larger, if the entrant has to incur huge sunk costs to enter the market. The higher sunk costs, the more costs cannot be recovered in the event of failure.[96]

Differing Economies of Scale in Both Markets

Entering two markets is more difficult than entering one, if the minimum efficient scale in both markets differs considerably. In this case, the entrant must choose between operating at an inefficiently small size in one market or at a larger than necessary scale in the other. Both strategies may significantly increase the operating costs of the entering firm.[97]

Uncertainty of Innovation

Given the uncertainty associated with the innovative process, the need to innovate successfully in two markets may decrease the probability of successful entry. To see this, assume that the probability of innovating successfully in one component is k. In this case, the chances of successful innovation in n components are k^n. Unless k is close to 1, this is considerably lower than k.[98] Thus, the probability of successful innovation in n components required to enter into n markets simultaneously is lower than the probability of successful innovation and successful entry

Guidelines].

96. For an argument along these lines, see Oliver E. Williamson, *Assessing Vertical Market Restrictions: Antitrust Ramifications of the Transaction Cost Approach*, 127 U. PA. L. REV. 953, 953-93 (1979); Non-Horizontal Merger Guidelines, *supra* note 95, § 4.212.

97. Non-Horizontal Merger Guidelines, *supra* note 95, § 4.212.

98. Carlton and Gertner, *supra* note 62, pp. 23-27; Choi & Stefanadis, *supra* note 77.

in one component market.

Existence of Indirect Network Effects

If the primary good is subject to indirect network effects[99] and any available complementary goods are offered exclusively with the monopolist's platform, an entrant into the primary market faces a "chicken and egg" problem: due to consumers' desire for variety in complementary products, consumers prefer a primary good that already offers a large number of complementary goods and services. At the same time, due to economies of scale and sunk costs in complementary product development, developers of complementary products prefer to develop products for primary goods that already have a large number of users. Thus, "[an entrant into the primary market] either has to offer consumers much lower value or has to incur large sunk costs to develop (or subsidize) a wide range of [complementary goods and services] before there is a large user base to purchase them." [100]

b) Application to the Internet

The conditions underlying this theory may well be present in the Internet context.

First, the exclusionary behavior in the complementary market must deprive a potential entrant to the market for Internet services of a source of complementary products.

By excluding rival producers of Internet portals, content and application from its network, the monopolist network provider may be able to drive its rivals from the nationwide market.

To deprive a potential entrant of a source of complementary products, the monopolist needs not only drive rival content and application producers from the market. He also needs to deny access to its own content and applications to consumers outside its network.[101] Otherwise,

---

99. For a definition of indirect network effects, see *supra* note 71, 351.

100. Richard J. Gilbert & Michael L. Katz, *An Economist's Guide to US v. Microsoft*, 15 J. ECON. PERSP. 25, 30 (2001) (referring to operating systems and application programs). Under the label "applications barrier to entry," this line of reasoning has featured prominently in the Microsoft case. *See, e.g.*, United States v. Microsoft Corp., 253 F.3d 34, 54-56 (D.C. Cir. 2001); Gilbert & Katz, *supra*, at 28-30.

101. In addition to offering its own content and applications to the customers of its Internet service, the monopolist may also "allow" independent producers of these products to offer their products to the customers of its Internet service, as long as they agree to offer their products exclusively to these customers. Stated differently, instead of depriving a potential entrant into the market for Internet services of a source of complementary products by driving rival content and application producers from the market, the monopolist could deprive the potential entrant of a source of complementary products by signing exclusive contracts with independent content and application producers. Whether a monopolist could profitably impose such an

customers of rival network providers could simply use the monopolist's content and applications with the rival's Internet service.[102] Hence, for a particular application or content, this strategy and the "primary good not essential" strategy are mutually exclusive.[103]

Thus, this theory is only applicable, if (a) an Internet service provider offers proprietary content and applications exclusively to customers of its Internet service,[104] and if (b) - potentially due to the exclusion of rivals from its customers - there are not enough remaining independent applications, content or portals available that could be used by customers of rival or newly entering network providers.[105] In this case, a new entrant into the market for Internet services needs to develop (or subsidize the development of) its own content or applications.

One may wonder whether the condition (b) may ever be fulfilled in

---

exclusivity provision, has been the subject of considerable debate. The Chicago school denied such a possibility, arguing that the other party to the exclusive contract would not agree to contracts that made it worse off, *e.g.*, BORK, *supra* note 39, at 309. More recent research has shown that this argument is incomplete: it does not consider the possibility that the exclusive contract imposes harm on third parties that are not parties to the contract, while not making the contracting parties worse off. In other words, the exclusive contract gives rise to a negative externality on third parties, and due to this externality, signing an exclusive contract is jointly optimal for the contracting parties. For a discussion of this question with pointers to the literature, see, for example, Gilbert & Katz , *supra* note 100, at 31-33; Michael D. Whinston, *Exclusivity and Tying in U.S. v. Microsoft: What We Know; and Don't Know*, 15 J. ECON. PERSP. 63, 66-70 (2001).

102. Usually, this theory is applied to cases, where the entrant's primary good is technically unable to take advantage of the set of applications developed for the monopolist's primary good. For example, software applications make use of a specific operating system's application programming interfaces and therefore run only on this operating system. As a result, customers of the entrant's operating system are technically unable to use applications developed for the incumbent's operating system. By contrast, as long as an application complies with the specifications of the Internet protocol, it can run over any physical network that supports the Internet protocol. As a result, applications adhering to that standard can be used by anyone connected to the Internet. Thus, from a technical point of view, the applications offered by the monopolist could be used by customers of a rival network provider as well. Therefore, the entrant's inability to use the monopolist's applications and content is not due to technical differences or incompatibility between the Internet services offered by the monopolist and a potential entrant, but results from the monopolist's business decision to offer its content and applications exclusively to customers of its own Internet service.

103. The strategy described here requires that the monopolist does not offer the content, application or portal to consumers outside its network; by contrast, in the "primary good not essential" strategy, the inability to earn monopoly profits on its sales to consumers outside its network is the reason that leads the monopolist to monopolize the complementary market as well. *See supra* Part II.B.3.1.

104. The potential anti-competitive implications of such a strategy are explored by, for example, MacKie-Mason, *supra* note 47, at 23-25; Rubinfeld & Singer, *supra* note 13, at 313-16.

105. Alternatively, the monopolist could reach the same result by allowing independent producers of applications, content and portals to offer their products to the customers of its Internet service, if they agree to provide the products exclusively to its customers. *See* the discussion *supra* note 101.

the Internet context: after all, there are a number of portals, content and applications that are available to anyone using the Internet today. The condition may be met in emerging markets such as the market for broadband Internet services, the market for Internet services for mobile phones or in emerging national markets in countries outside the United States. For example, there may be not enough independent applications or content that take advantage of broadband specific characteristics such as high transport speed or broadband's always on capacity.[106] Similarly, there may not be enough independent applications or content that are adapted to the specific limitations associated with using the Internet from mobile phones.[107] In a country that just started adopting the Internet, there may not be enough independent applications or content in the national language.

One may also imagine that consumers perceive certain applications and content as indispensable elements of Internet usage. If these applications and content are exclusively available with the incumbent's Internet service, consumers may not consider an entrant's Internet service an adequate alternative to the incumbent's Internet service, unless the entrant offers a similar set of applications and content itself. In this case, to deter entry to the market for Internet services, the incumbent does not need to drive all existing independent applications, portals and content from the market and offer all affiliated complementary products exclusively to customers of its Internet service. It suffices to restrict the exclusionary conduct to those applications and content that consumers view as essential. Although there are independent applications and content left that customers of a rival Internet service could use, the entrant will still be forced to enter the market for specific applications and content in order to be able to compete in the primary market.[108]

Second, simultaneously entering the market for Internet services and the market for content or applications must be more difficult or costly than entering the market for Internet services alone. This requirement is fulfilled as well. Simultaneous entry into both markets is more difficult or costly than entry into the market for Internet services alone if

---

106. Many broadband customers may simply use broadband Internet services to access narrowband offerings at higher speed. According to McKinsey, "so far, [. . .] faster and better access to the Internet is the sole killer application of broadband.," Beardsley et al., *supra* note 67, § "what happens next?" and Exhibit 6.

107. For example, compared to PCs, mobile handsets have small screens, limited keypads and not a lot of storage. *See, e.g.*, Francis Deprez et al., *Portals for All Platforms, in* MCKINSEY Q., Issue 1, 2002, at 92.

108. Finally, one may imagine a situation in which the nationwide market for Internet services consists of a collection of local monopolies who all bundle their content, portal and applications exclusively with their Internet service. In this case, a new entrant into the market for Internet services would have to enter the market for content, portals, or applications as well.

the two markets exhibit at least one of the four characteristics described above. In the Internet context, all four characteristics are present: first, entry to both markets requires very different capabilities, second, production in both markets is subject to differing economies of scale, third, success in the different markets is uncertain, and finally, due to the incumbent's exclusionary conduct, the provision of Internet service is subject to indirect network effects with respect to the individual provider's network.

First, developing software applications or interesting content requires very different capabilities than operating a network. As a result, a potential entrant to the market for Internet services may not necessarily have the capabilities required for entering the markets for applications or content.[109] In addition, most of the cost of entry into those complementary industries consists of the sunk costs of developing the offering that cannot be recovered in case of failure.[110] Due to these factors, the need to enter the complementary markets as well considerably increases the risks associated with entry to the primary market. Consequently, an entrant into both markets will most likely be charged higher rates for capital than an entrant to the primary market alone.

Second, the market for Internet services and the markets for complementary products are subject to very different economies of scale: for example, McKinsey estimates that assuming an average revenue per user of $ 18.00 to $ 22.50 a year in 2005, a broadband PC portal in Germany would need more than 8 million unique users to break even.[111] By contrast, the economies of building and operating physical networks over which IP services could be provided are much lower.[112]

Third, although network technology is undergoing rapid innovation, a new entrant into the market for Internet services can take advantage of existing technology and does not have to innovate itself. By contrast, the development of applications and content is subject to considerable uncertainty. If a potential entrant to the market for Internet services needs to develop several applications and services in order to be able to compete with the incumbent's Internet service, the uncertainty associated with each development reduces the likelihood of successful entry to the market for Internet services.

---

109. *See, e.g.*, Robert Niewijk et al., *Why European ISPs Need Partners, in* MCKINSEY Q., Issue 1, 2003, 98.

110. That the costs of capital may increase with the amount of entry costs that are sunk is discussed by W. KIP VISCUSI ET AL., ECONOMICS OF REGULATION AND ANTITRUST 157-58, 161 (3d ed. 2000).

111. Deprez, et al., *supra* note 107.

112. For example, as of June 30, 2001, the 10 largest providers in the market for broadband transport services in the United States had between 1,409,000 and 360,000 residential broadband customers, Yoo, *supra* note 31, at 256 tbl.7 (internal citations omitted).

Fourth, Internet service is subject to indirect network effects:[113] the more applications and content are available for users, the more valuable Internet service becomes. At the same time, the development of content and applications is subject to economies of scale.[114] As a larger number of users allows application and content developers to spread the fixed costs of development over a larger potential sales base, under free entry to these markets the variety of applications and content will be higher and their cost will be lower, the larger the number of users.

Technically, any application based on the Internet protocol can run over any network that is connected to the public Internet and supports the Internet protocol. As a result, from a technical point of view, the relevant network for indirect network effects is not an individual provider's network, but the global Internet. Thus, technically, Internet service providers compete under conditions of compatibility.

By excluding independent applications from its network and offering its own applications exclusively to its own Internet customers, an Internet service provider can reintroduce indirect network effects with respect to its own network.[115] Stated differently, as a result of this strategy, the benefits of adding a new user do not accrue to anyone connected to the Internet, but are limited to the customers of the new user's Internet service provider. Application and content developers now have to decide whether to offer their product to the customers of the Internet service provider with the "closed" network or to the customers of Internet service providers following an open system strategy. Due to economies of scale in the production of application and content, the developers will base their decision on the size of the different networks.

As a result, an entrant to the market for Internet services will have difficulties attracting application and content developers who write for its network instead of the incumbent's. Thus, due to the incumbent's strategy, the entrant faces the chicken and egg problem described above: consumers will not subscribe to its Internet service in the absence of an attractive amount of content and applications; application and content developers will not produce for its network in the absence of an attractive number of users.[116]

---

113. *E.g.*, Speta, *Handicapping*, *supra* note 5, at 83-84.

114. *See supra* Part II.B.2.3.b.

115. An Internet service provider could reach the same effect (i.e., reintroduce indirect network effects with respect to its own network) by using proprietary protocols inside its network, *see, e.g.*, COMPUTER SCI. & TELECOMM.S BD. & NAT'L RES. COUNCIL, THE INTERNET'S COMING OF AGE 147-49 (2001). An alternative strategy may be the provision of quality of service only within an Internet provider's network, *see, e.g.*, SHAPIRO & VARIAN, *supra* note 48, at 187.

116. COMPUTER SCI. & TELECOMM.S BD. & NAT'L RES. COUNCIL, *supra* note 115, at 147-49, describe a similar situation in the context of provider-specific indirect network effects

Thus, a monopolist provider of Internet services may be able to deter entry to the market for Internet services by excluding rival producers of applications, content and portals from the market and offering its own content and applications exclusively to the customers of its own Internet service.[117] This strategy may reduce consumers' valuation of its Internet service, as the exclusion of rival producers of applications, content and portals reduces the variety of complementary products available to customers of its Internet service. Thus, in deciding whether to employ such a strategy, the monopolist must trade off the loss in Internet service fees against the gains in future monopoly profits.

4. Profitability of Discrimination without Monopolization

In the network neutrality context, researchers commonly focus on the ability and incentive of a network provider to monopolize the market for selected complementary products. The previous sections have followed this approach. It is based on the implicit assumption that discrimination is only profitable, if the network provider manages to monopolize the complementary market. As the following section shows, this focus may be too narrow: A network provider may have an incentive to discriminate against an application even if the provider does not manage to drive it from the market.

Thus, researchers commonly underestimate the likelihood of discriminatory behavior by network providers: If discrimination requires the network provider to monopolize the complementary market to be a profitable strategy, discrimination will be restricted to those cases where the network provider can expect to drive its competitors from the complementary market. If, however, discrimination is a profitable strategy, even if the network provider does not manage to monopolize the complementary market, it is much more likely to occur.

The following analysis will cover four of the five exceptions outlined above.[118] It is based on the assumption that the exclusion of rivals from the network provider's Internet service customers increases the number of sales of the network provider's complementary product. At least some of the network provider's Internet service customers that would have used a rival's complementary product in the absence of ex-

---

due to the use of proprietary protocols inside the network.

117. As highlighted *supra* note 101, an alternative way of deterring entry would be to sign exclusive contracts with independent producers of applications, content and portals. Such a strategy would have the advantage that the monopolist does not have to bear losses with respect to its Internet service fees, as its customers would have access to all existing applications, content and portals.

118. The fifth exception, "monopoly preservation in the primary market," *supra* Part II.B.3.2, requires that rival producers of excluded complementary products are driven from the market.

clusion will use the network provider's offering instead. Thus, by excluding rival producers of applications or content from its network, the network provider gains additional sales from its Internet service customers at the expense of its rivals. If the complementary product is subject to economies of scale or network effects and the network provider offers its complementary product to customers nation-wide, the exclusion from the network provider's Internet service customers may force rivals to operate at an economically less efficient scale or with a smaller network of customers, putting the rivals at a competitive disadvantage in the rest of the market as well and potentially leading to even more additional sales for the network provider's complementary product.

Based on this assumption, the analysis will ask, whether a larger number of sales of the network provider's complementary product increase its profits, even if the network provider does not manage to monopolize the complementary market in question.

4.1. More Sales at Market Prices

In a perfectly competitive market subject to constant returns to scale, simply increasing the number of sales at the market price will not increase profits. In such an industry, long-run equilibrium prices equal marginal costs, resulting in zero profit per unit sold. As a result, a firm cannot increase its profits by making additional sales at the market price. Instead, it needs to gain a monopoly position that enables it to raise prices above marginal costs.

Markets for applications, content and portals are different: In these markets, the exclusionary conduct need not result in a monopoly to increase the network provider's profits; it suffices if it results in a larger number of sales.[119] This is due to the cost structure underlying the production of applications and content: the production of these goods is characterized by high fixed costs and very low marginal costs. While the costs of developing the first instance of an application or content may be significant, the costs of producing additional copies may be negligible. Due to the need to cover fixed costs, such products are priced signifi-

119. For an economic model demonstrating this effect in the context of tying, see Patrick DeGraba, *Why Lever into a Zero-Profit Industry: Tying, Foreclosure, and Exclusion*, 5 J. ECON. & MGMT. STRATEGY 433 (1996). In DeGraba's model, oligopolists sell a differentiated good (the primary good) and a homogenous good (the complementary good) that are used in fixed proportions to produce the final good. The homogenous good can be produced at constant marginal cost by any firm incurring a certain fixed cost. The homogenous market is characterized by free-entry Cournot competition. In such a market, the zero-profit price of the good is greater than the marginal cost. As a result, the oligopolist in DeGraba's model will tie in order to increase the sales of the complementary good. Note that this model does not require the complementary good to be a differentiated good.

cantly above marginal costs.[120]

If the market price is significantly above marginal costs, a firm does not need to be able to charge monopoly prices to increase its profits: Instead, making additional sales at the market price may be enough.[121] More sales enable the firm to spread the fixed costs of production over more units, resulting in lower average costs per unit and a higher profit margin at the same price. Put differently, once a firm has made enough sales to cover the fixed costs, any additional sale at the market price only adds to the profits. For example, given that gross margins of 80% or 90% are common in computer software,[122] any additional sale may lead to a significant increase in profits.

By excluding rival producers of complementary products from its network, the network provider gains additional sales. These additional sales increase the network provider's profit, even if the excluded rivals are not driven from the complementary market completely.

For example, this fact has important implications for the relevance of the exception "primary good not essential" outlined above.[123] Whether a network provider can monopolize the nation-wide complementary market in question by excluding its rivals from access to its Internet service customers, depends on a variety of factors such as the exact size of economies of scale with respect to the complementary product in question, the strength of any potential network effects and the size of both the monopolist's network and the remaining network. Ultimately, the cases in which monopolization is a realistic prospect may not be very common. As monopolization is not necessary to increase the network provider's profits, however, this restriction does not matter. As long as the exclusion of rivals from its Internet service customers enables the network provider to increase the number of sales of its complementary product and the additional profits resulting from more sales at the market price are larger than the costs of exclusion, exclusion will be a profitable strategy. Given how often the conditions underlying the "primary good not essential" exception[124] are met, this drastically increases the likelihood

---

120. If the price were equal to marginal costs, firms would not be able to cover their fixed costs and would earn negative profits. In the long run, firms would not operate in such a market. Thus, even if all firms earn zero profit per unit in long-run equilibrium, equilibrium prices are above marginal costs.

121. SHAPIRO & VARIAN, *supra* note 48, at 161. The importance of market share and number of units sold in knowledge-based products is also described by AFUAH & TUCCI, *supra* note 49, at 52-54. For an economic model demonstrating this effect in the context of tying, see DeGraba, *supra* note 119.

122. Katz & Shapiro, *supra* note 70.

123. *See supra* Part II.B.3.1.

124. As outlined *supra* Part II.B.3.1, these conditions are: The network provider has a monopoly in the primary market (i.e., the market for Internet services). The primary good is not essential (i.e., there are uses of the complementary product that do not require the primary

that exclusion may be a profitable strategy.

4.2. More Outside Revenue

As indicated above,[125] a network provider may have an incentive to monopolize the complementary market, if the complementary product is a source of outside revenue that cannot be extracted in the market for Internet services. For reasons outlined above, its outside revenue will be higher if it excludes its rivals and collects the outside revenue directly than if tries to capture some or all of its rivals' outside revenue by threatening exclusion.

This increase in profit, however, is not dependent on a monopolization of the complementary product market.

Although the network provider's revenue from outside sources will be highest if it manages to monopolize the market for access to its customers, increasing the number of customers who access the network provider's offering may still lead to higher profits than trying to extract the outside revenue from its rivals.

Evidence suggests that even without a monopoly, the relationship between the number of customers and advertising revenue is not a linear one: for example, MacKie-Mason reports that although Lycos had 72 percent as many unique visitors as Yahoo! in September 1999, it received only 36 percent as much advertising revenue.[126] This implies that selling access to one large group of customers as a whole may still yield substantially more revenue than selling access to subgroups of that group separately, even if the seller does not have a monopoly in the market for access to its customers.

In addition, through its billing relationship with customers of its Internet service, the network provider has data on customer demographics that enables it to charge higher advertising fees or commissions for online sales than many of its rivals in the market for Internet content, portals and applications.[127] Again, this ability is not dependent on a mo-

---

good). This condition is met when the Internet service provider offers its complementary product not only to its Internet service customers, but to customers nation-wide. The complementary market is subject to economies of scale or network effects, a condition that is met in most markets for applications, content or portals. The monopolist has a mechanism at its disposal that enables it to exclude its rivals from access to its primary good customers. In the Internet context, technology that enables the network provider to distinguish between applications running over its network and to control their execution provides the network provider with this capability.

125. *See supra* Part II.B.2.1.

126. MacKie-Mason, *supra* note 47, at 13.

127. Even if those rivals require consumers to register before using their product or service, they have no way to verify the information, unless they require payment; in this case, they can verify the information as part of the billing process, SHAPIRO & VARIAN, *supra* note 48, at 34-35; MacKie-Mason, *supra* note 47, at 11.

nopoly in the complementary market.

A similar argument applies to the variant of this exception described above.[128] In this variant, a network provider excludes Voice over IP (VoIP) providers from access to its Internet Service customers in order to preserve the outside revenue in the form of access charges associated with traditional long-distance calls. Such a strategy will also be profitable, if the network provider does not manage to exclude the VoIP providers from its customers completely: Access charges are per-call charges set by regulators; the ability to charge them is not dependent on keeping all long-distance customers. Every long-distance call lost to a VoIP provider reduces profits; the more conventional long-distance calls the network provider manages to keep, the higher its profits.

4.3. Monopoly Preservation in the Complementary Market

In the "monopoly preservation in the complementary market" exception outlined above,[129] the network provider excludes rival producers of a complementary product from access to its Internet service customers to preserve a legally acquired monopoly in the corresponding complementary market.

In the exception outlined above, the analysis assumes that the monopolist will be able to keep its rivals out of the nation-wide market by excluding them from access to its Internet service customers.

Even if the monopolist's footprint is not large enough to force its rivals to stay out of the market completely, exclusion from a part of the market may put them at a severe competitive disadvantage by forcing them to operate at a less efficient scale or with a smaller network. Compared to a world without exclusion, this may slow down the erosion of the network provider's monopoly in the complementary market, preserving its ability to charge monopoly profits for a longer time. Again, this may make exclusion a profitable strategy, even if the network provider does not manage to keep its rivals out of the market completely.[130]

C. Network Provider Faces Competition in the Market for Internet Services

Up this point, the analysis was based on the assumption that the network provider is at least a local monopolist in the market for Internet services. This assumption is in line with standard economic thinking on

---

128. *See supra* Part II.B.2.2.

129. *See supra* Part II.B.2.3.

130. *Cf.* POSNER, *supra* note 39, at 254 (making a similar argument with respect to the profitability of monopoly preservation through exclusionary conduct in new economy markets, if the monopoly is of intellectual property).

vertical exclusionary conduct in complementary markets: according to economic theory, an economic actor without monopoly power in the primary market will be incapable of excluding competitors in the complementary market using vertical practices such as tying, vertical mergers or exclusive dealing. A monopoly in the primary market is therefore considered to be an indispensable precondition for successful monopolization of the secondary market.[131]

Given this theory, it is not surprising that most of the literature on vertical exclusionary conduct in complementary product markets focuses on exclusionary conduct by monopolists: after all, the same conduct is unlikely to pose any significant anti-competitive threat, if the firm faces competition in the primary market.[132] This theory has also shaped the evaluation of existing firms' behavior in a complementary market: allegations of anti-competitive conduct in a secondary market are often countered by evidence that the accused firm does not have monopoly power in the primary market.[133] Alternatively, the analysis of the monopoly case is used as an argument "a maiore ad minus": if a monopolist in the primary market does not have the ability and incentive to impede competition in the secondary market, it is argued, then a competitive firm's conduct will pose even less of a threat.[134]

Based on this line of reasoning, most commentators believe that the threat of discrimination against independent providers of complementary products can be mitigated by competition in the market for Internet services. Stated differently, it is usually assumed that competition in the market for Internet services will restrict a network operator's ability and incentive to discriminate against independent content, portals or applications. This assumption forms the basis for two common policy proposals:

---

131. *E.g., id.* at 195; Yoo, *supra* note 31, at 188-91. Similarly, some sort of market power or political power is considered to be a prerequisite for strategies that raise rivals costs, *e.g.*, DENNIS W. CARLTON & JEFFREY M. PERLOFF, MODERN INDUSTRIAL ORGANIZATION 353 (3d ed. 2000).

132. For an important exception to this point, see the literature on the exercise of aftermarket power by a firm that faces competition in the foremarket. This literature focuses on the question whether primary market competition precludes anti-competitive aftermarket actions. For an analysis of these issues with pointers to the literature, *see* Jeffrey K. MacKie-Mason & John Metzler, *Links between Vertically Related Markets: Kodak, in* THE ANTITRUST REVOLUTION: ECONOMICS, COMPETITION, AND POLICY 386 (John E. Kwoka, Jr. & Lawrence J. White eds., 3d ed. 1999).

133. *See, e.g.*, Yoo, *supra* note 31, at 249-50, 253 in the context of the open access debate ("I conclude that the structure of the broadband industry renders it unlikely that such combinations will pose any significant anti-competitive threat. . ."); and Yoo, *Beyond Network Neutrality*, *supra* note 4, at 61 in the context of the network neutrality debate ("This suggests that for most of the country, competition should remain sufficiently robust to ameliorate concerns of anticompetitive effects.").

134. *E.g.*, Speta, *Vertical Dimension*, *supra* note 5, at 986 (discussing this notion in the context of the open access debate).

the first proposal assumes that fostering facilities-based competition (i.e., increased competition between operators of different physical networks) will mitigate a network provider's ability and incentive to discriminate.[135] The second proposal seeks to restore competition at the Internet service provider level by requiring the owners of broadband networks to allow independent Internet service providers to offer their services over these networks. This regulatory response is called "open access," "multiple access" or "forced access," depending on the point of view of the commentator.[136]

The following analysis shows that this assumption is not correct: a network provider may have the ability and incentive to exclude rival content, applications or portals from its network, even if it faces limited competition[137] in the market for Internet services.[138] Apart from increasing the number of cases in which unaffiliated providers of complementary products face a real threat of discrimination, this result also implies that neither facilities-based competition nor open access regulation are the appropriate tools to mitigate this threat.[139]

Three arguments drive this result: First, in the Internet context, the ability to exclude competitors from a complementary market (the markets for applications, content and portals) is not dependent on a monop-

---

135. *See, e.g.*, Yoo, *Mandating Network Neutrality*, *supra* note 4, at 67 ("On the other hand, regulators can adopt a more humble posture about their ability to distinguish anticompetitive from procompetitive behavior and attempt to resolve the problem by promoting entry by alternative broadband platforms. Once a sufficient number of alternative last mile providers exists, the danger of anticompetitive effects disappears, as any attempt to use an exclusivity arrangement to harm competition will simply induce consumers to obtain their services from another last mile provider").

136. An example of this line of reasoning can be found in the FCC memorandum and opinion in the AOL Time Warner merger proceeding. *AOL Memorandum Opinion & Order*, *supra* note 34, at 6594-95, ¶ 107 ("We believe that if unaffiliated ISPs receive nondiscriminatory access to Time Warner cable systems [. . .] the merged firm's incentives and ability to withhold unaffiliated content from its subscribers will be substantially mitigated."); *see id.* at 6596, ¶ 112; Lemley & Lessig, *Ex parte, supra* note 31.

137. The analysis assumes that the network provider competes with at least one other network provider. *See infra* note 140 and accompanying text.

138. *See also* Joseph Farrell, *Open Access Arguments: Why Confidence is Misplaced, in* NET NEUTRALITY OR NET NEUTERING: SHOULD BROADBAND INTERNET SERVICES BE REGULATED 195 (Thomas M. Lenard & Randoph J. May eds., 2006) (arguing that limited competition may not necessarily remove network providers' incentives to discriminate). For a similar argument in the context of the debate over censorship by private proxies, see Seth F. Kreimer, *Censorship by Proxy: the First Amendment, Internet Intermediaries, and the Problem of the Weakest Link*, 155 U. PA. L. REV. 11, 33-36 (2006) (arguing that competition between Internet service providers may not be sufficient to discipline Internet service providers that disable content needlessly based on arguments very similar to the ones advanced above).

139. There may be other reasons that justify these proposals, though. For example, according to Lemley & Lessig, *Ex parte, supra* note 31, at 21-25, ¶ 54-65, the reduction in application-level innovation by independent providers resulting from the threat of discrimination constitutes only one of three arguments in favor of open access.

oly position in the primary market (the market for Internet services). Instead, the power to exclude is conferred by network technology (Section 1). Second, realizing the benefits of exclusion (i.e., an increase in profits (or, sometimes, a preservation of current profits)) does not require a monopoly position in the primary market. The lack of monopoly in the primary market even increases the network provider's incentive to increase profits by engaging in exclusionary conduct in the complementary market, as the network provider cannot simply extract the available monopoly profit by charging higher prices in the primary market (Section 2). Third, due to various factors such as the existence of switching costs or the ability to use discrimination instead of exclusion, the exclusion of rivals will not necessarily cause the network provider's Internet service customers to switch to another provider, making the costs of exclusion lower than is commonly assumed (Section 3).

The following analysis assumes that the network provider competes with at least one other network provider.[140] In addition, the network provider may offer content or applications. A particular application or content may be offered to all consumers (affiliated product) or exclusively to the customers of its own Internet service (proprietary product).[141]

1. Ability to Exclude

Today, technology is available that enables network providers to distinguish between applications and content running over its network and to control their execution. This technology enables the network provider to exclude selected complementary products from its network or to slow down their execution.

This technology enables the network provider to exclude unaffiliated providers of complementary products from access to its Internet service customers, independent of a monopoly in the market for Internet services.

While the exclusionary power of the technology does not reach beyond the network provider's network, exclusion from the network provider's Internet service customers may suffice to drive rival producers of complementary products from the nation-wide market, if there are economies of scale or network effects in the complementary market.[142]

---

140. This assumption reflects the reality in the broadband market for residential customers in the US. According to a recent study by the United States Government Accountability Office, the median number of broadband providers available to residential users is two. United States Government Accountability Office, Report to Congressional Committees; Telecommunications; Broadband Deployment is Extensive throughout the United States, but It Is Difficult to Assess the Extent of Deployment Gaps in Rural Areas, at 18 (May 2006), *available at* http://www.gao.gov/new.items/d06426.pdf.

141. *See supra* Part II.A.

142. *See supra* Part II.B.3.1.

Whether this will happen, depends on the exact size of economies of scale with respect to the complementary product in question, on the strength of any potential network effects and on the nation-wide number of both the monopolist's Internet service customers and the customers of other network providers.[143] Thus, in this context, the ability to drive competitors from the nation-wide complementary market depends on the network provider's nation-wide market share in the market for Internet services. Again, a monopoly position in this market is not required.

2. Benefits of Exclusion

In a variety of cases, the exclusionary conduct will increase (or preserve) the network provider's profits in the complementary market. As the analysis will show, this increase is not dependent on a monopoly position in the market for Internet services; nor does it require the network provider to gain a monopoly in the complementary market.[144] Instead, the lack of monopoly in the primary market constrains the network provider's ability to extract profits in the market for Internet services, making the ability to realize profits in the complementary market even more attractive. As a result, there are many more cases in which exclusion may be profitable than is commonly assumed.

In general, by excluding rival producers of a specific complementary product from access to the network provider's Internet service customers, the network provider will increase the number of sales of its own complementary product.[145]

As set out in detail above, the increase in the number of sales will often lead to an increase in profits. In the cases outlined above, the increase in profits results from an increase in the number of sales, not from the ability to charge monopoly profits. Thus, to be profitable, the exclusionary conduct need not drive rivals from the complementary market completely.

In the cases described above, the network provider had a monopoly in the market for Internet services. As the following analysis will show, however, the increase in profits due to exclusion was not dependent on this monopoly position (Sections 2.1 – 2.3). In addition, it will highlight a variant of the "monopoly preservation in the primary market" exception outlined above: the network provider may exclude selected producers of complementary products from access to its customers to protect its

143. Even if the monopolist's footprint is not large enough to force its rivals to exit the market completely, exclusion from a part of the market may put them at a severe competitive disadvantage by forcing them to operate at a less efficient scale or with a smaller network. *See* the analysis *supra* Part II.B.4.1.

144. *See supra* Part II.B.4.

145. *See supra* Part II.B.4.

competitive position in the primary market (Section 2.4).

2.1. More Sales at Market Prices

In the exception "more sales at market prices,"[146] the increase in profits resulting from the higher number of sales in the complementary market was driven by the specific cost structure of the markets for applications, content or portals, which are characterized by high fixed costs and low marginal costs. This cost structure is not affected by the existence of market power in the market for Internet services.[147]

2.2. More Outside Revenue

In the exception "more outside revenue,"[148] the increase in profits resulted from the logic of pricing in the markets for advertising. This enabled the network provider to realize higher outside revenue by selling access to a large group of its Internet service customers directly, instead of letting rival producers of complementary products sell access to smaller groups of customers and extracting the outside revenue from them. Again, a monopoly in the market for Internet services is not required for this relationship to hold.

There is evidence that some Internet service providers (i.e., economic actors that face competition in the Internet service market) do in fact attempt to reduce the amount of time their customers spend on unaffiliated content or portal offerings. For example, in the AOL/Time Warner merger proceeding the FCC found that "[t]he record in this proceeding provides some evidence that AOL already seeks to limit its members' access to unaffiliated content on the World Wide Web. For example, AOL requires that content appearing on AOL web sites have only a limited number of hyperlinks to unaffiliated content." [References omitted][149]

In the variant of this exception,[150] the network provider was interested in excluding Voice over IP (VoIP) providers from access to its customers, because it could only charge access charges for long-distance calls placed using the conventional telephone service, not for long-

---

146. *See supra* Part II.B.3.1, and Part II.B.4.1.

147. DeGraba's model, DeGraba, *supra* note 119, which demonstrates this effect in the context of tying, supports this analysis. In the model, the producer of the primary good has an incentive to tie in order to increase the number of sales of the secondary good, although it competes with another producer in the primary market. Thus, in the model the incentive to exclude independent competitors from the secondary market is not dependent on a monopoly position in the primary market. The model is discussed in more detail *supra* note 119.

148. *See supra* Part II.B.2.1, and Part II.B.4.2.

149. *AOL Memorandum Opinion & Order*, *supra* note 34, at 6594 ¶ 106; *id.* at 6593-94, 104-06.

150. *See supra* Part II.B.2.2, and Part II.B.4.2.

distance calls using VoIP. Access charges are per-call charges set by regulation; they do not depend on a monopoly in the market for Internet services.

2.3. Monopoly Preservation in the Complementary Market

In the exception "monopoly preservation in the complementary market,"[151] the ability to preserve the monopoly in the complementary market depended on various factors such as the exact size of economies of scale with respect to the complementary product in question, on the strength of any potential network effects and on the nation-wide number of both the monopolist's Internet service customers and the customers of other network providers. A monopoly in the market for Internet services is not required.

2.4. Preserving Competitive Position in the Primary Market

The exclusion of rivals may protect the network provider's competitive position in the market for Internet services, even if it faces competition in this market. Such an incentive may occur, if an Internet transport provider offers proprietary content and applications exclusively to its transport customers. This is a common strategy, as it enables the transport provider to relax price competition in the market for Internet services by differentiating its transport service from rival offerings, to reduce customer turnover and increase profits by raising switching costs and to make additional profits by selling access to its customers to advertisers, content providers or online merchants.[152]

Independent content and applications that can be used from any provider threaten the success of this strategy:

First, they reduce the differentiation of a provider's offerings by providing comparable, but independent alternatives.

Second, independent offerings may reduce the switching costs of the network provider's Internet service customers. Switching costs are the costs a customer incurs when switching to a competitor.[153] For example, when switching from one dial up access provider to another, a consumer must reconfigure his or her Internet access program. When switching from broadband access over cable to DSL, a consumer also needs to buy and install new equipment such as a DSL modem. Switching costs reduce customer turnover: when considering whether to switch

151. *See supra* Part II.B.2.3, and Part II.B.4.2.

152. *See, e.g.*, MacKie-Mason, *supra* note 47, at 11.

153. *See, e.g.*, VARIAN, *supra* note 42, at 603-05 (providing overview of switching costs); *see also* SHAPIRO & VARIAN, *supra* note 48, chapters 5-6 (treating switching costs in the context of information goods).

to a competitor, a customer takes his switching costs into account. Switching costs also make demand more inelastic, enabling the seller to raise prices.[154]

Bundling Internet transport service with proprietary content and applications that are offered exclusively to transport customers is a common way to increase switching costs.[155] In this case, consumers loose access to their old provider's proprietary content and applications when they switch to another provider. As a result, they have to search for new ones and learn how to use them. If the new provider does not offer comparable content or applications, not being able to use the old provider's proprietary content or applications any more is itself a cost of switching. In addition, many proprietary offerings induce their customers to engage in nontransferable database creation and customization. For example, Internet service providers offer provider-specific e-mail addresses that cannot be transferred to another provider;[156] to take advantage of services like stock portfolio tracking, instant messaging or customized news pages, users have to enter nontransferable data as well. When switching providers, customers need to notify relevant parties of their new e-mail addresses or instant messaging IDs and loose their site-specific data.

Independent offerings may reduce the effectiveness of this strategy by reducing customers' switching costs: as the independent application or content is not tied to a specific provider of Internet services, consumers can continue to use it after switching providers. In addition, by creating site-specific data on independent offerings, customers can avoid becoming locked in to a specific access provider.[157]

Third, as has been set out above, independent alternatives may also reduce the time customers spend using proprietary offerings, reducing third party revenues such as advertising fees or commissions for online sales.

By excluding independent applications and content that compete with the network provider's proprietary offerings, the network provider may be able to prevent these problems.

3. Costs of Exclusion

Compared to the monopoly case, the existence of other, competing network providers may increase the costs of exclusionary behavior in the complementary market. Due to a variety of factors such as the existence of switching costs or the ability to use discrimination instead of exclu-

154. *E.g.*, VARIAN, *supra* note 42, at 604-05; Hausman et al., *Residential Demand, supra* note 31, at 164.

155. *See, e.g.*, MacKie-Mason, *supra* note 47, at 11.

156. *See, e.g.*, SHAPIRO & VARIAN, *supra* note 48, at 109-10.

157. They get locked in to the independent offering, though.

sion, the costs of exclusion will still be lower than is commonly assumed.

If the network provider is the only supplier of Internet services in a particular geographic area, consumers have no alternative way of accessing the excluded application or content. They either subscribe to the provider's Internet service or do not use Internet services at all. Thus, the costs of the exclusionary behavior are twofold: first, the price of Internet services will be lower due to the reduction in application and content variety.[158] Second, without being able to use the excluded application or content, some consumers may not value Internet services enough to pay the lower price.[159] Given that that the pricing of the service already reflects the reduced value, the number of lost transport customers will probably not be very high.

If the provider competes with at least one other network provider, consumers who desire access to the excluded application may switch to another provider. As these consumers do not have to forgo Internet services altogether, the number of lost transport customers will probably be higher than if the excluding network provider does not face competition. Thus, competition increases the costs of exclusionary behavior in the complementary market.[160]

Several factors may limit the costs of exclusionary behavior in spite of competition in the market for Internet services:

First, if the exclusionary conduct manages to drive the producers of the excluded application or content from the market, switching providers will not enable consumers to get access to the excluded product. As a result, fewer consumers will switch in response to the exclusion.[161]

Second, switching costs may prevent consumers from changing providers to get access to the excluded application.[162] This is the case, if the increased value from being able to use the excluded application is smaller than the costs of switching to another network provider. Thus, the higher switching costs, the lower the number of customers lost to other network providers.[163]

Third, and potentially most importantly, the network provider may be able to avoid this problem altogether by using discrimination instead

158. *See* Wu, *Network Neutrality*, *supra* note 1, at 153 (discussing the costs of a discriminatory pricing scheme that prohibits customers of a network provider's basic Internet service from using specific applications).

159. *See, e.g.*, Rubinfeld & Singer, *supra* note 13, at 310.

160. *See, e.g.*, *id.* at 310.

161. *See, e.g.*, *id.* at 312-13.

162. *See, e.g.*, Hausman et al., *Residential Demand*, *supra* note 31, at 164; Kreimer, *supra* note 138, at 34-35; NUECHTERLEIN & WEISER, *supra* note 9, at 156. For a discussion of switching costs in the market for Internet services, see *supra* notes 153-157 and accompanying text.

163. Switching costs do not protect the network provider from losing business from new customers.

of direct exclusion.[164] As today's network technology provides the ability to control the execution of applications running over the network, a network provider can negatively affect the execution of particular applications. For example, the network provider can slow down the transport of certain applications or the delivery of selected content. If a network provider discriminates against a rival's complementary product, consumers' use of the rival's product is less satisfactory than their use of the network provider's own offering, even if the rival's product is of higher quality.

Thus, discrimination works indirectly by changing consumers' perception of the quality of a rival's offering. As consumers are unable to detect the true cause of the lower quality, they may mistakenly attribute it to bad product design and use competing products whose use is more satisfactory. For example, a slow gaming experience may be due to bad application programming, insufficient server capacity at the gaming site or slow Internet transport. Similarly, long waiting times for pages from an online shop could result from bad programming of the underlying databases or insufficient server speed. If customers do not usually experience problems with network speed, they will be inclined to blame the online game or the online shop.

With discrimination, consumers have the option of choosing the rival's product, but prefer the network provider's product which they perceive to be of higher quality. Contrary to direct technical exclusion or tying, they will not feel that their choice has been restricted. As they do not wish to use the rival's product, the discrimination will neither reduce their valuation of the network provider's Internet services nor cause them to switch to a competing provider.

Thus, if the network provider discriminates against rival products instead of excluding them directly, competition in the market for Internet services does not increase the costs of the exclusionary conduct.

D. Conclusion

Although a network provider does not generally have an incentive to discriminate against independent providers of content, applications or content, the analysis has highlighted a variety of circumstances under which it may have such an incentive. Such an incentive may not only occur if it has a (local) monopoly in the market for Internet services, but also if it faces competition. Whether the conditions giving rise to such an incentive are present in a real life situation, is an empirical question. In most cases, however, the network provider need not be able to gain a monopoly in the complementary market to make exclusion a profitable

164. *See, e.g.*, Rubinfeld & Singer, *supra* note 13, at 310, 313.

strategy, making the threat of discrimination more relevant than commonly assumed.

In most cases, the network provider need not exclude all independent developers of complementary products from its network in order to increase its profits. Instead, it will often be profitable to exclude only those complementary products that directly compete with one of its own complementary products. This reduces the costs of exclusion, as the reduction in complementary goods variety is restricted to those products that are actually excluded.

Due to the specific characteristics of markets for applications and content such as the cost structure of information goods and (sometimes) the existence of network effects, the exclusion of rivals may lead to gains that are significantly higher than in traditional markets. As a result, it is more likely that the gains from exclusion exceed the associated costs, making it more likely that exclusion is a profitable strategy.[165]

II. IMPACT ON APPLICATION-LEVEL INNOVATION

The previous part has highlighted conditions under which a network provider may have an incentive to exclude independent producers of applications, content or portals from access to its Internet service customers. When these conditions are present, independent producers of complementary products face a real threat of discrimination.

The following section analyzes the impact of this threat on innovation in the markets for applications, content and portals ("application-level innovation"). It shows that the threat of discrimination reduces the amount of application-level innovation by independent producers of complementary products (Section A). While discrimination increases network providers' incentives to engage in application-level innovation, this increase cannot offset the reduction in innovation by independent producers (Section B). Thus, the threat of discrimination reduces the amount of application-level innovation.

A. Incentives of Independent Producers of Complementary Products

In the absence of network neutrality regulation, the threat of discrimination reduces the amount of application-level innovation by independent producers of complementary products in three ways.

First, when the conditions for profitable exclusion outlined above are present in a particular complementary market, a network provider will discriminate against rivals in this market. As indicated above, dis-

165. *Cf.* POSNER, *supra* note 39, at 254 (discussing the profitability of monopoly preservation through exclusionary conduct in new economy markets, if the monopoly is of intellectual property).

crimination will reduce their profits.[166] A potential innovator bases its decision to innovate on the expected costs and benefits of realizing the innovation. Facing the threat of discrimination, potential innovators in affected markets will expect lower profits. Thus, the threat of discrimination reduces their incentives to innovate.

Second, the profitability of exclusion depends on a large number of factors that may not be common knowledge for all market participants. As a result, an economic actor with an idea for a complementary product may not be able to decide whether the network provider will have an incentive to exclude the final product from the market.[167] As a result, potential innovators face a significant uncertainty with respect to their future competitive environment. This uncertainty may reduce a developer's incentive to innovate, even if the factual conditions for profitable exclusion are not present.

Third, the above analysis suggests that independent producers of complementary products need not be concerned about exclusion, if the network provider does not currently offer a competing product. This seems to imply that innovation will only be harmed where the network provider is already vertically integrated into one or more complementary markets. Economic theory shows that this is not correct: Even if the network provider does not currently offer a competing product, it may be tempted to imitate the entrant, exclude the entrant from its network and exploit the complementary market itself, once the entrant starts to make significant profits.

Economic models show that in the presence of demand uncertainty in a complementary market, a primary good monopolist with a selling advantage in this market may have an incentive to let an independent producer enter the complementary market first to let him "test the waters."[168] If the level of demand turns out to be large enough once the demand uncertainty is resolved, the primary good monopolist enters the

166. The exclusionary conduct hurts independent producers of excluded complementary products in several ways: first, they are excluded from the part of the complementary market that consists of the network provider's Internet service customers. As a result, they are unable to make any sales in that market. In addition, due to economies of scale and, potentially, network effects in the production of their products, the exclusion from a part of the market may put them at a competitive disadvantage in the rest of the market as well. In the worst case, they may be forced to exit the complementary market completely. If they had made at least some sales to the network provider's Internet service customers in the absence of the exclusionary conduct, the exclusion will reduce their profits.

167. Similarly, the network provider may fail to assess the situation correctly and discriminate against or exclude an independent provider of complementary products, even if none of the conditions under which this conduct would be profitable apply. Farrell & Weiser, *supra* note 13, at 114-17 (calling this problem "incompetent incumbents" and include it in their list of exceptions to their version of the "one monopoly rent" argument).

168. David A. Miller, *Invention under Uncertainty and the Threat of Ex Post Entry*, (Aug. 24, 2006), http://ssrn.com/abstract=319180.

market as well and uses its selling advantage to make most of the sales. Foreseeing this course of events, the independent producer refrains from entering the market. As a result, nobody enters the complementary market; there is a region of foregone invention where privately and socially beneficial innovations are not realized.

For this situation to occur, three conditions must be realized: First, there must be demand uncertainty in the complementary market. Second, in the presence of demand uncertainty, entry to the complementary market is attractive for the independent producer, but not for the primary good monopolist (e.g., due to cost heterogeneity). Third, the primary good monopolist has a selling advantage in the complementary market.

In the Internet context, these conditions will often be met: First, in markets for new applications or content, there is usually a considerable demand uncertainty. Second, the economics and business strategy literature highlights a variety of reasons, why an incumbent network provider may not have an incentive to enter a complementary market for a new product in the presence of demand uncertainty, while an independent producer may have such an incentive. For example, start-ups often have lower entry costs than an incumbent due to the different cost structure of incumbents and new entrants.[169] In addition, while a small level of demand may meet the growth needs of a small company, a large incumbent will need much higher levels of demand to meet its growth needs.[170] Similarly, even if the level of demand is too uncertain for the network provider to justify innovation, users may find it attractive to innovate to meet their own application needs.[171] Third, the ability to technically exclude a rival producer of complementary products from its network provides the network provider with a huge selling advantage in the complementary market.

Thus, the number of markets in which independent developers' incentives to innovate are reduced will be larger than implied by the exceptions outlined above.

B. Incentives of Network Providers

As the previous section has shown, the threat of discrimination reduces independent producers' incentives to innovate in the markets for applications, content or portals. This reduction is only relevant, if it is not offset by a corresponding increase in network providers' incentives to

169. *E.g.*, CLAYTON M. CHRISTENSEN, THE INNOVATOR'S DILEMMA; WHEN NEW TECHNOLOGIES CAUSE GREAT FIRMS TO FAIL 132 (rev., updated ed., Harper Business 2000) (1997).

170. *E.g.*, *id.* at 128-30.

171. *See* van Schewick, *supra note* 32, at 329-42 (providing pointers to the relevant literature).

innovate in these markets. While network providers' incentives to innovate at this level do rise due to the increase in profit under discrimination, this increase in a few network providers' incentives to innovate cannot compensate for the reduction in innovation by independent producers.

There are three reasons for this: First, the reduction in potential innovators results in less diverse approaches to innovation, with negative consequences for the amount and quality of innovation. Second, with respect to particular innovations, economic actors other than the network providers may have an incentive to innovate, while the network providers may lack such an incentive. This further reduces the amount of innovation. Third, there are specific benefits associated with specific types of independent innovators which a network provider cannot replicate.

First, while there are a large number of (potential) independent producers of complementary products, there are only a few network providers. Thus, by reducing the innovation incentives of a large number of independent developers, the threat of discrimination ultimately reduces the number of innovators at the application-level. In the presence of technological uncertainty, market uncertainty or consumer heterogeneity, this reduction negatively affects the amount and potentially the quality of application-level innovation.

Human beings and, consequently, the firms for which they work have different experiences, capabilities and organizations, a fact that is stressed by research in evolutionary economics and management strategy. Due to these differences, economic actors may react very differently when exposed to the same situation. The impact of these differences rises with technological uncertainty, market uncertainty or consumer heterogeneity. Under these conditions, an increase in the number of potential innovators will result in a more diverse set of approaches to innovation,[172] and a more diverse set of approaches will be socially beneficial:[173] It guarantees a more complete search of the space of potential complementary products and decreases the probability that beneficial uses of the platform remain undetected. It increases the expected quality of the resulting products and may increase the amount of heterogeneous consumer needs that are served.

Second, research in economics and management strategy has identified systematic differences in the nature and direction of innovative activity between different types of innovators. In particular, due to differences in history, economic position and capabilities, the same innovation may be attractive to one type of innovator, but not to another. This re-

---

172. *See id.* at 299-305 with pointers to the relevant literature.
173. *See* the discussion in *id.* at 305-10 with pointers to the relevant literature.

search suggests that there are a large number of cases in which economic actors other than the network provider may have an incentive to realize an innovative idea, while the network provider may lack such an incentive. For example, this has been shown for new entrants to a complementary market, for venture-capital backed firms and for users.[174] When independent producers loose their incentive to innovate, this innovation will be lost.

Third, there are specific benefits associated with specific types of independent innovators which a network provider cannot replicate. For example, research has shown that the participation of firms backed by venture capitalists may increase the amount and the quality of innovation. Enabling users to innovate, may leave less customer needs unserved. In addition, users often make their innovation freely available to others; as a result, such innovations will reach a higher level of diffusion than a similar innovation of comparable quality that is produced by a network provider which sells the innovation to make a profit.[175]

In the context of the Internet, technological and market uncertainty as well as user heterogeneity are high,[176] suggesting that the reduction in independent producers' incentives to innovate will have the detrimental impact on application-level innovation outlined above.

III. IMPACT ON SOCIAL WELFARE

Network neutrality rules prevent network providers from discriminating against independent producers of complementary products or excluding them from their network. In the absence of network neutrality regulation, there is a real threat of discrimination (see Part II). Regulatory intervention to remove this threat is only justified, if the social benefits of regulatory intervention are larger than the costs.

As Part III has shown, network neutrality regulation increases the

---

174. *Id.* at 311-24 (new entrants), 324-29 (venture capital backed firms), 329-42 (users), based on a discussion of the relevant literature.

175. *Id.* at 337-42, based on a discussion of the relevant literature. While it is difficult to quantify these benefits, there are indications that they may be significant. For example, surveys indicate that today's standard commercial products may on average leave between 46% and 54% of customer needs unserved. *See* Nikolaus Franke & Eric von Hippel, *Satisfying Heterogeneous User Needs Via Innovation Toolkits: The Case of Apache Security Software*, 32 RES. POL'Y 1199, 1201-02 (2003).

176. Both network technology as well as technologies for the development of applications are still evolving, creating considerable technological uncertainty. A large number of useful applications are still waiting to be identified; in these areas, market uncertainty is high. The more people and businesses get connected to the Internet, the higher the heterogeneity of Internet users will become. Ultimately, the heterogeneity of Internet users will mirror the heterogeneity of society. As a result, the heterogeneity of user needs is bound to be increasing, not decreasing.

amount of application-level innovation. The increase is only relevant, if it is socially beneficial (Section A). On the cost side, network neutrality rules reduce network providers' incentives to innovate at the network level and to deploy network infrastructure (Section B.1). While regulatory intervention has its own costs, these are not covered in detail (Section B.2). When deciding whether to introduce network neutrality regulation, regulators must trade-off the benefits against the costs (Section C).

The analysis shows that the increase in application-level innovation is socially beneficial and that these benefits are more important than the costs.

A. Benefits

Network neutrality rules increase the amount of application-level innovation. This increase is only relevant to public policy, if it increases social welfare. This question can be approached in several ways.

First, one may ask whether the amount of innovation is generally lower than the social optimum. In this case, an increase in the amount of innovation would be socially beneficial.

In dealing with such questions, economists often note that the link between innovation and social welfare is theoretically ambiguous:[177] on the one hand, some economic models highlight the possibility that in their desire to capture the rents from innovation, firms may increase the level of investment in research and development above the socially efficient amount.[178] On the other hand, the existence of uncompensated spillovers and other factors such as the inability of innovators to perfectly appropriate the increase in consumer surplus lead to the theoretical prediction that firms will not be able to completely appropriate the social gains from innovation, leading them to invest less than the socially optimal amount in innovation.[179]

A closer look at the underlying models indicates that under conditions of uncertainty this ambiguity may disappear, leading to the insight that the amount of innovation is usually too low, which makes any increase in innovation socially beneficial. In models where firms invest more than the socially efficient amount in innovation, the wedge

177. *See, e.g.*, Jennifer F. Reinganum, *The Timing of Innovation: Research, Development, and Diffusion*, *in* 1 HANDBOOK OF INDUSTRIAL ORGANIZATION 849 (Richard L. Schmalensee & Robert D. Willig eds., 1st ed. 1989); JEAN TIROLE, THE THEORY OF INDUSTRIAL ORGANIZATION 399-400 (1988); Michael L. Katz, *Intellectual Property Rights and Antitrust Policy: Four Principles for a Complex World*, 1 J. ON TELECOMM. & HIGH TECH. L. 325, 337 sec.C (2002).

178. For an overview of this literature, see Reinganum, *supra* note 177. For a particular example of such a model, see Partha Dasgupta & Eric Maskin, *The Simple Economics of Research Portfolios*, 97 ECON. J. 581 (1987).

179. *See, e.g.*, TIROLE, *supra* note 177, at 399-400.

between private and social benefits from innovation results from the argument that society does not care which firm is ultimately successful, whereas each individual firm wants to be the winner.[180] Thus, these models are based on the implicit assumption that similar approaches by different firms constitute a wasteful duplication of efforts that should better be avoided. As indicated above, such an assumption neglects differences in firm heterogeneity. Once firm heterogeneity is taken into account, having different firms approach a particular problem will often be socially beneficial.

These theoretical insights are supported by empirical studies. They indicate that there is indeed too little innovation, because private firms are typically unable to appropriate all social gains from the innovation.[181]

Second, one may ask whether in the specific case under analysis there is likely to be less innovation than the socially optimal amount. Innovation in platform products[182] and complementary products is subject to two types of externalities that are likely to reduce the amount of innovation below the social optimum:[183] while the first operates vertically between the platform product and each complementary product, the second externality operates horizontally between different complementary products.

Due to the complementarity between the platform product and complementary products, innovation in complementary products usually increases demand for the platform product and vice versa. If the platform product and the complementary product are developed by different economic actors, the innovator in a complementary component does not appropriate the positive effect on the platform product, and vice versa.[184]

Innovation in one complementary product usually increases demand for the platform product, which may in turn positively affect demand for other complementary products. If different economic actors pursue innovation in the different components, each actor does not appropriate the positive effect on the other components. As a result, each actor's incentives to innovate will be lower than the social optimum.

A common solution to the problems caused by such externalities is

---

180. *See, e.g.*, Dasgupta & Maskin, *supra* note 178, at 584-85.

181. *See, e.g.*, Edwin Mansfield et al., *Social and Private Rates of Return from Industrial Innovations*, 91 Q. J. ECON. 221 (1977); Charles I. Jones & John C. Williams, *Measuring the Social Return to R&D*, 113 Q.J. ECON. 1119 (1998).

182. A platform product is a product that may be used with a large number of complementary products. *See, e.g.*, Douglas G. Lichtman, *Property Rights in Emerging Platform Technologies*, 29 J. LEGAL STUD. 615 (2000).

183. This observation is made in two different contexts by Timothy F. Bresnahan & Manuel Trajtenberg, *General Purpose Technologies 'Engines of Growth,?'* 65 J. ECONOMETRICS 83 (1995) and Lichtman, *supra* note 182.

184. *See, e.g.*, Bresnahan & Trajtenberg, *supra* note 183, at 94; Farrell & Katz, *supra* note 41, at 414 and appendix.

integration by all affected parties. The integrated entity internalizes the externalities and has therefore higher incentives to innovate.[185] In the current context, this is not a feasible solution: no single economic actor will be able to identify and realize all beneficial uses of the Internet.[186]

Finally, any assessment of the benefits of additional application-level innovation needs to take account of the character of the Internet as a general purpose technology.[187]

As a general purpose technology, the Internet has the potential to significantly increase economic growth.[188] General-purpose technologies offer a generic functionality that can potentially be applied in a large number of sectors within the economy. As the use of a general-purpose technology spreads throughout the economy and increases productivity in the sectors in which it is applied, the promises for economic growth that this technology holds materialize. At the same time, new applications trigger new advances in the general-purpose technology itself; these advances may in turn spawn the adoption of the general-purpose technology in additional sectors of the economy or may lead to new or improved applications in sectors that already use the technology. Thus, the adoption of general-purpose technologies exhibits increasing returns to scale, leading to potentially enormous increases in economic growth.[189]

As the positive impact of a general purpose technology stems primarily from the productivity increases resulting from its adoption in more and more sectors of the economy, the existence of a general-purpose technology is not sufficient to positively impact economic growth. Instead, the rate at which a general purpose technology affects

---

185. See Farrell & Katz, *supra* note 41(discussing some important refinements to this statement). As Farrell & Katz demonstrate, integration between two firms that each are the sole supplier of a component that is complementary with the other does not necessarily increase the incentives to invest in socially valuable research and development. *(See id.* at appendix). In addition, they show that integration between a monopoly supplier of one component with one of several suppliers of a complementary component may inefficiently lower independent suppliers' incentives to innovate.

186. *See, e.g.*, Timothy F. Bresnahan & Shane Greenstein, *The Economic Contribution of Information Technology: Towards Comparative and User Studies*, 11 J. EVOLUTIONARY ECON. 95, 98 (2001); Lichtman, *supra* note 182.

187. *See* van Schewick, *supra note* 32, at 346-49 (providing a detailed exposition of the argument in the text with pointers to the literature).

188. On general-purpose technologies, see, e.g., Bresnahan & Trajtenberg, *supra* note 183; Bresnahan & Greenstein, *supra* note 186; and the collection of papers in GENERAL PURPOSE TECHNOLOGIES AND ECONOMIC GROWTH (Elhanan Helpman ed., 1998) [hereinafter GENERAL PURPOSE TECHNOLOGIES]. On the Internet as a general-purpose technology, see, *e.g.*, Richard G. Harris, *The Internet as a GPT. Factor Market Implications, in* GENERAL PURPOSE TECHNOLOGIES, *supra*, at 145.

189. *E.g.*, Bresnahan & Trajtenberg, *supra* note 183; Elhanan Helpman & Manuel Trajtenberg, *A Time to Sow and a Time to Reap; Growth Based on General Purpose Technologies, in* GENERAL PURPOSE TECHNOLOGIES, *supra* note 188, at 55.

economic growth depends on the rate of co-invention[190] (i.e., the rate at which potential uses of the technology are identified and realized).

With respect to the Internet, this analysis implies that identifying potential uses for the Internet and developing the corresponding applications is the prerequisite for realizing the enormous growth potential inherent in the Internet as a general-purpose technology.[191]

As a result, measures that reduce the amount of application-level innovation have the potential to significantly harm social welfare by significantly limiting economic growth.

B. Costs

On the cost side, network neutrality rules reduce network providers' incentives to innovate at the network level and to deploy network infrastructure (Section 1). Regulatory intervention also creates its own costs (Section 2); however, these are not covered in detail.

1. Impact on Incentives at the Network Level

As highlighted in Part II, there is a variety of cases in which discrimination increases (or preserves) network providers' profits. As network neutrality regulation prevents network providers from realizing these profits, network neutrality regulation reduces their profits. Due to the complementarity between applications, content and portals on the one hand and Internet services on the other hand, this reduction in profits also affects network providers' incentives to innovate at the network level and to deploy network infrastructure.[192]

---

190. The term "co-invention" denotes the innovative activity associated with identifying and realizing potential uses of the general purpose-technology in particular sectors of the economy, *e.g.*, Bresnahan & Trajtenberg, *supra* note 183, at 86-88; Bresnahan & Greenstein, *supra* note 186, at 95-97.

191. *See* ROBERT E. LITAN & ALICE M. RIVLIN, BEYOND THE DOT.COMS; THE ECONOMIC PROMISE OF THE INTERNET 104-07 (2001) (making a similar observation).

192. *See* THIERER, *supra* note 7, at 17-19; OWEN & ROSSTON, *supra* note 4, at 24-25. *See also* Yoo, *Beyond Network Neutrality*, *supra* note 4, at 27-37, 48-53 (arguing that network neutrality may increase concentration in the market for last-mile broadband access, based on a broader definition of network neutrality that includes mandating interconnection, non-discrimination, rate regulation and the adoption of standardized protocol interfaces such as TCP/IP). As Yoo's argument is based on the negative impact of measures such as mandating the adoption of standardized interfaces, which are not included in the definition of network neutrality used here, his arguments do not apply to the analysis of this paper. For a discussion of the differences in the usage of the term network neutrality, see *supra* notes 9-11 and accompanying text. *But see* Frischmann & van Schewick, *supra* note 1 (offering a critical reply to Yoo's argument).

2. Costs of Regulation

The costs of network neutrality regulation depend on the chosen form of implementation. While the costs of network neutrality regulation are not the focus of this article, existing literature suggests that the costs of regulation itself will not be significant. In particular, they will be significantly lower than the costs associated with implementing and overseeing an open access regime.[193]

C. Trade-Off

The social benefits and costs outlined above suggest that the introduction of network neutrality regulation requires a trade-off: On the one hand, network neutrality regulation increases the amount of application-level innovation, which is critically important for economic growth. On the other hand, it decreases network providers' incentives to innovate at the network level and to deploy network infrastructure. The following section analyzes the two trade-offs in turn.

1. Application-Level Innovation vs. Innovation at the Network Level

Research on information-technology based general-purpose technologies suggests that increasing co-invention[194] is more important than increasing innovation in the general-purpose technology itself. Applied to the Internet, this implies that increasing application-level innovation is relatively more important than increasing innovation at the network level.

In information technology-based general-purpose technologies the incentives to invest in advancing the general-purpose technology itself seem to be higher than the incentives to invest in co-invention,[195] making it relatively more important to foster co-invention. This difference is attributed to two factors: first, the science and engineering base of hardware technologies is more developed than the science base of software engineering and of finding attractive business uses. Second, due to their generality, general-purpose technologies have larger markets than the individual applications; after all, while not all users of a general-

---

193. *See, e.g.*, Weiser, *supra* note 1, at 79-80.

194. The term "co-invention" denotes the innovative activity associated with identifying and realizing potential uses of the general purpose-technology in particular sectors of the economy. *E.g.*, Bresnahan & Trajtenberg, *supra* note 183, at 86-88; Bresnahan & Greenstein, *supra* note 186, at 95-97.

195. Timothy F. Bresnahan, The Changing Structure of Innovation in Computing: Sources of and Threats to the Dominant U.S. Position 10, (July 21, 1998), http://www.stanford.edu/~tbres/research/step.pdf.

purpose technology need all applications, all users need the general-purpose technology.

These factors are also present in the context of the Internet, making it reasonable to assume that the imbalance between incentives to innovate found in information-based general-purpose technologies in general also exists in the context of the Internet: Network engineering has a more developed science base than the identification of uses and software engineering. Due to the generality of the networking infrastructure, the market for network technology itself is larger than the market for individual applications.

Thus, compared to the incentives to innovate at the application-level, incentives to innovate at the network level are higher. At the same time, application-level innovation is the main determinant of economic growth. This suggests that increasing the amount of application-level innovation is relatively more important than increasing innovation at the network level.

2. Application-Level Innovation vs. Deployment of Network Infrastructure

As indicated above, network neutrality regulation reduces network providers' profits. This reduction in profits will also affect their incentive to deploy network infrastructure. This causal relationship, however, does not say anything about the degree to which these incentives are reduced.

Thus, in determining the appropriate trade-off between infrastructure deployment and application-level innovation, two questions must be answered: First, will the reduction in profits reduce the incentive to deploy infrastructure below the necessary level? Second, even if this is the case, is allowing network providers to discriminate the appropriate solution to this problem?

First, it is an open question, whether network neutrality regulation will reduce incentives to deploy network infrastructure below the necessary level. Not surprisingly, network providers and their industry organizations have claimed that this is the case. There are several reasons to doubt this assessment, though: Network neutrality regulation does not forbid network providers to vertically integrate into complementary markets;[196] it only bans them from using discrimination to increase their sales at the expense of rivals. Thus, it does not prevent network providers from making profit in complementary markets; it just takes away the additional profits that could be realized due to discrimination.[197] It also

196. *E.g.*, Wu, *Broadband Debate*, *supra* note 1, at 89.

197. Whether and, if yes, what form of price discrimination should be forbidden under network neutrality regulation, is still an open question. *See supra* note 9.

does not prevent them from making profit in the market for Internet services. As a result, the remaining profit may still be sufficient to motivate them to deploy the necessary infrastructure. Moreover, new wireless technologies may ameliorate the problem by further reducing the costs of broadband infrastructure. Thus, it still needs to be proven that the reduction in profits caused by network neutrality regulation suffices to reduce network providers' incentives to deploy infrastructure so severely that it becomes relevant for public policy.

Second, even if network providers' incentives are too low to guarantee the necessary deployment of broadband infrastructure under network neutrality regulation, this does not necessarily imply that network providers should be allowed to discriminate.[198] As Michael Katz has put it, "In the antitrust – if not regulatory – context [. . .] U.S. policy rejects the notion that the otherwise illegal maintenance or acquisition of monopoly power in a market can be justified by 'good' use of the monopoly profits in that market or another one."[199] Following this line of reasoning, instead of allowing discrimination, regulators should contemplate other ways of ensuring a sufficient deployment of network infrastructure, if necessary. For example, in light of the severe consequences of stifling application-level innovation for economic growth, subsidizing the deployment of broadband infrastructure may be preferable to allowing network providers to discriminate.

Thus, in trading off application-level innovation against infrastructure deployment, it seems reasonable to opt for fostering application-level innovation in order to realize the enormous growth potential inherent in Internet technology, and to contemplate other ways of ensuring a sufficient deployment of network infrastructure, if necessary.

---

198. *But see* Yoo, *Beyond Network Neutrality*, *supra* note 4, Part II. Yoo argues that by focusing on competition in the market for applications and content, network neutrality proponents are focusing on the wrong policy problem. According to him, policy makers should focus on increasing competition in the market for last-mile broadband access, which is less competitive than the markets for applications and content, *id.* at Part II. In line with this assumption, he mainly rejects network neutrality proposals based on their negative impact on competition in last-mile broadband access. *Id.* at 27-37, 48-53. Apart from neglecting the different impact of innovation in these markets on economic growth, and accompanying text, this analysis fails to take account of the possibility to stimulate competition in the market for last-mile broadband access through other means. *See supra* notes 187-191, 194-195. In addition, his arguments about the negative impact of network neutrality on competition in the market for last-mile broadband access are based on a much broader definition of network neutrality than the one advocated by network neutrality proponents and used in this paper; as a result, his analysis does not carry over to the case of "pure" non-discrimination rules discussed here. *See also supra* notes 9-11 and accompanying text; *see supra* note 192 and accompanying text. For a critical appraisal of Yoo's work on network neutrality, see Frischmann & van Schewick, *supra* note 1; *see* Herman, *supra* note 1.

199. Katz, *supra* note 177, at 340.

CONCLUSION

This paper advances the debate over network neutrality by providing an economic framework within which calls for network neutrality regulation can be analyzed.

The analysis shows that calls for network neutrality regulation are justified: In the absence of network neutrality regulation, there is a real threat that network providers will discriminate against independent producers of applications, content or portals or exclude them from their network. This threat reduces the amount of innovation in the markets for applications, content and portals at significant costs to society.

While network neutrality rules remove this threat, they are not without costs: Apart from creating the costs of regulation itself, network neutrality rules reduce network providers' incentives to innovate at the network level and to deploy network infrastructure. Thus, regulators face a trade-off. As the paper shows, due to the potentially enormous benefits of application-level innovation for economic growth, increasing the amount of application-level innovation through network neutrality regulation is more important than the costs associated with it.

Before network neutrality regulation can be drafted, however, more research is needed. In particular, the open questions surrounding the scope of network neutrality regulation need to be resolved. In addition, the best way of implementing network neutrality rules still needs to be identified.[200]

The paper also contributes to the debate over "open access" and "facilities-based competition." As has been set out above, the proposals for "facilities-based competition" and "open access" are based on the assumption that competition in the market for Internet services will mitigate a network operator's ability and incentive to discriminate against or exclude independent portals, content and applications. The analysis has highlighted a variety of circumstances under which a network provider may have the ability and incentive to discriminate against unaffiliated producers of complementary products or exclude them from its network, even if it faces competition in the market for Internet services. Thus, neither increased facilities-based competition nor open access regulation are the appropriate tools to mitigate the threat of discrimination.

Finally, the paper shows that our intuitions regarding the profitability of exclusionary conduct that have been shaped by antitrust analysis of markets for conventional goods may be misleading in markets such as the markets for applications, portals and content that are characterized by high fixed costs, low marginal costs and, potentially, network effects, in

200. For an overview of open issues in these areas, see *supra* note 9 and accompanying text.

particular if the exclusionary conduct is based technological means.

From: Abraham, Magid
To: Josh Gottheimer
Sent: Wed Dec 01 22:10:20 2010
Subject: RE: Chairman Genachowski Remarks on Open Internet

Josh:
Thanks for sharing. This is a very sensible and balanced framework that will preserve and promote an open and vibrant Internet by preserving the consumer's right to get the content they want from any source, while respecting the access providers' needs to properly manage their network capacity and utilization.
Feel free to quote me if you like.
Best,
Magid

From: Josh Gottheimer
Sent: Wednesday, December 01, 2010 1:31 PM
To: Josh Gottheimer
Subject: Chairman Genachowski Remarks on Open Internet

Hi – attached and pasted below are the remarks Chairman Genachowski delivered this morning on preserving Internet freedom and openness. Please let me know if you have any questions.
Best, Josh

FEDERAL COMMUNICATIONS COMMISSION
CHAIRMAN JULIUS GENACHOWSKI
REMARKS ON PRESERVING INTERNET FREEDOM AND OPENNESS
WASHINGTON, DC
December 1, 2010

Good morning. After months of hard work at the FCC, in other parts of government, in the private sector, and in the public interest community, and after receiving more than 100,000 comments from citizens across America, we have reached an important milestone in our effort to protect Internet freedom and openness.

Yesterday, I circulated to my colleagues draft rules of the road to preserve the freedom and openness of the Internet. This framework, if adopted later this month, would advance a set of core goals: It would ensure that the Internet remains a powerful platform for innovation and job creation; it would empower consumers and entrepreneurs; it would protect free expression; it would increase certainty in the marketplace, and spur investment both at the edge and in the core of our broadband networks.

I am gratified by the broad support this proposal has already received this morning -- including from leading Internet and technology companies, founders and investors; consumer and public interest groups, unions, civil rights organizations, and broadband providers.

The proposed rules of the road are rooted in ideas first articulated by Republican Chairmen Michael Powell and Kevin Martin, and endorsed in a unanimous FCC policy statement in 2005. Similar proposals have been supported in Congress on a bipartisan basis. And they are consistent with President Obama's commitment to "keep the Internet as it should be – open and free."

Their adoption would culminate recent efforts to find common ground -- at the FCC, in Congress, and outside government, including approaches advanced by both Democrats and Republicans, and by stakeholders of differing perspectives. In particular, this proposal would build upon the strong and balanced framework developed by Chairman Henry Waxman, which garnered support from technology and telecommunications companies, big and small, as well as from consumer and public interest groups.

The animating force behind all of these efforts is a shared appreciation for the Internet's wondrous contributions to our economy and our way of life.

Millions of us depend on the Internet every day: at home, at work, in school -- and everywhere in between. The high-speed networks we call broadband are transforming health care, education, and energy usage for the better. It's hard to imagine life today without the Internet -- any more than we can imagine life without running water or electricity.

The Internet has been an unprecedented platform for speech and democratic engagement, and a place where the American spirit of innovation has flourished. We've seen new media tools like Twitter and YouTube used by democratic movements around the world.

Not only is the Internet becoming a central part of the daily lives of Americans, the Internet has been a strong engine of job creation and economic growth.

Internet companies have started as small start-ups, some of them famously in dorm rooms and garages with little more than a computer and access to the open Internet. Many have become large businesses, providing high-paying, high-tech jobs in communities across our country. It's the American dream at work.

Small businesses and start-ups have accounted for more than 22 million new American jobs over the last 15 years. And broadband has played a central part, enabling small businesses to start, to lower their costs, and to reach new customers in new markets around the country and, indeed, the globe.

Why has the Internet proved to be such a powerful engine for innovation, creativity and economic growth? A big part of the answer traces back to one key decision by the Internet's original architects: to make the Internet an open platform.

It is the Internet's openness and freedom -- the ability to speak, innovate, and engage in commerce without having to ask anyone's permission -- that has enabled the Internet's unparalleled success.

This openness is a quality -- a generative power -- that must be preserved and protected. And the record in the proceeding we've run over the past year, as well as history, shows that there are real risks to the Internet's continued freedom and openness. Broadband providers have natural business incentives to leverage their position as gatekeepers to the Internet. Even after the Commission announced open Internet principles in 2005, we have seen clear deviations from the Internet's openness -- instances when broadband providers have prevented consumers from using the applications of their choice without disclosing what they were doing.

The proposed open Internet framework is designed to guard against these risks, while recognizing the legitimate needs and interests of broadband providers.

In key respects, the interests of edge innovators – the entrepreneurs creating Internet content, applications and services -- broadband providers, and American consumers are aligned.
Innovation at the edge catalyzes consumer demand for broadband. Consumer demand spurs private investment in faster broadband networks. And faster networks spark ever-cooler innovation at the edge.

A central goal of the proposed open Internet framework is to foster this cycle of massive investment in both the edge and the core of broadband networks, to the benefit of consumers and our economy.

Protecting Internet freedom will drive the Internet job creation engine.

The crux of the proposal, which would establish open Internet rules for the first time, is straightforward:

First, consumers and innovators have a right to know basic information about broadband service, like how networks are being managed. The proposed framework therefore starts with a meaningful transparency obligation, so that consumers and innovators have the information they need to make smart choices about subscribing to or using a broadband network, or how to develop the next killer app. Sunshine can help solve problems early, reducing the number of issues that even come to the FCC.

Second, consumers and innovators have a right to send and receive lawful Internet traffic -- to go where they want and say what they want online, and to use the devices of their choice. Thus, the proposed framework would prohibit the blocking of lawful content, apps, services, and the connection of non-harmful devices to the network.

Third, consumers and innovators have a right to a level playing field. No central authority, public or private, should have the power to pick which ideas or companies win or lose on the Internet; that's the role of the market and the marketplace of ideas. And so the proposed framework includes a bar on unreasonable discrimination in transmitting lawful network traffic.

The proposed framework also recognizes that broadband providers must have the ability and investment incentives to build out and run their networks. Universal high-speed Internet access is a vital national goal that will require very substantial private sector investment in our 21st century digital infrastructure. For our global competitiveness, and to harness the opportunities of broadband for all Americans, we want world-leading broadband networks in the United States that are both the freest and the fastest in the world.

To this end, broadband providers need meaningful flexibility to manage their networks -- for example, to deal with traffic that's harmful to the network or unwanted by users, and to address the effects of congestion. Reasonable network management is an important part of the proposal, recognizing that what is reasonable will take account of the network technology and architecture involved.

Our work has also demonstrated the importance of business innovation to promote network investment and efficient use of networks, including measures to match price to cost such as usage-based pricing.

The record in our proceeding reflects both the importance of openness principles to mobile broadband, and the appropriateness of recognizing differences between fixed and mobile broadband. This is not a new point, but one that I've made consistently since the beginning of this proceeding. For example, mobile broadband is at an earlier stage of development than fixed broadband, and is evolving rapidly.

Accordingly, the proposal takes important but measured steps in this area -- including transparency and a basic no blocking rule. Under the framework, the FCC would closely monitor the development of the mobile broadband market and be prepared to step in to further address anti-competitive or anti-consumer conduct as appropriate.

The work of the FCC staff on this proceeding has been exceptional, no more so than in connection with the complex legal issues. Informed by the staff's additional legal analysis and the extensive comments on this issue over the past year, the proposal is grounded in a variety of provisions of the communications laws, but would not reclassify broadband as a Title II telecommunications service. I am satisfied that we have a sound legal basis for this approach.

I want to emphasize that moving this item to a vote at the Commission is not designed or intended to preclude action by Congress. As always, I welcome the opportunity for the Commission to serve as a resource to Congress.

The Commission itself has a duty and an obligation to fulfill -- a duty to address important open proceedings based on the record, and an obligation to be a cop on the beat to protect broadband consumers and foster innovation, investment, and competition. I believe the proposed framework advances this mission, and that its adoption will provide increased certainty and benefits to the American public.

I look forward to ongoing work with my Commission colleagues on this and other issues. We have very important work to do for the American people in the months ahead, as we strive to harness the opportunities of broadband and communications for the benefit of our economy and for all Americans.

Thank you.

ARTICLE

Studios lick their lips over new-look Netflix

By Paul Bond
Mon Aug 16, 2010 12:38am EDT

LOS ANGELES (Hollywood Reporter) - Perhaps it's fitting that while the stock market was tanking last week, shares of Netflix were touching all-time highs. What else to expect from a company that is remaking Hollywood's home-distribution model?

That seems to be the popular consensus, anyway, since it was revealed last week that Netflix will pay almost $1 billion to Epix over five years for online-streaming rights for movies from Paramount, Lionsgate and MGM.

Before that, online streaming might have been considered more a necessary annoyance than an actual business with revenue and profit. But with the stroke of a pen, Netflix has upped the ante so significantly that movie studios have no doubt begun to ratchet up expectations for licensing fees they can charge going forward.

"It certainly proves that on-demand streaming rights have value," said Jim Packer, co-president of MGM Worldwide Television.

"We're the only ones that have put real money in play," Netflix chief content officer Ted Sarandos said.

In fact, Netflix probably paid more money for digital rights than traditional pay TV has been willing to pay for television rights in recent years. That's a sea change.

"For studios, the fear was the Internet was going to turn dollars into pennies," Sarandos said. "But there's a real economic model online. We didn't get any big discount because we're on the Internet. We're showing that everything that happens online is not going to be stolen or given away."

Richard Greenfield, an analyst with BTIG Research, figures that when Paramount, Lionsgate and MGM ditched Showtime to create Epix, Showtime was offering $175 million a year for their combined content, whereas Netflix values that same content, delivered over the Internet, at $200 million a year.

Netflix had already signaled a willingness to outbid more traditional outlets in July when it agreed to pay Relativity Media about $100 million a year for rights to movies that otherwise might have gone to HBO or Showtime.

"We're definitely another buyer in this market," Sarandos said.

That's great for studios, especially at a time when DVD revenue is dwindling, but it's bad for the likes of Showtime and HBO. Aristotle Munarriz of the Motley Fool went so far as to predict HBO is "toast" because of the Epix-Netflix deal.

After all, through video game consoles, Blu-ray players, TiVo and other boxes, millions of Netflix users are streaming movies onto their TV screens where they would otherwise be watching HBO. And it's costing Netflix users only $8.99 a month for unlimited access to thousands of titles, whereas cablers charge up to $15 a month for a premium movie channel.

Sarandos notes that the movies for which Netflix has streaming rights have accounted for 46% of the year's box office so far this year, one percentage point more than HBO, which has films from Universal, Fox, Warners and DreamWorks.

Nevertheless, Sarandos bristles at the suggestion Netflix is a competitor to Showtime and HBO and, in fact, said he wouldn't mind striking Epix-style licensing deals with them. He noted Netflix has no plans to follow HBO and Showtime into original programing and that Netflix won't stream movies until 90 days after Epix first shows them on television, so they have created a window separate from that of pay TV.

Munarriz predicts Netflix will eventually reserve streaming for its higher-tiered subscribers. Though Sarandos says there's no plans for that now, it makes sense that Netflix would want to pass along the cost of big licensing deals to its customers.

1 2 NEXT

TECHNOLOGY MEDIA

Comments (0)

This discussion is now closed. We welcome comments on our articles for a limited period after their publication.

Related News

Studios, cable closer to home movie deals: sources
Wed, Aug 11 2010

UPDATE 4-Netflix, Epix strike programming deal
Tue, Aug 10 2010

Hollywood glad kids willing to pay for content
Thu, Aug 5 2010

UPDATE 3-ITV moves into pay-TV with new strategy
Tue, Aug 3 2010

UPDATE 2-Comcast, CBS reach 10-year programming agreement
Mon, Aug 2 2010

Related Topics

Technology »
Media »

READ

1 WikiLeaks rival Openleaks "coming soon": website
8:48am EST

2 UK mosque denounced Stockholm bomber for militancy
9:59am EST

3 Murdered Alabama children were tortured: documents
09 Dec 2010

4 "Black Swan" leads Critics' Choice award nominees
6:05am EST

5 Infants all released unharmed from French nursery
7:01am EST

DISCUSSED

168 Obama unveils tax deal with Republicans

87 WikiLeaks founder Assange arrested in Britain

69 Special Report: STD fears sparked case against WikiLeaks boss

WATCHED

Row over vote orgasm video
Fri, Nov 19 2010

Bejeweled bra exposed in NY
Thu, Oct 21 2010

Bill Clinton backs Obama tax plan
Fri, Dec 10 2010

Login or **register**

Latest from My Topics

MORE FROM REUTERS

Judge rejects key part of Obama healthcare law

WASHINGTON (Reuters) - A U.S. judge in Virginia on Monday declared unconstitutional a key part of President Barack Obama's landmark healthcare law in the first major setback on an issue that will likely end up at the Supreme Court.

CONTINUE READING

HAPPENING NOW

Swedish bombings probed

A Middle East-born man who died in a blast in Stockholm was wearing an explosives belt and likely aimed to attack a crowded train station or department store when the device went off prematurely, an official said. **Video | Full Article**

Q&A: A "lone wolf" attack?

Analysis: WikiLeaks raises Yemen risk

WORLD

TOP VIDEOS

Berlusconi vote risk for Italy

Fast train links Finland and Russia

Chinese fakes plague US retailers

» More Top Videos

MARKETS

US Indices

DOW
49.20
11,459.52
+0.43%

NASDAQ
2.91
2,640.45
+0.11%

S&P 500
4.78
1,245.18
+0.39%

TR US INDEX
0.43
113.81
+0.38%

Int'l Indices

NIKKEI
10,293.89

HANG SENG
23,317.61

» Markets

TOP NEWS

Tax deal seen passing despite cost concerns

GE to buy UK oil pipemaker Wellstream for $1.3 billion

WikiLeaks rival Openleaks "coming soon": website

BofA looks to sell toxic mortgages worth $1 billion: report

Exclusive: China extends special RRR hike for top banks

Midwest AND mid-Atlantic face icy temperatures | VIDEO

WikiLeaks founder Assange says Pentagon plans prosecution

» More Top News

ANALYSIS & OPINION

Is Obama endangering Social Security?
Felix Salmon

Should DSK run for president of France?
Katharine Herrup

Sarah Palin's huge 2012 move
James Pethokoukis

» More Analysis & Opinion

TODAY IN PICTURES

Editor's choice

A selection of our best photos from the past 24 hours.

View Slideshow

REUTERS

Editorial Editions:

Click here to download plugin.

REUTERS
- Contact Us
- Advertise With Us
- Help
- Journalism Handbook
- Archive
- Site Index
- Video Index
- Reader Feedback
- Reuters on Facebook
- Mobile
- Newsletters
- RSS
- Podcasts
- Widgets
- Your View
- Analyst Research

THOMSON REUTERS
- Copyright
- Disclaimer
- Privacy
- Professional Products
- Professional Products Support
- Financial Products
- About Thomson Reuters
- Careers

ONLINE PRODUCTS
- Acquisitions Monthly
- Buyouts
- Venture Capital Journal
- International Financing Review
- Project Finance International
- PEhub.com
- PE Week
- FindLaw

Login or **register**

Latest from My Topics

Business

Posted on Monday, 10.18.10

Showtime for Univisión

A new studio and fresh programming take center stage as the company scripts its future.

The long-running talk show of Cristina Saralegui was canceled. PEDRO PORTAL / STAFF

- Photos

- Friday Business Report: Media, Artsy Souls and Disaster Cleanup

Related Content

- Twists and turns mark Univisión's plot line
- Network pulled plug, Cristina says of show

Univisión at a glance

• Operates through television, radio, Internet

• Operates three networks: Broadcast networks Univisión and TeleFutura and cable network Galavisión

• Owns 60 Univisión and TeleFutura stations

• There are 80 affiliate-owned Univisión and TeleFutura stations

• Has 68 radio stations in 16 of the top 25 U.S. Hispanic markets and 5 stations in Puerto Rico

• Runs 18 national and 72 local online and mobile websites, which got more than 30 million visits in September

• Reaches 80 percent of the U.S. Hispanic population across all its platforms

• Cesar Conde was appointed president of

BY BRIDGET CAREY AND GLENN GARVIN
BCAREY@MIAMIHERALD.COM

When Nielsen Media Research sent out its weekly television ratings report for the first week in September, TV programmers and marketing directors across America blinked, rubbed their eyes, then blinked again. Could it really be? Spanish-language Univisión finishing first? Not first in the Spanish rankings, not first in Miami or Los Angeles: First in the whole United States, ahead of Fox, CBS, ABC and NBC.

It could be. It was. And, says Univisión Networks President Cesar Conde, it will be again.

``Our goal is to be the No. 1 network in this country regardless of language,'' says Conde, from his office in Doral. ``[If] we continue to perform the way we have -- and you have to have some macroeconomic trends continue -- that is feasible, that Univisión, a Spanish-language network, can be the No. 1 network in this country regardless of language, within the next five years.''

It's time to insert a couple of asterisks here. Univisión was No. 1 not in total viewers but in the 18-to-49 age demographic that TV advertisers covet. And its victory came during a lull in English-language programming, a week between the end of the summer broadcast season and the beginning of fall's roll-out of new programs when the schedule consists almost entirely of reruns and cheap reality shows.

But it also capped a year-long show of Nielsen strength when Univisión frequently finished among the top two or three broadcast networks in overall viewers, a performance that threatened to erase the traditional distinction between Spanish-language stations and what the industry refers to as the ``general market,'' the four big broadcast networks.

``They're competing head to head with English-language networks,'' says Miami Hispanic-media consultant Adam Jacobsen. ``Forget the Spanish part. They're America's fifth network, period.''

MORE BUSINESS

- **Back in the holiday spirit: Caterers see resurgence in office parties**
- **Caviar is out; comfort food is in**
- **HIS MISSION: FAIR TRADE**
- **85 years later, Gunster still 'focused on Florida'**
- **6 must-have, brand-boosting apps for small firms**
- **Survival of the arts is in our hands**

More

VIDEOS

Univisión Networks one year ago, and has made Fortune Magazine's 40 Under 40 list for two consecutive years.

• Univisión's 2010 World Cup Final averaged 8.8 million total viewers, the most-watched final in the network's history

Ratings watch

Univisión has increased its share of viewers in the key 18-to-34 and 18-to-49 age demographics more than any other broadcast network except CBS over the past year.

In adults, ages 18-34, the primetime average audience percent change since last year:

CBS +13% Univisión +9% NBC +2% CW -4% ABC -11% Fox -15% In adults, ages 18-49

Univisión +10% CBS +7% NBC +5% CW +4% ABC -13% Fox -14% SOURCES: Univisión, Nielsen

Just four years ago, when Los Angeles billionaire Jerry Perenchio sold Univisión to a consortium of five private equity firms for $12.3 billion, the network's future seemed uncertain. The stream of hit telenovelas that allowed it to dominate the Spanish-language market was about to dry up after a bitter and expensive spat with its supplier, Mexican media conglomerate Televisa.

The challenge from Spanish competitor Telemundo, powered by an influx of cash and expertise from new corporate master NBC Universal, was growing fiercer. And the eccentric and secretive policies the company had inherited from Perenchio seemed increasing out of step with the modern media environment. ``You weren't supposed to talk about *anything*,'' recalls one former Univisión executive. ``If your name turned up in a newspaper, you were dead meat.''

Since then, Univisión has emerged as a force to be reckoned with in that media environment. Moving beyond its immigrant roots, the network began targeting younger second- and third-generation U.S. Hispanics by broadening its news and public affairs programming beyond its traditional focus on Latin America, as well as experimenting with reality shows and other entertainment beyond the traditional telenovela format.

Result: The network has increased its share of viewers in the key 18-to-34 and 18-to-49 age demographics more than any other broadcast network except CBS over the past year. Univisión, once content to let cable systems air the signals of its affiliate stations for free, since 2007 has been asking fees as high as $1 per subscriber, a practice that brought in more than $175 million last year.

Full Story 1 | 2 | 3 | 4 | Next »

Join the discussion

The Miami Herald is pleased to provide this opportunity to share information, experiences and observations about what's in the news. Some of the comments may be reprinted elsewhere in the site or in the newspaper. We encourage lively, open debate on the issues of the day, and ask that you refrain from profanity, hate speech, personal comments and remarks that are off point. Thank you for taking the time to offer your thoughts.

We have introduced a new commenting system called Disqus for our articles. This allows readers the option of signing in using their Facebook, Twitter, Disqus or existing MiamiHerald.com username and password.

Read more about the new commenting system on MiamiHerald.com.

0 Comments

ADD NEW COMMENT

Required: Please login below to comment.

The Miami Herald

Post as ...

SHOWING 0 COMMENTS

Sort by Oldest first Subscribe by email Subscribe by RSS

Real-time updating is **enabled**. (Pause)

Recommendations

Sign Up Create an account or **log in** to see what your friends are recommending.

Kentucky Republicans propose Arizona-style immigration law - Politics AP
45 people shared this.

Path to the future: Pedestrians, cyclists get urban greenway - Miami-Dade
32 people shared this.

Back in the holiday spirit: Caterers see resurgence in office parties - Business Monday - MiamiHeral
15 people shared this.

Facebook social plugin

QUICK JOB SEARCH

Enter Keyword(s)

Enter City

Select a State - All United States -

Select a Category - All Job Categories -

Search

Advanced Job Search
Search by Category
powered by careerbuilder.com

Caution Concerning Forward-Looking Statements

Caution Concerning Forward-Looking Statements

The information on this web site may contain "forward-looking statements" within the meaning of the Private Securities Litigation Reform Act of 1995, particularly statements anticipating future growth in revenues, Operating Income (Loss) before Depreciation and Amortization, cash provided by operating activities and other financial measures. Words such as "anticipates," "estimates," "expects," "projects," "intends," "plans," "believes" and words and terms of similar substance used in connection with any discussion of future operating or financial performance identify forward-looking statements. These forward-looking statements are based on management's current expectations and beliefs about future events. As with any projection or forecast, they are susceptible to uncertainty and changes in circumstances.

The Company operates in a highly competitive, consumer and technology driven and rapidly changing business that is affected by government regulation and economic, strategic, political and social conditions. Various factors could adversely affect the operations, business or financial results of TWC in the future and cause TWC's actual results to differ materially from those contained in the forward-looking statements, including those factors discussed in detail in Item 1A, "Risk Factors," in TWC's Annual Report on Form 10-K for the year ended December 31, 2009, and in TWC's subsequent filings made from time to time with the SEC. In addition, important factors that could cause the Company's actual results to differ materially from those in its forward-looking statements include:

- increased competition from video, high-speed data and voice providers, particularly direct broadcast satellite operators, incumbent local telephone companies, companies that deliver programming over broadband Internet connections, and wireless broadband and phone providers;
- the Company's ability to deal effectively with the current economic slowdown or further deterioration in the economy, which may negatively impact customers' demand for the Company's services and also result in a reduction in the Company's advertising revenues;
- the Company's continued ability to exploit new and existing technologies that appeal to residential and commercial customers;
- changes in the regulatory and tax environments in which the Company operates, including, among others, regulation of broadband Internet services under Title II of the Communications Act of 1934, as amended, "net neutrality" legislation or regulation and federal, state and local taxation;
- increased difficulty negotiating programming and retransmission agreements on favorable terms, resulting in increased costs to the Company and/or the loss of popular programming; and
- changes in the Company's plans, initiatives and strategies.

Any forward-looking statements made by the Company on this web site speak only as of the date on which they were made. The Company is under no obligation to, and expressly disclaims any obligation to, update or alter its forward looking statements whether as a result of changes in circumstances, new information, subsequent events or otherwise.

Last updated: August 16, 2010

Close Window

Chicago Tribune

September 30, 2010 Thursday
Chicagoland Final Edition

HOW TO DROP THE BOX;
(AND SURVIVE)

BYLINE: By Stephen Cavendish, Tribune Newspapers

SECTION: LIVE! ; ZONE C; Pg. 1

LENGTH: 874 words

So, you want to ditch your cable box. ...

Viva la 21st century! Everything is on the Interwebs, right?

Before you switch, however, you need to do a little self-evaluating. You, as an average American, will spend about four hours a day watching television. How important is it in your life? Do you need to be able to talk about the latest episode of "Mad Men" around the water cooler the morning after it airs, or is waiting a few days OK? Is it important to watch in high-definition, or is standard just fine? Do you need the warm glow of a 50-inch screen, or will a laptop satisfy you? Are you a news or sports junkie?

As much as this is an entertainment decision, getting rid of cable is also a lifestyle choice that will involve rearranging a few things in your life, and maybe investing in a little bit of hardware. Here are two paths to cable independence:

If all you care about is the programming, not when or how you see it, fire up the laptop. You're a realist who's looking to save money on your cable bill. Here are three things you'll need to start:

YOUR PC AS YOUR TV

Hulu Plus: A joint venture between Fox, NBC and ABC, Hulu.com is a free, commercial-based Web site that streams the five most recent episodes of most broadcast network shows, except CBS and The CW. Hulu Plus, launched earlier this year, allows access to Hulu's complete catalog (more than 2,600 titles strong), including previous seasons of existing shows and some that are no longer on air. Want to watch "Lost" from the beginning? Big fan of "Arrested Development"? It has them. One caveat: New shows on Hulu and Hulu Plus are available the day after they air on TV. Cost: $9.99 a month

Netflix streaming: Netflix originally launched as a mail-based DVD service to compete with Blockbuster and other brick-and-mortar video outlets, but has hit big with its online-based streaming service. For a monthly fee, you have instant access to about 20 percent of the Netflix catalog, mostly in the form of new-release movies and TV shows available on DVD (the catch being that you have to wait for the shows' release on DVD). Last week Netflix signed a deal to stream NBC's shows on a next-day basis. Cost: $8.99 a month

iTunes: You might already have iTunes for music, but there's a lot of video available. For $1.99 an episode, you can cherry-pick the entertainment options not covered by

Hulu and Netflix (CBS, niche networks like TLC, History Channel and the Travel Channel), and most TV is available the day after it airs. Cost: Download player for free.

What you?re not getting

First, nothing is available live. Tuesday night, you're not watching "Glee" without an antenna. Second, with streaming services the quality can vary based on your download speeds, so a mediocre broadband connection can mean a subpar picture. Heavy sports- and news-channel watchers are going to have to get creative. Soap operas? Syndicated daily talk shows? Sorry, you'll need cable for that.

YOUR TV WITH HELP

Let's be honest: Sitting on your couch with a big flat-screen HDTV is a superior way to experience television. Here are some hardware options to help you do that:

Apple TV: Apple CEO Steve Jobs announced an intriguing revamp to Apple TV earlier this month: For $99, you get a small box that hooks your TV into the iTunes store via your broadband connection. But unlike previous incarnations, TV shows would now be available to rent for 99 cents. Jobs promised Fox and ABC shows to start and is hopeful that other studios will follow. If you pair that with the movie offerings already available, that's an awful lot of content. Plus, Apple has said it will have access to Netflix.

Roku: For $59, the Roku player is cheaper and connects your TV to a number of different subscription services, including Netflix and Amazon Video on Demand. Unlike other services, Roku also has sports options (via subscription) such as MLB.tv and the UFC.

Your game system: That Wii, PlayStation 3 or Xbox 360 already attached to your TV? It already has Netflix access. Just connect it to the Internet.

What you're not getting

Again, you're missing out on news and sports channels. Heavy TV watchers, particularly of movie channels, should price out their current viewing habits, as even cheap rentals can add up. And though a lot of prime-time content is out there, daytime fare is hard to come by.

WHAT ABOUT SPORTS?

Sports content is where things get a little tricky. Your cable and satellite providers know that sports programming is a big plus for them. On a Saturday afternoon, even a half-decent cable sports package gives you tons of college football across the country. Meanwhile, regional sports networks hold the rights to most MLB, NBA and NHL games. And DirecTV?s deal to show all NFL games is worth billions. Ditching your box means paying up.

MLB: $120 a season.

NHL: $169 a season.

NBA: $149.95 a season.

NFL: $350 a season, available only to those who can't or aren't allowed to set up a DirecTV dish.

ESPN3: Free with your broadband subscription (depending on your Internet service provider), ESPN3 streams the majority of the Worldwide Leader's live games. It offers a substantial amount of coverage, particularly of college sports, and includes a wide array of international as well as domestic sports.

LOAD-DATE: September 30, 2010

LANGUAGE: ENGLISH

GRAPHIC: Photo (color): (Cable box)\ Photo (color): (hulu logo)\ Photo (color): (Netflix logo)\ Photo (color): iITunes logo)\ Photo (color): (Apple TV)\ Photo (color): (Ruku)\ Photo (color): (Your game system)\ Photo (color): (Football)
Photo(s) Graphic(s)

PUBLICATION-TYPE: Newspaper

10-K 1 a2196834z10-k.htm FORM 10-K

Use these links to rapidly review the document
TABLE OF CONTENTS

Table of Contents

UNITED STATES
SECURITIES AND EXCHANGE COMMISSION

Washington, D.C. 20549

FORM 10-K

(Mark One)

ý **ANNUAL REPORT PURSUANT TO SECTION 13 or 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934**

For the fiscal year ended December 31, 2009

OR

o **TRANSITION REPORT PURSUANT TO SECTION 13 or 15(d) OF THE SECURITIES EXCHANGE ACT OF 1934**

For the transition period from to

Commission file number 333-106529

DIRECTV HOLDINGS LLC
DIRECTV FINANCING CO., INC.

(Exact name of registrant as specified in its charter)

DIRECTV Holdings LLC—Delaware **DIRECTV Financing Co., Inc.—Delaware** (State or other jurisdiction of incorporation or organization)	**25-1902628** **59-3772785** (I.R.S. Employer Identification No.)
2230 East Imperial Highway, El Segundo, California (Address of Principal Executive Offices)	**90245** (Zip Code)

Registrant's telephone number, including area code: (310) 964-5000
Securities registered pursuant to Section 12(b) of the Act: None
Securities registered pursuant to Section 12(g) of the Act: None

Indicate by check mark if the registrant is a well-known seasoned issuer, as defined in Rule 405 of the Securities Act. Yes o No ý

Indicate by check mark if the registrant is not required to file reports pursuant to Section 13 or Section 15(d) of the Act. Yes o No ý

Indicate by check mark whether the registrant: (1) has filed all reports required to be filed by Section 13 or 15(d) of the Securities Exchange Act of 1934 during the preceding 12 months (or for such shorter period that the registrant was required to file such reports), and (2) has been subject to such filing requirements for the past 90 days. Yes ý No o

Indicate by check mark whether the registrant has submitted electronically and posted on its corporate Web site, if any, every Interactive Data File required to be submitted and posted pursuant to Rule 405 of Regulation S-T (§232.405 of this chapter) during the preceding 12 months (or for such shorter period that the registrant was required to submit and post such files). Yes ý No o

Indicate by check mark if disclosure of delinquent filers pursuant to Item 405 of Regulation S-K is not contained herein, and will not be contained, to the best of registrant's knowledge, in definitive proxy or information statements incorporated by reference in Part III of this Form 10-K or any amendment to this Form 10-K. ý

Indicate by check mark whether the registrant is a large accelerated filer, an accelerated filer, a non-accelerated filer, or a smaller reporting company. See definition of "large accelerated filer," "accelerated filer" and "smaller reporting company" in Rule 12b-2 of the Exchange Act. (Check one):

Large accelerated filer o	Accelerated filer o
Non-accelerated filer ý (Do not check if a smaller reporting company)	Smaller reporting company o

Indicate by check mark whether the registrant is a shell company (as defined in Rule 12b-2 of the Exchange Act). Yes o No ý

State the aggregate market value of the voting and non-voting common equity held by non-affiliates computed by reference to the price at which the common equity was last sold, or the average bid and asked price of such common equity, as of the last business day of the registrant's most recently completed second fiscal quarter. None.

The registrant has met the conditions set forth in General Instruction I (1) (a) and (b) of Form 10-K and is therefore filing this Annual Report on Form 10-K with the reduced disclosure format.

TABLE OF CONTENTS

Table of Contents

DIRECTV HOLDINGS LLC

CAUTIONARY STATEMENT FOR PURPOSE OF THE "SAFE HARBOR" PROVISIONS OF THE PRIVATE SECURITIES LITIGATION REFORM ACT OF 1995

This Annual Report on Form 10-K may contain certain statements that we believe are, or may be considered to be, "forward-looking statements" within the meaning of various provisions of the Securities Act of 1933 and of the Securities Exchange Act of 1934. These forward-looking statements generally can be identified by use of statements that include phrases such as we "believe," "expect," "estimate," "anticipate," "intend," "plan," "foresee," "project" or other similar words or phrases. Similarly, statements that describe our objectives, plans or goals also are forward-looking statements. All of these forward-looking statements are subject to certain risks and uncertainties, including, without limitation, risk factors discussed in more detail in Item 1A of this Annual Report, which could cause our actual results to differ materially from historical results or from those expressed or implied by the relevant forward-looking statement. The forward-looking statements included in this Annual Report are made only as of the date of this Annual Report and we undertake no obligation to publicly update these forward-looking statements to reflect subsequent events or circumstances.

PART I

ITEM 1. BUSINESS

DIRECTV Holdings LLC is a wholly-owned subsidiary of DIRECTV and consists of DIRECTV Enterprises, LLC and its wholly-owned subsidiaries and DIRECTV Financing Co., Inc. We sometimes refer to DIRECTV Holdings LLC as DIRECTV Holdings, DIRECTV U.S, we or us and sometimes refer to DIRECTV as our Parent.

On November 19, 2009, The DIRECTV Group, Inc., or DIRECTV Group, and Liberty Media Corporation, which we refer to as Liberty or Liberty Media, obtained shareholder approval of and closed a series of related transactions which we refer to collectively as the Liberty Transaction. The Liberty Transaction included the split-off of certain of the assets of the Liberty Entertainment group into Liberty Entertainment, Inc., or LEI, which was then split-off from Liberty. Following the split-off, DIRECTV Group and LEI merged with subsidiaries of DIRECTV. As a result of the Liberty Transaction, DIRECTV Group, which is comprised of the DIRECTV U.S. and DIRECTV Latin America businesses, and LEI, which held Liberty's 57% interest in DIRECTV Group, a 100% interest in three regional sports networks, a 65% interest in Game Show Network LLC, or GSN, approximately $120 million in cash and cash equivalents and approximately $2.1 billion of indebtedness and a series of related equity collars became wholly-owned subsidiaries of DIRECTV. DIRECTV Holdings remained a direct subsidiary of DIRECTV Group and became an indirect subsidiary of DIRECTV.

We provide over 18.5 million subscribers with access to hundreds of channels of digital-quality video pictures and CD-quality audio programming that we transmit directly to subscribers' homes or businesses via high-powered geosynchronous satellites.

We believe we provide one of the most extensive collections of programming available in the multi-channel video programming distribution, or MVPD, industry. As of December 31, 2009, we distributed more than 2,000 digital video and audio channels, including about 200 basic entertainment and music channels, 40 premium movie channels, over 50 regional and specialty sports networks, over 120 Spanish and other foreign language special interest channels, over 31 pay-per-view movie and event choices, and over 130 national high-definition, or HD, television channels. Although we distribute more than 1,500 local channels—over 500 in high-definition—a subscriber generally receives only the local channels in the subscriber's home market. In addition, we offer an on demand service named DIRECTV on DEMAND which, as of the end of 2009, provided a selection of about 6,000 movie and television programs to our subscribers who have a broadband connection to their set-top receiver. As of

Table of Contents

December 31, 2009, we provided local channel coverage in standard definition to markets covering about 95% of U.S. television households. In addition, we provided HD local channels to markets representing approximately 92% of U.S. TV households. In the second quarter of 2010, we expect to further expand our offering of HD channels when the recently launched DIRECTV 12 satellite begins operations.

We also provide premium professional and collegiate sports programming such as the NFL SUNDAY TICKET™ package, which allows subscribers to view the largest selection of NFL games available each Sunday during the regular season. Under our contract with the NFL, we have exclusive rights to provide this service through the 2014 season, including rights to provide related broadband, HD, interactive and mobile services.

To subscribe to the DIRECTV® service, subscribers acquire receiving equipment from either us, our national retailers, independent satellite television retailers or dealers, or regional telephone companies, which we refer to as telcos. Most set-top receivers provided to new and existing subscribers are leased subsequent to the introduction of a lease program on March 1, 2006.The receiving equipment consists of a small receiving satellite dish antenna, a digital set-top receiver and a remote control, which we refer to as a DIRECTV® System. After acquiring and installing a DIRECTV System, subscribers activate the DIRECTV service by contacting us and subscribing to one of our programming packages.

Key Strengths

- ***Large Subscriber Base.*** We are the largest provider of direct-to-home, or DTH, digital television services and the second largest MVPD provider in the United States, in each case based on the number of subscribers. We believe that our large subscriber base provides us with the opportunity to obtain programming on favorable terms and secure unique and exclusive programming. We also believe that our large subscriber base contributes to achieving other economies of scale in areas such as DIRECTV System equipment purchasing, customer service, installation and repair service, broadcast operations and general and administrative services.

- ***Leading Brand Name.*** Results from a study we commissioned in 2009 indicated that 96% of consumers in the United States recognized the DIRECTV brand name. We believe the strength of our brand name is an important factor in our ability to attract new subscribers. In addition, we believe our recognized brand name enhances our ability to secure strategic alliances with programmers, distributors and other technology and service providers.

- ***Substantial Channel Capacity and Programming Content.*** As a result of our significant channel capacity, we believe we are able to deliver to our subscribers one of the widest selections of local and national programming available today in the United States, including exclusive programming such as the NFL SUNDAY TICKET package and international programming. In addition, we have a substantial amount of capacity in the Ka-Band spectrum which enables us to provide one of the most extensive national HD offerings currently available in the industry.

- ***High-Quality Digital Picture and Sound, Including HD Programming.*** Our video and audio programming is 100% digitally delivered, providing subscribers with digital-quality video and CD-quality sound. We believe this compares favorably with most cable providers that frequently offer popular programming in an analog format and offer a selection of digital channels for an additional fee. In addition, we believe we currently offer one of the nation's most comprehensive selections of HD channels, including the largest choice of 1080p movies.

- ***Strong Customer Service.*** We have attained top rankings in customer satisfaction studies for our industry. For example, we have been rated ahead of every major cable company in customer service for nine consecutive years in the American Customer Satisfaction Index™. We believe that providing high-quality customer service is an important element in minimizing subscriber disconnection, or churn, and attracting new subscribers.

- ***Valuable Orbital Slots and Satellite-Based Technology.*** We believe our regulatory authorization to use desirable orbital slots and broadcast spectrum helps sustain our position as one of the leading companies in the MVPD industry. The Federal Communications Commission, or FCC, has designated three direct broadcast satellite, or DBS, orbital slots in the Ku-Band spectrum that provide full coverage across the 48 contiguous states of the United States, often referred to as CONUS coverage. Within these three orbital slots, there are 96 assigned DBS frequencies. We hold licenses to broadcast our services from 46 of these 96 DBS frequencies. The FCC is currently considering licensing additional DBS slots for satellites that are sometimes referred to as "tweeners" which would provide CONUS coverage. See "Government Regulation—FCC Regulation Under the Communications Act and Related Acts" and "Risk Factors—The ability to maintain FCC licenses and other regulatory approvals is critical to our business" for more information related to these types of slots and satellites.

  In addition, we hold licenses in three orbital slots (99° west longitude, or WL, 101° WL, and 103° WL) in the Ka-Band spectrum. The satellites that have been launched into these orbital slots have substantially increased our channel capacity, allowing us to provide one of the most extensive HD channel offerings currently available across the United States. We also have obtained approval from the FCC to transmit our signal in the Ku-Band from one of our satellites that has been stationed at a temporary orbital location at 72.5° WL and from leased capacity on a satellite at 95° WL.

  Our satellite-based service provides us with many advantages over ground-based cable television services. We have the ability to distribute hundreds of channels to millions of recipients nationwide with minimal incremental infrastructure cost per additional subscriber. In addition, we have comprehensive coverage to areas with low population density in the United States and the ability to quickly introduce new services to a large number of subscribers.

Business Strategy

Our vision is to provide customers with the best video experience in the United States both inside and outside of the home by offering subscribers unique, differentiated and compelling programming through leadership in content, technology and customer service.

- ***Offer Differentiated and Exclusive Content and Services.*** To fulfill our goal, we believe we must provide the most extensive collection of valuable programming and interactive services to our customers.

  - *Provide the Most Extensive Collection of Valuable Programming.* We believe that we currently have one of the most extensive collections of programming in the MVPD industry and our strategy is to continue improving our offering. For example, we offer content which is not offered by other MVPD providers such as NFL SUNDAY TICKET where subscribers can watch up to 14 games each week, most of which are offered in HD. We have also signed agreements to be the exclusive MVPD provider of NCAA® MEGA MARCH MADNESS®. In addition, we offer our customers The 101® Network, a free premium channel dedicated to the broadcast of unique and exclusive content including series such as Friday Night Lights, Deadwood®, Sleeper Cell® and The Nine™ as well as concert

Table of Contents

performances by top-rated artists. In 2009, we also launched new shows on The 101® Network such as The Dan Patrick Show™ and Trailer Park Boys®.

We also believe we currently have one of the most extensive national HD channel offerings as well as the largest lineup of 1080p movies in the MVPD industry. Additionally, when our DIRECTV 12 satellite is put into service in the second quarter of 2010, we expect to have the capacity to broadcast approximately 200 national HD channels to nearly all U.S. television households. As part of this rollout, we plan on offering local channels in HD to 19 additional markets, bringing the total number of HD local channel markets to 157—covering over 95% of TV homes. Subscribers receiving local HD channels will generally only receive the channels broadcast in their home market. Additionally, in 2010 we plan on being one of the first MVPD providers to offer dedicated 3D programming by introducing three 3D channels to our HD customers who have purchased 3D television sets.

We also expect to expand our DIRECTV on DEMAND, a video-on-demand, or VOD, service for subscribers that have the DIRECTV Plus® digital video recorder, or DVR, or DIRECTV Plus® HD DVR set-top receivers. As of year end 2009, DIRECTV on DEMAND offered about 6,000 titles providing thousands of hours of top programming from the major broadcast and cable networks, as well as popular movies. Most of the titles are offered free of charge and are downloaded from the Internet through a broadband connection for those subscribers with a DIRECTV Plus HD DVR. In addition, we download top movies via our satellites to a customer's DVR hard drive. In 2010, we expect to introduce a new movie service, DIRECTV Cinema™, which will substantially increase the number of new release movies available for our customers to view and purchase from either their television, laptop computer or mobile telephone.

- *Expand and Enhance Interactive Services.* We believe that enhanced and interactive services play an important role in the subscriber experience. For example, NFL SUNDAY TICKET subscribers can view a mix-channel with up to 8 games on one screen while the SUPERCAST™ service lets viewers access games and interactive statistics online through a personal computer. We also offer interactive services for many major tennis and golf events, including the Masters® golf tournament, where we dedicate several extra channels of event coverage, interactive scoreboards and a mix channel, all of which no other MVPD operator provides. In 2009, we launched our free TV Apps service, which are applications that appear on the TV, including Flickr®, weather forecasts and other user generated programs for customers who connect their DIRECTV Plus HD DVR to their broadband router. We have added interactive applications for the 2010 Winter Olympics, such as medal count pages and special USA team coverage.

- ***Technology Leadership.*** We believe that technological leadership has been and will continue to be important to our ability to introduce services that are easy to use and subscriber-friendly, while also reducing costs. We believe that advancements in technology will drive subscriber demand for enhanced DVRs, and HD equipment, VOD, a whole-house entertainment solution, mobile and portable devices.

    - *Introduce Multi-Room, Whole-House, Mobile and Portable Services.* We believe that it is important for our subscribers to have multiple ways to access DIRECTV® programming throughout the home and on devices outside of the home. Accordingly, in 2010 we will be introducing a multi-room viewing service. This service will enable customers with the proper equipment to share content and recorded shows around the house. In the second half of 2010, we expect to introduce a home media center that will provide HD, DVR and standard-definition video functionality throughout the home and allow customers to access stored content, including video, photos and music, seamlessly from any connected television in a home. We also intend to make DIRECTV programming more ubiquitous by offering it on portable and mobile devices, including cell phones. For example, in 2009, subscribers to our NFL SUNDAY TICKET™ SuperFan® package were able to stream live NFL games to their mobile phones. In addition, we believe that our ongoing marketing relationships with the major wireless telephony providers such as AT&T and Verizon provide us a unique opportunity to develop compelling applications for our customers.

- *Enhance/Improve User Interface/Guide.* We are constantly striving to improve our guide and user interface because it is important that our subscribers are able to access the many offerings we provide in as easy and intuitive a manner as possible. For example, in January 2010 we introduced Smart Search which helps customers find what they are looking for on TV faster and easier as well as providing significantly more information about the television programs and actors they are interested in watching.

  DIRECTV has led the industry in the application of remote DVR scheduling technology as over 2 million of our customers have scheduled over 12 million recordings remotely to DIRECTV DVRs through 2009. We also introduced ScoreGuide™ in 2009 which, at the press of a button, enables customers to easily track scores and start times of major sporting events, see a list of channels carrying each event and tune directly to those channels. In 2010, we expect to further expand ScoreGuide, to include Olympic and soccer coverage. Another example of our improved user interface is GameSearch™ which automatically recognizes when a customer has tuned to a channel with a blacked-out sports game and immediately looks to see if the game is on another channel. It then provides the customer a message telling them where to find the game or if it is unavailable.

- ***Enhance Sales and Marketing; Focus on High Quality Subscribers; Improve Customer Service, Distribution and Installation.*** We expect to continue to grow our subscriber base and maintain relatively low churn levels by focusing on acquiring higher quality subscribers as well as improving our customer service, distribution and installation.

  - *Enhance Sales and Marketing.* We expect to continue growing our subscriber base through marketing programs that capitalize on the strength of our brand and extensive programming. In addition, we expect that our expanded national and local HD programming, as well as many of our new services including multi-room viewing, DIRECTV Cinema and the home media center will increase sales from customers purchasing these services. We also intend to continue focusing on local advertising and marketing to ensure that our competitive strengths are effectively targeted based on competitive factors, demographics and geography.

  - *Maintain Low Levels of Churn by Attracting High Quality Subscribers.* We believe that in order to maintain churn at relatively low levels, we must continue to improve the overall quality of our subscriber base by regularly refining our credit and identification policies to properly reflect the changing competitive and economic landscape. However, we believe it is also important to balance churn levels by adjusting our upgrade and retention policies and costs to help assure appropriate financial returns.

  - *Improve Customer Service, Distribution and Installation.* We strive to attain the gold standard in customer service throughout a customer's lifecycle. We expect to improve customer service, distribution and installation services while also improving operational efficiencies. For example, in 2008 and 2009 we entered into several transactions which resulted in a substantial portion of our previously outsourced service and installation network technicians becoming DIRECTV U.S. employees. By having these technicians as employees of DIRECTV U.S., we have reduced turnover and improved the overall customer experience, and performance of the remaining outsourced technicians has also generally improved. We have also improved the quality and usage of our web-based customer service capabilities, improved the tools that our customer service representatives have at their disposal, and simplified our customer bills. In addition, we have implemented a new work order management system that has improved the scheduling and tracking of our installation and service calls including the use of wireless handheld devices so that our service technicians can improve the efficiency of their daily work orders. In 2009, we have seen substantial improvements in many of our customer service and installation metrics and we expect to make further improvements in 2010.

- *Improve and Expand Relationships with Telcos.* In February 2009, AT&T began marketing a bundle of broadband Internet, telephone services and DIRECTV video service to new and existing customers. AT&T's territories include 22 states and cover approximately 44 million households. With this relationship, we now have agreements with the three major telecommunications companies—AT&T, Verizon and Qwest—covering approximately 90 million homes in the United States. We are also working with the telcos to develop new services including more integrated bundles and wireless applications.

Infrastructure

Satellites. We currently have a fleet of twelve geosynchronous satellites, including eleven owned satellites and one leased satellite. We have seven Ku-Band satellites at the following orbital locations: 101° WL (three), 110° WL (one), 119° WL (one), 72.5° WL (one), and 95° WL (one-leased). We also have five Ka-Band satellites at our 99° WL (two) and 103° WL (three) orbital locations. The 72.5° WL orbital location is used pursuant to an arrangement with Telesat Canada and Bell ExpressVu.

We are currently evaluating whether to begin construction of an additional satellite to provide additional services as well as backup capacity. If we do decide to acquire such a satellite, we expect that it would be launched and go into service in 2013.

Satellite Risk Management. At times, we use launch and in-orbit insurance to mitigate the potential financial impact of satellite fleet launch and in-orbit failures unless the premium costs are considered to be uneconomical relative to the risk of satellite failure. The insurance generally does not compensate for business interruption or loss of future revenues or subscribers. We rely on in-orbit spare satellites and excess transponder capacity at key orbital slots to mitigate the impact of a potential satellite failure on our ability to provide service. However, programming continuity cannot be assured in all instances or in the event of multiple satellite losses.

Launch insurance typically covers the time frame from ignition of the launch vehicle through separation of the satellite from the launch vehicle. In the past, we have launched satellites without insurance. As of December 31, 2009, the net book value of our in-orbit satellites was $1,516 million, none of which is insured.

Digital Broadcast Centers. To gather programming content, ensure its digital quality, and transmit content to our satellites, we have built two digital broadcast centers, located in Castle Rock, Colorado and Los Angeles, California. These facilities provide the majority of our national and local standard-definition and HD programming. We have also built five uplink facilities which are used to provide HD local channels. Our broadcast centers receive programming from content providers via satellite, fiber optic cable and/or special tape. Most satellite-delivered programming is then digitized, encoded and transmitted to our satellites. We designed each broadcast center and uplink facility with redundant systems to minimize service interruptions.

Installation Network. The DIRECTV home service provider, or HSP, installation and service network performs installation, upgrades and other service call work for us. In 2008 and 2009, we entered into several transactions that brought a significant portion of this HSP network activity in-house. We now directly employ nearly 4,000 technicians and utilize an additional 11,000 technicians from seven outsourced companies around the United States. The combined workforce completed approximately 93% of all in-home visits in 2009. We set the standards for the quality of installation and service, perform quality control, manage inventory and monitor the overall service network performance for nearly all of the third-party installation network.

Customer Service Centers. As of December 31, 2009, we used 36 customer service centers employing over 16,000 customer service representatives. Most of these customer service centers are operated by Convergys Customer Management Group, Inc., Precision Response Corporation, Sitel Operating Corporation, N.E.W. Customer Service Companies, Inc., VXI Global Solutions, Inc. and Teleperformance. We currently own and operate six customer service centers located in: Boise, Idaho; Tulsa, Oklahoma; Huntsville, Alabama; Missoula, Montana; Huntington, West Virginia and Denver, Colorado that employ approximately 5,000 customer service representatives. Potential and existing subscribers can call a single telephone number 24 hours a day, seven days a week, to request assistance for hardware, programming, installation, technical and other support. We continue to increase the functionality of telephone-based and web-based self-care features in order to better manage customer service costs and improve service levels.

Competition

We face substantial competition in the MVPD industry and from emerging digital media distribution providers. Our competition includes companies that offer video, audio, interactive programming, telephony, data and other entertainment services, including cable television, other DTH companies, telcos, wireless companies and companies that are developing new technologies. Many of our competitors have access to substantially greater financial and marketing resources. We believe our brand, the quality and variety of video, audio and interactive programming, quality of picture, access to service, availability of HD and DVR services, customer service and price are the key elements for attaining and retaining subscribers. Our over 18.5 million subscribers represent approximately 19% of MVPD subscribers at December 31, 2009.

- *Cable Television.* We encounter substantial competition in the MVPD industry from cable television companies. According to the National Cable & Telecommunications Association's 2008 Industry Overview, 96% of the 128.6 million U.S. housing units are passed by cable. Most cable television operators have a large, established customer base, and many have significant investments in companies that provide programming content. Approximately 100 million households subscribe to an MVPD service and approximately 62% of MVPD subscribers receive their programming from a cable operator. In addition, most cable providers have completed network upgrades that allow for enhanced service offerings such as digital cable, HD channels, broadband Internet access and telephony services. Cable companies bundle these services, offering discounts and providing one bill to the consumer.

- *Telephone Companies.* Several telcos have upgraded a significant portion of their infrastructure by replacing their older copper wire telephone lines with high-speed fiber optic lines. These fiber lines provide the telcos with significantly greater capacity enabling them to offer new and enhanced services, such as Internet access at much greater speeds and digital-quality video. For example, Verizon announced that at the end of 2009, it had the capability to serve 15 million homes with fiber optic lines with the goal of having the capability to serve 18 million homes by the end of 2010. In addition, AT&T has begun deploying fiber optic lines to neighborhoods and expects to have the capability to serve approximately 30 million of its customers by the end of 2011. As of year end 2009, Verizon had nearly 3 million video subscribers and AT&T had approximately 2 million subscribers. Similar to the cable companies, the telcos expect to offer their customers multiple services at a discount on one bill.

- *Other Direct Broadcast Satellite and Direct-To-Home Satellite System Operators.* We also compete with DISH Network Corporation, or DISH Network, which had over 14 million subscribers at the end of 2009, representing approximately 14% of MVPD subscribers. Other domestic and foreign satellite operators also have proposed to offer DTH satellite service to U.S. customers using U.S.-licensed satellite frequencies or foreign-licensed frequencies that have the ability of covering the United States.

- *Video via the Internet.* With the large increase in the number of consumers with broadband service, a significant amount of video content has become available on the Internet for users to download and view on their personal computers, televisions and other devices. For example, Apple™ offers two hundred television shows and 400 movies for rental or purchase, some in high-definition, on the online iTunes® Store. In addition, Hulu™ is an online video service website which provides free movies and TV shows from over 190 content providers including Fox, Disney, NBC Universal, MGM Studios, Sony Pictures and Warner Bros. This content can be accessed on demand through its website and those of its partners—AOL, MSN, MySpace and Yahoo. In addition, several companies, such as Netflix, Blockbuster and Amazon.com, have begun selling and renting movies via Internet download. For example, Netflix has a library of 17,000 movies and TV shows available for download to its over 12 million subscribers. There are also several similar initiatives by companies such as Intel, Microsoft and Sony to make it easier to view Internet-based video on television and personal computer screens. Many television models, Blu-Ray Disc® players and gaming consoles like the Xbox® can be directly connected to the Internet and have the capacity to stream video to the television.

- *Mobile Video.* Many companies are beginning to offer mobile applications for video allowing consumers to watch video on the go. For example, AT&T offers mobile TV which provides users the ability to watch full length TV shows from ABC, CBS, ESPN and other programmers on their cell phones. Verizon Wireless offers V Cast™ which allows subscribers to watch many of the top TV shows including college football and basketball on their mobile phone for a modest fee. In addition other mobile applications and services are becoming available, such as FLO TV from Qualcomm. FLO TV™ Service provides portable TV with full length shows from programmers such as Fox, CBS, Fox News Channel and ESPN on a portable device. Other cable and satellite distributors are also focused on distributing their content to their customers on the go.

- *Small and Rural Telephone Companies.* Other telephone companies are also finding ways to deliver video programming services over their wireline facilities or in a bundle with other MVPD providers. For example, DISH Network has agreements with Embarq, CenturyTel, Windstream, TDS, and Frontier to bundle their individual DSL and telephony services with DISH Network's video service.

- *Local Broadcasters.* Most areas of the United States can receive traditional digital television broadcasts of between three and ten channels. These broadcasters are often low to medium power operators with a limited coverage area and provide local, network and syndicated programming typically free of charge. There are over 2,000 TV broadcast stations in the U.S. split among 210 TV markets.

ACQUISITIONS, STRATEGIC ALLIANCES AND DIVESTITURES

We review our competitive position on an ongoing basis and, from time to time, consider various acquisitions, strategic alliances and divestitures, including potential wireless broadband investments or alliances, in order to continue to compete effectively, improve our financial results, grow our business and allocate our resources efficiently. We also consider periodically making equity investments in companies with which we can jointly provide services to our subscribers.

GOVERNMENT REGULATION

We are subject to government regulation in the United States, primarily by the FCC and, to a certain extent, by the legislative branches, other federal agencies and state and local authorities. We are also subject to the rules and procedures of the International Telecommunications Union, or ITU, a specialized agency of the United Nations within which governments and the private sector coordinate global telecommunications networks and services. Depending upon the circumstances, noncompliance with legislation or regulations promulgated by these entities could result in the suspension or revocation of our licenses or registrations, the termination or loss of contracts or the imposition of contractual damages, civil fines or criminal penalties.

This section sets forth a summary of regulatory issues pertaining to our operations in the United States and is not intended to describe all present and proposed government regulation and legislation affecting the MVPD industry or our business.

FCC Regulation Under the Communications Act and Related Acts. The Communications Act and other related acts give the FCC broad authority to regulate the operations of our company.

The ownership and operation of our DBS/DTH system is regulated by the FCC primarily for:

- the licensing of DBS and DTH satellites, earth stations and ancillary authorizations;
- the assignment of frequencies and orbital slots, the relocation of satellites to different orbital locations or the replacement of an existing satellite with a new satellite;
- compliance with the terms and conditions of assignments and authorizations, including required timetables for construction and operation of satellites;
- avoidance of interference by and to DBS/DTH operations with operations of other entities that make use of the radio spectrum; and
- compliance with the Communications Act and FCC rules governing U.S.- licensed DBS and DTH systems.

The FCC grants authorizations to satellite operators that meet its legal, technical and financial qualification requirements. The FCC conditions such authorizations on satisfaction of ongoing due diligence, construction, reporting and related obligations.

All of our satellites and earth stations are or have been licensed by the FCC. Currently, two of our satellites are licensed by the government of Canada. While the FCC generally issues DTH space station licenses for a fifteen-year term, DBS space station and earth station licenses are generally issued for a

Table of Contents

ten-year term, which is less than the useful life of a healthy direct broadcast satellite. Upon expiration of the initial license term, the FCC has the option to renew a satellite operator's license or authorize an operator to operate for a period of time on special temporary authority, or decline to renew the license. If the FCC declines to renew the operator's license, the operator is required to cease operations and the frequencies it was previously authorized to use would revert to the FCC.

Currently, we have several applications pending before the FCC, including applications to launch and operate future satellites to support DIRECTV's services. In general, the FCC's approval of these applications is required for us to continue to expand our range of service offerings while increasing the robustness of our satellite fleet. We may not obtain these approvals in a timely fashion or at all.

As a DBS/DTH licensee and operator we are subject to a variety of Communications Act requirements, FCC regulations and copyright laws that could materially affect our business. They include the following:

- ***Local-into-Local Service and Limitation on Retransmission of Distant Broadcast Television Signals.*** The Satellite Home Viewer Improvement Act, or SHVIA, allows satellite carriers to retransmit the signals of local broadcast television stations in the stations' local markets without obtaining authorization from the holders of copyrights in the individual programs carried by those stations. Another portion of SHVIA, as amended by the Satellite Home Viewer Extension and Reauthorization Act of 2004, or SHVERA, also permits satellite retransmission of distant network stations (those that originate outside of a satellite subscriber's local television market) only to "unserved households." A subscriber qualifies as an "unserved household" if he or she cannot receive, over the air, a signal of sufficient intensity from a local station affiliated with the same network, or falls into one of a few other very limited exceptions. SHVERA also prohibits satellite carriers from signing up a new subscriber to distant analog or digital signals if that subscriber lives in a local market where the satellite carrier provides local analog or local digital signals, respectively. SHVERA imposes a number of notice and reporting requirements, and also permits satellite retransmission of distant stations in neighboring markets where they are determined by the FCC to be "significantly viewed." In implementing SHVIA, the FCC has required satellite carriers to delete certain programming, including sports programming, from the signals of certain distant stations. In addition, the FCC's continuing interpretation, implementation and enforcement of other provisions of SHVIA and SHVERA, as well as judicial decisions interpreting and enforcing these laws, could hamper our ability to retransmit local and distant network and superstation signals, reduce the number of our existing or future subscribers that can qualify for receipt of these signals, impose costs on us in connection with the process of complying with the rules, or subject us to fines, monetary damages or injunctions. Also, the FCC's sports blackout requirements, which apply to all distant network signals, may require costly upgrades to our system. Further, an FCC order interpreting the requirement that satellite carriers retransmit local digital signals with "equivalent bandwidth" of significantly viewed digital signals may constrain our ability to deliver such significantly viewed digital signals. The distant-signal provisions of SHVERA were set to expire at the end of 2009, but Congress has extended that deadline to February 28, 2010. Congress may decline to renew those provisions, which could severely restrict our ability to retransmit distant signals. Congress could also adopt amendments to SHVERA with respect to local or distant signals, including limiting the provision of distant signals. In particular, Congress is considering for the first time making subscribers ineligible for distant signals where they can receive local digital multicast signals over the air. This could adversely affect our ability to deliver distant signals to our existing or future subscribers.

- ***Must Carry Requirement.*** SHVIA also imposes a must carry obligation on satellite carriers. This must carry obligation requires satellite carriers that choose to take advantage of the statutory copyright license in a local market to carry upon request the signals of all qualifying television broadcast stations within that local market, subject to certain limited exceptions. The FCC has implemented SHVIA's must carry requirement and adopted further detailed must carry rules covering our carriage of both commercial and non-commercial broadcast television stations. These rules generally require us to carry all of the local broadcast stations requesting carriage in a timely and appropriate manner in markets in which we choose to retransmit the signals of local broadcast stations. We have limited capacity, and the projected number of markets in which we can deliver local broadcast programming will continue to be constrained because of the must carry requirement and may be reduced depending on the FCC's interpretation of its rules in pending and future rulemaking and complaint proceedings, as well as judicial decisions interpreting must carry requirements. For example, the FCC issued an order requiring mandatory carriage of high-definition digital signals in an increasing number of markets each year, requiring so-called "HD carry-one, carry-all" in all local markets served by 2013. We may not be able to comply with these must carry rules, or compliance may mean that we will be required to use capacity that could otherwise be used for new or additional local or national programming services. Moreover, Congress may amend the must carry rules when it considers SHVERA reauthorization. For example, Congress has in the past proposed legislation and may in the future enact legislation that would require us to provide local channels via satellite in all markets in the United States. We currently provide local channel coverage to approximately 155 markets representing approximately 95% of U.S. television households. If such legislation were enacted, we would be required to provide local channel coverage to an additional 55 markets representing about 5% of U.S. television households on an accelerated timetable. We believe that the capital expenditures and ongoing costs to provide this coverage would not be covered by the incremental revenue from the additional subscribers we could potentially gain in these markets. Moreover, depending upon the timetable imposed, we may not be able to comply in a timely manner.

- ***Public Interest Requirement.*** Under a requirement of the Communications Act, the FCC has imposed certain public interest obligations on DBS operators, including a requirement that such providers set aside four percent of channel capacity exclusively for noncommercial programming of an educational or informational nature, for which we must charge programmers below-market rates and for which we may not impose additional charges on subscribers. FCC rules also require us to comply with a number of political broadcasting requirements to which broadcasters are subject under the Communications Act, as well as limits on the commercialization of children's programming applicable to cable operators. We believe that we are in compliance with all of these requirements, but some of them require our interpretations, which we believe are reasonable and consistent with industry practice. However, if we are challenged, the FCC may not agree with our interpretations. In addition, the FCC could, in the future, attempt to impose additional public interest or content requirements on us, for example, by seeking to impose rules on indecent programming.

- ***Emergency Alert System.*** The Emergency Alert System, or EAS, requires participants to interrupt programming during nationally-declared emergencies and to pass through emergency-related information. The FCC has adopted rules that require satellite carriers to participate in the "national" portion of EAS. It is also considering whether to mandate that satellite carriers also interrupt programming for local emergencies and weather events. We believe that any such requirement would be very difficult to implement, would require costly changes to our DBS/DTH system, and, depending on how it is implemented, could inconvenience or confuse our

Table of Contents

viewers. The FCC is also considering whether to require that EAS alerts be provided in multiple languages or via text messages, which could also prove difficult and costly to implement depending upon the nature of any such requirement adopted.

- ***Spectrum Allocation and License Assignment Rules.*** We depend upon the FCC's allocation of sufficient DBS frequencies and assignment of DBS licenses in order to operate our business. DBS frequencies and available DBS orbital locations capable of supporting our business have become increasingly scarce. While we have obtained additional DTH service capacity and continue to explore new sources of DBS/DTH capacity, there can be no assurance that we will obtain further capacity. In addition, the FCC had adopted a system of competitive bidding to assign licenses for additional DBS frequencies. On June 21, 2005, the United States Court of Appeals for the D.C. Circuit held that such an auction process was not authorized by statute. The FCC subsequently voided the previous auction and implemented a freeze on applications for authority to provide DBS service in the United States using new frequencies or new orbital locations not assigned to the United States in the ITU Region 2 Broadcasting Satellite Service, or BSS, Plan. On August 18, 2006, the FCC began a proceeding to identify a new system for assigning DBS authorizations. There can be no assurance that we will be able to obtain additional DBS capacity under whatever system the FCC implements in the future.

  In 2007, the FCC adopted new service and licensing rules for the BSS in the 17.3-17.8 GHz and 24.75-25.25 GHz bands, or 17/24 GHz BSS. This spectrum, also known as the "reverse band" (in that transmissions from these satellites to consumers would occur in spectrum currently used for uplinking programming to traditional DBS satellites), could provide a new source of additional DTH capacity. Among other things, the FCC established a licensing procedure under which the four parties with applications then pending—including DIRECTV—would be allowed to amend their applications to conform to the new rules and would be entitled to have those applications processed on a co-equal basis with one another before any new applications would be accepted. On July 28, 2009, the FCC granted four DIRECTV satellite applications in this band. However, foreign operators who may have international priority have indicated an interest in using slots that may conflict with some or all of these licenses. One foreign licensed operator, Spectrum Five LLC, has filed a petition seeking reconsideration of one of DIRECTV's licenses at an orbital location where Spectrum Five also proposes to operate, and that petition remains pending.

- ***Rules Governing Co-Existence With Other Satellite and Terrestrial Services and Service Providers in the MVPD Industry.*** The FCC has adopted rules to allow non-geostationary orbit fixed satellite services to operate on a co-primary basis in the same frequency band as the one used by direct broadcast satellite and Ku-Band-based fixed satellite services. In the same proceeding, the FCC concluded that multi-channel video and data distribution services, or MVDDS, can share spectrum with DBS operators on a non-interference basis, and adopted rules and a method for assigning licenses in that service, as well. While the FCC has established service and technical rules to govern the non-geostationary orbit and MVDDS services to protect DBS operations from harmful interference, these rules may not be sufficient to prevent such interference, and the introduction of such services into spectrum used by us for DBS service may have a material adverse impact on our operations. A number of aspects of these rules remain subject to judicial review. In addition, one MVDDS operator recently requested a waiver of the applicable rules so that it could operate systems at substantially higher power levels in 80 markets where it holds MVDDS licenses. If granted, such a waiver may have a material adverse impact on our operation in the affected markets. Although we have opposed that waiver request, there can be no assurance that the FCC will deny it.

Table of Contents

On August 18, 2006, the FCC released a notice of proposed rulemaking regarding the possible operation of "tweener" or "short spaced" satellites—satellites that would operate in the same DBS uplink and downlink frequency bands as us, from orbital positions located in between those now assigned to the DBS service. This rulemaking follows applications by SES and Spectrum Five LLC to operate tweener satellites. Under rules that the FCC is considering, a provider could, by complying with certain technical restrictions, operate a satellite in between two orbital locations where we have already positioned our satellites without completing coordination of its operations with us and without demonstrating that such operations would not "affect" us as that term is defined by the ITU. We have opposed this proposal, and believe that tweener satellites as proposed by applicants would cause interference to our current and planned operations and impose a significant constraint on the further growth of our DBS service. We cannot predict what if any action the FCC may take or the effect of such a proceeding on our business.

On November 29, 2006, despite the pendency of the tweener satellite rulemaking and over our opposition, the FCC's International Bureau granted Spectrum Five's application to operate a tweener satellite at the 114.5° WL orbital location, only 4.5° away from our DBS satellites operating at the 110° WL and 119° WL orbital locations. While the Bureau limited Spectrum Five's operations to levels below those at which the ITU deems one DBS system to "affect" another in the absence of agreement from all affected DBS operators (including us), the Bureau's grant of Spectrum Five's application prior to coordination could ultimately permit Spectrum Five to operate at levels that would cause interference to our operations. On February 1, 2008, the full FCC denied reconsideration of the International Bureau's order, but clarified that, if Spectrum Five is unable to coordinate its tweener satellite, it must file for a modification of its authorization and demonstrate that its proposed operational parameters would not exceed the ITU trigger for coordination. To date, Spectrum Five has neither contacted us to attempt coordination of its tweener system nor filed for modification of its authorization as directed by the FCC.

The FCC has also adopted rules that require satellite operators to take certain measures to mitigate the dangers of collision and orbital debris. Among other things, these rules impose certain requirements for satellite design and end-of-life disposal maneuvers for all satellites launched after March 18, 2002, which apply to eight of our in-orbit satellites. We believe that we are in compliance with all of these requirements and expect that we will continue to be able to comply with them going forward, but the requirements for end-of-life disposal could result in a slight reduction in the operational life of each new satellite.

- ***Geographic Service Rules.*** The FCC requires DBS licensees to comply with certain geographic service obligations intended to foster the provision of DBS service to subscribers residing in the states of Alaska and Hawaii. We believe that we are in compliance with these rules although, in the past, some have argued otherwise to the FCC. The FCC has not acted on petitions filed several years ago by the State of Hawaii and an Alaska satellite television dealer. We cannot be sure that the FCC will agree with our view that we are in compliance with the agency's geographic services rules, or that the FCC will not require us to make potentially cumbersome and costly changes to our offerings. The FCC has also adopted similar rules for the 17/24 GHz BSS service.

- ***FCC Conditions Imposed In Connection With the Liberty and News Corporation Transactions.*** In approving Liberty's 2008 acquisition of News Corporation's equity investment in DIRECTV, the FCC imposed a number of regulatory conditions on us and Liberty, some of which directly or indirectly affected our business. In granting authority for the merger of Liberty

DIRECTV HOLDINGS LLC

Entertainment, Inc. and DIRECTV in 2009, the FCC conditioned its approval of the transaction on continued compliance with those conditions. Accordingly, the FCC has imposed on us program carriage conditions intended to prevent discrimination against all forms of unaffiliated programming; and certain program access conditions intended to ensure non-discriminatory access to much of the programming carried on the DIRECTV service. In particular, we may be required to submit to "baseball style" arbitration if we cannot arrive at terms for carriage of our regional sports network programming with an MVPD. We cannot predict what effect our compliance with or the FCC's enforcement of these conditions will have on our business.

International Telecommunications Union Rules. We are required by international rules to coordinate the use of the frequencies on our satellites with other satellite operators who may interfere with us or who may suffer interference from our operations.

Other Legal and Regulatory Requirements. DBS/DTH providers are subject to other federal and state regulatory requirements, such as Federal Trade Commission, FCC and state telemarketing and advertising rules, and subscriber privacy rules similar to those governing other MVPDs. We have agreed with the Federal Trade Commission to (1) review and monitor compliance with telemarketing laws by any companies we authorize to do telemarketing as well as by independent retailers, (2) investigate and respond to complaints about alleged improper telemarketing and (3) terminate our relationship with marketers or retailers found in violation. Similarly, we have agreed with certain state attorneys general to comply with advertising disclosure requirements and monitor compliance by independent retailers.

In addition, although Congress has granted the FCC exclusive jurisdiction over the provision of DTH satellite services, aspects of DBS/DTH service remain regulated at the state and local level. For example, the FCC has promulgated rules prohibiting restrictions by local government agencies, such as zoning commissions and private organizations, such as homeowners associations, on the placement of DBS receiving antennas. Local governments and homeowners associations, however, may continue to regulate the placement of such antennas if necessary to accomplish a clearly defined public safety objective or to preserve a recognized historic district, and may also apply to the FCC for a waiver of FCC rules if there are other local concerns of a special or unusual nature. In addition, a number of state and local governments have attempted to impose consumer protection, customer service and other types of regulation on DBS operators. Also, while Congress has prohibited local taxation of the provision of DBS service, taxation at the state level is permissible, and many states have imposed such taxes, and additional states have attempted to do so recently. Incident to conducting a consumer directed business, we occasionally receive inquiries or complaints from authorities such as state attorneys general and state consumer protection offices. These matters are generally resolved in the ordinary course of business.

INTELLECTUAL PROPERTY

All DIRECTV companies maintain active programs for identifying and protecting our important intellectual property. With the exception of certain U.S. trademark registrations held by us, pursuant to trademark license agreements and various intellectual property licensed from third parties, DIRECTV Group owns all of the intellectual property for the benefit of the company and its subsidiaries.

We believe that our growing portfolio of pending and issued patents are important assets. We presently hold over 1,950 issued patents worldwide relating to our past and present businesses, including over 450 patents developed by, or otherwise relating to, the businesses of DIRECTV U.S. We hold a worldwide portfolio of over 1,100 trademarks in over 130 countries related to the DIRECTV brand, the Cyclone Design and DIRECTV products and services. In particular, DIRECTV U.S. holds trademark registrations relating to its business, including registrations of the primary "DIRECTV" and Cyclone Design trademarks. In many instances, these trademarks are licensed royalty-free to third

Table of Contents

parties for use in support of the DIRECTV U.S. business. We actively protect our important patents, trademarks and other intellectual property rights against unauthorized or improper use by third parties.

DIRECTV, DIRECTV Plus, SuperFan, SUPERCAST, ScoreGuide, DIRECTV Cinema, GameSearch, The 101 Network, and the DIRECTV Cyclone Design are trademarks of The DIRECTV Group, Inc. and/or its related entities. Other trademarks, service marks and trade names appearing in this Annual Report are the property of their respective holders.

ENVIRONMENTAL REGULATION

We are subject to the requirements of federal, state, local and foreign environmental laws and regulations. These include laws regulating air emissions, water discharge and universal and hazardous waste management activities. We have an environmental management function designed to track, facilitate and support our compliance with these requirements and attempt to maintain compliance with all such requirements. We have made and will continue to make, as necessary, capital and other expenditures to comply with environmental requirements. We do not, however, expect capital or other expenditures for environmental compliance to be material in 2010. In addition, we periodically review environmental stewardship concepts (such as green initiatives and energy conservation strategies) and implement these whenever feasible. Environmental requirements are complex, change frequently and have become more stringent over time. Accordingly, we cannot provide assurance that these requirements will not change or become more stringent in the future in a manner that could have a material adverse effect on our business.

We are also subject to environmental laws requiring the investigation and cleanup of environmental contamination at facilities we formerly owned or operated or currently own or operate or to which we sent hazardous wastes, including specified universal wastes, for treatment, service, disposal or recycling. We are aware of contamination at one of our former sites. We are in the process of complying with the requirements stipulated by the government agency overseeing the site clean up and have allocated the funds to achieve the decontamination goals.

EMPLOYEES

As of December 31, 2009, we had approximately 15,900 full-time and 300 part-time employees.

***

ITEM 1A. RISK FACTORS

You should carefully consider the following risk factors, as well as the more detailed descriptions of our business elsewhere in this Annual Report. The risks described below are not the only ones facing our company. Additional risks not presently known to us or that we currently deem immaterial may also adversely affect our business, financial condition or results of operations.

Our business, financial condition or results of operations could be materially and adversely affected by the following:

We compete with other MVPDs, some of whom have greater resources than we do and levels of competition are increasing.

We compete in the MVPD industry against cable television, telcos, and wireless companies and other land-based and satellite-based system operators with service offerings including video, audio and interactive programming, data and other entertainment services and telephony service. Some of these competitors have greater financial, marketing and other resources than we do.

Table of Contents

Some cable television operators have large, established customer bases and many cable operators have significant investments in, and access to, programming. According to the National Cable & Telecommunications Association's 2008 Industry Overview, 96% of the 128.6 million U.S. housing units are passed by cable. Of the 128.6 million U.S. housing units, approximately 97.6 million subscribe to an MVPD service and approximately 62% of MVPD subscribers receive their programming from a cable operator. Cable television operators have advantages relative to us, including or as a result of:

- being the incumbent MVPD operator with an established subscriber base in the territories in which we compete;
- bundling their analog video service with expanded digital video services delivered terrestrially or via satellite, or with efficient two-way high-speed Internet access or telephone service on upgraded cable systems;
- having the ability to provide certain local and other programming, including HD programming, in geographic areas where we do not currently provide local or local HD programming; and
- having legacy arrangements for exclusivity in certain multiple dwelling units and planned communities.

In addition, cable television operators have grown their subscriber bases through mergers and acquisitions, and a recent federal appeals court decision invalidating the cap on the number of subscribers a single cable operator may allow them additional avenues for growth. Moreover, mergers, joint ventures and alliances among franchise, wireless or private cable television operators, telcos, broadband service providers and others may result in providers capable of offering bundled television, data and telecommunications services in competition with our services.

We do not currently offer local channel coverage to markets covering approximately five percent of U.S. television households, which places us at a competitive disadvantage in those markets. We also have been unable to secure certain international programming, due to exclusive arrangements of programming providers with certain competitors, which has constrained our ability to compete for subscribers who wish to obtain such programming. And as discussed below, certain cable-affiliated programmers have withheld their programming from us in certain markets, which has further constrained our ability to compete for subscribers in those markets.

In the United States, various telcos and broadband service providers have deployed fiber optic lines directly to customers' homes or neighborhoods to deliver video services, which compete with the DIRECTV service. It is uncertain whether we will be able to increase our satellite capacity, offer a significant level of new services in existing markets in which we compete or expand to additional markets as may be necessary to compete effectively. Some of these various telcos and broadband service providers also sell the DIRECTV service as part of a bundle with their voice and data services. A new broadly-deployed network with the capability of providing video, voice and data services could present a significant competitive challenge and, in the case of the telcos currently selling the DIRECTV service, could result in such companies focusing less effort and resources selling the DIRECTV service or declining to sell it at all. We may be unable to develop other distribution methods to make up for lost sales through the telcos.

As a result of these and other factors, we may not be able to continue to expand our subscriber base or compete effectively against cable television or other MVPD operators in the future.

Emerging digital media competition could materially adversely affect us.

Our business is focused on television, and we face emerging competition from other providers of digital media, some of which have greater financial, marketing and other resources than we do. In particular, programming offered over the Internet has become more prevalent as broadband networks have improved their speed and quality of service. Significant changes in consumer behavior with regard to the means by which they obtain video entertainment and information in response to this emerging digital media competition could materially adversely affect our revenues and earnings or otherwise disrupt our business.

We depend on others to produce programming and programming costs are increasing.

We depend on third parties to provide us with almost all of our programming services, including third parties who are our affiliates and third parties controlled by competitors. As discussed below, a limited number of cable-affiliated programmers have in the past denied us access to their programming. Our ability to compete successfully will depend on our ability to continue to obtain desirable programming and deliver it to our subscribers at competitive prices. Our programming agreements generally have remaining terms ranging from less than one to up to ten years and contain various renewal and cancellation provisions. We may not be able to renew these agreements on favorable terms, or at all, or these agreements may be canceled prior to expiration of their original terms. If we are unable to renew any of these agreements or the other parties cancel the agreements, we may not be able to obtain substitute programming, or if we are able to obtain such substitute programming, it may not be comparable in quality or cost to our existing programming.

In addition, many of our programming agreements are long term agreements and contain fixed annual price increases. When offering new programming, or upon expiration of existing contracts, programming suppliers have historically increased the rates they charge us for programming, increasing our costs. We expect this practice to continue. Increases in programming costs could cause us to increase the rates that we charge our subscribers, which could in turn, especially in a difficult economic environment, cause subscribers to terminate their subscriptions or potential new subscribers to refrain from subscribing to our service. Furthermore, due to the economy and other factors, we may be unable to pass programming cost increases on to our subscribers, which could have a material adverse effect on our earnings or cash flow.

Increased subscriber churn or subscriber upgrade and retention costs could materially adversely affect our financial performance.

Turnover of subscribers in the form of subscriber service cancellations, or churn, has a significant financial impact on the results of operations of any subscription television provider, including us, as does the cost of upgrading and retaining subscribers. Any increase in our upgrade and retention costs for our existing subscribers may adversely affect our financial performance or cause us to increase our subscription rates, which could increase churn. Churn may also increase due to factors beyond our control, including churn by subscribers who are unable to pay their monthly subscription fees, a slowing economy, significant signal theft, consumer fraud, a maturing subscriber base and competitive offers. Any of the risks described in this Annual Report that could potentially have a material adverse impact on our cost or service quality or that could result in higher prices for our subscribers could also, in turn, cause an increase in churn and consequently have a material adverse effect on our earnings and financial performance.

Our subscriber acquisition costs could materially increase.

We incur costs relating to subscribers acquired by us and subscribers acquired through third parties. These costs are known as subscriber acquisition costs. For instance, we provide installation incentives to our retailers to enable them to offer standard professional installation as part of the subscriber's purchase or lease of a DIRECTV System. In addition, we pay commissions to retailers for their efforts in offering a DIRECTV System at a lower cost to consumers. Our subscriber acquisition costs may materially increase to the extent we continue or expand current sales promotion activities or introduce other more aggressive promotions, or due to increased competition. Any material increase in subscriber acquisition costs from current levels would negatively impact our earnings and could materially adversely affect our financial performance.

Results are impacted by the effect of, and changes in, United States economic conditions and weakening economic conditions may reduce subscriber spending and our rate of growth of subscriber additions and may increase subscriber churn.

Our business may be affected by factors in the United States that are beyond our control, such as downturns in economic activity, or in the MVPD industry. Factors such as interest rates and the health of the housing market may impact our business. A substantial portion of our revenues comes from residential customers whose spending patterns may be affected by prevailing economic conditions. Our market share in multiple dwelling units such as apartment buildings is lower than that of many of our competitors. If unemployment and foreclosures of single family residences increase, our earnings and financial performance could be negatively affected more than those of our competitors. In addition, if our customers seek alternative means to obtain video entertainment, they may choose to purchase fewer services from us. Due to the economic and competitive environment, we may need to spend more to acquire and retain customers who in turn spend less on our services. If our average monthly revenue per subscriber, or ARPU, decreases, our margins could become compressed and the long term value of a customer would then decrease. The weak economy may affect our net subscriber additions and reduce subscriber spending and, if these economic conditions continue or deteriorate further, our subscriber growth could decline and our churn rate could increase which would have a material adverse effect on our earnings and financial performance.

Our ability to keep pace with technological developments is uncertain.

In the video industry, changes occur rapidly as new technologies are developed, which could cause our services and products that deliver our services to become obsolete. We may not be able to keep pace with technological developments. If the new technologies on which we intend to focus our investments fail to achieve acceptance in the marketplace or our technology does not work and requires significant cost to replace or fix, we could suffer a material adverse effect on our future competitive position, which could cause a reduction in our revenues and earnings. For example, our competitors could be the first to obtain proprietary technologies that are perceived by the market as being superior. Further, after incurring substantial costs, one or more of the technologies under development by us or any of our strategic partners could become obsolete prior to its introduction.

In addition, technological innovation depends, to a significant extent, on the work of technically skilled employees. Competition for the services of these employees has been vigorous. We cannot assure you that we will be able to continue to attract and retain these employees.

To access technologies and provide products that are necessary for us to remain competitive, particularly in the area of broadband services, we may make future acquisitions and investments and may enter into strategic partnerships with other companies. Such investments may require a commitment of significant capital and human and other resources. The value of such acquisitions, investments and partnerships and the technology accessed may be highly speculative. Arrangements with third parties can lead to contractual and other disputes and dependence on the development and delivery of necessary technology on third parties that we may not be able to control or influence. These relationships may commit us to technologies that are rendered obsolete by other developments or preclude the pursuit of other technologies which may prove to be superior.

New technologies could also create new competitors for us. Entities such as telcos are supporting digital video delivery over existing telephone lines and building out fiber optic lines to enhance their capabilities to deliver programming services. Satellite operators such as SES have begun offering turn-key packages of digital programming on a wholesale basis for distribution by rural telcos. In addition, programming services offered over the Internet have become more prevalent as broadband networks have improved their speed and quality of service. We may not be able to compete successfully with new entrants in the market for video services.

Our business relies on intellectual property, some of which is owned by third parties, and we may inadvertently infringe patents and proprietary rights of others.

Many entities, including some of our competitors, have or may in the future obtain patents and other intellectual property rights that cover or affect products or services related to those that we currently offer or may offer in the future. In general, if a court determines that one or more of our services or the products used to transmit or receive our services infringes on intellectual property owned by others, we and the applicable manufacturers or vendors may be required to cease developing or marketing those services and products, to obtain licenses from the owners of the intellectual property or to redesign those services and products in such a way as to avoid infringing the intellectual property rights. If a third party holds intellectual property rights, it may not allow us or the applicable manufacturers to use its intellectual property at any price, which could materially adversely affect our competitive position.

We may not be aware of all intellectual property rights that our services or the products used to transmit or receive our services may potentially infringe. In addition, patent applications in the United States are confidential until the Patent and Trademark Office issues a patent. Therefore, we cannot evaluate the extent to which our services or the products used to transmit or receive our services may infringe claims contained in pending patent applications. Further, without lengthy litigation, it is often not possible to determine definitively whether a claim of infringement is valid.

We cannot estimate the extent to which we may be required in the future to obtain intellectual property licenses or the availability and cost of any such licenses. Those costs, and their impact on our earnings, could be material. Damages in patent infringement cases may also include treble damages in certain circumstances. To the extent that we are required to pay royalties to third parties to whom we are not currently making payments, these increased costs of doing business could materially adversely affect our operating results. We are currently being sued in patent infringement actions related to use of technologies in our DTH business. There can be no assurance that the courts will conclude that our services or the products used to transmit or receive our services do not infringe on the rights of third parties, that we or the manufacturers would be able to obtain licenses from these persons on commercially reasonable terms or, if we were unable to obtain such licenses, that we or the manufacturers would be able to redesign our services or the products used to transmit or receive our services to avoid infringement. The final disposition of these claims is not expected to have a material adverse effect on our consolidated financial position, but could possibly be material to our consolidated results of operations for any one period. Further, no assurance can be given that any adverse outcome would not be material to our consolidated financial position.

See "Legal Proceedings—Intellectual Property Litigation" in Part I, Item 3 of this Annual Report.

John C. Malone has significant influence over actions requiring stockholder approval and his interests may differ from ours.

The Chairman of our Parent's Board of Directors, John Malone, is also Chairman and Chief Executive of Liberty Media, Chairman of Liberty Global, Inc., and owns significant voting interests in each of DIRECTV, Liberty Media, Liberty Global, and Discovery Communications, Inc. Mr. Malone, his wife and certain trusts for the benefit of their children own shares of DIRECTV common stock, which represent approximately 24.3% of the total voting power of the outstanding shares of DIRECTV as of December 31, 2009. DIRECTV has two classes of common stock, the Class A common stock entitling holders to one vote per share and the Class B common stock entitling holders to 15 votes per share. The shares of DIRECTV Class B common stock also have certain limited consent rights with respect to certain share distributions and certain amendments to the DIRECTV Amended and Restated Certificate of Incorporation. By virtue of such rights as well as Mr. Malone's position as DIRECTV's Chairman, Mr. Malone may have significant influence over the outcome of any corporate transaction or other matters submitted to DIRECTV stockholders for approval, including the election of directors, mergers, consolidations and the sale of all or substantially all of DIRECTV's assets.

We rely on key personnel.

We believe that our future success will depend to a significant extent upon the performance of certain of our key executives. The loss of certain of our key executives could have a material adverse effect on our business, financial condition and results of operations.

Construction or launch delays on satellites could materially adversely affect our revenues and earnings.

A key component of our business strategy is our ability to expand our offering of new programming and services, including increased local and HD programming. In order to accomplish this goal, we need to construct and launch new satellites. The construction and launch of satellites are often subject to delays, including satellite and launch vehicle construction delays, periodic unavailability of reliable launch opportunities due to competition for launch slots, weather and also due to general delays that result when a launch provider experiences a launch failure, and delays in obtaining regulatory approvals. A significant delay in the future delivery of any satellite would materially adversely affect the use of the satellite and thus could materially adversely affect our anticipated revenues and earnings. If satellite construction schedules are not met, there can be no assurance that a launch opportunity will be available at the time a satellite is ready to be launched. Certain delays in satellite construction could also jeopardize a satellite authorization that is conditioned on timely construction and launch of the satellite.

Our satellites are subject to significant launch and operational risks.

Satellites are subject to significant operational risks relating to launch and while in orbit. Launch and operational risks include launch failure, incorrect orbital placement or improper commercial operation. Launch failures result in significant delays in the deployment of satellites because of the need both to construct replacement satellites, which can take up to 36 months, and obtain other launch opportunities. We estimate the overall historical loss rate for all launches of commercial satellites in the last seven years to be approximately 5% but it may be higher. Any significant delays or failures in successfully launching and deploying our satellites could materially adversely affect our ability to generate revenues. While we have traditionally purchased insurance covering the launch and, in limited cases, operation of our satellites, such policies typically cover the loss of the satellite itself or a portion thereof, and not the business interruption or other associated direct and indirect costs. For example, we purchased launch insurance covering a portion of our DIRECTV 12 satellite, which we launched at the end of 2009, and launch vehicle costs in the event of a total loss of the satellite prior to separation from the launch vehicle, but did not purchase in-orbit insurance for it.

In-orbit risks include malfunctions, commonly referred to as anomalies, and collisions with meteoroids, other spacecraft or other space debris. Anomalies occur as a result of various factors, such as satellite manufacturing errors, problems with the power systems or control systems of the satellites and general failures resulting from operating satellites in the harsh space environment. We work closely with our satellite manufacturers to determine and eliminate the potential causes of anomalies in new satellites and provide for redundancies of critical components in the satellites as well as having backup satellite capacity. However, we cannot assure you that we will not experience anomalies in the future, nor can we assure you that our backup satellite capacity will be sufficient for our business purposes. Any single anomaly or series of anomalies could materially adversely affect our operations and revenues and our relationships with our subscribers, as well as our ability to attract new subscribers for our services. Anomalies may also reduce the expected useful life of a satellite, thereby creating additional expenses due to the need to provide replacement or backup satellites and potentially reducing revenues if service is interrupted. Finally, the occurrence of anomalies may materially adversely affect our ability to insure our satellites at commercially reasonable premiums, if at all. While some anomalies are currently covered by existing insurance policies, others are not now covered or may not be covered in the future.

Our ability to earn revenue also depends on the usefulness of our satellites. Each satellite has a limited useful life. A number of factors affect the useful life of a satellite, including, among other things:

- the design;
- the quality of its construction;
- the durability of its component parts;
- the launch vehicle's insertion of the satellite into orbit;
- any required movement, temporary or permanent, of the satellite;
- the ability to continue to maintain proper orbit and control over the satellite's functions; and
- the remaining on-board fuel following orbit insertion.

Generally, the minimum design life of the satellites in our fleet is between 12 and 16 years. The actual useful lives of the satellites may be shorter or longer, in some cases significantly. Our operating results could be adversely affected if the useful life of any of our satellites were significantly shorter than 12 years from the date of launch.

In the event of a failure or loss of any of our satellites, we may relocate another satellite and use it as a replacement for the failed or lost satellite. In the event of a complete satellite failure, our services provided via that satellite could be unavailable for several days or longer while backup in-orbit satellites are repositioned and services are moved. We are not insured for any resultant lost revenues. The use of backup satellite capacity for our programming may require us to discontinue some programming services due to potentially reduced capacity on the backup satellite. Any relocation of our satellites would require prior FCC approval and, among other things, a demonstration to the FCC that the replacement satellite would not cause additional interference compared to the failed or lost satellite. Such FCC approval may not be obtained. We believe we have or will have in 2010, in-orbit satellite capacity to expeditiously recover transmission of most our programming in the event one of

Table of Contents

our in-orbit satellites fails. However, programming continuity cannot be assured in the event of multiple satellite losses.

The cost of commercial insurance coverage on our satellites or the loss of a satellite that is not insured could materially adversely affect our earnings.

We use in-orbit and launch insurance to mitigate the potential financial impact of satellite fleet in-orbit and launch failures unless the premium costs are considered uneconomic relative to the risk of satellite failure. When insurance is obtained, it generally covers all or a portion of the unamortized book value of covered satellites. Although the insurance does not compensate for business interruption or loss of future revenues or subscribers, we rely on in-orbit spare satellites and excess transponder capacity at key orbital slots to mitigate the impact that a satellite failure may have on our ability to provide service.

The price, terms and availability of insurance fluctuate significantly. Launch and in-orbit policies on satellites may not continue to be available on commercially reasonable terms or at all. In addition to higher premiums, insurance policies may provide for higher deductibles, shorter coverage periods and satellite health-related policy exclusions.

Any launch vehicle failure, or loss or destruction of any of our satellites, even if insured, could have a material adverse effect on our financial condition and results of operations, our ability to comply with FCC regulatory obligations and our ability to fund the construction or acquisition of replacement satellites in a timely fashion, or at all. At December 31, 2009, the net book value of in-orbit satellites was $1,516 million, none of which was insured.

We depend on the Communications Act for access to cable-affiliated programming and changes impacting that access could materially adversely affect us.

We purchase a substantial percentage of our programming from programmers that are affiliated with cable system operators, including key regional sports networks, or RSNs. Currently, under certain provisions of the Communications Act governing access to programming, cable-affiliated programmers generally must sell and deliver their programming services to all MVPDs on non-discriminatory terms and conditions. The Communications Act and the FCC rules also prohibit certain types of exclusive programming contracts involving programming from cable-affiliated programmers.

Any change in the Communications Act or the FCC's rules that would permit programmers that are affiliated with cable system operators to refuse to provide such programming or to impose discriminatory terms or conditions could materially adversely affect our ability to acquire programming on a cost-effective basis, or at all. The Communications Act prohibitions on certain cable industry exclusive contracting practices with cable-affiliated programmers were extended by the FCC through October 2012, though it is currently considering proposals that could shorten the term of this extension if a cable operator could show that competition from new entrant MVPDs had reached a sufficient penetration level in the relevant marketing area.

In addition, certain cable providers have denied us and other MVPDs access to a limited number of channels created by programmers with which the cable providers are affiliated. In other cases, such programmers have denied MVPDs high definition feeds of such programming. The cable providers have asserted that they are not required to provide such programming (or resolution) due to the manner in which that programming is distributed, which they argue is not covered by the program access provisions of the Communications Act. The FCC recently adopted new rules under which such programming would also be subject to certain non-exclusivity and non-discrimination requirements. These rules have not yet gone into effect, and likely will be challenged in court. In addition, they will

Table of Contents

require a further evidentiary showing by an MVPD seeking access to such programming. If these new rules are successfully challenged in court or we cannot make the required evidentiary showing, we may continue to be precluded from obtaining such programming, which in turn could materially adversely affect our ability to compete in regions serviced by those cable providers. Although the FCC also addressed some of these issues in a limited fashion by placing access conditions on certain regional sports networks affiliated with Time Warner Cable, Inc. and Comcast Corporation, it is not clear that we will be able to assure continued access to this programming on fair and nondiscriminatory terms.

DIRECTV itself is subject to similar restrictions with respect to certain programmers affiliated with us. The FCC imposed a number of conditions on its approval of Liberty Media's acquisition of News Corporation's interest in DIRECTV in 2007. Among other things, those conditions require DIRECTV to offer national and regional programming services it controls to all MVPDs on non-exclusive and non-discriminatory terms and conditions, and prohibits DIRECTV from entering into exclusive arrangements with affiliated programmers or unduly influencing such programmers in their dealings with other MVPDs. The conditions also require DIRECTV to engage in "baseball style" arbitration if elected by an MVPD where the parties cannot agree on terms and conditions for carriage of RSN programming owned, managed or controlled by DIRECTV. This condition currently applies to the three RSNs DIRECTV acquired from Liberty Media in 2009.

Changes to and implementation of statutory copyright license requirements may negatively affect our ability to deliver local and distant broadcast stations, as well as other aspects of our business.

We carry the signals of distant broadcast stations pursuant to statutory copyright licenses contained in the Satellite Home Viewer Improvement Act (SHVIA) and its successors, including the Satellite Home Viewer Extension and Reauthorization Act of 2004 (SHVERA). Critical provisions of SHVERA related to distant signals were due to expire at the end of 2009, but Congress has extended that deadline to February 28, 2010. Expiration of or changes to SHVERA, the FCC's interpretation, implementation and enforcement of provisions of SHVIA and SHVERA, as well as judicial decisions interpreting and enforcing these laws, could hamper our ability to retransmit distant network and superstation signals, reduce the number of our existing or future subscribers that can qualify for receipt of these signals, impose costs on us in connection with the process of complying with the rules, or subject us to fines, monetary damages or injunctions.

SHVERA, related laws, and FCC implementing rules also govern our provision of local broadcast signals. While those provisions of SHVERA do not expire on February 28, 2010, they may be changed by Congress. Such changes could limit our ability to deliver local broadcast signals. More generally, we have limited capacity, and the projected number of markets in which we can deliver local broadcast programming will continue to be constrained because of SHVERA's "carry-one, carry-all" requirement and may be reduced depending on changes to that requirement, the FCC's interpretation of its rules in pending and future rulemaking and complaint proceedings, as well as judicial decisions interpreting must carry requirements. We may not be able to comply with these must carry rules, or compliance may mean that we are not able to use capacity that could otherwise be used for new or additional local or national programming services. In addition, the FCC has issued an increasing obligation for carriage of local digital broadcast transmissions in HD format. We may be unable to comply with this requirement in markets where we currently carry such signals without ceasing HD local service entirely in some markets, and would be precluded from launching additional markets currently planned.

In addition, the FCC has adopted rules requiring us to negotiate in good faith with broadcast stations seeking carriage outside of the mandatory carriage regime described elsewhere. The rules for "retransmission consent" negotiations, which are similar to those that have applied to broadcast stations for years, require us to comply with certain indicia of good faith negotiation, as well as to

Table of Contents

demonstrate good faith under a "totality of the circumstances" test. Failure to comply with these rules could subject us to administrative sanctions and other penalties.

Satellite programming signals have been stolen and may be stolen in the future, which could result in lost revenues and would cause us to incur incremental operating costs that do not result in subscriber acquisition.

The delivery of subscription programming requires the use of conditional access technology to limit access to programming to only those who subscribe and are authorized to view it. The conditional access system uses, among other things, encryption technology to protect the transmitted signal from unauthorized access. It is illegal to create, sell or otherwise distribute software or devices to circumvent that conditional access technology. However, theft of cable and satellite programming has been widely reported, and the access cards used in our conditional access system have been compromised in the past and could be compromised in the future.

We have undertaken various initiatives with respect to our conditional access system to further enhance the security of the DIRECTV signal. To help combat signal theft, we provide our subscribers with more advanced access cards that we believe significantly enhance the security of our signal. Currently, we believe these access cards have not been compromised. However, we cannot guarantee that those advanced access cards will prevent the theft of our satellite programming signals in the future. Furthermore, there can be no assurance that we will succeed in developing the technology we need to effectively restrict or eliminate signal theft. If our current access cards are compromised, our revenue and our ability to contract for video and audio services provided by programmers could be materially adversely affected. In addition, our operating costs could increase if we attempt to implement additional measures to combat signal theft.

The ability to maintain FCC licenses and other regulatory approvals is critical to our business.

If we do not obtain all requisite U.S. regulatory approvals for the construction, launch and operation of any of our existing or future satellites for the use of frequencies at the orbital locations planned for these satellites or for the provision of service, or the licenses obtained impose operational restrictions on us, our ability to generate revenue and profits could be materially adversely affected. In addition, under certain circumstances, existing licenses are subject to revocation or modification and upon expiration, renewal may not be granted. If existing licenses are not renewed, or are revoked or materially modified, our ability to generate revenue could be materially adversely affected.

In certain cases, satellite system operators are obligated by governmental regulation and procedures of the ITU to coordinate the operation of their systems with other users of the radio spectrum in order to avoid causing interference to those other users. Coordination may require a satellite system operator to reduce power, avoid operating on certain frequencies, relocate its satellite to another orbital location and/or otherwise modify planned or existing operations. For example, the FCC has conditionally granted Spectrum Five authority to provide DBS service using frequencies assigned to it by the Government of the Netherlands from an orbital slot located halfway between slots at which we currently operate. Other operators have filed similar requests. We believe this closer proximity, if ultimately implemented, would significantly increase the risk of interference which could adversely affect the quality of service provided to our subscribers. We may not be able to successfully coordinate our satellites to the extent we are required to do so, and any modifications we make in the course of coordination, or any inability to successfully coordinate, may materially adversely affect our ability to generate revenue. In addition, the FCC is currently conducting a rulemaking proceeding to consider, among other things, the adoption of operating parameters under which such "tweener" systems would be automatically deemed coordinated.

Table of Contents

Other regulatory risks include, among others:

- the relocation of satellites to different orbital locations if the FCC determines that relocation is in the public interest;
- the denial by the FCC of an application to replace an existing satellite with a new satellite, or to operate a satellite beyond the term of its current authorization, or to operate an earth station to communicate with such satellite;
- the loss of authorizations to operate satellites on certain frequencies at certain locations if we do not construct, launch and operate satellites for those locations by certain dates; and
- the authorization by the United States or foreign governments of the use of frequencies by third party satellite or terrestrial facilities that have the potential to interfere with communication to or from our satellites, which could interfere with our contractual obligations or services to subscribers or other business operations.

All of our FCC satellite authorizations are subject to conditions imposed by the FCC in addition to the FCC's general authority to modify, cancel or revoke those authorizations. Use of FCC licenses and other authorizations are often subject to conditions, including technical requirements and implementation deadlines. Failure to comply with such requirements, or comply in a timely manner, could lead to the loss of authorizations and could have a material adverse effect on our ability to generate revenue. For example, loss of an authorization could potentially reduce the amount of programming and other services available to our subscribers. The materiality of such a loss of authorization would vary based upon, among other things, the orbital location at which the frequencies may be used.

Moreover, some of our authorizations and future applications may be subject to petitions and oppositions, and there can be no assurance that our authorizations will not be canceled, revoked or modified or that our applications will not be denied. Moreover, the FCC has adopted new rules for licensing satellites that may limit our ability to file applications and secure licenses in the future.

Congress has continued to shape the scope of the FCC's regulatory authority and enact legislation that affects our business. In addition, FCC proceedings to implement legislation and enact additional regulations are ongoing. The outcomes of these legislative or regulatory proceedings or their effect on our business cannot be predicted.

We control a substantial portion of interaction with our customers and we may not be as efficient or effective as our outsourced providers resulting in higher costs.

We have a number of insourced call centers and installation service providers to handle customer service calls, installations and repairs. We may not be as efficient or effective as our outsourced providers resulting in higher costs. Also, there is a risk that our customer satisfaction could be impacted, which may lead to higher subscriber churn and an inability to attract new subscribers. In addition, our outsourced providers could encounter financial difficulties, which may disrupt our ability to make installation service calls or to provide a level of customer service we expect, and which also may lead to higher subscriber churn and an inability to attract new subscribers.

We have significant debt.

We have debt totaling $6.8 billion as of December 31, 2009. If we do not have sufficient income or other sources of cash, it could affect our ability to service debt and pay other obligations.

Table of Contents

We face risks arising from possible union legislation in the United States.

There is a possibility that the proposed Employee Free Choice Act, or EFCA, may be enacted. The EFCA, also referred to as the "card check" bill, if passed in its current form could significantly change the nature of labor relations in the United States, specifically, how union elections and contract negotiations are conducted. With respect to our owned and operated home service provider installation business, it would be easier for unions to win elections and we could face arbitrator-imposed labor scheduling, costs and standards. Therefore, the EFCA could impose more labor relations requirements and union activity on our business, thereby potentially increasing our costs, and could have a material adverse effect on our overall competitive position. Currently, neither we nor most of our outsourced home service provider installation vendors have any unions.

We may not be able to obtain or retain certain foreign regulatory approvals.

There can be no assurance that any current regulatory approvals held by us are, or will remain, sufficient in the view of foreign regulatory authorities, or that any additional necessary approvals will be granted on a timely basis or at all, in all jurisdictions in which we operate, or that applicable restrictions in those jurisdictions will not be unduly burdensome. The failure to obtain the authorizations necessary to operate satellites or provide satellite service internationally could have a material adverse effect on our ability to generate revenue and our overall competitive position.

We may have a significant indemnity obligation to Liberty Media, which is not limited in amount or subject to any cap, if parts of the merger transactions are treated as a taxable transaction.

Despite obtaining a private letter ruling from the IRS and an opinion of legal counsel to the effect that parts of the merger transactions with Liberty Media qualified as a tax-free distribution for U.S. federal income tax purposes, the continuing validity of such ruling and opinion is subject to the accuracy of factual representations and certain assumptions. Any inaccuracy in such representations could invalidate the ruling or failure to comply with any undertakings made in connection with such tax opinion, could alter the conclusions reached in such opinion. Even if parts of the merger transactions otherwise qualify for tax-free treatment, it would result in a significant U.S. federal income tax liability to Liberty Media if one or more persons acquire a 50% or greater interest in the DIRECTV common stock as part of a plan or series of related transactions that includes the merger transactions within a certain time frame. The process for determining whether an acquisition is part of a plan under these rules is complex, inherently factual and subject to interpretation of the facts and circumstances of a particular case. Liberty Media or DIRECTV might inadvertently cause or permit a prohibited change in the ownership of DIRECTV to occur, thereby triggering a tax liability to Liberty Media.

In addition, Liberty Media entered into a tax matters agreement with News Corporation in connection with the News/Liberty transaction in 2008, pursuant to which Liberty Media agreed, among other things, to indemnify News Corporation and certain related persons for taxes resulting from actions taken by Liberty Media or its affiliates that cause the News/Liberty transaction (or related restructuring transactions) not to qualify as tax-free transactions. Liberty Media's indemnification obligations to News Corporation and certain related persons are not limited in amount or subject to any cap.

Under a Tax Sharing Agreement between Liberty Media and our parent, DIRECTV is obligated to indemnify Liberty Media and certain related persons for any losses and taxes resulting from the failure of the merger transactions to be tax-free transactions in certain circumstances and from any losses resulting from Liberty Media's indemnity obligations to News Corporation under the tax matters agreement between News Corporation and Liberty. If DIRECTV is required to indemnify Liberty

Media or certain related persons under the circumstances set forth in the Tax Sharing Agreement, we may be subject to substantial liabilities not limited in amount or subject to any cap. In such a circumstance, we may be required to make payments or dividends to satisfy such liabilities that could either breach covenants in our credit facilities and bond indentures or require additional or accelerated payments, which could materially adversely affect our financial position and short term operating results.

We may be required to forgo certain transactions in order to avoid the risk of incurring significant tax-related liabilities.

We might be required to forgo certain transactions that might have otherwise been advantageous in order to preserve the tax-free treatment of the News/Liberty transaction. In particular, we might be required to forgo certain transactions, including asset dispositions or other strategic transactions for some period of time following the Liberty Transaction so as not to trigger any liability under the tax indemnification obligations.

We face risks arising from the outcome of various legal proceedings.

We are involved in various legal proceedings, including those arising in the ordinary course of business and those described under the caption "Legal Proceedings" in Item 3. Such matters include investigations and legal actions by state attorneys general where regulators may seek monetary damages and may also seek to require or prohibit certain actions by the Company with regard to its current or potential customers. While we do not believe that any of these proceedings alone or in the aggregate will have a material effect on our consolidated financial position, an adverse outcome in one or more of these matters or the imposition of conditions by regulators on the conduct of our business could be material to our consolidated results of operations and cash flows for any one period. Further, no assurance can be given that any adverse outcome would not be material to our consolidated financial position.

We may face other risks described from time to time in periodic reports filed by us with the SEC.

We urge you to consider the above risk factors carefully in evaluating forward-looking statements contained in this Annual Report. The forward-looking statements included in this Annual Report are made only as of the date of this Annual Report and we undertake no obligation to publicly update these forward- looking statements to reflect subsequent events or circumstances.

ITEM 1B. UNRESOLVED STAFF COMMENTS

None.

***

ITEM 2. PROPERTIES

As of December 31, 2009, we had approximately 190 owned and leased locations operating in the United States. The major locations include eight administrative offices, two broadcast centers and six call centers.

***

DIRECTV HOLDINGS LLC

ITEM 3. LEGAL PROCEEDINGS

(a) Material pending legal proceedings, other than ordinary routine litigation incidental to the business, to which we became or were a party during the year ended December 31, 2009 or subsequent thereto, but before the filing of the report, are summarized below:

Intellectual Property Litigation. We are a defendant in several unrelated lawsuits claiming infringement of various patents relating to various aspects of our businesses. In certain of these cases other industry participants are also defendants, and also in certain of these cases we expect that any potential liability would be the responsibility of our equipment vendors pursuant to applicable contractual indemnification provisions. To the extent that the allegations in these lawsuits can be analyzed by us at this stage of their proceedings, we believe the claims are without merit and intend to defend the actions vigorously. The final disposition of these claims is not expected to have a material adverse effect on our consolidated financial position, but could possibly be material to our consolidated results of operations of any one period. No assurance can be given that any adverse outcome would not be material to our consolidated financial position.

Finisar Corporation. As previously reported, we were successful in 2008 getting the jury verdict in the Finisar case vacated on appeal. The original verdict found the patent to be valid and willfully infringed, and the jury awarded approximately $79 million in damages. The trial court increased the damages award by $25 million because of the jury finding of willful infringement and awarded pre-judgment interest of $13 million. DIRECTV was also ordered to pay into escrow $1.60 per new set-top receiver manufactured for use with the DIRECTV system beginning June 17, 2006 and continuing until the patent expires in 2012 or was otherwise found to be invalid. On April 18, 2008, the Court of Appeals reversed the verdict of the district court in part, vacated the findings of infringement, and remanded for further proceedings on the remaining issues finding that the district court had applied erroneous interpretations of certain terms of the claims. On remand, we sought and obtained summary judgment on the invalidity of all remaining claims, and the case against DIRECTV was dismissed on May 19, 2009. Finisar filed a Notice of Appeal with the Fifth Circuit Court of Appeals, and oral argument on the appeal was held on January 6, 2010. On January 8, 2010, the Court of Appeals affirmed per curiam the grant of summary judgment on all claims. This case is now resolved and there will be no further proceedings in this matter.

Early Cancellation Fees. In 2008, a number of plaintiffs filed putative class action lawsuits in state and federal courts challenging the early cancellation fees we assess our customers when they do not fulfill their programming commitments. Several of these lawsuits are pending—some in California state court purporting to represent statewide classes, and some in federal courts purporting to represent nationwide classes. The lawsuits seek both monetary and injunctive relief. While the theories of liability vary, the lawsuits generally challenge these fees under state consumer protection laws as both unfair and inadequately disclosed to customers. Each of the lawsuits is at an early stage. Where possible, we are moving to compel these cases to arbitration in accordance with our Customer Agreement, but in states such as California where the enforceability of the arbitration provision is limited, we intend to defend against these allegations in court. We believe that our early cancellation fees are adequately disclosed, and represent reasonable estimates of the costs we incur when customers cancel service before fulfilling their programming commitments.

From time to time, we receive investigative inquiries or subpoenas from state authorities with respect to alleged violations of state statutes. These inquiries may lead to legal proceedings in some cases. Currently, DIRECTV U.S. is the subject of an investigation by a multistate group of state attorneys general regarding alleged violations of their respective state consumer protection statutes. The state of Washington, originally a part of the multistate group, filed an action in Washington state court in December 2009 seeking injunctive relief and civil penalties of up to $2,000 per violation of Washington's Consumer Protection Act. The multistate investigation and the Washington lawsuit allege a variety of purported violations of the statutes, but primarily allege that we do not adequately disclose the terms and conditions of consumer offers, including subscriber commitments and early cancellation fees. We are cooperating with the multistate group by providing information about our sales and marketing practices and customer complaints. We are defending the Washington lawsuit.

Other. We are subject to other legal proceedings and claims that arise in the ordinary course of our business. The amount of ultimate liability with respect to such actions is not expected to materially affect our financial position, results of operations or liquidity.

(b) The following previously reported legal proceedings were terminated during the fourth quarter ended December 31, 2009: None.

ITEM 4. SUBMISSION OF MATTERS TO A VOTE OF SECURITY HOLDERS

None.

***

Table of Contents

DIRECTV HOLDINGS LLC

PART II

ITEM 5. MARKET FOR THE REGISTRANT'S COMMON EQUITY, RELATED STOCKHOLDER MATTERS AND ISSUER PURCHASES OF EQUITY SECURITIES

All of DIRECTV Holdings LLC's equity is indirectly owned by DIRECTV. All of DIRECTV Financing Co., Inc.'s common equity is owned by DIRECTV Holdings LLC. There is no established public trading market for our equity. Dividends on equity will be paid when and if declared by our Boards of Directors. None of our equity is subject to outstanding options or warrants.

***

ITEM 6. SELECTED FINANCIAL DATA

Omitted.

***

Table of Contents

DIRECTV HOLDINGS LLC

ITEM 7. MANAGEMENT'S DISCUSSION AND ANALYSIS OF FINANCIAL CONDITION AND RESULTS OF OPERATIONS

The following is a discussion of our results of operations and financial condition. This discussion should be read in conjunction with the consolidated financial statements and related notes included elsewhere in this Annual Report. Information in this section is organized as follows:

- Summary Results of Operations and Financial Condition
- Significant Transactions Affecting the Comparability of the Results of Operations
- Key Terminology
- Executive Overview and Outlook
- Results of Operations
- Liquidity and Capital Resources
- Contractual Obligations
- Off-Balance Sheet Arrangements
- Contingencies
- Certain Relationships and Related Party Transactions
- Critical Accounting Estimates
- Accounting Changes and New Accounting Pronouncements
- Security Ratings

Table of Contents

DIRECTV HOLDINGS LLC

SUMMARY RESULTS OF OPERATIONS AND FINANCIAL CONDITION

	Years Ended December 31,		
	2009	2008	2007
	(Dollars in Millions)		
Consolidated Statements of Operations Data:			
Revenues	$ 18,671	$ 17,310	$ 15,527
Total operating costs and expenses	16,261	14,980	13,125
Operating profit	2,410	2,330	2,402
Interest income	4	37	69
Interest expense	(348)	(315)	(216)
Other, net	(17)	5	(5)
Income before income taxes	2,049	2,057	2,250
Income tax expense	(794)	(807)	(891)
Net income	$ 1,255	$ 1,250	$ 1,359

	December 31,	
	2009	2008
	(Dollars in Millions)	
Consolidated Balance Sheet Data:		
Cash and cash equivalents	$ 1,716	$ 1,149
Total current assets	3,560	2,946
Total assets	12,408	12,546
Total current liabilities	3,388	3,006
Long-term debt	6,500	5,725
Total owner's equity	1,451	2,647

Table of Contents

DIRECTV HOLDINGS LLC

	Years Ended December 31, 2009	2008	2007
	(Dollars in Millions)		
Other Data:			
Operating profit before depreciation and amortization[1]			
Operating profit	$ 2,410	$ 2,330	$ 2,402
Add: Depreciation and amortization expense	2,275	2,061	1,448
Operating profit before depreciation and amortization	$ 4,685	$ 4,391	$ 3,850
Operating profit before depreciation and amortization margin	25.1%	25.4%	24.8%
Cash flow information			
Net cash provided by operating activities	$ 3,691	$ 3,277	$ 2,909
Net cash used in investing activities	(1,496)	(1,857)	(2,335)
Net cash used in financing activities	(1,628)	(1,073)	(1,128)
Free cash flow[2]			
Net cash provided by operating activities	$ 3,691	$ 3,277	$ 2,909
Less: Cash paid for property and equipment	(443)	(501)	(621)
Less: Cash paid for subscriber leased equipment—subscriber acquisitions	(564)	(599)	(762)
Less: Cash paid for subscriber leased equipment—upgrade and retention	(419)	(537)	(774)
Less: Cash paid for satellites	(59)	(128)	(169)
Free cash flow	$ 2,206	$ 1,512	$ 583

Reference should be made to the notes to the Consolidated Financial Statements.

(1) Operating profit before depreciation and amortization, which is a financial measure that is not determined in accordance with accounting principles generally accepted in the United States of America, or GAAP, can be calculated by adding amounts under the caption "Depreciation and amortization expense" to "Operating profit." This measure should be used in conjunction with GAAP financial measures and is not presented as an alternative measure of operating results, as determined in accordance with GAAP. Our management and DIRECTV use operating profit before depreciation and amortization to evaluate the operating performance of our company and our business segments and to allocate resources and capital to business segments. This metric is also used as a measure of performance for incentive compensation purposes and to measure income generated from operations that could be used to fund capital expenditures, service debt or pay taxes. Depreciation and amortization expense primarily represents an allocation to current expense of the cost of historical capital expenditures and for acquired intangible assets resulting from prior business acquisitions. To compensate for the exclusion of depreciation and amortization expense from operating profit, our management and DIRECTV separately measure and budget for capital expenditures and business acquisitions.

We believe this measure is useful to investors, along with GAAP measures (such as revenues, operating profit and net income), to compare our operating performance to other communications, entertainment and media service providers. We believe that investors use current and projected operating profit before depreciation and amortization and similar measures to estimate our current or prospective enterprise value and make investment decisions. This metric provides investors with a means to compare operating results exclusive of depreciation and amortization expense. Our management believes this is useful given the significant variation in depreciation and amortization expense that can result from the timing of capital expenditures, the capitalization of intangible assets, potential variations in expected useful lives when compared to other companies and periodic changes to estimated useful lives.

Operating profit before depreciation and amortization margin is calculated by dividing operating profit before depreciation and amortization by Revenues.

Table of Contents

(2) Free cash flow, which is a financial measure that is not determined in accordance with GAAP, can be calculated by deducting amounts under the captions "Cash paid for property and equipment," "Cash paid for subscriber leased equipment—subscriber acquisitions," "Cash paid for subscriber leased equipment—upgrade and retention" and "Cash paid for satellites" from "Net cash provided by operating activities" from the Consolidated Statements of Cash Flows. This financial measure should be used in conjunction with other GAAP financial measures and is not presented as an alternative measure of cash flows from operating activities, as determined in accordance with GAAP. Our management and DIRECTV use free cash flow to evaluate the cash generated by our current subscriber base, net of capital expenditures, for the purpose of allocating resources to activities such as adding new subscribers, retaining and upgrading existing subscribers, for additional capital expenditures and other capital investments or transactions and as a measure of performance for incentive compensation purposes. We believe this measure is useful to investors, along with other GAAP measures (such as cash flows from operating and investing activities), to compare our operating performance to other communications, entertainment and media companies. We believe that investors also use current and projected free cash flow to determine the ability of revenues from our current and projected subscriber base to fund required and discretionary spending and to help determine our financial value.

SIGNIFICANT TRANSACTIONS AFFECTING THE COMPARABILITY OF THE RESULTS OF OPERATIONS

Acquisitions

180 Connect. In July 2008, we acquired 100% of 180 Connect's outstanding common stock and exchangeable shares. Simultaneously, in a separate transaction, UniTek USA, LLC acquired 100% of 180 Connect's cable service operating unit and operations in certain of our installation services markets in exchange for satellite installation operations in certain markets and $7 million in cash. These transactions provide us with control over a significant portion of our home service provider network. We paid $91 million in cash, net of the $7 million we received from UniTek USA, for the acquisition, including the equity purchase price, repayment of assumed debt and related transaction costs.

Other Developments

In addition to the item described above, the following items had a significant effect on the comparability of our operating results and financial position as of and for the years ended December 31, 2009, 2008 and 2007:

Lease Program. On March 1, 2006, we introduced a new set-top receiver lease program. Prior to March 1, 2006, we expensed most set-top receivers provided to new and existing subscribers upon activation as a subscriber acquisition or upgrade and retention cost in the Consolidated Statements of Operations. Subsequent to the introduction of our lease program, we lease most set-top receivers provided to new and existing subscribers, and therefore capitalize the set-top receivers in "Property and equipment, net" in the Consolidated Balance Sheets.

The following table sets forth the amount of set-top receivers we capitalized, and depreciation expense we recorded under the lease program for the years ended December 31:

	2009	2008	2007
	(Dollars in Millions)		
Capitalized subscriber leased equipment:			
Cash paid for subscriber leased equipment—subscriber acquisitions	$ 564	$ 599	$ 762
Cash paid for subscriber leased equipment—upgrade and retention	419	537	774
Total subscriber leased equipment capitalized	$ 983	$ 1,136	$ 1,536
Depreciation expense—subscriber leased equipment	$ 1,333	$ 1,100	$ 645

DIRECTV HOLDINGS LLC

Financing Transactions. On September 22, 2009, we issued $1 billion in five year 4.750% senior notes due in 2014 at a 0.3% discount resulting in $997 million of proceeds. We also issued $1 billion in 10 year 5.875% senior notes due in 2019 at a 0.7% discount resulting in $993 million of proceeds.

On September 22, 2009, we purchased, pursuant to a tender offer, $583 million of our then outstanding $910 million 8.375% senior notes at a price of 103.125% plus accrued and unpaid interest, for a total of $603 million. On September 23, 2009, we exercised its right to redeem the remaining $327 million of the 8.375% senior notes at a price of 102.792% plus accrued and unpaid interest. We redeemed the remaining 8.375% senior notes on October 23, 2009 for a total of $339 million.

The purchase of our 8.375% senior notes resulted in a 2009 pre-tax charge of $34 million, $21 million after tax, of which $29 million resulted from the premium paid for redemption of our 8.375% senior notes and $5 million resulted from the write-off of deferred debt issuance costs and other transaction costs. The charge was recorded in "Other, net" in our Consolidated Statements of Operations.

In May 2008, we issued $1.5 billion in senior notes and amended our senior secured credit facility to include a new $1.0 billion Term Loan C. The senior notes bear interest at a rate of 7.625% and the principal balance is due in May 2016. The Term Loan C currently bears interest at a rate of 5.25% and was issued at a 1% discount. Principal payments on the Term Loan C began on September 30, 2008. The principal is payable in installments with the final installment due in April 2013.

KEY TERMINOLOGY

Revenues. We earn revenues mostly from monthly fees we charge subscribers for subscriptions to basic and premium channel programming, HD programming and access fees, pay-per-view programming, and seasonal and live sporting events. We also earn revenues from monthly fees that we charge subscribers with multiple non-leased set-top receivers (which we refer to as mirroring fees), monthly fees we charge subscribers for leased set-top receivers, monthly fees we charge subscribers for digital video recorder, or DVR, service, hardware revenues from subscribers who lease or purchase set-top receivers from us, our published programming guide, warranty service fees and advertising services.

Broadcast Programming and Other These costs primarily include license fees for subscription service programming, pay-per-view programming, live sports and other events. Other costs include expenses associated with the publication and distribution of our programming guide, continuing service fees paid to third parties for active subscribers, warranty service costs and production costs for on-air advertisements we sell to third parties.*Subscriber Service Expenses.* Subscriber service expenses include the costs of customer call centers, billing, remittance processing and certain home services expenses, such as in-home repair costs.

Broadcast Operations Expenses. These expenses include broadcast center operating costs, signal transmission expenses (including costs of collecting signals for our local channel offerings), and costs of monitoring, maintaining and insuring our satellites. Also included are engineering expenses associated with deterring theft of our signal.

Subscriber Acquisition Costs. These costs include the cost of set-top receivers and other equipment, commissions we pay to national retailers, independent satellite television retailers, dealers, telcos, and the cost of installation, advertising, marketing and customer call center expenses associated with the acquisition of new subscribers. Set-top receivers leased to new subscribers are capitalized in "Property and equipment, net" in the Consolidated Balance Sheets and depreciated over their useful lives. The amount of set-top receivers capitalized each period for subscriber acquisitions is included in

Table of Contents

"Cash paid for subscriber leased equipment- subscriber acquisitions" in the Consolidated Statements of Cash Flows.

Upgrade and Retention Costs. Upgrade and retention costs are associated with upgrade efforts for existing subscribers that we believe will result in higher average monthly revenue per subscriber, or ARPU, and lower churn. Our upgrade efforts include subscriber equipment upgrade programs for DVR, HD and HD DVR receivers and local channels, our multiple set-top receiver offer and similar initiatives. Retention costs also include the costs of installing and providing hardware under our movers program for subscribers relocating to a new residence. Set-top receivers leased to existing subscribers under upgrade and retention programs are capitalized in "Property and equipment, net" in the Consolidated Balance Sheets and depreciated over their useful lives. The amount of set-top receivers capitalized each period for upgrade and retention programs is included in "Cash paid for subscriber leased equipment- upgrade and retention" in the Consolidated Statements of Cash Flows.

General and Administrative Expenses. General and administrative expenses include departmental costs for legal, administrative services, finance, marketing and information technology. These costs also include expenses for bad debt and other operating expenses, such as legal settlements, and gains or losses from the sale or disposal of fixed assets.

Average Monthly Revenue Per Subscriber. We calculate ARPU by dividing average monthly revenues for the period (total revenues during the period divided by the number of months in the period) by average subscribers for the period. We calculate average subscribers for the period by adding the number of subscribers as of the beginning of the period and for each quarter end in the current year or period and dividing by the sum of the number of quarters in the period plus one.

Average Monthly Subscriber Churn. Average monthly subscriber churn represents the number of subscribers whose service is disconnected, expressed as a percentage of the average total number of subscribers. We calculate average monthly subscriber churn by dividing the average monthly number of disconnected subscribers for the period (total subscribers disconnected, net of reconnects, during the period divided by the number of months in the period) by average subscribers for the period.

Subscriber Count. The total number of subscribers represents the total number of subscribers actively subscribing to our service, including seasonal subscribers, subscribers who are in the process of relocating and commercial equivalent viewing units. In March 2008, we implemented a change in our commercial pricing and packaging to increase our competitiveness. As a result, during the first quarter of 2008, we made a one-time downward adjustment to the subscriber count of approximately 71,000 subscribers related to commercial equivalent viewing units.

SAC. We calculate SAC, which represents total subscriber acquisition costs stated on a per subscriber basis, by dividing total subscriber acquisition costs for the period by the number of gross new subscribers acquired during the period. We calculate total subscriber acquisition costs for the period by adding together "Subscriber acquisition costs" expensed during the period and the amount of cash paid for equipment leased to new subscribers during the period.

EXECUTIVE OVERVIEW AND OUTLOOK

The United States is continuing to undergo a period of substantial economic uncertainty. As discussed in "Competition" in Item 1, in addition to cable and satellite system operators, we are experiencing increasing competition from telcos and other emerging digital media distribution providers. A more severe downturn in economic activity or further competitive pressures could have a detrimental impact on our forecasted revenue, operating margins, net subscriber additions, free cash flow and net income. Please refer to "Risk Factors" in Item 1A for a further discussion of risks which may affect forecasted results or our business generally.

Our revenue growth is generated by both increases in the average monthly rates we earn from subscribers, or ARPU, and increases in the total number of subscribers. In 2010, we expect revenue growth in the mid-to-high single digit percentage range. We anticipate higher ARPU growth in 2010, offset by a reduction in revenue growth from net subscriber additions compared to 2009.

In 2010, as a result of the anticipated growth in revenues, the economies of scale in our business, and lower gross subscriber additions, we expect operating profit before depreciation and amortization growth in the low-teens percentage range, an anticipated improvement compared to the 6.7% operating profit before depreciation and amortization growth achieved in 2009.

In 2010, we expect capital expenditures to approximate capital expenditures reported for 2009.

Table of Contents

DIRECTV HOLDINGS LLC

Year Ended December 31, 2009 Compared with the Year Ended December 31, 2008

The following table provides operating results and a summary of key subscriber data:

	2009	2008	Change $	Change %
	(Dollars in Millions, Except Per Subscriber Amounts)			
Revenues	$ 18,671	$ 17,310	$ 1,361	7.9%
Operating costs and expenses				
Costs of revenues, exclusive of depreciation and amortization expense				
Broadcast programming and other	8,027	7,424	603	8.1%
Subscriber service expenses	1,268	1,139	129	11.3%
Broadcast operations expenses	274	265	9	3.4%
Selling, general and administrative expenses, exclusive of depreciation and amortization expense				
Subscriber acquisition costs	2,478	2,191	287	13.1%
Upgrade and retention costs	1,045	1,027	18	1.8%
General and administrative expenses	894	873	21	2.4%
Depreciation and amortization expense	2,275	2,061	214	10.4%
Total operating costs and expenses	16,261	14,980	1,281	8.6%
Operating profit	$ 2,410	$ 2,330	$ 80	3.4%
Other data:				
Operating profit before depreciation and amortization	$ 4,685	$ 4,391	$ 294	6.7%
Total number of subscribers (000's)(1)	18,560	17,621	939	5.3%
ARPU	$ 85.48	$ 83.90	$ 1.58	1.9%
Average monthly subscriber churn %	1.53%	1.47%	—	4.1%
Gross subscriber additions (000's)	4,273	3,904	369	9.5%
Subscriber disconnections (000's)	3,334	3,043	291	9.6%
Net subscriber additions (000's)	939	861	78	9.1%
Average subscriber acquisition costs—per subscriber (SAC)	$ 712	$ 715	$ (3)	(0.4)%

(1) As discussed above in "Key Terminology," during 2008, we had a one-time downward adjustment to our subscriber count of approximately 71,000 subscribers related to commercial equivalent viewing units. This adjustment did not affect our revenue, operating profit, cash flows, net subscriber additions or average monthly subscriber churn.

Subscribers. In 2009, gross subscriber additions increased primarily due to more aggressive promotions, marketing of the AT&T/DIRECTV bundle which began in February 2009, higher demand for advanced services and the impact of the transition to digital programming by broadcasters in the first half of 2009. Net subscriber additions increased from 2008 primarily due to the increase in gross additions, partially offset by higher subscriber disconnections due to a higher average monthly churn rate on a larger subscriber base. Average monthly subscriber churn increased primarily due to stricter upgrade and retention policies for existing customers as well as more aggressive competitor promotions combined with a weaker economy.

Revenues. Our revenues increased as a result of the larger subscriber base and higher ARPU. The increase in ARPU resulted primarily from price increases on programming packages, higher HD and

DVR product penetration, partially offset by more competitive customer promotions, the elimination of satellite lease revenue and lower premium movie package buy rates.

Operating profit before depreciation and amortization. The improvement of operating profit before depreciation and amortization was primarily due to the gross profit generated from the higher revenues, partially offset by higher subscriber acquisition costs principally related to the increase in gross subscriber additions.

Broadcast programming and other costs increased due to the larger number of subscribers in 2009 and annual program supplier rate increases. Subscriber service expenses increased primarily due to a larger subscriber base in 2009 and costs associated with service quality improvement initiatives.

Subscriber acquisition costs increased primarily due to an increase in gross subscriber additions compared to 2008 and increased marketing and advertising costs. SAC per subscriber, which includes the cost of capitalized set-top receivers, decreased primarily due to lower set-top receiver costs and greater savings related to the increased usage of refurbished set-top receivers through our lease program.

Upgrade and retention costs increased in 2009 primarily due to the larger subscriber base, partially offset by decreased installation costs and decreased spending on other programs due to stricter spending policies.

General and administrative expenses increased in 2009 primarily due to increased labor and benefit expense from the increase in headcount within our owned and operated home service provider installation business, partially offset by a $14 million charge in 2008 for the write-off of accounts receivable for equipment and other costs incurred to effect the orderly transition of services from one of our home service providers that ceased operations.

Operating profit. The increase in operating profit was primarily due to higher operating profit before depreciation and amortization, partially offset by higher depreciation and amortization expense in 2009 resulting from the capitalization of set-top receivers under the lease program.

Interest income and expense. The decrease in interest income to $4 million in 2009 from $37 million in 2008 was due to lower interest rates, partially offset by a higher average cash balance. The increase in interest expense to $348 million in 2009 from $315 million in 2008 was due to an increase in the average debt balance compared to 2008, partially offset by decreased interest rates.

Income Tax Expense. We recognized income tax expense of $794 million in 2009 compared to $807 million in 2008. The lower income tax expense in 2009 is primarily attributable to the lower income before income taxes

Table of Contents

DIRECTV HOLDINGS LLC

Year Ended December 31, 2008 Compared with the Year Ended December 31, 2007

The following table provides operating results and a summary of key subscriber data:

	2008	2007	Change $	Change %
	(Dollars in Millions, Except Per Subscriber Amounts)			
Revenues	$ 17,310	$ 15,527	$ 1,783	11.5%
Operating costs and expenses				
Costs of revenues, exclusive of depreciation and amortization expense				
Broadcast programming and other	7,424	6,681	743	11.1%
Subscriber service expenses	1,139	1,137	2	0.2%
Broadcast operations expenses	265	216	49	22.7%
Selling, general and administrative expenses, exclusive of depreciation and amortization expense				
Subscriber acquisition costs	2,191	1,901	290	15.3%
Upgrade and retention costs	1,027	958	69	7.2%
General and administrative expenses	873	784	89	11.4%
Depreciation and amortization expense	2,061	1,448	613	42.3%
Total operating costs and expenses	14,980	13,125	1,855	14.1%
Operating profit	$ 2,330	$ 2,402	$ (72)	(3.0)%
Other data:				
Operating profit before depreciation and amortization	$ 4,391	$ 3,850	$ 541	14.1%
Total number of subscribers (000's)[(1)]	17,621	16,831	790	4.7%
ARPU	$ 83.90	$ 79.05	$ 4.85	6.1%
Average monthly subscriber churn %	1.47%	1.51%	—	(2.6)%
Gross subscriber additions (000's)	3,904	3,847	57	1.5%
Subscriber disconnections (000's)	3,043	2,969	74	2.5%
Net subscriber additions (000's)	861	878	(17)	(1.9)%
Average subscriber acquisition costs—per subscriber (SAC)	$ 715	$ 692	$ 23	3.3%

(1) As discussed above in "Key Terminology," during 2008, we had a one-time downward adjustment to our subscriber count of approximately 71,000 subscribers related to commercial equivalent viewing units. This adjustment did not affect our revenue, operating profit, cash flows, net subscriber additions or average monthly subscriber churn.

Subscribers. In 2008, gross subscriber additions increased primarily due to growth in the direct sales and retail distribution channels due in large part to more attractive promotions and higher demand for HD and DVR services, partially offset by the loss of a distribution relationship with BellSouth during the last three quarters of 2008. Average monthly subscriber churn decreased primarily due to increased sales of HD and DVR services as well as from lower involuntary churn associated with the continued effect of stringent credit policies. Net subscriber additions declined slightly from 2007 as the increase in gross additions was offset by higher subscriber disconnections.

Revenues. Revenues increased as a result of higher ARPU and the larger subscriber base. The increase in ARPU resulted primarily from price increases on programming packages, higher HD and DVR service fees, and an increase in lease fees due to higher average number of receivers per subscriber, partially offset by more competitive customer promotions and slightly lower upfront equipment fees.

Operating profit before depreciation and amortization. The improvement of operating profit before depreciation and amortization was primarily due to the gross profit generated from the higher revenues, partially offset by higher subscriber acquisition, upgrade and retention costs for the increased number of new and existing customers adding HD and DVR services, as well as increased general and administrative costs.

Broadcast programming and other costs increased due to annual program supplier rate increases and the larger number of subscribers in 2008. Subscriber service expenses remained essentially flat with a larger subscriber base in 2008 due to the cost savings from a decline in customer call volume and a lower call handle time. Broadcast operations expense increased in 2008 due primarily to costs to support advanced services, HD enhancements and VOD.

Subscriber acquisition costs increased due to higher sales, marketing and advertising costs and higher costs associated with the acquisition of higher quality and advanced product customers. SAC per subscriber, which includes the cost of capitalized set-top receivers, increased due to higher sales, marketing and advertising costs and higher costs associated with the acquisition of higher quality and advanced product customers, partially offset by lower set-top receiver costs.

Upgrade and retention costs increased in 2008 due to an increase in the movers program and other marketing programs.

General and administrative expenses increased in 2008 primarily due to a $25 million one-time gain recognized in the second quarter of 2007 related to hurricane insurance recoveries, a $14 million charge in 2008 for the write-off of accounts receivable for equipment and other costs incurred to effect the orderly transition of services from one of our home service providers that ceased operations, $24 million in charges associated with the settlement of multiple legal proceedings and an increase in labor and benefit costs.

Operating profit. The increase in operating profit was primarily due to higher operating profit before depreciation and amortization, partially offset by higher depreciation and amortization expense in 2008 resulting from the capitalization of set-top receivers under the lease program.

Interest income and expense. The decrease in interest income to $37 million in 2008 from $69 million in 2007 was due to lower interest rates and lower average cash balances due mostly from the remittance of $3.4 billion in dividends to our Parent. The increase in interest expense to $315 million in 2008 from $216 million in 2007 was due to an increase in the average debt balance compared to 2007 and lower capitalization of interest cost in 2008. We capitalized $18 million of interest costs in 2008 and $51 million in 2007. The reduction in the capitalization of interest costs was due to the successful completion and launch of two satellites.

Income Tax Expense. We recognized income tax expense of $807 million in 2008 compared to $891 million in 2007. The lower income tax expense in 2008 is primarily attributable to the lower income before income taxes.

LIQUIDITY AND CAPITAL RESOURCES

Our principal sources of liquidity are our cash, cash equivalents and the cash flow that we generate from our operations. From 2007 to 2009 we experienced significant growth in net cash provided by operating activities and free cash flow. We expect net cash provided by operating activities and free cash flow to continue to grow and believe that our existing cash balances and cash provided by operations will be sufficient to fund our existing business plan. Additionally, as of December 31, 2009, we had the ability to borrow up to $500 million under our existing credit facility, which is available until 2011. Borrowings under this facility may be required to fund strategic investment opportunities should they arise.

At December 31, 2009, our cash and cash equivalents totaled $1.7 billion compared with $1.1 billion at December 31, 2008.

As a measure of liquidity, the current ratio (ratio of current assets to current liabilities) was 1.05 at December 31, 2009 and 0.98 at December 31, 2008. Working capital increased by $232 million to $172 million at December 31, 2009 from a working capital deficit of $60 million at December 31, 2008. The increase during the period was mostly due to the increase in our cash and cash equivalent balances resulting from the changes discussed below.

Summary Cash Flow Information

	Years Ended December 31,		
	2009	2008	2007
	(Dollars in Millions)		
Net cash provided by operating activities	$ 3,691	$ 3,277	$ 2,909
Net cash used in investing activities	(1,496)	(1,857)	(2,335)
Net cash used in financing activities	(1,628)	(1,073)	(1,128)
Free cash flow:			
Net cash provided by operating activities	$ 3,691	$ 3,277	$ 2,909
Less: Cash paid for property and equipment	(443)	(501)	(621)
Less: Cash paid for subscriber leased equipment—subscriber acquisitions	(564)	(599)	(762)
Less: Cash paid for subscriber leased equipment—upgrade and retention	(419)	(537)	(774)
Less: Cash paid for satellites	(59)	(128)	(169)
Free cash flow	$ 2,206	$ 1,512	$ 583

Cash Flows Provided By Operating Activities

The increases in net cash provided by operating activities in 2009 and 2008 were primarily due to our higher operating profit before depreciation and amortization, which resulted from the higher gross profit generated from an increase in revenues and lower payments for income taxes in 2009.

Cash Flows Used In Investing Activities

During both 2008 and 2009, we experienced a reduction in set-top receiver costs and benefited from the use of refurbished set-top receivers from our lease program, which resulted in a reduction in capital expenditures for property and equipment in 2008 and 2009.

Also during 2007, 2008 and 2009, we were in the process of constructing three satellites. We have completed and placed two of these satellites into service, which resulted in decreasing satellite capital expenditures over the three year period. We expect to place the last of these satellites in service in the second quarter of 2010. Additionally, our capital expenditures for broadcast facilities and equipment to support our HD programming has decreased from 2007 to 2009 as we have largely completed the build out of the infrastructure necessary to launch HD programming both locally and nationally.

Additionally, we paid $11 million in 2009 and $97 million in 2008 for investments, net of cash acquired, in various companies. Significant acquisistions are described in Note 12 of the Notes to the

Consolidated Financial Statements in Part II, Item 8 of this Annual Report. Our cash spending on investments is discretionary and we may fund strategic investment opportunities should they arise in the future.

Cash Flows Used in Financing Activities

During 2009, we had $1,990 million of net cash proceeds from the issuance of senior notes which were completed in September 2009. We also repaid $1,018 million of our long-term debt. During 2008, we had $2,490 million of net cash proceeds from the issuance of senior notes and borrowings under our senior secured credit facility which were completed in May 2008 as described in Note 7 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report. Also, we paid cash dividends to our Parent in the amounts of $2,500 million in 2009, $3,400 million in 2008 and $1,050 million in 2007. We have provided $1,500 million in additional dividends to our Parent in 2010 and will likely continue to provide additional dividends to our Parent to fund its cash requirements, including the $3.5 billion share repurchase program announced February 18, 2010 or other distributions to shareholders, or to fund strategic investment opportunities should they arise. We may use available cash and cash equivalents, cash from operations, or incur additional borrowings, which may include borrowings under our $500 million revolving credit facility, to fund such dividends.

Free Cash Flow

Free cash flow increased from 2007 to 2009 due to an increase in net cash provided by operating activities described above, and the decrease in capital expenditures. The decrease in capital expenditures resulted from lower costs for set-top receivers capitalized under our lease program and lower capital expenditures for satellite and broadcast facilities and equipment to support HD programming.

During 2010, we expect continued free cash flow growth primarily as a result of the anticipated increase in operating profit before depreciation and amortization.

Debt

At December 31, 2009, we had $6,808 million in total outstanding borrowings, bearing a weighted average interest rate of 5.2%. Our outstanding borrowings primarily consist of notes payable and amounts borrowed under our senior secured credit facility as more fully described in Note 7 of the Notes to the Consolidated Financial Statements in Item 8, Part II of this Annual Report, which we incorporate herein by reference.

Our notes payable and senior secured credit facility mature as follows: $308 million in 2010, $108 million in 2011, $20 million in 2012, $1,887 million in 2013, $1,000 million in 2014 and $3,500 million thereafter. These amounts do not reflect potential prepayments that may be required under our senior secured credit facility, which could result from a computation that we are required to make at each year end under the credit agreement. We were not required to make a prepayment for the years ended December 31, 2009 and 2008.

Covenants and Restrictions. The senior secured credit facility requires us to comply with certain financial covenants. The senior notes and the senior secured credit facility also include covenants that restrict our ability to, among other things, (i) incur additional indebtedness, (ii) incur liens, (iii) pay dividends or make certain other restricted payments, investments or acquisitions, (iv) enter into certain transactions with affiliates, (v) merge or consolidate with another entity, (vi) sell, assign, lease or otherwise dispose of all or substantially all of its assets, and (vii) make voluntary prepayments of certain debt, in each case subject to exceptions as provided in the credit agreement and senior notes

Table of Contents

indentures. Our 4.750% and 5.875% senior notes are rated as investment grade and have fewer covenants and restrictions than our other senior notes. Should we fail to comply with these covenants, all or a portion of its borrowings under the senior notes and senior secured credit facility could become immediately payable and its revolving credit facility could be terminated. At December 31, 2009, we were in compliance with all such covenants and we expect to continue to be in compliance with all covenants in 2010.

Contingencies

Several factors may affect our ability to fund our operations and commitments that we discuss in "Contractual Obligations", "Off-Balance Sheet Arrangements" and "Contingencies" below. In addition, our future cash flows may be reduced if we experience, among other things, significantly higher subscriber additions than planned, increased subscriber churn or upgrade and retention costs, higher than planned capital expenditures for satellites and broadcast equipment, satellite anomalies or signal theft or if we are required to make a prepayment on our term loans under our senior secured credit facility. Additionally, our ability to borrow under the senior secured credit facility is contingent upon us meeting financial and other covenants associated with its facility as more fully described above.

CONTRACTUAL OBLIGATIONS

The following table sets forth our contractual obligations as of December 31, 2009, including the future periods in which payments are expected. Additional details regarding these obligations are provided in the Notes to the Consolidated Financial Statements in Part II, Item 8 referenced in the table.

	Payments due by period				
Contractual Obligations	Total	Less than 1 year	1-3 years	3-5 years	More than 5 years
	(Dollars in Millions)				
Long-term debt obligations (Note 7)[a]	$ 9,102	$ 665	$ 830	$ 3,473	$ 4,134
Purchase obligations (Note 13)[b]	8,297	1,671	3,450	2,539	637
Operating lease obligations (Note 13)[c]	232	41	67	46	78
Other long-term liabilities reflected on the Consolidated Balance Sheets under GAAP (Notes 13)[(d),(e)]	198	104	71	11	12
Total	$ 17,829	$ 2,481	$ 4,418	$ 6,069	$ 4,861

(a) Long-term debt obligations include interest calculated based on the rates in effect at December 31, 2009, however, the obligations do not reflect potential prepayments that may be required under our senior secured credit facility, if any, or permitted under our indentures.

(b) Purchase obligations consist primarily of broadcast programming commitments, regional professional team rights agreements, service contract commitments and satellite launch contracts. Broadcast programming commitments include guaranteed minimum contractual commitments that are typically based on a flat fee or a minimum number of required subscribers subscribing to the related programming. Actual payments may exceed the minimum payment requirements if the actual number of subscribers subscribing to the related programming exceeds the minimum amounts. Service contract commitments include minimum commitments for the purchase of services that have been outsourced to third parties, such as billing services, telemetry, tracking and control services and broadcast center services. In most cases, actual payments, which are typically based on volume, usually exceed these minimum amounts.

(c) Certain of our operating leases contain escalation clauses and renewal or purchase options, which we do not consider in the amounts disclosed.

(d) Other long-term liabilities consist of the amounts we owe to National Rural Telecommunications Cooperative, or NRTC, for the purchase of distribution rights and to the NRTC members that elected the long-term payment option resulting from the NRTC acquisition transactions in 2004, capital lease obligations, including interest, and satellite contracts.

(e) Payments due by period for other long-term liabilities reflected on the Consolidated Balance Sheet under GAAP do not include payments that could be made related to our net unrecognized tax benefits liability, which amounted to $45 million as of December 31, 2009. The timing and amount of any future payments is not reasonably estimable, as such payments are dependent on the completion and resolution of examinations with tax authorities. We do not expect a significant payment related to these obligations within the next twelve months.

OFF-BALANCE SHEET ARRANGEMENTS

As of December 31, 2009, we were contingently liable under standby letters of credit and bonds in the aggregate amount of $2 million.

CONTINGENCIES

For a discussion of "Contingencies", see Note 13 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report, which we incorporate herein by reference.

CERTAIN RELATIONSHIPS AND RELATED-PARTY TRANSACTIONS

For a discussion of "Certain Relationships and Related-Party Transactions," see Note 11 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report, which we incorporate herein by reference.

CRITICAL ACCOUNTING ESTIMATES

The preparation of the consolidated financial statements in conformity with accounting principles generally accepted in the United States requires management to make estimates, judgments and assumptions that affect amounts reported. Management bases its estimates, judgments and assumptions on historical experience and on various other factors that are believed to be reasonable under the circumstances. Due to the inherent uncertainty involved in making estimates, actual results reported for future periods may be affected by changes in those estimates. The following represents what we believe are the critical accounting policies that may involve a higher degree of estimation, judgment and complexity. For a summary of our significant accounting policies, including those discussed below, see Note 2 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report.

Multi-Year Programming Contracts for Live Sporting Events. We charge the cost of multi-year programming contracts for live sporting events with minimum guarantee payments, such as our agreement with the NFL, based on the contractual rates in the contract per season, unless the contractual rates are inconsistent with the relative value of the programming from season to season, in which case we record the expense based on the ratio of each period's sports programming package revenues to the estimated total package revenues to be earned over the contract period. Management evaluates estimated total programming package revenues at least annually. Estimates of forecasted revenues rely on assumptions regarding the number of subscribers to a given sporting events package and the estimated package price throughout the contract. While we base our estimates on past experience and other relevant factors, actual results could differ from our estimates. If actual results were to significantly vary from forecasted amounts, the profit recorded on such contracts in a future period could vary from current rates and the resulting change in profits recorded could be material to our consolidated results of operations.

DIRECTV HOLDINGS LLC

Income Taxes. We must make certain estimates and judgments in determining provisions for income taxes. These estimates and judgments occur in the calculation of tax credits, tax benefits and deductions, and in the calculation of certain tax assets and liabilities, which arise from differences in the timing of recognition of revenue and expense for tax and financial statement purposes.

We assess the recoverability of deferred tax assets at each reporting date and where applicable, record a valuation allowance to reduce the total deferred tax asset to an amount that will, more-likely-than-not, be realized in the future. Our assessment includes an analysis of whether deferred tax assets will be realized in the ordinary course of operations based on the available positive and negative evidence, including the scheduling of deferred tax liabilities and forecasted income from operating activities. The underlying assumptions we use in forecasting future taxable income require significant judgment. In the event that actual income from operating activities differs from forecasted amounts, or if we change our estimates of forecasted income from operating activities, we could record additional charges or reduce allowances in order to adjust the carrying value of deferred tax assets to their realizable amount. Such adjustments could be material to our consolidated financial statements.

In addition, the recognition of a tax benefit for tax positions involves dealing with uncertainties in the application of complex tax regulations. Judgment is required in assessing the future tax consequences of events that have been recognized in our financial statements or tax returns. We provide for taxes for uncertain tax positions where assessments have not been received. We believe such tax reserves are adequate in relation to the potential for additional assessments. Once established, we adjust these amounts only when more information is available or when an event occurs necessitating a change to the reserves. Future events such as changes in the facts or law, judicial decisions regarding the application of existing law or a favorable audit outcome will result in changes to the amounts provided.

Contingent Matters. Determining when, or if, an accrual should be recorded for a contingent matter, including but not limited to legal and tax issues, and the amount of such accrual, if any, requires a significant amount of management judgment and estimation. We develop our judgments and estimates in consultation with outside counsel based on an analysis of potential outcomes. Due to the uncertainty of determining the likelihood of a future event occurring and the potential financial statement impact of such an event, it is possible that upon further development or resolution of a contingent matter, we could record a charge in a future period that would be material to our consolidated financial statements.

Depreciable Lives of Leased Set-Top Receivers. We currently lease most set-top receivers provided to new and existing subscribers and therefore capitalize the cost of those set-top receivers. We depreciate capitalized set-top receivers over a three year estimated useful life, which is based on, among other things, management's judgment of the risk of technological obsolescence. Changes in the estimated useful lives of set-top receivers capitalized could result in significant changes to the amounts recorded as depreciation expense. Based on the book value of the set-top receivers capitalized as of December 31, 2009, if we extended the depreciable life of the set-top receivers by one half of a year, it would result in an approximately $200 million reduction in annual depreciation expense.

Valuation of Long-Lived Assets. We evaluate the carrying value of long-lived assets to be held and used, other than goodwill and intangible assets with indefinite lives, when events and circumstances warrant such a review. We consider the carrying value of a long-lived asset impaired when the anticipated undiscounted future cash flow from such asset is separately identifiable and is less than its carrying value. In that event, we recognize a loss based on the amount by which the carrying value exceeds the fair value of the long-lived asset. We determine fair value primarily using the estimated future cash flows associated with the asset under review, discounted at a rate commensurate with the risk involved, and other valuation techniques. We determine losses on long-lived assets to be disposed of in a similar manner, except that we reduce the fair value for the cost of disposal. Changes in estimates of future cash flows could result in a write-down of the asset in a future period.

Valuation of Goodwill and Intangible Assets with Indefinite Lives. We evaluate the carrying value of goodwill and intangible assets with indefinite lives annually in the fourth quarter or more frequently when events and circumstances change that would more likely than not result in an impairment loss. We completed our annual impairment testing during the fourth quarter of 2009, and determined that there was no impairment of goodwill or intangible assets with indefinite lives. As of December 31, 2009, the fair value of the Company and our intangible assets with indefinite lives significantly exceed their carrying values. See Note 5 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report, which we incorporate herein by reference.

The goodwill evaluation requires the estimation of our business as a whole. We determine fair values primarily using estimated cash flows discounted at a rate commensurate with the risk involved, when appropriate. Estimation of future cash flows requires significant judgment about future operating results, and can vary significantly from one evaluation to the next. Risk adjusted discount rates are not fixed and are subject to change over time. As a result, changes in estimated future cash flows and/or changes in discount rates could result in a write-down of goodwill or intangible assets with indefinite lives in a future period, which could be material to our consolidated financial statements.

ACCOUNTING CHANGES AND NEW ACCOUNTING PRONOUNCEMENTS

For a discussion of accounting changes and new accounting pronouncements see Note 2 of the Notes to the Consolidated Financial Statements in Part II, Item 8 of this Annual Report, which we incorporate herein by reference.

SECURITY RATINGS

Debt ratings by the various rating agencies reflect each agency's opinion of the ability of issuers to repay debt obligations as they come due and expected loss given a default. Ratings in the Baa range for Moody's Investors Service, and the BBB range for Standard & Poor's Ratings Services, or S&P, and the BBB range for Fitch Ratings, generally indicate adequate current protection of interest payments and principal security, with certain protective elements lacking. Ratings in the Ba range for Moody's and the BB range for S&P and Fitch, generally indicate moderate protection of interest and principal payments, potentially outweighed by exposure to uncertainties or adverse conditions. In general, lower ratings result in higher borrowing costs. A security rating is not a recommendation to buy, sell, or hold securities and may be subject to revision or withdrawal at any time by the assigning rating organization.

Currently, we have the following security ratings:

	Senior Secured	Senior Unsecured	Corporate	Outlook
Standard & Poor's	BBB-	BBB-	BBB-	Stable
Moody's	Baa2	Ba2	Ba1	Stable
Fitch	BBB	BBB-	BBB-	Stable

***

Table of Contents

DIRECTV HOLDINGS LLC

ITEM 7A. QUANTITATIVE AND QUALITATIVE DISCLOSURES ABOUT MARKET RISK

The following discussion and the estimated amounts generated from the sensitivity analyses referred to below include forward-looking statements of market risk which assume for analytical purposes that certain adverse market conditions may occur. Actual future market conditions may differ materially from such assumptions and the amounts noted below are the result of analyses used for the purpose of assessing possible risks and the mitigation thereof. Accordingly, you should not consider the forward-looking statements as our projections of future events or losses.

Interest Rate Risk

We are subject to fluctuating interest rates, which may adversely impact our consolidated results of operations and cash flows. We had outstanding debt of $6,808 million at December 31, 2009, which consisted of fixed rate borrowings of $4,490 million and variable rate borrowings of $2,316 million. As of December 31, 2009, a hypothetical one percentage point increase in interest rates related to our outstanding variable rate debt would have increased our annual interest expense by approximately $23 million.

***

DIRECTV HOLDINGS LLC

ITEM 8. FINANCIAL STATEMENTS AND SUPPLEMENTARY DATA

REPORT OF INDEPENDENT REGISTERED PUBLIC ACCOUNTING FIRM

To the Board of Directors of
DIRECTV Holdings LLC
El Segundo, California

We have audited the accompanying consolidated balance sheets of DIRECTV Holdings LLC (the "Company") as of December 31, 2009 and 2008, and the related consolidated statements of operations, changes in owner's equity, and cash flows for each of the three years in the period ended December 31, 2009. Our audits also included the financial statement schedule listed in the Index at Item 15. These financial statements and financial statement schedule are the responsibility of the Company's management. Our responsibility is to express an opinion on the financial statements and financial statement schedule based on our audits.

We conducted our audits in accordance with the standards of the Public Company Accounting Oversight Board (United States). Those standards require that we plan and perform the audit to obtain reasonable assurance about whether the financial statements are free of material misstatement. An audit includes examining, on a test basis, evidence supporting the amounts and disclosures in the financial statements. An audit also includes assessing the accounting principles used and significant estimates made by management, as well as evaluating the overall financial statement presentation. We believe that our audits provide a reasonable basis for our opinion.

In our opinion, such consolidated financial statements present fairly, in all material respects, the financial position of DIRECTV Holdings LLC at December 31, 2009 and 2008, and the results of its operations and its cash flows for each of the three years in the period ended December 31, 2009, in conformity with accounting principles generally accepted in the United States of America. Also, in our opinion, such financial statement schedule, when considered in relation to the basic consolidated financial statements taken as a whole, present fairly, in all material respects, the information set forth therein.

We have also audited, in accordance with the standards of the Public Company Accounting Oversight Board (United States), the Company's internal control over financial reporting as of December 31, 2009, based on the criteria established in *Internal Control —Integrated Framework* issued by the Committee of Sponsoring Organizations of the Treadway Commission and our report dated February 25, 2010 expressed an unqualified opinion on the Company's internal control over financial reporting.

/s/ DELOITTE & TOUCHE LLP

Los Angeles, California
February 25, 2010

Table of Contents

DIRECTV HOLDINGS LLC

CONSOLIDATED STATEMENTS OF OPERATIONS

	Years Ended December 31,		
	2009	2008	2007
	(Dollars in Millions)		
Revenues	$ 18,671	$ 17,310	$ 15,527
Operating costs and expenses			
Costs of revenues, exclusive of depreciation and amortization expense			
Broadcast programming and other	8,027	7,424	6,681
Subscriber service expenses	1,268	1,139	1,137
Broadcast operations expenses	274	265	216
Selling, general and administrative expenses, exclusive of depreciation and amortization expense			
Subscriber acquisition costs	2,478	2,191	1,901
Upgrade and retention costs	1,045	1,027	958
General and administrative expenses	894	873	784
Depreciation and amortization expense	2,275	2,061	1,448
Total operating costs and expenses	16,261	14,980	13,125
Operating profit	2,410	2,330	2,402
Interest income	4	37	69
Interest expense	(348)	(315)	(216)
Other, net	(17)	5	(5)
Income before income taxes	2,049	2,057	2,250
Income tax expense	(794)	(807)	(891)
Net income	$ 1,255	$ 1,250	$ 1,359

The accompanying notes are an integral part of these Consolidated Financial Statements.

Table of Contents

DIRECTV HOLDINGS LLC

CONSOLIDATED BALANCE SHEETS

	December 31, 2009	December 31, 2008
	(Dollars in Millions)	
ASSETS		
Current assets		
Cash and cash equivalents	$ 1,716	$ 1,149
Accounts receivable, net	1,421	1,308
Inventories	200	182
Deferred income taxes	60	46
Prepaid expenses and other	163	261
Total current assets	3,560	2,946
Satellites, net	1,870	1,980
Property and equipment, net	2,998	3,348
Goodwill	3,167	3,189
Intangible assets, net	582	871
Other assets	231	212
Total assets	$ 12,408	$ 12,546
LIABILITIES AND OWNER'S EQUITY		
Current liabilities		
Accounts payable and accrued liabilities	$ 2,727	$ 2,582
Unearned subscriber revenues and deferred credits	353	316
Current portion of long-term debt	308	108
Total current liabilities	3,388	3,006
Long-term debt	6,500	5,725
Deferred income taxes	559	405
Other liabilities and deferred credits	510	763
Commitments and contingencies		
Owner's equity		
Capital stock and additional paid-in capital	1,076	2,403
Retained earnings	375	244
Total owner's equity	1,451	2,647
Total liabilities and owner's equity	$ 12,408	$ 12,546

The accompanying notes are an integral part of these Consolidated Financial Statements.

Table of Contents

DIRECTV HOLDINGS LLC

CONSOLIDATED STATEMENTS OF CHANGES IN OWNER'S EQUITY

	Capital stock and additional paid-in capital	Retained earnings	Total owner's equity
	(Dollars in Millions)		
Balance at January 1, 2007	$ 3,786	$ 602	$ 4,388
Net income		1,359	1,359
Dividend to Parent	(1,050)		(1,050)
Capital contribution from Parent	41		41
Adjustment to initially record cumulative effect of adopting accounting standard for uncertainty in income taxes, net of tax		(3)	(3)
Other	5		5
Balance at December 31, 2007	2,782	1,958	4,740
Net income		1,250	1,250
Dividend to Parent	(436)	(2,964)	(3,400)
Capital contribution from Parent	43		43
Other	14		14
Balance at December 31, 2008	2,403	244	2,647
Net income		1,255	1,255
Dividend to Parent	(1,376)	(1,124)	(2,500)
Capital contribution from Parent	44		44
Other	5		5
Balance at December 31, 2009	$ 1,076	$ 375	$ 1,451

The accompanying notes are an integral part of these Consolidated Financial Statements.

Table of Contents

DIRECTV HOLDINGS LLC

CONSOLIDATED STATEMENTS OF CASH FLOWS

	Years Ended December 31, 2009	2008	2007
	(Dollars in Millions)		
Cash Flows From Operating Activities			
Net income	$ 1,255	$ 1,250	$ 1,359
Adjustments to reconcile income from continuing operations to net cash provided by operating activities:			
Depreciation and amortization	2,275	2,061	1,448
Amortization of deferred revenues and deferred credits	(48)	(104)	(98)
Share-based compensation expense	44	43	41
Deferred income taxes	229	52	162
Other	20	11	11
Change in other operating assets and liabilities			
Accounts receivable, net	(121)	93	(120)
Inventories	(10)	22	(47)
Prepaid expenses and other	98	(109)	3
Accounts payable and accrued liabilities	(76)	(47)	138
Unearned subscriber revenues and deferred credits	33	(4)	60
Other, net	(8)	9	(48)
Net cash provided by operating activities	3,691	3,277	2,909
Cash Flows From Investing Activities			
Cash paid for property and equipment	(443)	(501)	(621)
Cash paid for subscriber leased equipment—subscriber acquisitions	(564)	(599)	(762)
Cash paid for subscriber leased equipment—upgrade and retention	(419)	(537)	(774)
Cash paid for satellites	(59)	(128)	(169)
Investment in companies, net of cash acquired	(11)	(97)	—
Other	—	5	(9)
Net cash used in investing activities	(1,496)	(1,857)	(2,335)
Cash Flows From Financing Activities			
Cash proceeds from debt issuance	1,990	2,490	—
Debt issuance costs	(14)	(19)	—
Repayment of long-term debt	(1,018)	(53)	(10)
Repayment of other long-term obligations	(90)	(98)	(72)
Cash dividends to Parent	(2,500)	(3,400)	(1,050)
Excess tax benefit from share-based compensation	4	7	4
Net cash used in financing activities	(1,628)	(1,073)	(1,128)
Net increase (decrease) in cash and cash equivalents	567	347	(554)
Cash and cash equivalents at beginning of the year	1,149	802	1,356
Cash and cash equivalents at end of the year	$ 1,716	$ 1,149	$ 802
Supplemental Cash Flow Information			
Cash paid for interest	$ 341	$ 289	$ 211
Cash paid for income taxes	529	753	730

The accompanying notes are an integral part of these Consolidated Financial Statements.

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS

Note 1: Description of Business

DIRECTV Holdings LLC is an indirect, wholly-owned subsidiary of DIRECTV and consists of DIRECTV Enterprises, LLC and its wholly-owned subsidiaries and DIRECTV Financing Co., Inc. We sometimes refer to DIRECTV Holdings LLC as DIRECTV Holdings, DIRECTV U.S., we or us and sometimes refer to DIRECTV as our Parent. We are the largest provider of direct-to-home, or DTH, digital television services and the second largest provider in the multi-channel video programming distribution, or MVPD, industry in the United States.

On November 19, 2009, The DIRECTV Group, Inc., or DIRECTV Group, and Liberty Media Corporation, which we refer to as Liberty or Liberty Media, obtained shareholder approval of and closed a series of related transactions which we refer to collectively as the Liberty Transaction. The Liberty Transaction included the split-off of certain of the assets of the Liberty Entertainment group into Liberty Entertainment, Inc., or LEI, which was then split-off from Liberty. Following the split-off, DIRECTV Group and LEI merged with subsidiaries of DIRECTV. As a result of the Liberty Transaction, DIRECTV Group, which is comprised of the DIRECTV U.S. and DIRECTV Latin America businesses, and LEI, which held Liberty's 57% interest in DIRECTV Group, a 100% interest in three regional sports networks, a 65% interest in GSN, approximately $120 million in cash and cash equivalents and approximately $2.1 billion of indebtedness and a related series of equity collars became wholly-owned subsidiaries of DIRECTV. DIRECTV Holdings remained a direct subsidiary of DIRECTV Group and became an indirect subsidiary of DIRECTV.

Note 2: Basis of Presentation and Summary of Significant Accounting Policies

Principles of Consolidation

We include the accounts of DIRECTV Holdings and our majority owned subsidiaries, after elimination of intercompany accounts and transactions, in the accompanying consolidated financial statements.

Use of Estimates in the Preparation of the Consolidated Financial Statements

We prepare our consolidated financial statements in conformity with accounting principles generally accepted in the United States of America, which requires us to make estimates and assumptions that affect amounts reported herein. We base our estimates and assumptions on historical experience and on various other factors that we believe to be reasonable under the circumstances. Due to the inherent uncertainty involved in making estimates, our actual results reported in future periods may be affected by changes in those estimates.

Revenue Recognition

We recognize subscription and pay-per-view revenues when programming is broadcast to subscribers. We recognize subscriber fees for multiple set-top receivers, our published programming guide, warranty services and equipment rental as revenue, as earned. We recognize advertising revenues when the related services are performed. We defer programming payments received from subscribers in advance of the broadcast as "Unearned subscriber revenues and deferred credits" in the Consolidated Balance Sheets until earned. We recognize revenues to be received under contractual commitments on a straight-line basis over the minimum contractual period.

Table of Contents

Broadcast Programming and Other

We recognize the costs of television programming distribution rights when we distribute the related programming. We recognize the costs of television programming rights to distribute live sporting events for a season or tournament to expense using the straight-line method over the course of the season or tournament. However, we charge the cost of multi-year programming contracts for live sporting events with minimum guarantee payments, such as DIRECTV U.S.' agreement with the NFL, based on the contractual rates in the contract per season, unless the contractual rates are inconsistent with the relative value of the programming from season to season, in which case we record the expense based on the ratio of each period's sports programming package revenues to the estimated total package revenues to be earned over the contract period. We evaluate estimated total contract revenues at least annually.

We defer advance payments in the form of cash and equity instruments from programming content providers for carriage of their signal and recognize them as a reduction of "Broadcast programming and other" in the Consolidated Statements of Operations on a straight-line basis over the related contract term. We record equity instruments at fair value based on quoted market prices or values determined by management.

Subscriber Acquisition Costs

Subscriber acquisition costs consist of costs we incur to acquire new subscribers. We include the cost of set-top receivers and other equipment, commissions we pay to national retailers, independent satellite television retailers, dealers, telephone communication companies and the cost of installation, advertising, marketing and customer call center expenses associated with the acquisition of new subscribers in subscriber acquisition costs. We expense these costs as incurred, or when subscribers activate the DIRECTV® service, as appropriate, except for the cost of set-top receivers leased to new subscribers which we capitalize in "Property and equipment, net" in the Consolidated Balance Sheets. Although paid in advance, the retailer or dealer earns substantially all commissions paid for customer acquisitions over 12 months from the date of subscriber activation. Should the subscriber cancel our service during the 12 month service period, we are reimbursed for the unearned portion of the commission by the retailer or dealer and record a decrease to subscriber acquisition costs. We include the amount of our set-top receivers capitalized each period for subscriber acquisition activities in the Consolidated Statements of Cash Flows under the caption "Cash paid for subscriber leased equipment—subscriber acquisitions." See Note 4 below for additional information.

Upgrade and Retention Costs

Upgrade and retention costs consist primarily of costs we incur for loyalty programs offered to existing subscribers. The costs for loyalty programs include the costs of installing or providing hardware under our movers program (for subscribers relocating to a new residence), multiple set-top receiver offers, digital video recorder, or DVR, high-definition, or HD, local channel upgrade programs and other similar initiatives, and third party commissions we incur for the sale of additional set-top receivers to existing subscribers. We expense these costs as incurred, except for the cost of set-top receivers leased to existing subscribers which we capitalize in "Property and equipment, net" in the Consolidated Balance Sheets. We include the amount of our set-top receivers capitalized each period for upgrade and retention activities in the Consolidated Statements of Cash Flows under the caption "Cash paid for subscriber leased equipment—upgrade and retention." See Note 4 below for additional information.

Cash and Cash Equivalents

Cash and cash equivalents consist of highly liquid investments we purchase with original maturities of three months or less.

Inventories

We state inventories at the lower of average cost or market. Inventories consist of finished goods for DIRECTV System equipment and DIRECTV System access cards.

Property and Equipment, Satellites and Depreciation

We carry property and equipment, and satellites at cost, net of accumulated depreciation. The amounts we capitalize for satellites currently being constructed and those that have been successfully launched include the costs of construction, launch, launch insurance, incentive obligations and related capitalized interest. We generally compute depreciation using the straight-line method over the estimated useful lives of the assets. We amortize leasehold improvements over the lesser of the life of the asset or term of the lease.

Goodwill and Intangible Assets

Goodwill and intangible assets with indefinite lives are carried at historical cost and are subject to write-down, as needed, based upon an impairment analysis that we must perform at least annually, or sooner if an event occurs or circumstances change that would more likely than not result in an impairment loss. We perform our annual impairment analysis in the fourth quarter of each year. If an impairment loss results from the annual impairment test, we would record the loss as a pre-tax charge to operating income.

We amortize other intangible assets using the straight-line method over their estimated useful lives, which range from 5 to 15 years.

Valuation of Long-Lived Assets

We evaluate the carrying value of long-lived assets to be held and used, other than goodwill and intangible assets with indefinite lives, when events and circumstances warrant such a review. We consider the carrying value of a long-lived asset impaired when the anticipated undiscounted future cash flow from such asset is separately identifiable and is less than its carrying value. In that event, we would recognize a loss based on the amount by which the carrying value exceeds the fair value of the long-lived asset. We determine fair value primarily using estimated future cash flows associated with the asset under review, discounted at a rate commensurate with the risk involved, or other valuation techniques. We determine losses on long-lived assets to be disposed of in a similar manner, except that we reduce the fair value for the cost of disposal.

Investments and Financial Instruments

We account for investments in which we own at least 20% of the voting securities or have significant influence under the equity method of accounting. We record equity method investments at cost and adjust for the appropriate share of the net earnings or losses of the investee. We record investee losses up to the amount of the investment plus advances and loans made to the investee, and financial guarantees made on behalf of the investee.

The carrying value of cash and cash equivalents, accounts receivable, investments and other assets, accounts payable, and amounts included in accrued liabilities and other meeting the definition of a financial instrument approximated their fair values at December 31, 2009 and 2008.

Debt Issuance Costs

We defer costs we incur to issue debt and amortize these costs to interest expense using the straight-line method over the term of the respective obligation.

Share-Based Payment

DIRECTV grants restricted stock units and common stock options to our employees.

We record compensation expense equal to the fair value of stock-based awards at the date approved on a straight-line basis over the requisite service period of up to three years, reduced for estimated forfeitures and adjusted for anticipated payout percentages related to the achievement of performance targets.

Income Taxes

We join in the filing of DIRECTV's consolidated U.S. federal income tax return. We determine our income taxes based upon our tax sharing agreement with our Parent, which generally provides that the current income tax liability or receivable be computed as if we were a separate taxpayer.

We determine deferred tax assets and liabilities based on the difference between the financial statement and tax basis of assets and liabilities, using enacted tax rates in effect for the year in which we expect the differences to reverse. We must make certain estimates and judgments in determining income tax provisions, assessing the likelihood of recovering our deferred tax assets, and evaluating tax positions.

We recognize a benefit in "Income tax expense" in the Consolidated Statements of Operations for uncertain tax positions that are more-likely-than-not to be sustained upon examination, measured at the largest amount that has a greater than 50% likelihood of being realized upon settlement. Unrecognized tax benefits represent tax benefits taken or expected to be taken in income tax returns, for which the benefit has not yet been recognized in "Income tax expense" in the Consolidated Statements of Operations due to the uncertainty of whether such benefits will be ultimately realized. We recognize interest and penalties accrued related to unrecognized tax benefits in "Income tax expense" in the Consolidated Statements of Operations. Unrecognized tax benefits are recorded in "Income tax expense" in the Consolidated Statement of Operations at such time that the benefit is effectively settled.

Advertising Costs

We expense advertising costs primarily in "Subscriber acquisition costs" in the Consolidated Statements of Operations as incurred. Advertising expenses, net of payments received from programming content providers for marketing support, were $222 million in 2009, $211 million in 2008 and $178 million in 2007.

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Market Concentrations and Credit Risk

We sell programming services and extend credit, in amounts generally not exceeding $200 each, to a large number of individual residential subscribers throughout the United States. As applicable, we maintain allowances for anticipated losses.

Accounting Changes

Noncontrolling interests. On January 1, 2009 we adopted new accounting standards for the accounting and reporting of noncontrolling interests in subsidiaries, also known as minority interests, in consolidated financial statements. The new standards also provide guidance on accounting for changes in the parent's ownership interest in a subsidiary and establishes standards of accounting for the deconsolidation of a subsidiary due to the loss of control. Reporting entities must now present certain noncontrolling interests as a component of equity and present net income and consolidated comprehensive income attributable to the parent and the noncontrolling interest separately in the consolidated financial statements. These new standards are required to be applied prospectively, except for the presentation and disclosure requirements, which must be applied retrospectively for all periods presented. Our adoption of these changes did not have any effect on our consolidated financial statements.

Business Combinations. On January 1, 2009 we adopted a new business combination accounting standard that requires the acquiring entity in a business combination to record 100% of all assets and liabilities acquired, including goodwill and any non-controlling interest, generally at their fair values for all business combinations, whether partial, full or step acquisitions. Under the new standard, certain contingent assets and liabilities, as well as contingent consideration, are also required to be recognized at fair value on the date of acquisition and acquisition-related transaction and restructuring costs will be expensed. Additionally, disclosures are required describing the nature and financial effect of the business combination and the standard also changes the accounting for certain income tax assets recorded in purchase accounting. The adoption of the new accounting requirements as required, on January 1, 2009, changed the way we account for adjustments to deferred tax asset valuation allowances recorded in purchase accounting for prior business combinations so that adjustments to these deferred tax asset valuation allowances will no longer be recorded to goodwill but rather adjustments will be recorded in "Income tax expense" in the Consolidated Statements of Operations. Additionally, the adoption of the new accounting guidance changed the accounting for all business combinations we consummate after January 1, 2009.

Fair Value Recognition, Measurement and Disclosure. On January 1, 2008 we adopted new accounting standards which permit, but do not require, companies to report at fair value the majority of recognized financial assets, financial liabilities and firm commitments. Under this standard, unrealized gains and losses on items for which the fair value option is elected are reported in earnings at each subsequent reporting date. Our adoption of these accounting standards did not have any effect on our consolidated financial statements, as we have not elected to report subject instruments at fair value.

On January 1, 2008 we adopted new accounting standards for fair value measurements which defines fair value, sets out a framework for measuring fair value under accounting principles generally accepted in the United States of America, or GAAP, and expands disclosures about fair value measurements of assets and liabilities to include disclosure about inputs used in the determination of fair value using the following three categories:

Level 1: Quoted market prices in active markets for identical assets or liabilities.
Level 2: Observable market based inputs or unobservable inputs that are corroborated by market data.
Level 3: Unobservable inputs that are not corroborated by market data.

The new accounting standards apply under other accounting pronouncements previously issued by the Financial Accounting Standards Board, or FASB, that require or permit fair value measurements. Our adoption of the new accounting standards did not have any effect on our consolidated financial statements.

Payments to Manufacturers and Resellers. On January 1, 2008 we adopted new accounting standards which provide guidance to service providers regarding the proper reporting of consideration given to manufacturers or resellers of equipment necessary for an end-customer to receive its services. Depending on the circumstances, such consideration is reported as either an expense or a reduction of revenues. Our adoption of the new accounting standards did not have any effect on our consolidated financial statements.

Uncertain Tax Positions. We adopted accounting standards for accounting for uncertainty in income taxes on January 1, 2007, the cumulative effect of which resulted in a $3 million decrease to "Retained earnings" in the Consolidated Balance Sheets. As of the date of adoption, our unrecognized tax benefits and accrued interest totaled $26 million, including $3 million of tax positions the recognition of which would affect the annual effective income tax rate. As of the date of adoption, we have accrued $1 million in interest and penalties as part of our liability for unrecognized tax benefits. See Note 8 for additional information regarding unrecognized tax benefits.

Pensions. On December 31, 2007 our Parent adopted new accounting standards that requires the measurement of plan assets and benefit obligations as of the date of our fiscal year end and accordingly resulted in a change in our measurement date, which was previously November 30.

New Accounting Standards

Consolidation of Variable Interest Entities. In June 2009, the FASB, issued revisions to consolidation accounting standards for variable interest entities, or VIEs. The new standard replaces the quantitative-based risks and rewards calculation for determining which enterprise, if any, has a controlling financial interest in a variable interest entity. Instead, the new approach is qualitative and focused on identifying which enterprise has the power to direct the activities of a VIE that most significantly impact the entity's performance and (1) the obligation to absorb the losses of an entity or (2) the right to receive benefits from the entity. As a result of the changed requirements, it is possible that an entity's previous assessment of a VIE will change, and the standard now requires ongoing reassessments of whether an enterprise is the primary beneficiary of a VIE. Disclosure requirements under the new standard have been enhanced, and now include disclosure of the method the entity used to determine whether they are the primary beneficiary of the VIE. We do not expect the adoption of these changes to have an effect on our consolidated results of operations and financial position, when adopted, as required, on January 1, 2010.

Multiple Element Revenue Arrangements. In September 2009, the FASB approved a revised standard for revenue arrangements with multiple deliverables. Under the revised standard, the criteria

Table of Contents

for determining whether a deliverable should be considered a separate unit of accounting has changed to remove a limitation for separation to only items with objective and reliable evidence of fair value. Instead, the revised standard allows entities to use the "best estimate of selling price" in addition to third-party evidence or actual selling prices for determining the fair value of a deliverable. The standard also includes additional disclosure requirements for revenue arrangements for multiple deliverables. We currently do not expect the adoption of the revised standard to have an effect on our consolidated results of operations and financial position, when adopted, as required, on January 1, 2011.

Note 3: Accounts Receivable, Net

The following table sets forth the amounts recorded for "Accounts receivable, net" in our Consolidated Balance Sheets as of December 31:

	2009	2008
	(Dollars in Millions)	
Subscriber	$ 887	$ 807
Trade and other	563	533
Subtotal	1,450	1,340
Less: Allowance for doubtful accounts	(29)	(32)
Accounts receivable, net	$ 1,421	$ 1,308

Note 4: Satellites, Net and Property and Equipment, Net

The following table sets forth the amounts recorded for "Satellites, net" and "Property and equipment, net" in our Consolidated Balance Sheets at December 31:

	Estimated Useful Lives (years)	2009	2008
		(Dollars in Millions)	
Satellites	10-15	$ 2,355	$ 2,483
Satellites under construction		354	292
Total		2,709	2,775
Less: Accumulated depreciation		(839)	(795)
Satellites, net		$ 1,870	$ 1,980
Land and improvements	—	$ 26	$ 26
Buildings and leasehold improvements	6-30	301	287
Machinery and equipment	2-23	3,015	2,911
Subscriber leased set-top receivers	3-5	3,995	3,577
Construction in progress	—	330	235
Total		7,667	7,036
Less: Accumulated depreciation		(4,669)	(3,688)
Property and equipment, net		$ 2,998	$ 3,348

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

We capitalized interest costs of $18 million in 2009, $18 million in 2008 and $51 million in 2007, as part of the cost of our property and satellites under construction. Depreciation expense was $1,984 million in 2009, $1,707 million in 2008 and $1,094 million in 2007.

On March 1, 2006, we introduced a set-top receiver lease program. Prior to March 1, 2006, most set-top receivers provided to new and existing subscribers were immediately expensed upon activation as a subscriber acquisition or upgrade and retention cost in the Consolidated Statements of Operations. Subsequent to the introduction of the lease program, we lease most set-top receivers provided to new and existing subscribers, and therefore capitalize the set-top receivers in "Property and equipment, net" in the Consolidated Balance Sheets. We depreciate capitalized set-top receivers over a three year estimated useful life and include the amount of set-top receivers capitalized each period in the Consolidated Statements of Cash Flows.

The following table sets forth the amount of set-top receivers we capitalized, and depreciation expense we recorded under the lease program for each of the periods presented:

	Years ended December 31,		
Capitalized subscriber leased equipment:	**2009**	**2008**	**2007**
	(Dollars in Millions)		
Subscriber leased equipment—subscriber acquisitions	$ 564	$ 599	$ 762
Subscriber leased equipment—upgrade and retention	419	537	774
Total subscriber leased equipment capitalized	$ 983	$ 1,136	$ 1,536
Depreciation expense—subscriber leased equipment	$ 1,333	$ 1,100	$ 645

Note 5: Goodwill and Intangible Assets

The following table sets forth the changes in the carrying amounts of "Goodwill" in the Consolidated Balance Sheets for the years ended December 31, 2009 and 2008:

	(Dollars in Millions)
Balance as of January 1, 2008	$ 3,032
Acquisition related to home service provider business	157
Balance as of December 31, 2008	$ 3,189
Purchase or acquisition accounting adjustments:	
New acquisitions	24
Finalization of prior acquisitions	(46)
Balance as of December 31, 2009	$ 3,167

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

The following table sets forth the components for "Intangible assets, net" in the Consolidated Balance Sheets at:

	Estimated Useful Lives (years)	December 31, 2009 Gross Amount	December 31, 2009 Accumulated Amortization	December 31, 2009 Net Amount	December 31, 2008 Gross Amount	December 31, 2008 Accumulated Amortization	December 31, 2008 Net Amount
				(Dollars in Millions)			
Orbital slots	Indefinite	$ 432		$ 432	$ 432		$ 432
72.5° WL orbital license	5	219	$ 219	—	219	$ 180	39
Subscriber related	5-10	1,348	1,309	39	1,348	1,116	232
Dealer network	15	130	90	40	130	79	51
Distribution rights	7	334	263	71	334	217	117
Total intangible assets		$ 2,463	$ 1,881	$ 582	$ 2,463	$ 1,592	$ 871

Amortization expense for intangible assets was $289 million in 2009, $352 million in 2008 and $353 million in 2007.

Estimated amortization expense for intangible assets in each of the next five years and thereafter is as follows: $90 million in 2010, $34 million in 2011, $10 million in 2012, $10 million in 2013, $5 million in 2014 and $1 million thereafter.

We performed our annual impairment tests for goodwill and orbital slots in the fourth quarters of 2009, 2008, and 2007. The estimated fair values for each reporting unit and the orbital slots exceeded our carrying values, and accordingly, no impairment losses were recorded during 2009, 2008, or 2007.

Note 6: Accounts Payable and Accrued Liabilities; Other Liabilities and Deferred Credits

The following represent significant components of "Accounts payable and accrued liabilities" in our Consolidated Balance Sheets as of December 31:

	2009	2008
	(Dollars in Millions)	
Programming costs	$ 1,679	$ 1,532
Accounts payable	302	325
Payroll and employee benefits	122	96
Interest payable	47	45
Property and income taxes	45	35
Other	532	549
Total accounts payable and accrued liabilities	$ 2,727	$ 2,582

Table of Contents

The significant components of "Other liabilities and deferred credits" in our Consolidated Balance Sheets are as follows as of December 31:

	2009	2008
	(Dollars in Millions)	
Other accrued taxes	$ 229	$ 202
Deferred credits	78	122
Programming costs	76	251
Other	127	188
Total other liabilities and deferred credits	$ 510	$ 763

Note 7: Debt

The following table sets forth our outstanding debt:

	December 31,	
	2009	2008
	(Dollars in Millions)	
Senior notes	$ 4,490	$ 3,410
Senior secured credit facility, net of unamortized discount of $7 million as of December 31, 2009 and $9 million as of December 31, 2008	2,316	2,421
Unamortized bond premium	2	2
Total debt	6,808	5,833
Less: Current portion of long-term debt	(308)	(108)
Long-term debt	$ 6,500	$ 5,725

All of the senior notes and the senior secured credit facility were issued by us. The senior secured credit facility is secured by substantially all of our assets.

2009 Financing Transactions

On September 22, 2009, we issued $1,000 million in five-year 4.750% senior notes due in 2014 at a 0.3% discount resulting in $997 million of proceeds and $1,000 million in 10 year 5.875% senior notes due in 2019 at a 0.7% discount resulting in $993 million of proceeds in private placement transactions. Principal on these senior notes is payable upon maturity, while interest is payable semi-annually commencing April 1, 2010. We incurred $14 million of debt issuance costs in connection with these transactions. The senior notes have been fully and unconditionally guaranteed, jointly and severally, by substantially all of our current and certain of its future domestic subsidiaries on a senior unsecured basis. Pursuant to a registration rights agreement with the initial purchasers of the senior notes, we have filed a registration statement, whereby all holders of the original notes can elect to exchange their existing notes for registered notes with identical terms, except that the registered notes will be registered under the Securities Act of 1933, as amended and will not bear the legends restricting their transfer. We expect to complete the registration and exchange of these senior notes within the first half of 2010.

On September 22, 2009, we purchased, pursuant to a tender offer, $583 million of our then outstanding $910 million 8.375% senior notes at a price of 103.125% plus accrued and unpaid interest, for a total of $603 million. On September 23, 2009, we exercised our right to redeem the remaining

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

$327 million of the 8.375% senior notes at a price of 102.792% plus accrued and unpaid interest. On October 23, 2009, we redeemed the remaining $327 million of our 8.375% senior notes at a price of 102.792% plus accrued and unpaid interest for a total of $339 million.

The redemption of our 8.375% senior notes resulted in a 2009 pre-tax charge of $34 million, $21 million after tax, of which $27 million resulted from the premium paid for redemption of our 8.375% senior notes and $7 million resulted from the write-off of deferred debt issuance and other transaction costs. The charge was recorded in "Other, net" in our Consolidated Statements of Operations.

2008 Financing Transactions

In May 2008, we completed financing transactions that included the issuance of senior notes and an amendment to our existing senior secured credit facility as discussed below. We incurred $19 million of debt issuance costs in connection with these transactions.

We issued $1,500 million in senior notes due in 2016 in a private placement transaction. The eight-year notes bear interest at 7.625%. Principal on the senior notes is payable upon maturity, while interest is payable semi-annually commencing November 15, 2008. The senior notes have been fully and unconditionally guaranteed, jointly and severally, by substantially all of our current and certain of its future domestic subsidiaries on a senior unsecured basis. On November 11, 2008, we completed an exchange offer in which holders of substantially all of the outstanding principal amount of the senior notes exchanged the original senior notes for registered notes with identical terms, except that the registered notes are registered under the Securities Act of 1933, as amended, and do not bear the legends restricting their transfer.

We also amended its senior secured credit facility to include a new $1,000 million Term Loan C, which was issued at a 1% discount, resulting in $990 million of proceeds. Initially, borrowings under Term Loan C bear interest at 5.25%, however the rate is variable based on changes in the London InterBank Offered Rate, or LIBOR. The interest rate may be increased or decreased under certain conditions. The Term Loan C has a final maturity of April 13, 2013, and we began making quarterly principal payments totaling 1% annually on September 30, 2008. The senior secured credit facility is secured by substantially all of our assets and the assets of our current and certain of our future domestic subsidiaries and is fully and unconditionally guaranteed, jointly and severally, by substantially all of our current and certain of our future domestic subsidiaries.

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Senior Notes. The following table sets forth our outstanding senior notes balance and fair value as of December 31:

	Outstanding Balance		Fair value	
	2009	2008	2009	2008
	(Dollars in millions)			
8.375% senior notes due in 2013	$ —	$ 910	$ —	$ 904
4.750% senior notes due in 2014, net of unamortized discount of $3 million as of December 31, 2009	997	—	1,017	—
6.375% senior notes due in 2015	1,000	1,000	1,038	911
7.625% senior notes due in 2016	1,500	1,500	1,642	1,451
5.875% senior notes due in 2019, net of unamortized discount of $7 million as of December 31, 2009	993	—	1,016	—
Total senior notes	$ 4,490	$ 3,410	$ 4,713	$ 3,266

We calculated the fair values based on quoted market prices of our senior notes, which is a Level 1 input under the accounting guidance.

All of our senior notes were issued by us and have been, or in the case of the 4.750% and 5.875% senior notes are in the process of being, registered under the Securities Act of 1933, as amended. All of our senior notes are unsecured and have been fully and unconditionally guaranteed, jointly and severally, by substantially all of our assets and additionally, our 4.750% and 5.875% senior notes are rated as investment grade. Principal on the senior notes is payable upon maturity, while interest is payable semi-annually.

Credit Facility. At December 31, 2009, our senior secured credit facility consisted of a $375 million six-year Term Loan A, a $962 million eight-year Term Loan B, a $979 million five-year Term Loan C and a $500 million undrawn six-year revolving credit facility. The Term Loan A, Term Loan B and Term Loan C components of the senior secured credit facility currently bear interest at a rate equal to the LIBOR plus 0.75%, 1.50% and 2.25%, respectively. The weighted average interest for the senior secured credit facilities at December 31, 2009 was 3.143%. In addition, we pay a commitment fee of 0.175% per year for the unused commitment under the revolving credit facility. The interest rate and commitment fee may be increased or decreased under certain conditions. The senior secured credit facility is secured by substantially all of our assets and is fully and unconditionally guaranteed, jointly and severally by substantially all of our material domestic subsidiaries.

Our notes payable and credit facility mature as follows: $308 million in 2010, $108 million in 2011, $20 million in 2012, $1,887 million in 2013, $1,000 million in 2014 and $3,500 million thereafter. These amounts do not reflect potential prepayments that may be required under our senior secured credit facility, which could result from a computation of excess cash flows that we may be required to make at each year end under the credit agreement. We were not required to make a prepayment for the years ended December 31, 2009, 2008, or 2007. The amount of interest accrued related to our outstanding debt was $47 million at December 31, 2009 and $45 million at December 31, 2008.

Covenants and Restrictions. The senior secured credit facility requires us to comply with certain financial covenants. The senior notes and the senior secured credit facility also include covenants that restrict our ability to, among other things, (i) incur additional indebtedness, (ii) incur liens, (iii) pay dividends or make certain other restricted payments, investments or acquisitions, (iv) enter into certain

Table of Contents

transactions with affiliates, (v) merge or consolidate with another entity, (vi) sell, assign, lease or otherwise dispose of all or substantially all of its assets, and (vii) make voluntary prepayments of certain debt, in each case subject to exceptions as provided in the credit agreement and senior notes indentures. Should we fail to comply with these covenants, all or a portion of our borrowings under the senior notes and senior secured credit facility could become immediately payable and our revolving credit facility could be terminated. At December 31, 2009, we were in compliance with all such covenants. The senior notes and senior secured credit facility also provide that the borrowings may be required to be prepaid if certain change-in-control events occur. In September 2008, Liberty Media became the majority owner of the DIRECTV Group's outstanding common stock. There was no ratings decline for the senior notes associated with that event, and we were not required either to offer to redeem any of the senior notes pursuant to their respective indentures or to prepay any of the borrowings under the senior secured credit facility.

Restricted Cash. Restricted cash of $2 million as of December 31, 2009 and $8 million as of December 31, 2008 was included as part of "Prepaid expenses and other" in our Consolidated Balance Sheets. These amounts secure our letter of credit obligations. Restrictions on the cash will be removed as the letters of credit expire.

Note 8: Income Taxes

We base our income tax expense or benefit on reported "Income before income taxes." Deferred income tax assets and liabilities reflect the impact of temporary differences between the amounts of assets and liabilities recognized for financial reporting purposes and such amounts recognized for tax purposes, which are available to us pursuant to our tax sharing agreement with our Parent and as measured by applying currently enacted tax laws.

Our income tax expense consisted of the following for the years ended December 31:

	2009	2008	2007
	(Dollars in Millions)		
Current tax expense:			
U.S. federal	$ (532)	$ (691)	$ (633)
State and local	(69)	(64)	(156)
Total	(601)	(755)	(789)
Deferred tax expense:			
U.S. federal	(165)	(12)	(82)
State and local	(28)	(40)	(20)
Total	(193)	(52)	(102)
Total income tax expense	$ (794)	$ (807)	$ (891)

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Our income tax expense was different than the amount computed using the U.S. federal statutory income tax rate for the reasons set forth in the following table for the years ended December 31:

	2009	2008	2007
	(Dollars in Millions)		
Expected expense at U.S. federal statutory income tax rate	$ (717)	$ (720)	$ (788)
U.S. state and local income tax expense, net of federal benefit	(64)	(77)	(102)
Tax credits and other	(13)	(10)	(1)
Total income tax expense	$ (794)	$ (807)	$ (891)

Temporary differences and carryforwards that gave rise to deferred tax assets and liabilities at December 31 were as follows:

	2009		2008	
	Deferred Tax Assets	Deferred Tax Liabilities	Deferred Tax Assets	Deferred Tax Liabilities
	(Dollars in Millions)			
Depreciation and amortization	$ —	$ 637	$ —	$ 503
Accruals and advances	120	85	90	50
Programming contract liabilities	108	—	143	—
Prepaid expenses	—	19	—	26
State taxes	—	12	—	25
Net operating loss and tax credit carryforwards	29	—	17	—
Other temporary differences	8	9	3	6
Subtotal	265	762	253	610
Valuation allowance	(2)	—	(2)	—
Total deferred taxes	$ 263	$ 762	$ 251	$ 610

We assessed the deferred tax assets for the respective periods for recoverability and, where applicable, we recorded a valuation allowance to reduce the total deferred tax assets to an amount that will, more likely than not, be realized in the future.

The valuation allowance balances of $2 million at December 31, 2009 and $2 million at December 31, 2008, are attributable to unused capital losses which are available for carry-forward.

As of December 31, 2009, we have $35 million of federal net operating loss carryforward which expire between 2027 and 2028. The utilization of the federal net operating loss carryforward is subject to an annual limitation under Section 382 of the Internal Revenue Code, however we believe that we will have sufficient taxable income during the limitation period to utilize all of the carryforward.

A reconciliation of the beginning and ending balances of the total amounts of gross unrecognized tax benefits is as follows:

	Unrecognized tax benefits
	(Dollars in Millions)
Gross unrecognized tax benefits at January 1, 2007	$ 25
Increases in tax positions for prior years	31
Increases in tax positions for the current year	32
Gross unrecognized tax benefits at December 31, 2007	88
Increases in tax positions for prior years	76
Increases in tax positions for the current year	21
Settlements with taxing authorities	1
Gross unrecognized tax benefits at December 31, 2008	186
Increases in tax positions for prior years	22
Increases in tax positions for the current year	9
Settlements with taxing authorities	(2)
Gross unrecognized tax benefits at December 31, 2009	$ 215

As of December 31, 2009, our unrecognized tax benefits totaled $215 million, including $22 million of tax positions the recognition of which would affect the annual effective income tax rate.

We recognize interest and penalties accrued related to unrecognized tax benefits in "Income tax expense" in the Consolidated Statements of Operations. As of December 31, 2009, we have accrued $23 million in interest and penalties as part of our liability for unrecognized tax benefits.

We file numerous consolidated and separate income tax returns in the U.S. federal jurisdiction and in many state jurisdictions. For U.S. federal tax purposes, the tax years 2007 through 2009 remain open to examination. The California tax years 1994 through 2009 remain open to examination and the income tax returns in the other state tax jurisdictions in which we have operations are generally subject to examination for a period of 3 to 5 years after filing of the respective return.

We do not anticipate changes to the total unrecognized tax benefits in the next twelve months which will have a significant effect on our results of operations or financial position.

Note 9: Pension and Other Postretirement Benefit Plans

Most of our employees are eligible to participate in our Parent's funded non-contributory defined benefit pension plan, which provides defined benefits based on either years of service and final average salary, or eligible compensation while employed by us. We have not separately determined the accumulated benefit obligation and net assets available for benefits for our employees and do not include these items in our Consolidated Balance Sheets. In addition to pension benefits, DIRECTV charges us for the cost of certain other post-retirement benefits. The accumulated other post-retirement benefit obligation related to our employees has not been separately determined and is not included in the accompanying Consolidated Balance Sheets. We also participate in other health and welfare plans of DIRECTV. Our portion of the cost of these benefit plans, allocated from DIRECTV, amounted to $19 million in 2009, $16 million in 2008 and $16 million in 2007.

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Note 10: Share-Based Payment

Under The DIRECTV Group, Inc. Amended and Restated 2004 Stock Plan, or the DIRECTV Plan, as approved by our Parent's stockholders on June 5, 2007, shares, rights or options to acquire up to 21 million shares of common stock plus the number of shares that were granted under a former plan but which, after December 22, 2003 are forfeited, expire or are cancelled without the delivery of shares of common stock or otherwise result in the return of such shares to us, were authorized for grant through June 4, 2017, subject to the approval of the Compensation Committee of our Parent's Board of Directors. Under the DIRECTV Plan, our Parent issues new shares of its Class A common stock when restricted stock units are earned and when stock options are exercised.

Restricted Stock Units. The Compensation Committee of DIRECTV has granted restricted stock units under our Parent's stock plans to certain of our employees and executives. Annual awards are mostly performance-based, with final payments in shares of our Parent's Class A common stock. Final payment can be reduced from the target award amounts based on our Parent company's performance over a three-year performance period in comparison with pre-established targets.

During the year ended December 31, 2009, our employees were granted 2.4 million restricted stock units with a weighted average grant-date fair value of $21.26 per share. During the year ended December 31, 2008, our employees were granted 2.3 million restricted stock units with a weighted average grant-date fair value of approximately $23.15 per share. During the year ended December 31, 2007, our employees were granted 2.5 million restricted stock units with a weighted average grant-date fair value of approximately $23.69 per share. The grant date fair value of restricted stock units is based on the closing stock price of our Parent's Class A common stock on the date of grant.

Stock Options. DIRECTV's Compensation Committee has also granted stock options to acquire our Parent's Class A common stock under our Parent's stock plans to certain of our employees and executives. The exercise price of options granted is equal to at least 100% of the fair market value of the common stock on the date the options were granted. These nonqualified options generally vest over one to five years, expire ten years from date of grant and are subject to earlier termination under certain conditions. During the year ended December 31, 2007, our employees were granted 1.2 million stock options with a grant-date fair value of approximately $8.27 per share. No stock options were granted to our employees during 2008 or 2009. The grant date fair value of common stock options is determined by our Parent using the Black-Scholes valuation model.

The following table presents amounts recorded related to share-based compensation for the years ended December 31:

	2009	2008	2007
	(Dollars in Millions)		
Share-based compensation expense recognized	$ 44	$ 43	$ 41
Tax benefits associated with share-based compensation expense	17	16	16
Actual tax benefits realized for the deduction of share-based compensation expense	24	30	19

As of December 31, 2009, there was $50 million of unrecognized compensation costs related to unvested restricted stock units, which we expect to recognize as follows: $33 million in 2010 and $17 million in 2011.

As of December 31, 2009, our employees held 6.4 million stock options and 6.2 million restricted stock units.

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Note 11: Related Party Transactions

In the ordinary course of our operations, we enter into transactions with related parties as discussed below.

DIRECTV and affiliates

We determine our income taxes based upon our tax sharing agreement with our Parent, which generally provides that the current income tax liability or receivable be computed as if we were a separate taxpayer. Payments made to our Parent under this tax sharing arrangement were $502 million in 2009, $721 million in 2008 and $709 million in 2007. We also receive an allocation of employee benefit expenses from DIRECTV. We believe that our consolidated financial statements reflect our cost of doing business in accordance with SEC Staff Accounting Bulletin No. 55, "Allocation of Expenses and Related Disclosures in Financial Statements of Subsidiaries, Divisions or Lesser Business Components of Another Entity."

We paid dividends to our Parent in the amounts of $2,500 million in 2009, $3,400 million in 2008 and $1,050 million in 2007 from available cash and cash equivalents. In addition, we paid $1,500 million in dividends to our Parent during January and February 2010.

Beginning November 19, 2009, transactions with the regional sports networks which were acquired by DIRECTV on that date are also included as transactions with DIRECTV and affiliates.

Liberty Media, Liberty Global and Discovery Communications

Beginning with Liberty's acquisition of its ownership interest in DIRECTV Group from News Corporation on February 27, 2008, transactions with Liberty Media Corporation, or Liberty Media, and its affiliates, including its equity method investees, may be considered to be related party transactions. Our transactions with Liberty Media and its affiliates consist primarily of the purchase of programming.

Although as a result of the Liberty Transaction, Liberty no longer has any equity interest in DIRECTV, John Malone, Chairman of the Board of Directors of DIRECTV and of Liberty Media, has an approximate 24% voting interest in DIRECTV, an approximate 31% voting interest in Discovery Communications, Inc., or Discovery Communications, and an approximate 40% voting interest in Liberty Global Inc., or Liberty Global, and serves as Chairman of Liberty Global, and certain of Liberty Media's management and directors also serve as directors of Discovery Communications or Liberty Global. As a result of this common ownership and management, transactions with Discovery Communications and Liberty Global, and their subsidiaries or equity method investees may be considered to be related party transactions. Our transactions with Discovery Communications and Liberty Global consist primarily of purchases of programming created, owned or distributed by Discovery Communications and its subsidiaries and investees.

News Corporation and affiliates

News Corporation and its affiliates were considered related parties until February 27, 2008, when News Corporation transferred its 41% interest in our Parent's common stock to Liberty Media. Accordingly, the following contractual arrangements with News Corporation and its affiliates are considered related party transactions and reported through February 27, 2008: purchase of programming, products and advertising; license of certain intellectual property, including patents; purchase of system access products, set-top receiver software and support services; sale of advertising space; purchase of employee services; and use of facilities.

Table of Contents

The majority of payments under contractual arrangements with Liberty Media, Discovery Communications, Liberty Global and News Corporation entities relate to multi-year programming contracts. Payments under these contracts are typically subject to annual rate increases and are based on the number of subscribers receiving the related programming.

Other

Companies in which we hold equity method investments are also considered related parties.

Beginning November 19, 2009, transactions with the Game Show Network, which our Parent holds an equity method investment in, are also included as transactions with Other.

The following table summarizes sales and purchase transactions with related parties:

	2009	2008	2007
	(Dollars in Millions)		
Sales:			
Liberty Media and affiliates	$ 56	$ 36	$ —
Discovery Communications, Liberty Global and affiliates	11	10	—
News Corporation and affiliates	—	2	17
DIRECTV and affiliates	7	—	—
Other	1	2	—
Total	$ 75	$ 50	$ 17
Purchases:			
Liberty Media and affiliates	$ 355	$ 267	$ —
Discovery Communications, Liberty Global and affiliates	219	164	—
News Corporation and affiliates	—	157	835
DIRECTV and affiliates	7	—	1
Other	69	35	23
Total	$ 650	$ 623	$ 859

The following table sets forth the amount of accounts receivable from and accounts payable to related parties as of December 31:

	2009	2008
	(Dollars in Millions)	
Accounts receivable	$ 23	$ 29
Accounts payable	166	156

The accounts receivable and accounts payable balances as of December 31, 2009 and 2008 are primarily related to affiliates of Liberty Media.

Note 12: Acquisitions

Home Services Providers

180 Connect. In July 2008, we acquired 100% of 180 Connect Inc.'s outstanding common stock and exchangeable shares. Simultaneously, in a separate transaction, UniTek USA, LLC acquired 100% of 180 Connect's cable service operating unit and operations in certain of our installation services markets in exchange for satellite installation operations in certain markets and $7 million in cash. These transactions provide us with control over a significant portion of our home service provider network. We paid $91 million in cash, net of the $7 million we received from UniTek USA, for the acquisition, including the equity purchase price, repayment of assumed debt and related transaction costs.

We accounted for the 180 Connect acquisition using the purchase method of accounting, and began consolidating the results from the date of acquisition. The December 31, 2009 consolidated financial statements reflect the final allocation of the $91 million net purchase price to assets acquired and the liabilities assumed based on their estimated fair values at the date of acquisition using information currently available. The assets acquired included approximately $5 million in cash. The excess of the purchase price over the estimated fair values of the net assets has been recorded as goodwill, $28 million of which will be deductible for tax purposes.

The following table sets forth the final allocation of the purchase price to the 180 Connect net assets acquired in July 2008 (dollars in millions):

Total current assets	$ 18
Property and equipment	16
Goodwill	97
Investments and other assets	51
Total assets acquired	$ 182
Total current liabilities	$ 83
Other liabilities	8
Total liabilities assumed	$ 91
Net assets acquired	$ 91

The following selected unaudited pro forma information is being provided to present a summary of the combined results of us and 180 Connect for 2008 as if the acquisition had occurred as of the beginning of the period, giving effect to purchase accounting adjustments. The pro forma data is presented for informational purposes only and may not necessarily reflect the results of our operations had 180 Connect operated as part of us for the period presented, nor are they necessarily indicative of the results of future operations. The pro forma information excludes the effect of non-recurring charges.

	Year Ended December 31, 2008 (Dollars in Millions)
Revenues	$ 17,310
Net income	1,208

Note 13: Commitments and Contingencies

Commitments

At December 31, 2009, minimum future commitments under noncancelable operating leases having lease terms in excess of one year were primarily for satellite transponder leases and real property and aggregated $232 million, payable as follows: $41 million in 2010, $35 million in 2011, $32 million in 2012, $31 million in 2013, $15 million in 2014 and $78 million thereafter. Certain of these leases contain escalation clauses and renewal or purchase options, which we have not considered in the amounts disclosed. Rental expenses under operating leases were $62 million in 2009, $53 million in 2008 and $52 million in 2007.

Table of Contents

At December 31, 2009, our minimum payments under agreements to purchase broadcast programming, regional professional team rights and the purchase of services that we have outsourced to third parties, such as billing services, and satellite telemetry, tracking and control, satellite launch contracts and broadcast center services aggregated $8,297 million, payable as follows: $1,671 million in 2010, $1,663 million in 2011, $1,787 million in 2012, $1,392 million in 2013, $1,147million in 2014 and $637 million thereafter.

At December 31, 2009, other long-term obligations totaling $183 million are payable approximately as follows: $97 million in 2010, $58 million in 2011, $9 million in 2012, $7 million in 2013, $3 million in 2014 and $9 million thereafter. These amounts are recorded in "Accounts payable and accrued liabilities" and "Other liabilities and deferred credits" in the Consolidated Balance Sheets.

Contingencies

Litigation. Litigation is subject to uncertainties and the outcome of individual litigated matters is not predictable with assurance. Various legal actions, claims and proceedings are pending against us arising in the ordinary course of business. We have established loss provisions for matters in which losses are probable and can be reasonably estimated. Some of the matters may involve compensatory, punitive, or treble damage claims, or demands that if granted, could require us to pay damages or make other expenditures in amounts that could not be estimated at December 31, 2009. After discussion with counsel representing us in those actions, it is the opinion of management that such litigation is not expected to have a material adverse effect on our consolidated financial statements.

Finisar Corporation. As previously reported, we were successful in 2008 getting the jury verdict in the Finisar case vacated on appeal. The original verdict found the patent to be valid and willfully infringed, and the jury awarded approximately $79 million in damages. The trial court increased the damages award by $25 million because of the jury finding of willful infringement and awarded pre-judgment interest of $13 million. DIRECTV was also ordered to pay into escrow $1.60 per new set-top receiver manufactured for use with the DIRECTV system beginning June 17, 2006 and continuing until the patent expires in 2012 or was otherwise found to be invalid. On April 18, 2008, the Court of Appeals reversed the verdict of the district court in part, vacated the findings of infringement, and remanded for further proceedings on the remaining issues finding that the district court had applied erroneous interpretations of certain terms of the claims. On remand, we sought and obtained summary judgment on invalidity of all remaining claims, and the case against DIRECTV was dismissed on May 19, 2009. Finisar filed a Notice of Appeal, and oral argument on the appeal was held on January 6, 2010. On January 8, 2010, the Court of Appeals affirmed per curiam the grant of summary judgment on all claims. This case is now resolved and there will be no further proceedings in this matter.

Satellites We may purchase in-orbit and launch insurance to mitigate the potential financial impact of satellite launch and in-orbit failures if the premium costs are considered economic relative to the risk of satellite failure. The insurance generally covers the unamortized book value of covered satellites. We do not insure against lost revenues in the event of a total or partial loss of the capacity of a satellite. We generally rely on in-orbit spare satellites and excess transponder capacity at key orbital slots to mitigate the impact a satellite failure could have on our ability to provide service. At December 31, 2009, the net book value of in-orbit satellites was $1,516 million, all of which was uninsured.

Table of Contents

Other We are contingently liable under standby letters of credit and bonds in the aggregate amount of $2 million at December 31, 2009.

Note 14: Condensed Consolidating Financial Statements

The following presents the condensed consolidating statements of operations for the years ended December 31, 2009, 2008 and 2007, the condensed consolidating balance sheets as of December 31, 2009 and 2008, and the condensed consolidating statements of cash flows for the years ended December 31, 2009, 2008 and 2007 of DIRECTV Holdings together with DIRECTV Financing Co., Inc., or the Co-Issuers, and each of DIRECTV Holdings' material subsidiaries (other than DIRECTV Financing), or the Guarantor Subsidiaries, and the eliminations necessary to present DIRECTV Holdings' financial statements on a consolidated basis. These condensed consolidating financial statements should be read in conjunction with the accompanying consolidated financial statements of DIRECTV Holdings.

Condensed Consolidating Statement of Operations
For the Year Ended December 31, 2009

	Co-Issuers	Guarantor Subsidiaries	Eliminations	DIRECTV Holdings Consolidated
	(Dollars in Millions)			
Revenues	$ 332	$ 18,671	$ (332)	$ 18,671
Operating costs and expenses				
Costs of revenues, exclusive of depreciation and amortization expense				
Broadcast programming and other	—	8,027	—	8,027
Subscriber service expenses	—	1,268	—	1,268
Broadcast operations expenses	—	274	—	274
Selling, general and administrative expenses, exclusive of depreciation and amortization expense				
Subscriber acquisition costs	—	2,478	—	2,478
Upgrade and retention costs	—	1,045	—	1,045
General and administrative expenses	—	1,226	(332)	894
Depreciation and amortization expense	—	2,275	—	2,275
Total operating costs and expenses	—	16,593	(332)	16,261
Operating profit	332	2,078	—	2,410
Equity in income of consolidated subsidiaries	1,275	—	(1,275)	—
Interest income	4	—	—	4
Interest expense	(335)	(13)	—	(348)
Other, net	(34)	17	—	(17)
Income before income taxes	1,242	2,082	(1,275)	2,049
Income tax benefit (expense)	13	(807)	—	(794)
Net income	$ 1,255	$ 1,275	$ (1,275)	$ 1,255

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Statement of Operations
For the Year Ended December 31, 2008

	Co-Issuers	Guarantor Subsidiaries	Eliminations	DIRECTV Holdings Consolidated
		(Dollars in Millions)		
Revenues	$ 287	$ 17,310	$ (287)	$ 17,310
Operating costs and expenses				
Costs of revenues, exclusive of depreciation and amortization expense				
Broadcast programming and other	—	7,424	—	7,424
Subscriber service expenses	—	1,139	—	1,139
Broadcast operations expenses	—	265	—	265
Selling, general and administrative expenses, exclusive of depreciation and amortization expense				
Subscriber acquisition costs	—	2,191	—	2,191
Upgrade and retention costs	—	1,027	—	1,027
General and administrative expenses	—	1,160	(287)	873
Depreciation and amortization expense	—	2,061	—	2,061
Total operating costs and expenses	—	15,267	(287)	14,980
Operating profit	287	2,043	—	2,330
Equity in income of consolidated subsidiaries	1,235	—	(1,235)	—
Interest income	36	1	—	37
Interest expense	(298)	(17)	—	(315)
Other, net	—	5	—	5
Income before income taxes	1,260	2,032	(1,235)	2,057
Income tax expense	(10)	(797)	—	(807)
Net income	$ 1,250	$ 1,235	$ (1,235)	$ 1,250

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Statement of Operations
For the Year Ended December 31, 2007

	Co-Issuers	Guarantor Subsidiaries	Eliminations	DIRECTV Holdings Consolidated
		(Dollars in Millions)		
Revenues	$ 67	$ 15,527	$ (67)	$ 15,527
Operating costs and expenses				
Costs of revenues, exclusive of depreciation and amortization expense				
Broadcast programming and other	—	6,681	—	6,681
Subscriber service expenses	—	1,137	—	1,137
Broadcast operations expenses	—	216	—	216
Selling, general and administrative expenses, exclusive of depreciation and amortization expense				
Subscriber acquisition costs	—	1,901	—	1,901
Upgrade and retention costs	—	958	—	958
General and administrative expenses	—	851	(67)	784
Depreciation and amortization expense	—	1,448	—	1,448
Total operating costs and expenses	—	13,192	(67)	13,125
Operating profit	67	2,335	—	2,402
Equity in income of consolidated subsidiaries	1,394	—	(1,394)	—
Interest income	69	—	—	69
Interest expense	(193)	(23)	—	(216)
Other, net	—	(5)	—	(5)
Income before income taxes	1,337	2,307	(1,394)	2,250
Income tax benefit (expense)	22	(913)	—	(891)
Net income	$ 1,359	$ 1,394	$ (1,394)	$ 1,359

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Balance Sheet
As of December 31, 2009

	Co-Issuers	Guarantor Subsidiaries	Eliminations	DIRECTV Holdings Consolidated
	(Dollars in Millions)			
ASSETS				
Total current assets	$ 1,756	$ 1,854	$ (50)	$ 3,560
Satellites, net	—	1,870	—	1,870
Property and equipment, net	—	2,998	—	2,998
Goodwill	1,828	1,339	—	3,167
Intangible assets, net	—	582	—	582
Other assets	10,228	3,873	(13,870)	231
Total assets	$ 13,812	$ 12,516	$ (13,920)	$ 12,408
LIABILITIES AND OWNER'S EQUITY				
Total current liabilities	$ 380	$ 3,057	$ (49)	$ 3,388
Long-term debt	6,500	—	—	6,500
Deferred income taxes	—	775	(216)	559
Other liabilities and deferred credits	5,481	510	(5,481)	510
Owner's equity				
Capital stock and additional paid-in capital	1,076	4,526	(4,526)	1,076
Retained earnings	375	3,648	(3,648)	375
Total owner's equity	1,451	8,174	(8,174)	1,451
Total liabilities and owner's equity	$ 13,812	$ 12,516	$ (13,920)	$ 12,408

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Balance Sheet
As of December 31, 2008

	Co-Issuers	Guarantor Subsidiaries	Eliminations	DIRECTV Holdings Consolidated
		(Dollars in Millions)		
ASSETS				
Total current assets	$ 1,221	$ 1,821	$ (96)	$ 2,946
Satellites, net	—	1,980	—	1,980
Property and equipment, net	—	3,348	—	3,348
Goodwill	1,827	1,362	—	3,189
Intangible assets, net	—	871	—	871
Other assets	8,070	1,739	(9,597)	212
Total assets	$ 11,118	$ 11,121	$ (9,693)	$ 12,546
LIABILITIES AND OWNER'S EQUITY				
Total current liabilities	$ 216	$ 2,888	$ (98)	$ 3,006
Long-term debt	5,725	—	—	5,725
Deferred income taxes	—	621	(216)	405
Other liabilities and deferred credits	2,530	763	(2,530)	763
Owner's equity				
Capital stock and additional paid-in capital	2,403	4,476	(4,476)	2,403
Retained earnings	244	2,373	(2,373)	244
Total owner's equity	2,647	6,849	(6,849)	2,647
Total liabilities and owner's equity	$ 11,118	$ 11,121	$ (9,693)	$ 12,546

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Statement of Cash Flows
For the Year Ended December 31, 2009

	Co-Issuers	Guarantor Subsidiaries	DIRECTV Holdings Consolidated
	(Dollars in Millions)		
Cash flows from operating activities			
Net cash provided by operating activities	$ 2,108	$ 1,583	$ 3,691
Cash flows from investing activities			
Cash paid for property and equipment	—	(443)	(443)
Cash paid for subscriber leased equipment—subscriber acquisitions	—	(564)	(564)
Cash paid for subscriber leased equipment—upgrade and retention	—	(419)	(419)
Cash paid for satellites	—	(59)	(59)
Investment in companies, net of cash acquired	—	(11)	(11)
Net cash used in investing activities	—	(1,496)	(1,496)
Cash flows from financing activities			
Cash proceeds from debt issuance	1,990	—	1,990
Debt issuance costs	(14)	—	(14)
Repayment of long-term debt	(1,018)	—	(1,018)
Repayment of other long-term obligations	—	(90)	(90)
Cash dividends to Parent	(2,500)	—	(2,500)
Excess tax benefit from share-based compensation	—	4	4
Net cash used in financing activities	(1,542)	(86)	(1,628)
Net increase in cash and cash equivalents	566	1	567
Cash and cash equivalents at beginning of the period	1,143	6	1,149
Cash and cash equivalents at the end of the period	$ 1,709	$ 7	$ 1,716

Table of Contents

DIRECTV HOLDINGS LLC

NOTES TO THE CONSOLIDATED FINANCIAL STATEMENTS—(Continued)

Condensed Consolidating Statement of Cash Flows
For the Year Ended December 31, 2008

	Co-Issuers	Guarantor Subsidiaries	DIRECTV Holdings Consolidated
		(Dollars in Millions)	
Cash flows from operating activities			
Net cash provided by operating activities	$ 1,335	$ 1,942	$ 3,277
Cash flows from investing activities			
Cash paid for property and equipment	—	(501)	(501)
Cash paid for subscriber leased equipment—subscriber acquisitions	—	(599)	(599)
Cash paid for subscriber leased equipment—upgrade and retention	—	(537)	(537)
Cash paid for satellites	—	(128)	(128)
Investment in companies, net of cash acquired	—	(97)	(97)
Other	—	5	5
Net cash used in investing activities	—	(1,857)	(1,857)
Cash flows from financing activities			
Cash proceeds from debt issuance	2,490	—	2,490
Debt issuance costs	(19)	—	(19)
Repayment of long-term debt	(53)	—	(53)
Repayment of other long-term obligations	—	(98)	(98)
Cash dividends to Parent	(3,400)	—	(3,400)
Excess tax benefit from share-based compensation	—	7	7
Net cash used in financing activities	(982)	(91)	(1,073)
Net increase (decrease) in cash and cash equivalents	353	(6)	347
Cash and cash equivalents at beginning of the period	790	12	802
Cash and cash equivalents at the end of the period	$ 1,143	$ 6	$ 1,149

Table of Contents

Condensed Consolidating Statement of Cash Flows
For the Year Ended December 31, 2007

	Co-Issuers	Guarantor Subsidiaries	DIRECTV Holdings Consolidated
		(Dollars in Millions)	
Cash flows from operating activities			
Net cash provided by operating activities	$ 496	$ 2,413	$ 2,909
Cash flows from investing activities			
Cash paid for property and equipment	—	(621)	(621)
Cash paid for subscriber leased equipment-subscriber acquisitions	—	(762)	(762)
Cash paid for subscriber leased equipment-upgrade and retention	—	(774)	(774)
Cash paid for satellites	—	(169)	(169)
Other	—	(9)	(9)
Net cash used in investing activities	—	(2,335)	(2,335)
Cash flows from financing activities			
Repayment of long-term debt	(10)	—	(10)
Repayment of other long-term obligations	—	(72)	(72)
Cash dividend to Parent	(1,050)	—	(1,050)
Excess tax benefit from share-based compensation	—	4	4
Net cash used in financing activities	(1,060)	(68)	(1,128)
Net (decrease) increase in cash and cash equivalents	(564)	10	(554)
Cash and cash equivalents at beginning of the year	1,354	2	1,356
Cash and cash equivalents at the end of the year	$ 790	$ 12	$ 802

Note 15: Selected Quarterly Data (Unaudited)

The following table presents unaudited selected quarterly data for 2009 and 2008:

	1st	2nd	3rd	4th
		(Dollars in Millions)		
2009 Quarters				
Revenues	$ 4,303	$ 4,539	$ 4,703	$ 5,126
Operating profit	397	652	611	750
Net income	197	350	311	397
2008 Quarters				
Revenues	$ 4,049	$ 4,196	$ 4,324	$ 4,741
Operating profit	593	717	532	488
Net income	332	402	266	250

***

DIRECTV HOLDINGS LLC

ITEM 9. CHANGES IN AND DISAGREEMENTS WITH ACCOUNTANTS ON ACCOUNTING AND FINANCIAL DISCLOSURE

None.

***

ITEM 9A. CONTROLS AND PROCEDURES

Disclosure Controls and Procedures

We carried out an evaluation as of the end of the year covered by this Annual Report on Form 10-K under the supervision and with the participation of management, including our principal executive officers and financial officers, of the effectiveness of our disclosure controls and procedures (as defined in Rules 13a-15(e) and 15d-15(e) under the Securities Exchange Act of 1934, as amended, or the Exchange Act). Based on the evaluation, our principal executive officers and our financial officers concluded that our disclosure controls and procedures were effective as of December 31, 2009.

There has been no change in our internal control over financial reporting (as defined in Rules 13a-15(f) and 15d-15(f) under the Exchange Act) that occurred during our fiscal quarter ended December 31, 2009, that has materially affected, or is reasonably likely to materially affect, our internal control over financial reporting.

Internal Control Over Financial Reporting

Management's Report on Internal Control Over Financial Reporting

Our management is responsible for establishing and maintaining adequate internal control over financial reporting as defined in Rule 13a-15(f) or 15d-15(f) promulgated under the Securities Exchange Act of 1934. Those rules define internal control over financial reporting as a process designed to provide reasonable assurance regarding the reliability of financial reporting and the preparation of financial statements for external purposes in accordance with accounting principles generally accepted in the United States of America, or GAAP, and includes those policies and procedures that:

- pertain to the maintenance of records that in reasonable detail accurately and fairly reflect the transactions and dispositions of the assets of the company;
- provide reasonable assurance that transactions are recorded as necessary to permit preparation of financial statements in accordance with GAAP, and that receipts and expenditures of the company are being made only in accordance with authorizations of management and directors of the company; and
- provide reasonable assurance regarding prevention or timely detection of unauthorized acquisition, use or disposition of the company's assets that could have a material effect on the financial statements.

Because of its inherent limitations, internal control over financial reporting may not prevent or detect misstatements. Projections of any evaluation of effectiveness to future periods are subject to the risk that controls may become inadequate because of changes in conditions, or that the degree of compliance with the policies or procedures may deteriorate.

Our management assessed the effectiveness of our internal control over financial reporting as of December 31, 2009. In making this assessment, our management used the criteria established in

Table of Contents

Internal Control-Integrated Framework issued by the Committee of Sponsoring Organizations of the Treadway Commission (COSO). Based on their assessment and those criteria, management believes that, as of December 31, 2009, our internal control over financial reporting is effective.

Our independent registered public accounting firm has issued an audit report on internal control over financial reporting, which appears below.

Table of Contents

DIRECTV HOLDINGS LLC

REPORT OF INDEPENDENT REGISTERED PUBLIC ACCOUNTING FIRM

To the Board of Directors of DIRECTV Holdings LLC
El Segundo, California

We have audited the internal control over financial reporting of DIRECTV Holdings LLC (the "Company") as of December 31, 2009, based on criteria established in *Internal Control—Integrated Framework* issued by the Committee of Sponsoring Organizations of the Treadway Commission. The Company's management is responsible for maintaining effective internal control over financial reporting and for its assessment of the effectiveness of internal control over financial reporting, included in the accompanying Management's Report on Internal Control over Financial Reporting. Our responsibility is to express an opinion on the Company's internal control over financial reporting based on our audit.

We conducted our audit in accordance with the standards of the Public Company Accounting Oversight Board (United States). Those standards require that we plan and perform the audit to obtain reasonable assurance about whether effective internal control over financial reporting was maintained in all material respects. Our audit included obtaining an understanding of internal control over financial reporting, assessing the risk that a material weakness exists, testing and evaluating the design and operating effectiveness of internal control based on the assessed risk, and performing such other procedures as we considered necessary in the circumstances. We believe that our audit provides a reasonable basis for our opinion.

A company's internal control over financial reporting is a process designed by, or under the supervision of, the company's principal executive and principal financial officers, or persons performing similar functions, and effected by the company's board of directors, management, and other personnel to provide reasonable assurance regarding the reliability of financial reporting and the preparation of financial statements for external purposes in accordance with generally accepted accounting principles. A company's internal control over financial reporting includes those policies and procedures that (1) pertain to the maintenance of records that, in reasonable detail, accurately and fairly reflect the transactions and dispositions of the assets of the company; (2) provide reasonable assurance that transactions are recorded as necessary to permit preparation of financial statements in accordance with generally accepted accounting principles, and that receipts and expenditures of the company are being made only in accordance with authorizations of management and directors of the company; and (3) provide reasonable assurance regarding prevention or timely detection of unauthorized acquisition, use, or disposition of the company's assets that could have a material effect on the financial statements.

Because of the inherent limitations of internal control over financial reporting, including the possibility of collusion or improper management override of controls, material misstatements due to error or fraud may not be prevented or detected on a timely basis. Also, projections of any evaluation of the effectiveness of the internal control over financial reporting to future periods are subject to the risk that the controls may become inadequate because of changes in conditions, or that the degree of compliance with the policies or procedures may deteriorate.

In our opinion, the Company maintained, in all material respects, effective internal control over financial reporting as of December 31, 2009, based on the criteria established in *Internal Control—Integrated Framework* issued by the Committee of Sponsoring Organizations of the Treadway Commission.

We have also audited, in accordance with the standards of the Public Company Accounting Oversight Board (United States), the consolidated financial statements and financial statement schedule

DIRECTV HOLDINGS LLC

as of and for the year ended December 31, 2009 of the Company and our report dated February 25, 2010 expressed an unqualified opinion on those financial statements and financial statement schedule.

/s/ DELOITTE & TOUCHE LLP

Los Angeles, California
February 25, 2010

Table of Contents

DIRECTV HOLDINGS LLC

ITEM 9B. OTHER INFORMATION

None.

***

PART III

ITEM 10. DIRECTORS AND EXECUTIVE OFFICERS OF THE REGISTRANT

Omitted.

***

ITEM 11. EXECUTIVE COMPENSATION

Omitted.

***

ITEM 12. SECURITY OWNERSHIP OF CERTAIN BENEFICIAL OWNERS AND MANAGEMENT AND RELATED STOCKHOLDER MATTERS

Omitted

***

ITEM 13. CERTAIN RELATIONSHIPS AND RELATED TRANSACTIONS

Omitted.

***

ITEM 14. PRINCIPAL ACCOUNTING FEES AND SERVICES

The aggregate fees Deloitte & Touche LLP billed for professional services in 2009 and 2008 were:

Type of Fees	2009	2008
	(Dollars in Millions)	
Audit Fees and Audit-Related Services	$ 3	$ 3

"Audit Fees" are fees Deloitte & Touche LLP bills us for professional services for the audit of our consolidated financial statements included in Form 10-K and review of our consolidated financial statements included in Form 10-Qs. Deloitte & Touche LLP bills us for "Audit-Related Services," which are principally for accounting consultations and assurance and related services associated with our financing transactions. DIRECTV engages our accountant on our behalf to render audit and non-audit services for us.

***

Table of Contents

DIRECTV HOLDINGS LLC

PART IV

ITEM 15. EXHIBITS AND FINANCIAL STATEMENTS SCHEDULES

		Page Number
1.	All Financial Statements	See Part II
2.	Financial Statement Schedule II-Valuation and Qualifying Accounts for the Years Ended December 31, 2009, 2008 and 2007	
3.	Exhibits (Including Those Incorporated By Reference)	

Exhibit Number	Exhibit Name
*3.1	Certificate of Formation of DIRECTV Holdings LLC dated as of June 11, 2002 (incorporated by reference to Exhibit 3.1 to the Form S-4 of DIRECTV Holdings LLC filed June 26, 2003 (SEC File No. 333-106529)).
*3.2	Certificate of Incorporation of DIRECTV Financing Co., Inc. dated as of February 5, 2003 (incorporated by reference to Exhibit 3.2 to the Form S-4 of DIRECTV Holdings LLC filed June 26, 2003 (SEC File No. 333-106529)).
*3.3	Limited Liability Company Agreement of DIRECTV Holdings LLC dated as of June 11, 2002 (incorporated by reference to Exhibit 3.9 to the Form S-4 of DIRECTV Holdings LLC filed June 26, 2003 (SEC File No. 333-106529)).
*3.4	Amended and Restated By-laws of DIRECTV Financing Co., Inc. (incorporated by reference to Exhibit 3.9 of the Form S-4 of DIRECTV Holdings, LLC filed on February 5, 2010 (SEC File No. 333-106529)).
*4.1	Indenture, dated as of June 15, 2005, by and among DIRECTV Holdings LLC, DIRECTV Financing Co., Inc., the Guarantors signatory thereto and The Bank of New York, as trustee (incorporated by reference to Exhibit 10.1 to the Current Report on Form 8-K of the DIRECTV Holdings LLC filed on June 20, 2005 (SEC File No. 333-106529).
*4.2	Form of $6^3/_8$% Senior Notes due 2015 (incorporated by reference to Exhibit 10.1 to the Current Report on Form 8-K of the DIRECTV Holdings LLC filed on June 20, 2005 (SEC File No. 333-106529).
*4.3	Supplemental Indenture dated as of April 28, 2006 by and among LABC Productions, LLC, DIRECTV Holdings LLC, DIRECTV Financing Co, Inc., the Guarantors signatory thereto and The Bank of New York, as trustee (incorporated by reference to Exhibit 4.8 to the Form 10-K of The DIRECTV Group, Inc. filed March 1, 2007 (SEC File No. 1-31945))
*4.4	Indenture, dated as of May 14, 2008, by and among DIRECTV Holdings LLC, DIRECTV Financing Co, Inc., the Guarantors signatory thereto and The Bank of New York, as trustee (incorporated by reference to Exhibit 10.1 to the Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed May 16, 2008 (SEC File No. 333-106529))

Table of Contents

DIRECTV HOLDINGS LLC

Exhibit Number	Exhibit Name
*4.5	Form of $7^{5}/_{8}$% Senior Notes due 2016 (incorporated by reference to Exhibit 4.10 to the Form 10-K of The DIRECTV Group, Inc. filed February 26, 2009 (SEC File No. 1-31945))
*4.6	Indenture, dated as of September 22, 2009, by and among DIRECTV Holdings LLC, DIRECTV Financing Co, Inc., the Guarantors signatory thereto and The Bank of New York Mellon Trust Company, N.A., as trustee (incorporated by reference to Exhibit 10.1 of the Form 8-K of DIRECTV Holdings, LLC filed on September 25, 2009 (SEC File No. 333-106529)).
*4.7	Form of $4^{3}/_{4}$% Senior Notes due 2014 (included in Exhibit 4.6).
*4.8	Form of $5^{7}/_{8}$% Senior Notes due 2019 (included in Exhibit 4.6)
*10.1	Intellectual Property License Agreement dated as of February 10, 2003, between HEC and DIRECTV Enterprises, LLC, as licensee (incorporated by reference to Exhibit 10.16 of the Form S-4 of DIRECTV Holdings LLC filed June 26, 2003 (SEC File No. 333-106529)).
*10.2	Credit Agreement dated as of April 13, 2005 by and among DIRECTV Holdings LLC, Bank of America, N.A., as Administrative Agent and Collateral Agent, the lenders party to the Credit Agreement, certain subsidiaries of the DIRECTV Holdings LLC, as guarantors, JP Morgan Chase Bank, N.A., as Syndication Agent, Credit Suisse First Boston, Goldman Sachs Credit Partners, L.P. and Citicorp North America, Inc. as Co-Documentation Agents, and Banc of America Securities LLC and J.P. Morgan Securities Inc., as Co-Lead Arrangers and Co-Book Managers (incorporated by reference to Exhibit 10.1 to the Current Report on Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed April 13, 2005 (SEC File No. 333-106529)).
*10.3	Security Agreement, dated as of April 13, 2005, by and among DIRECTV Holdings LLC, its subsidiaries named therein as grantors and Bank of America, N.A., as Collateral Agent (incorporated by reference to Exhibit 10.2 to the Current Report on Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed April 13, 2005 (SEC File No. 333-106529)).
*10.4	Pledge Agreement, dated as of April 13, 2005, by and among DIRECTV Holdings LLC, its subsidiaries named therein as pledgors and Bank of America, N.A., as Collateral Agent (incorporated by reference to Exhibit 10.3 to the Current Report on Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed April 13, 2005 (SEC File No. 333-106529)).
*10.5	Amendment No. 1, dated as of May 14, 2008, by and among DIRECTV Holdings LLC, the Guarantors and Lenders signatory thereto and Bank of America, N.A. as Administrative Agent (incorporated by reference to Exhibit 10.3 of the Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed May 16, 2008 (SEC File No. 333-106529)).

Table of Contents

DIRECTV HOLDINGS LLC

Exhibit Number	Exhibit Name
*10.6	Tranche C Term Loan Joinder Agreement, dated as of May 14, 2008, by and among DIRECTV Holdings LLC and Bank of America, N.A., as Administrative Agent and Collateral Agent (incorporated by reference to Exhibit 10.4 of the Form 8-K of DIRECTV Holdings LLC and DIRECTV Financing Co., Inc. filed May 16, 2008 (SEC File No. 333-106529)).
*10.7	Registration Rights Agreement dated as of September 22, 2009, by and among DIRECTV Holdings LLC, DIRECTV Financing Co., Inc., the Guarantors signatory thereto and the Initial Purchasers named therein (incorporated by reference to Exhibit 10.2 of the Form 8-K of DIRECTV Holdings LLC filed on September 25, 2009 (SEC File No. 333-106529))
***31.1	Certification of Chief Executive Officer of DIRECTV Holdings LLC pursuant to Section 302 of the Sarbanes-Oxley Act of 2002 ("Section 302").
***31.2	Certification of Chief Financial Officer of DIRECTV Holdings LLC pursuant to Section 302.
***31.3	Certification of Chief Executive Officer of DIRECTV Financing Co., Inc. pursuant to Section 302.
***31.4	Certification of Chief Financial Officer of DIRECTV Financing Co., Inc. pursuant to Section 302.
***32.1	Certification of the Chief Executive Officer of DIRECTV Holdings LLC pursuant to 18 U.S.C. Section 1350, As Adopted Pursuant to Section 906 of the Sarbanes-Oxley Act of 2002 ("Section 906").
***32.2	Certification of the Chief Financial Officer of DIRECTV Holdings LLC pursuant to Section 906.
***32.3	Certification of the Chief Executive Officer of DIRECTV Financing Co., Inc. pursuant to Section 906.
***32.4	Certification of the Chief Financial Officer of DIRECTV Financing Co., Inc. pursuant to Section 906.
****101.INS	XBRL Instance Document
****101.SCH	XBRL Taxonomy Extension Schema Document
****101.CAL	XBRL Taxonomy Extension Calculation Linkbase Document
****101.DEF	XBRL Taxonomy Extension Definition Linkbase Document
****101.LAB	XBRL Taxonomy Extension Label Linkbase Document
****101.PRE	XBRL Taxonomy Extension Presentation Linkbase Document

Table of Contents

DIRECTV HOLDINGS LLC

---

* Incorporated by reference.

** Filed herewith.

*** Furnished not filed.

**** Pursuant to Rule 406T of Regulation S-T, these interactive data files are deemed not filed or part of a registration statement or prospectus for purposes of Sections 11 or 12 of the Securities Act of 1933 or Section 18 of the Securities Exchange Act of 1934 and otherwise are not subject to liability.

A copy of any of the exhibits included in this Annual Report on Form 10-K, other than those as to which confidential treatment has been granted by the Securities and Exchange Commission, upon payment of a fee to cover the reasonable expenses of furnishing such exhibits, may be obtained by written request to us at the address set forth on the front cover, attention General Counsel.

***

Table of Contents

DIRECTV HOLDINGS LLC

SCHEDULE II—VALUATION AND QUALIFYING ACCOUNTS

Description	Balance at Beginning of year	Additions Charged to costs and expenses	Additions Charged to other accounts	Deductions	Balance at end of year
			(Dollars in Millions)		
For the Year Ended December 31, 2009					
Allowances Deducted from Assets					
Accounts receivable	$ (32)	$ (200)	$ (238)[(a)]	$ 441[(b)]	$ (29)
For the Year Ended December 31, 2008					
Allowances Deducted from Assets					
Accounts receivable	$ (39)	$ (181)	$ (192)[(a)]	$ 380[(b)]	$ (32)
For the Year Ended December 31, 2007					
Allowances Deducted from Assets					
Accounts receivable	$ (39)	$ (172)	$ (158)[(a)]	$ 330[(b)]	$ (39)

(a) Primarily reflects the recovery of accounts previously written-off.

(b) Primarily relates to accounts written-off.

Reference should be made to the Notes to the Consolidated Financial Statements.

Table of Contents

DIRECTV HOLDINGS LLC

SIGNATURES

Pursuant to the requirements of Section 13 or 15(d) of the Securities Exchange Act of 1934, the Registrant has duly caused this report to be signed on its behalf by the undersigned, hereunto duly authorized.

DIRECTV HOLDINGS LLC
(Registrant)

By: /s/ PATRICK T. DOYLE

Patrick T. Doyle
Executive Vice President and
Chief Financial Officer

Date: February 25, 2010

Pursuant to the requirements of the Securities Exchange Act of 1934, this report has been signed below on this 25th day of February 2010 by the following persons on behalf of the Registrant and in the capacities indicated.

/s/ MICHAEL D. WHITE Michael D. White	President, Chief Executive Officer and Director	}Principal Executive Officer
/s/ PATRICK T. DOYLE Patrick T. Doyle	Executive Vice President and Chief Financial Officer	}Principal Financial Officer
/s/ JOHN F. MURPHY John F. Murphy	Senior Vice President, Controller, Chief Accounting Officer and Director	}Principal Accounting Officer
/s/ LARRY D. HUNTER Larry D. Hunter	Executive Vice President, General Counsel and Director	

Table of Contents

DIRECTV HOLDINGS LLC

SIGNATURES

Pursuant to the requirements of Section 13 or 15(d) of the Securities Exchange Act of 1934, the registrant has duly caused this report to be signed on its behalf by the undersigned, hereunto duly authorized.

DIRECTV FINANCING CO., INC.
(Registrant)

By: /s/ PATRICK T. DOYLE

Patrick T. Doyle
Executive Vice President and
Chief Financial Officer

Date: February 25, 2010

Pursuant to the requirements of the Securities Exchange Act of 1934, this report has been signed below on this 25th day of February 2010 by the following persons on behalf of the registrant and in the capacities indicated.

/s/ MICHAEL D. WHITE Michael D. White	President and Chief Executive Officer	}Principal Executive Officer
/s/ PATRICK T. DOYLE Patrick T. Doyle	Executive Vice President and Chief Financial Officer	}Principal Financial Officer
/s/ JOHN F. MURPHY John F. Murphy	Senior Vice President, Controller and Chief Accounting Officer	}Principal Accounting Officer
/s/ LARRY D. HUNTER Larry D. Hunter	Executive Vice President, General Counsel and Director	

Table of Contents

DIRECTV HOLDINGS LLC

EXHIBIT INDEX

Exhibit Number	Exhibit Name
31.1	Certification of Chief Executive Officer of DIRECTV Holdings LLC pursuant to Section 302 of the Sarbanes-Oxley Act of 2002 ("Section 302").
31.2	Certification of Chief Financial Officer of DIRECTV Holdings LLC pursuant to Section 302.
31.3	Certification of Chief Executive Officer of DIRECTV Financing Co., Inc. pursuant to Section 302.
31.4	Certification of Chief Financial Officer of DIRECTV Financing Co., Inc. pursuant to Section 302.
32.1	Certification of the Chief Executive Officer of DIRECTV Holdings LLC pursuant to 18 U.S.C. Section 1350, As Adopted Pursuant to Section 906 of the Sarbanes-Oxley Act of 2002 ("Section 906").
32.2	Certification of the Chief Financial Officer of DIRECTV Holdings LLC pursuant to Section 906.
32.3	Certification of the Chief Executive Officer of DIRECTV Financing Co., Inc. pursuant to Section 906.
32.4	Certification of the Chief Financial Officer of DIRECTV Financing Co., Inc. pursuant to Section 906.
101.INS	XBRL Instance Document
101.SCH	XBRL Taxonomy Extension Schema Document
101.CAL	XBRL Taxonomy Extension Calculation Linkbase Document
101.DEF	XBRL Taxonomy Extension Definition Linkbase Document
101.LAB	XBRL Taxonomy Extension Label Linkbase Document
101.PRE	XBRL Taxonomy Extension Presentation Linkbase Document

A copy of any of the exhibits included in this Annual Report on Form 10-K, other than those as to which confidential treatment has been granted by the Securities and Exchange Commission, upon payment of a fee to cover the reasonable expenses of furnishing such exhibits, may be obtained by written request to us at the address set forth on the front cover, attention General Counsel.

Broadband decisions:
What drives consumers to switch – or stick with – their broadband Internet provider

FCC Working Paper
December 2010

SUMMARY OF FINDINGS

Summary of Findings

The Federal Communications Commission's April 2010 survey sought to understand the context surrounding people's decisions about their home broadband service. Specifically, the survey explored the considerations people have in mind when contemplating or making a change in broadband service. This involved asking people who have not switched service providers whether they considered doing so and what influenced their decision to stick with their provider. For those who have switched service, the survey inquired about why they did so and how the process of changing service went for them.

The survey found that, at least in the prior three years, a minority of home broadband subscribers switched service providers. The survey also found that there are things that inhibit users from switching service, such as the need to pay set-up fees for new service and the basic hassle of making a switch. Those who have switched broadband providers are typically looking for better price or performance, and very few switched because they want additional features from their providers such as more email accounts or online storage.

Here are the survey's main findings:

Just over one-third of Internet users changed their service provider in the prior three years.

- When asked whether they had switched service in the prior three years, 36% of Internet users had done so, while 62% had not.
  - 23% of online users switched once in this time frame.
  - 13% switched more than once.

- Of those who had switched, 43% also moved residences within the prior three years.

Among those who have not switched Internet service providers (ISPs) in the prior three years, few have considered making a change in ISP service.

- 30% of those who have not switched ISPs have considered switching, 13% *very seriously* and 17% *somewhat seriously*.

- 69% of those who have not switched ISPs say they have not considered it; 54% have *not at all seriously* considered changing ISPs and 15% have *not too seriously* considered it.

Among broadband users who say they have a choice of Internet service providers, a minority would seriously consider switching to a provider in their area and a majority thought it would be at least somewhat easy to do.

- Just 21% of broadband adopters with a choice of more than one provider say they would seriously consider presently changing their ISP.

- 63% of broadband adopters with a choice of multiple providers said it would be easy to switch providers, with 33% saying it would be *very easy* and 30% saying it would be *somewhat easy*.

Financial and non-financial factors, such as installation fees or the hassle of getting new service, can inhibit consumers from changing service.

When asked about things that might keep them from switching service, broadband users with the choice of more than one provider said the following:

- 50% said paying set-up or installation fees were *major* factors in keeping service.
- 43% said dealing with the hassle of getting new service installed was a *major* reason they have kept service.
- 40% said putting down a deposit for new service was a *major* reason for keeping their service.
- 39% said that having to change their current bundle of Internet, TV, and phone service was a *major* reason for keeping service.
- 34% said having to give up their current email address from their ISP was a *major* reason for not changing service.
- 32% said paying termination fees to their current ISP was a *major* reason for keeping service.

Internet users who have switched service in the past three years cite a desire for better service or price as a reason for their change. Very few of those who changed service said it was because they wanted more features from their provider, such as more email accounts or online storage.

Among the 36% of Internet users who have switched service in the past three years, here is what they say when asked to identify the *major reason* for the change:

- 49% said the desire for a faster or higher-performance Internet connection was a *major* reason for the switch.
- 47% said getting a better price on service was the *major* reason behind the change.
- 39% said getting a bundle of Internet, TV, and phone services from a single company was the *major* reason for the switch.
- 28% said poor customer service from their old ISP was the *major* reason for the change.
- 9% said getting more features such as more email accounts or online storage was the *major* reason for the switch.

Introduction

The Federal Communications Commission's April 2010 survey sought to explore the context in which consumers make decisions about broadband service at home.[1] In doing so, the survey asked home broadband users what factors are important to them in choosing a provider; whether they have considered switching their home broadband providers; and whether, in the past three years, they have in fact changed their home Internet service provider.

The National Broadband Plan (NBP) suggests the reasons why undertaking such a survey is important. Although the analysis underpinning the NBP went to great lengths to explain why some one-third of adult Americans do not have broadband at home, the NBP also emphasized the need for policymakers to understand better what shapes adoption choices. The NBP observes that many fixed broadband users "have little information about the actual speed and performance of the service they purchase" and goes on to note research gaps in understanding "price and service terms and conditions."[2]

The survey results reported here, in conjunction with earlier reports on users' perspectives on broadband speed and on bill shock and early termination fees for cell phones, fulfill the NBP's promise that the FCC would field a survey on these issues and produce a report analyzing results.[3]

Overview of adoption

The April 2010 survey contained standard questions that seek to determine whether the respondent is an Internet user, whether he or she uses the Internet at home and, if so, whether the home connection type is broadband.[4] The survey found that 69% of adult Americans are Internet users, with 57% of adults having a high-speed Internet connection at home and just 5% having a dial-up connection. Some 6% of adults are online users but do not access the Internet from home.

These figures differ from those contained in the FCC's October-November 2009 survey, which found that 78% of adults were Internet users, with 65% of adults having home high-speed connections. This drop in home broadband connections is puzzling, though not inconsistent with other research. The Pew Research Center's Internet & American Life Project's December 2009 survey found that 60% of Americans had broadband at home, a slight decline from the 63% figure registered in April 2009. Pew had also found in April 2009 that some Americans (7%) had cut back on home Internet service costs in

---

[1] The FCC survey was conducted between April 19 and May 2, 2010 and interviewed 3,005 adults in English or, if the respondent chose, Spanish. The margin of error is plus or minus 2 percentage points for the entire sample.

[2] *Connecting America: The National Broadband Plan*. Chapter 4, p.44.

[3] *Ibid*., p. 64, footnote 53.

[4] See John B. Horrigan, *Broadband Adoption and Use in America.* OBI Working Paper Series No. 1. See page 14 for discussion of how connection type at home was determined in the survey.

the prior year. The persistence of the economic recession may have heightened that phenomenon and is, perhaps, reflected in the FCC's latest survey.

More recently, the Pew Internet Project found that broadband adoption changed little from 2009 to 2010, with 66% of adults reporting that they had broadband at home in an April 2010 survey.[5] The different findings for broadband adoption from the FCC in April 2010 (57%) and the Pew Internet Project at the same time (66%) are significant. The FCC and Pew frame questions differently to determine broadband adoption and it is possible some of the difference is attributable to that fact. The Pew sample also did not conduct interviews in Spanish; because Hispanics who opt to take a survey in Spanish have lower broadband adoption rates than those who do not, the results from the Pew survey are likely to be several percentage points higher than would be the case with a Spanish-language option.[6] Whatever the reasons for the differences – the economy, question wording, or sample – it seems clear that the days of rapid broadband adoption growth are, for now at least, past.

Frequency of switching home Internet service

The survey asked adults whether they have switched Internet service providers (ISPs) in the past three years and, if they have not switched, whether they have considered switching. Overall, 62% of home Internet users have *not* switched ISPs in the past three years, while 36% have done so one or more times.[7] Among home broadband users, these figures are essentially the same, with 62% having not switched and 37% having done so in the past three years.[8] Specifically:

- 23% of home Internet users have switched ISPs once in the past three years;
- 10% have switched twice; and
- 3% have switched more than twice.

The survey also inquired about whether the respondent had changed residences in the past three years, which would indicate whether the switch was prompted by moving to a new home or some other reason. Some 29% of respondents said they had changed residences in the past three years and, of those who moved, 50% also changed their Internet service provider. This means that those who moved and switched service account for 43% of all those who switched their ISP in the past three years.

Assuming an even distribution of switching in the three-year time horizon respondents thought about when asked the question, this means roughly 17% switch ISPs in a given

[5] Aaron Smith, *Home Broadband Adoption 2010*. Pew Internet & American Life Project, August 2010. Available online at: http://www.pewInternet.org/Reports/2010/Home-Broadband-2010.aspx
[6] In the October-November 2009 survey conducted for the National Broadband Plan, which did include a Spanish-language option, 65% of the entire sample had broadband at home. For English-speakers, 67% had broadband at home, while 20% of those who took the survey in Spanish had broadband at home.
[7] Roughly 2 percent of home Internet users did not know or refused to answer the question.
[8] The discrepancies between findings for broadband users and all users are small because the overwhelming majority of home Internet users (90%) have broadband connections.

year, with roughly 7% have switched and changed their residence at the same time.[9] A later section of this report discusses the reasons why users switch ISPs.

To put this rate of service churn in context, 19% of cell phone users have changed service in the prior three years according to this survey. However, the FCC's *Mobile Competition Report* notes that monthly churn rates run between 1.5% and 3.3% per month, indicating annual churn could run between 18% and approximately 40%.[10]

Thinking about switching service

This section focuses on the 62% of respondents who have not switched ISPs in the past three years. For this group, the survey sought to understand the context for that answer – how seriously the respondent considered switching, whether the respondent believes there was a choice of provider in his or her area, and what factors were important in a choice not to switch.

As to how seriously home online users considered switching service, some two-thirds (69%) of those who have not switched ISPs in the past three years said they had not seriously considered switching. That breaks down to 54% who "not at all seriously" considered switching and 15% who "not too seriously" considered switching. That leaves, among those who have not switched ISPs in the prior three years, nearly one-third (30%) who did consider switching. In this group, 13% "very seriously" considered switching and 17% "somewhat seriously" considered switching. The figures for the narrower set of home broadband users do not differ from the ones for all home Internet users.

Shifting from the past to the present, the survey asked respondents whether they might *now* seriously consider switching to another ISP. To get at this question, the survey walked respondents through questions that sought to determine whether more than one broadband provider serves their area.

Those who said they had *more* than one provider where they lived were then asked if there was a broadband provider in their service area to which they would presently *seriously* consider switching. Among broadband adopters who believed they had more than one provider available – and that came to 71% of all home broadband subscribers – 21% said they would consider switching ISPs, while 75% said they would not.

The past switching behavior of those who say they would today consider a switch in ISP (again, among those who say they have more than one provider) is in line with that of the overall population of broadband users. Among the 21% of broadband users with multiple service providers in their area who would presently seriously consider switching, 35% said they switched their ISP in the past three years. That compares with the 36% average for all home Internet users.

---

[9] This roughly takes into account those users who report switching ISPs multiple times over the three-year time horizon.
[10] FCC 10-81, available online at: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-81A1.pdf, paragraph 245.

Approximately one-third of broadband users with a choice of service providers thought it would be difficult or impossible to switch, while approximately two-thirds thought it would be somewhat or very easy to switch. Specifically:

- 33% said it would be *very easy* to switch ISPs;
- 30% said it would be *somewhat easy*;
- 21% said it would be *somewhat difficult*;
- 6% said it would be *very difficult*; and
- 4% said they *could not do it or it would be impossible*.

There are differences among broadband users in response to this question depending on how seriously they might consider switching. Those who say they would consider switching are somewhat less likely to say it would be easy to switch. The following table shows the comparison.

Table 1. Anticipated ease of switching broadband provider

	Would consider switching broadband provider	Would ***not*** consider switching broadband provider
Very Easy	25%	36%
Somewhat Easy	33	30
Somewhat Difficult	27	19
Very Difficult	7	6
Could not do it/impossible	3	5

On its face, it may seem paradoxical that those who would consider switching providers are more likely to find switching difficult. However, it is possible that those who have considered switching have looked into it more closely than those who have not – and as a result have found it to be a more involved process than those with less information.

Overall, broadband users do not exhibit a high rate of churn (about 17% per year), nor do non-switchers indicate that they are likely switchers. About one-third (30%) say they have considered switching in the past three years. Among those with a choice of broadband service providers, 21% say they would seriously consider switching today.

Any number of factors could be behind respondents' reasons for sticking with their provider. The survey probed this further by asking about financial factors that might inhibit switching and non-monetary costs associated with switching. The following table shows the responses to the question among broadband users with the choice of more than one provider:

Table 2. Factors in the decision to stay with current provider among broadband users with a choice of provider

	Major	Minor	Not	DK	Ref.
Paying set-up or installation fees to get new service	50	25	24	1	*
Dealing with the hassle of getting a new service installed	43	28	28	1	*
Putting down a deposit to get a new service	40	29	29	2	*
Having to change your current bundle of Internet, TV and phone services	39	22	35	3	*
Having to give up your current email address from your Internet provider	34	19	46	1	*
Paying termination fees to your current Internet company	32	26	38	3	*

The typical broadband user who has a choice of more than one provider cites two of the reasons listed above as major factors in sticking with their current provider; 28% cite four or more reasons.

There are some differences in responses depending on whether broadband users say they have seriously considered switching providers or not.

Table 3. Factors in the decision to stay with current provider among broadband users with a choice of provider, by consumer type (% citing issue as "major" reason)

	Would consider switching broadband provider	Would ***not*** consider switching broadband provider
	% citing issue as "major" reason	
Paying set-up or installation fees to get new service	53	48
Dealing with the hassle of getting a new service installed	47	41
Putting down a deposit to get a new service	43	39
Having to change your current bundle of Internet, TV and phone services	44	38
Having to give up your current email address from your Internet provider	33	34
Paying termination fees to your current Internet company	40	30

As Table 3 shows, people who have considered switching, but have not switched, are generally more likely to perceive barriers to switching. For them, financial reasons loom large, as they are more likely than other respondents to worry about paying a termination fee and set-up or installation costs. Non-monetary factors also come into play, the hassles of dealing with installation and changing bundles are greater issues for them.

People who have switched service

As noted, 36% of broadband users have switched service in the past three years, with 43% of these switchers having done so in conjunction with a change in residence. When asked to think about the reasons for their last ISP switch, here is what all respondents said:

Table 4. Reasons why people switched ISPs (among those who have switched in the past three years)

	Major	Minor	Not	DK	Ref.
Getting a faster or higher performance Internet connection	49	20	29	3	0
Getting a better price for Internet service	47	16	34	3	0
Getting a bundle of Internet, TV and phone services from a single company	39	15	44	2	*
Any other MAJOR reason that I have not mentioned	15	0	14	65	6
Poor customer service from your old Internet provider	28	12	57	3	0
Getting more features such as added email accounts or online storage	9	18	71	2	*

Price and speed are the most prominent reasons switchers cite for changing their service, with nearly half saying this. In fact, two-thirds (67%) of switchers cite either price or speed as a major reason behind their decision to change ISPs. Many (39%) say getting a bundle is a motivation for changing ISPs. Few switchers – just 9% – say that getting added features from their ISP prompted the change. Poor customer service is a major reason for 28% of switchers, with another 12% saying it was a minor one.

Reasons given differ significantly depending on whether the respondent's service switch was accompanied by a change in residence or not.

Table 5. Reasons why people switched ISPs, among those who have switched in the past three years (% citing issue as "major" reason)

	Change in ISP **did not** involve change in residence	Change in ISP **did** involve change in residence
Getting a faster or higher performance Internet connection	55%	40%
Getting a better price for Internet service	54	39
Getting a bundle of Internet, TV and phone services from a single company	44	31
Any other MAJOR reason that I have not mentioned	15	17
Poor customer service from your old Internet provider	31	24
Getting more features such as added email accounts or online storage	9	8

Performance and price are also leading reasons for switching for cell phone users. Among cell phone users who have changed providers in the past three years:

- 49% said they wanted to get a better signal in places they use their cell phone;
- 47% wanted to pay less per month for service;
- 39% wanted to get a new cell phone;
- 32% said they received poor customer service from their old provider; and
- 10% switched so they could add Internet access to their cell phone.

It is worth noting that the reasons for switching match up reasonably well with the reasons all broadband users (*i.e.*, not just those who have switched) cite for their choice of provider. When home broadband users were asked about the reasons they chose their current ISP:

- 50% said the monthly quoted price was the *major* reason for the choice.
- 43% said the advertised connection speed was the *major* reason behind the choice.
- 42% said the bundle of Internet, TV, and phone service was the *major* reason for the choice.

For the most part, switchers found doing so easy, with 56% saying it was "very easy" and 30% saying it was "somewhat easy," with 10% finding it "somewhat difficult" and 3% "very difficult". These figures are very much in line with figures for those who have switched cell phone providers in the past three years; among that group, 56% said switching was "very easy" and 28% said it was "somewhat easy."

When asked about the process of changing service providers, those who have switched in the past three years said the following:

- 49% said they had to pay a set-up, installation, or equipment fee to their new company;
- 37% said that they or someone else had to spend considerable time waiting at home for the installation;
- 29% said they had to wait more than a week before new service was installed;
- 9% said they had to put down a deposit to qualify for service from the new company; and
- 9% said they had to pay a termination fee to the old company.

Among broadband users who had to pay a termination fee, only two-thirds could identify the fee's amount. The picture is a bit different for set-up or installation fees. Many (43%) either were not subject to a fee or, if they were, did not know its level; 25% fell into that category. For those who switched ISPs in the past three years:

- 12% said their set-up fee was between $1 and $49;
- 11% said it was between $50 and $99;
- 6% said it was between $100 and $149;

- 1% said it was between $150 and $199; and
- 2% said it was in excess of $200.

I/S: A JOURNAL OF LAW AND POLICY FOR THE INFORMATION SOCIETY

"Net Neutrality," Non-Discrimination and Digital Distribution of Content Through the Internet*

NICHOLAS ECONOMIDES**

Abstract: The vast majority of U.S. residential consumers face a monopoly or duopoly in broadband Internet access. Until now, the Internet has been characterized by a regime of "net neutrality," which means there has been no discrimination between the price of transmitting packets based on the identity of either the transmitter or the identity of the receiver, based on the application, or the type of content the packet contains. Providers of DSL or cable modem Internet access in the United States are taking advantage of a recent regulatory change that effectively abolishes "net neutrality" and non-discrimination protections. Due to their market power, these service providers are considering a variety of discriminatory pricing schemes. This article discusses and evaluates the effect a number of these schemes would have on the prices and profitability of network access, as well as the effect on complementary application and content providers. This article also discusses an assortment of anti-competitive effects created by price discrimination and evaluates the possibility of "net neutrality" being imposed by law.

---

* In an earlier draft of this article, I benefited from comments by the following individuals: Carla Bulford, Richard Clarke, David Gabel, Bill Sharkey, Peter Shane, Scott Shenker, Brian Viard, and Glenn Woroch. I gratefully acknowledge financial support from the Newhouse Foundation and the Entertainment, Media, and Technology program of the Stern School of Business.

** Professor of Economics, Stern School of Business, N.Y.U., 44 West 4th Street, New York, NY 10012, (212) 998-0864, economides@stern.nyu.edu, http://www.stern.nyu.edu/networks/, and Executive Director, NET Institute, http://www.NETinst.org.

I. INTRODUCTION

The Internet is a global, interconnected network of computers that allows data transfers and provides a variety of interactive, real-time and time-delayed telecommunications services. Internet communications are based on common, public protocols. Hundreds of millions of computers are connected to the Internet at any moment. The vast majority of computers connect to the Internet through commercial Internet Service Providers ("ISP"s).[1] Users connect to the Internet through ISP dial-ups, cable modems connections, residential Digital Subscriber Lines ("DSL"), or through corporate networks (Local Area Networks ("LAN"s)). Ninety-eight percent of domestic residential broadband customers access the Internet through DSL or a cable modem.[2] Only about half of residential consumers have a choice between even two providers. Typically, the routers and switches owned by the ISP send the caller's packets to a local Point of Presence ("POP") on the Internet. In dial-up, cable modem, and DSL, the access POPs, as well as corporate networks dedicated access circuits, connect to high-speed hubs. Generally, access POPs (which serve dial-up, cable modem and DSL connections) and corporate networks with dedicated access circuits connect to high-speed hubs. High-speed circuits, leased from or owned by telephone companies, connect the high-speed hubs, forming an Internet Backbone Network ("IBN").

The Internet is the primary global network for digital communications. A number of different services are provided on the Internet, including, among numerous others, e-mail servers, browser interfaces (using Internet Explorer, Firefox, Opera, or others), Peer-to-Peer file exchange services, and Internet telephony (Voice over Internet Protocol ("VOIP")). A number of software applications run on top of the Internet browser, including information services (Google, Yahoo, MSN), image displays, video transmissions and others. Since the advent of Mosaic, the first Internet browser, in 1993, the Internet has evolved beyond text-based interface to support images, sound, and video transmitted in digital format. Even full-length movies are regularly downloaded, rented, or sold through

---

[1] Educational institutions and government departments are also connected to the Internet but do not offer commercial ISP services.

[2] See Senate Committee on Commerce, Science, and Transportation, *Hearing on "Network Neutrality"* (testimony of Vinton G. Cerf), 109th Cong., 1st sess., 2006, http://commerce.senate.gov/pdf/cerf-020706.pdf (accessed April 10, 2008).

commercial services over the Internet and viewed on personal computers or television sets.

As video services and the digital distribution of content over the Internet grow, Internet broadband access providers including AT&T, Verizon, and a number of cable TV companies, have recently demanded additional compensation for carrying digital services. Ed Whitacre, the Chief Executive Officer of AT&T, expressed his company's dislike of existing regulatory structures: "Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it."[3]

The claim that consumers, content providers, or applications providers use the Internet for free is certainly incorrect.[4] Currently, users pay ISPs for access to the Internet. Similarly, ISPs pay fees to Internet backbones for access to the Internet.[5] ISPs pay per month for

---

[3] "Online Extra: At SBC, It's All About 'Scale and Scope,'" *BusinessWeek*, November 7, 2005, http://www.businessweek.com/@@n34h*IUQu7KtOwgA/magazine/content/05_45/b3958092.htm (accessed April 10, 2008).

Interview of Ed Whitacre:

> Q. How concerned are you about Internet upstarts like Google (GOOG), MSN, Vonage, and others?
>
> A. How do you think they're going to get to customers? Through a broadband pipe. Cable companies have them. We have them. Now what they would like to do is use my pipes free, but I ain't going to let them do that because we have spent this capital and we have to have a return on it. So there's going to have to be some mechanism for these people who use these pipes to pay for the portion they're using. Why should they be allowed to use my pipes?
> The Internet can't be free in that sense, because we and the cable companies have made an investment and for a Google or Yahoo! (YHOO) or Vonage or anybody to expect to use these pipes [for] free is nuts!

[4] Of course, the categories of consumers, content providers and applications providers intersect since a consumer could also be providing content to some extent. In making the distinction between these three categories of Internet participants I define them by their primary function.

[5] This service is called "transit." See Nicholas Economides, "The Economics of the Internet Backbone," in *Handbook of Telecommunications*, ed. S. Majumder, et al., 379–381 (New York, NY: Elsevier B.V. 2005), http://www.stern.nyu.edu/networks/Economides_ECONOMICS_OF_THE_INTERNET_BACKBONE.pdf (accessed April 10, 2008); Nicholas Economides, "The Economics of the Internet," in *The New Palgrave Dictionary of Economics* (forthcoming), http://www.stern.nyu.edu/networks/Economides_

a virtual "pipe" of a certain bandwidth, according to their expected use.[6] When digital content (or information packets of any service) is downloaded by consumer *A* from provider *B*, both *A* and *B* pay. *A* pays his ISP through his monthly subscription, and *B* pays similarly. In turn, ISPs pay their respective backbones through their monthly subscriptions. Unlike a traditional telephone call arrangement in which only the calling party pays, Internet backbones collect from both sides of a communication.

So, what change would AT&T's CEO like to see in the pricing and industry structure? He desires the abolition of "net neutrality," the regime that does not distinguish in terms of price between bits or information packets according to the services that they provide, and additionally fails to distinguish in price based on the identities of the uploader and downloader. This pricing regime has prevailed since the inception of the commercial Internet.[7] Presently, an information packet used for VOIPs, email, images, or video is priced equally as a part of the large number of packets that correspond to the subscription services of the originating and terminating ISPs.

In addition to content neutrality, there is no distinction made according to the identities of the uploader and downloader. AT&T, Verizon, and cable Internet access providers would like to abolish the regime of "net neutrality" and in its place substitute a pricing schedule that charges both the final customer for his or her basic transmission service and the transmission's originating party (such as Google, etc.) for the provision of content. An access network, for example AT&T, wants to charge fees to an originating party even when the originating party does not connect to the Internet using AT&T and therefore does not have any contractual relationship with AT&T. Access network operators have also reserved the right to charge differently based on the identity of the provider even for the same type of packets; for example, an ISP may charge Google more than Yahoo for the same transmission. The proposed Internet model, without "net neutrality," would more closely mirror the traditional pre-Internet

---

Economics_of_the_Internet_for_Palgrave.pdf (accessed April 8, 2008). In addition to transit service, Internet backbones of comparable size "peer" with each other, which means that they agree not to exchange money for exchanged traffic.

[6] See Economides, *The Economics of the Internet Backbone*, Table 5.

[7] We disregard pricing issues in the pre-commercial Internet when it was first primarily a network among military contractors and later a network among primarily academic communities.

telecommunications model in which customers pay *per service*.[8] This would be a very sharp departure from the way the Internet was designed to operate and how it has run since its inception (that is, pricing without reference to particular services or functions of the transmitted information packets).

After the acquisition of AT&T by Southwestern Bell ("SBC")[9] and of Microwave Communications Inc. (MCI) by Verizon, enabled by a change in regulatory rules by the Federal Communications Commission, the resulting consolidated companies (AT&T and Verizon) now advocate price discrimination according to the type of application and the provider used to transmit the content.[10] AT&T, Verizon, and cable TV companies would like to abolish the regime of "net neutrality" and substitute a complex pricing schedule where, besides the basic charge for transmission of bits, there will also be additional charges by the Internet access operator applied to the originating party (such as Google, Yahoo, or MSN). These charges would apply even when the application provider is not directly connected to AT&T or Verizon, that is, even when Google's ISP is not AT&T or Verizon.[11]

The broadband Internet access providers' new pricing scheme will most likely impose price discrimination on the provider side of the market and not on the subscriber. That is, the change will implement two-sided pricing. This is uniquely possible for firms operating within a network structure. Outside of traditional networks, such two-sided pricing is also made possible by the intermediaries operating between trading parties in exchange networks (such as the exchanges themselves).[12] There is presently considerable debate over the

---

[8] See Nicholas Economides, *Telecommunications Regulation: An Introduction*, in *The Limits and Complexity of Organizations*, ed. Richard R. Nelson, 48–76 (New York, NY: Russell Sage Foundation Press, 2005), http://www.stern.nyu.edu/networks/Economides_Telecommunications_Regulation.pdf (accessed April 10, 2008). A discussion of the differences between the Internet and earlier digital data networks, and an exposition of traditional telecommunications regulation.

[9] SBC changed its name to AT&T after it acquired AT&T.

[10] Recently, Deutsche Telecom and Telecom Italia have made similar proposals.

[11] See Economides, "Telecommunications Regulation: An Introduction." The proposed Internet model without "net neutrality" would be closer to the traditional pre-Internet telecommunications model where customers pay per service.

[12] See Nicholas Economides, "Competition Policy in Network Industries: An Introduction," in *The New Economy and Beyond: Past, Present and Future*, ed. Dennis Jansen, 112–13 (London: Edward Elgar, 2006), http://www.stern.nyu.edu/networks/Economides_

legality, as well as the efficiency, of the implementation of the proposed changes. There is additional concern due to the considerable market power of such firms.

II. ABOLITION OF NON-DISCRIMINATION REQUIREMENTS

Electronic networks are created by a number of different, complementary levels of necessary operation. The Internet is supported by low-level sets of protocols, primarily Transmission Control Protocol/Internet Protocol ("TCP/IP"). These protocols define three basic levels of functions in the network: (1) the hardware/electronics level of the physical network, (2) the (logical) network level where basic communication and interoperability is established, and (3) the applications/services level.[13] The Internet separates the network interoperability level from the applications/services level. This means that, unlike earlier centralized digital electronic communications networks, such as CompuServe, AT&T Mail, Prodigy, and early AOL, the Internet allows a large variety of applications and services to be run "at the edge" of the network and *not* centrally. This means that users have a tremendous amount of choice: if a user elects to download video, he can do so without asking permission from a central authority in the network. For example, if a user elects to run a spyware-stopper, he may do so according to his preference; the network does not select security software for him.

The tremendous degree of choice of applications and content on the Internet is a direct consequence of its design, in which intelligence, applications, services, and content live "at the edge" of the network and are only dependent on the network for connectivity. A key consequence of "net neutrality" pricing has been successful innovation resulting, for example, in Google, Yahoo, and MSN as well as the large number of applications developed by companies that do not own any network infrastructure. Many companies have been able to innovate at the edge of the network. These innovations include new

---

Competition_Policy.pdf (accessed April 10, 2008), for a discussion of two-sided pricing in a network.

[13] See Richard S. Whitt, "A Horizontal Leap Forward: Formulating a New Communications Public Policy Framework Based on the Network Layers Model," *Federal Communications Law Journal* 56 (May 2004): 587–672; Senate Committee, *Hearing on "Network Neutrality."*

methods of content distribution (both news and entertainment),[14] the distribution and modification of applications (including patching and updates), and the creation of many new applications such as interactive advertising.

Since the beginning of the commercial Internet, Internet pricing did not discriminate with respect to the identity of those receiving information packets, those sending them, or the nature of the information packets and the function they served. The content of the packets and the frequency of interactions are all irrelevant. Networks simply set different prices according to the bandwidth required for transfers. Transmitters and receivers of Internet information packets are charged according to the amount of bandwidth they subscribe to. For example, a residential DSL customer may buy from his ISP a 384Kb per second bandwidth pipe, while a business customer can buy a multiple of the same. Similarly, ISPs are charged—by Internet backbones—subscription fees according to the bandwidth they require/use.

Typically, Internet transmissions are carried over infrastructure owned by telecommunications companies, cable TV companies, and terrestrial satellites. Following the regulatory tradition of the United States, until the summer of 2005, telecommunication-facility-based Internet transmissions were subject to common carrier regulation that included non-discrimination requirements. Other Internet transmissions, those not telecommunication-facility-based, were not subject to common carrier regulation. Thus, DSL service was considered a common carrier service, and therefore subject to non-discrimination provisions. Cable modem service, in contrast, was not considered common carrier service, and therefore did not have to abide by such provisions.

In the summer of 2005, the Federal Communications Commission changed the classification of Internet transmissions from "telecommunications services" to "information services."[15] This implied that there were no longer "non-discrimination" restrictions on Internet service pricing. The remarks of the president of SBC (now AT&T after SBC acquired AT&T in 2005–2006), and similar

---

[14] There are significant changes in many industries because of the Internet. For example, dissemination of news through the Internet has cut radically into the circulation of newspapers and has resulted in a round of consolidations among newspapers.

[15] In mid-2005 the FCC reclassified Internet service to no longer be subject to non-discrimination rules. See *Nat'l Cable & Telecomm. Ass'n v. Brand X Internet Servs*, 125 S. Ct. 2688 (2005).

expressions by Verizon and cable TV companies, underscore the concerns of network infrastructure operators who are keen to extract more of the value generated by the information packets they transport. This value accrues to both final consumers as consumers' surplus[16] and to application or content providers as profits.

It is widely believed that an additional reason for the proposed change is the increasing introduction of video services by AT&T and Verizon. It is expected that video services will congest "last mile" broadband Internet access as it is presently sold. Therefore, AT&T and Verizon would like to set up pricing differentiation so that consumers will buy the content generated by their service provider rather than the content offered by the service provider's competitors. However, broadband access providers have not committed to any restriction on their ability to extract additional surplus from their consumers and content or application providers. In addition, broadband access providers have not committed to restrictions on the use of price discrimination instruments. Industry lobbyists have proposed congressional bills that legalize the ability of an access provider to impose *any* price discrimination scheme it chooses. Presently, residential consumers pay at most $24 billion a year for broadband Internet access, as shown in Section IV. The combination of the consumers' surplus and the profits generated by Internet-distributed complementary applications and Internet-distributed content are a very large multiple of the current cost of residential broadband service. Thus, changes in fee structure proposed by access providers have the potential to seriously disrupt the current distribution of wealth between content, applications, and transmission service providers.

To put the proposed change in perspective, it is useful to understand what unrestricted discriminatory pricing would mean in the context of a traditional telecommunications network. If a telephone company were free from legal restrictions on price discrimination the company could, for example, routinely charge more for phone calls between investment bankers. This additional charge may be "justified" by the company because such phone calls are more likely to generate value than the average phone call. If phone companies were unregulated with respect to price discrimination, they could charge more for fax telephone calls than for other calls, since fax transmissions are likely to be more valuable on average than phone calls. Similarly, a telephone company without a non-discrimination

[16] Consumers' surplus is the difference between what consumers are willing to pay and what they actually pay.

requirement could charge a high price for 911 emergency calls because the willingness to pay for these calls is obviously high.

As discussed above, the Internet under the "net neutrality" model separated the network layer from the applications/services layer. This allowed firms to innovate "at the edge of the network" without seeking approval from network operators.[17] The decentralization of the Internet based on "net neutrality" facilitated innovation resulting in successes such as the creation of the World Wide Web, Google, MSN, Skype, Yahoo, etc. "Net neutrality" also increased competition among the applications and services that operate "at the edge of the network," which did not need to own a network in order to compete. The existence of network effects (the increase in value that each user experiences as more users are added to the network) on the Internet implies that efficient prices to users on both sides (consumers and applications) are lower than they would be in a market without network effects.[18] A departure from "net neutrality" is likely to increase prices, which will reduce network effects and hamper innovation.

III. Detailed Examination of Anti-Competitive Concerns Arising from the Abolition of "Net Neutrality"

A. Horizontal Concerns

The abolition of "net neutrality" raises both horizontal and vertical antitrust and public interest issues. In addition to the pricing issues, there are concerns that network operators will discriminate against certain types of content and political opinions.[19]

---

[17] Vint Cerf, one of the "fathers of the Internet," has called this environment "innovation without permission" of the network. Senate Committee, *Hearing on "Network Neutrality,"* (testimony of Vinton G. Cerf).

[18] See Nicholas Economides, "The Economics of Networks," *International Journal of Industrial Organization* 14 (1996): 675–99, http://www.stern.nyu.edu/networks/Economides_Economics_of_Networks.pdf (accessed April 10, 2008).

[19] See, for example, House Committee on the Judiciary, *Hearing on "Network Neutrality: Competition, Innovation, and Nondiscriminatory Access*," 109th Cong., 2nd sess., 2006 (testimony of Tim Wu), at http://judiciary.house.gov/media/pdfs/wu042506.pdf (accessed April 10, 2008). Wu discusses how Western Union, in the 1860s, when it had a telegraph monopoly, wrote an exclusive contract with the Associated Press that discriminated in price against other news organizations, and that resulted in a near monopoly for the Associated Press.

This section starts with a discussion of the horizontal antitrust concerns. Carriers in the "last mile" to the home have significant market power. Residential retail customers may have difficulty changing ISPs in response to price or quality changes. For 98% of residential consumers in the United States, there are only one or two choices for broadband Internet access: either DSL or cable modem access.[20]

Cable TV broadband Internet service is available to 92% of U.S. households but market penetration is significantly lower.[21] Most cable TV companies offer broadband Internet access only in conjunction with a digital cable TV package.[22] Due to technical limitations, DSL is offered only to households that are close to a local telephone company switch; the capabilities of the connection diminish as the distance from the switch increases. The vast majority of U.S. households cannot buy DSL service (so-called "naked DSL") without at the same time subscribing to voice telephone service on the same line.[23] Even where naked DSL is available, its price often significantly exceeds the price of DSL service that includes voice provision on the same line.

Due to coverage and bundling issues, and the very limited number of residential broadband providers, existing providers, typically AT&T, Verizon, or a cable TV company, have significant market power. The complications of changing equipment, configuration, email addresses, etc., imply significant switching costs for customers. Such costs add to the market power of existing local access providers. Finally, residential customers are affected by bundling of broadband Internet access with other services, such as telecommunications and cable television. However, despite the significant market power and high concentration in the Internet broadband access market, carriers are unable to effectively discriminate in price between monopoly and

[20] Senate Committee on Commerce, Science, and Transportation, *Hearing on "Network Neutrality,"* (testimony of Vinton G. Cerf).

[21] See National Cable and Telecommunications Association, http://www.ncta.com/Statistic/Statistic/ResidentialCableHighSpeedDataSubscribers.aspx (accessed April 10, 2008).

[22] Even when broadband Internet access is offered by itself, it is typically offered at the full price of the bundle of Internet access and digital cable TV combined.

[23] There is no technical requirement for this, and the EU has mandated unbundling of the fixed local telecommunications network that allows DSL to be provided separately from voice service, as well as in its absence.

duopoly customers. Marketing through mass channels constrains carriers by forcing them to set prices for large regions, typically covering multiple states. Some carriers have nationwide pricing. Thus, access carriers with significant market power are unable to extract value from consumers to an extent proportional with their market power.

Carriers have much less market power upstream on the Internet backbone because, despite some concentration, there is a much more egalitarian distribution of market share on the backbone than in the residential access market. Market share of national backbones are listed in Table 1 based on 1999 data and projections. In papers filed in support of the merger of SBC and AT&T, as well as the merger of Verizon with MCI, there was mention of two recent traffic studies by Ryan Hankin Kent Research ("RHK"). These studies, showing traffic for 2004, are summarized in Table 2. The data demonstrate a dramatic change in the ranking of the networks, with AT&T ranked first and MCI fourth in 2004. They also show that a much larger share of traffic (over 40%) is now carried by smaller networks. These latest traffic studies show that earlier concerns, expressed in the European Union ("EU") and by the United States Department of Justice, that the Internet backbone market would tilt to create monopoly situations, have proven overstated.[24]

Table 1. Market Shares of National Internet Backbones[25]

Company	1997	1999	2001 (projected in 1999)	2003 (projected in 1999)
MCI WorldCom	43%	38%	35%	32%
GTE-BBN	13%	15%	16%	17%
AT&T	12%	11%	14%	19%
Sprint	12%	9%	8%	7%
Cable &Wireless	9%	6%	6%	6%
All Other	11%	21%	22%	19%
Total	100%	100%	100%	100%

[24] See Economides, *Competition Policy in Network Industries* for a more detailed discussion of the EU and DOJ concerns regarding the WorldCom-MCI and MCI-Sprint mergers.

[25] Senate Committee on the Judiciary, *Hearing on the MCI WorldCom-Sprint Merger*, 106th Cong., 1st sess., 1999 27–38 (testimony of Tod A. Jacobs, Senior Telecommunications Analyst, Sanford C. Bernstein & Co., Inc.); Bernstein Research, MCI WorldCom (Bernstein Report, March 1999), 51.

Table 2. Carrier Traffic in Petabytes per Month in 2004[26]

Company	Traffic				Market share among all networks
	1Q2004	2Q2004	3Q2004	4Q2004	4Q2004
A (AT&T)	37.19	38.66	44.54	52.33	12.58%
B	36.48	36.50	41.41	51.31	12.33%
C	34.11	35.60	36.75	45.89	11.03%
D (MCI)	24.71	25.81	26.86	30.87	7.42%
E	18.04	18.89	21.08	25.46	6.12%
F	16.33	17.78	17.47	19.33	4.65%
G	16.67	15.04	14.93	15.19	3.65%
Total traffic Top 7 networks	183.53	188.28	203.04	240.38	57.78%
Total traffic all networks	313	313	353	416	100%

As shown in the above tables, concentration in the Internet backbone market is lower than in the broadband access market and has decreased in the last five years. Additionally, both firms and ISPs can connect with multiple suppliers. This practice, "multi-homing," is engaged in by many ISPs as well as many of their business customers for two reasons: first, ISPs and large business customers multi-home on various backbones to avoid outages; second, both ISPs and customers multi-home to place additional competitive pressure on their service providers. In contrast to the residential customer, who must often select among a small group of broadband access providers, business customers, especially large business customers, have many choices. The fact that the Internet access market is more competitive for large business customers is reflected in the significantly lower price per unit of bandwidth that large business customers pay, both in comparison to the prices residential customers pay and to the prices small business customers pay.

I first consider two-sided pricing by a monopolist who charges both final consumers and applications or content providers. I then

[26] Data from *RHK Traffic Analysis–Methodology and Results*, May 2005, as reported in Declaration of Marius Schwartz to the FCC in the SBC-AT&T merger. The identities of all networks are not provided, but it is likely that B, C, E and F are Level 3, Quest, Sprint, and SBC in unknown order.

discuss general price discrimination strategies by a monopolist. I follow up with the price discrimination issues in an oligopoly situation.

1. TWO-SIDED PRICING MODEL

I model the two-sided network as follows. Consider the strategic interactions between a network access monopolist A_0, an applications or content company B_1 (selling a complementary good to network access) and the final consumers of content when the network can charge a fee to both consumers and applications providers.[27] In the mathematical part of the text, for brevity I will be using the word "application" to mean both applications and content. The network access firm sells an Internet connection subscription to end users at price p_0. The application provider sells the application to end users at price p_1. The application provider also pays the network a per unit access fee s, which the network has set.

Assuming a linear demand structure, let the demand function of network access service be $q_0 = a_0 - b_0p_0 - dp_1$, and the demand of the application B_1 be $q_1 = a_1 - b_1p_1 - dp_0$.[28] In this model, the quantity intercept a_0 of the network access demand (representing actual sales when all prices are zero) depends on the inherent quality and function of the network and the variety (number) of applications that are transported on the network.[29] In the demand function, the parameter d measures the strength of the complementarity between the network and the application.[30,31] The profit function of the access network is

---

[27] The mathematical structure of this model is similar to Nicholas Economides and Evangelos Katsamakas, *Two-sided Competition of Proprietary vs. Open Source Technology Platforms and the Implications for the Software Industry*, 52 MGMT. SCI. 1057, 1071 (2006), http://www.stern.nyu.edu/networks/Economides_Katsamakas_Two-sided.pdf.

[28] Ibid. This demand system can be generated by a population of users with differing willingness to pay. For example, it can be generated by a population of users of uniformly distributed types, each with a unit demand. This demand system can also be generated by a representative consumer with quadratic utility function.

[29] Ibid. The maximum sales of the network, a_0, may be larger than the maximum sales of the application, a_1, *i.e.*, $a_1 \leq a_0$.

[30] The degree of complementarity between two goods measures the extent to which two goods are used together.

[31] I assume b_0, $b_1 > d$, *i.e.*, that the own-price effect for each product dominates the cross-price effect. To create a benchmark, I assume zero cost.

$\pi_0 = \pi_{0u} + \pi_{0a}$, where $\pi_{0u} = p_0 q_0$ is the network profit from users, and $\pi_{0a} = sq_1$ is the network profit from the application access fees. The profit function of the application provider is $\pi_1 = (p_1 - s)q_1$.

I assume that network access firms and applications firms set prices in a two-stage game. In stage one, the access network sets the access fee *s* paid by the application provider. In stage two, the network access and the application provider set the price the end-user pays, p_0, p_1 simultaneously. We assume a non-cooperative game and we find and characterize the subgame-perfect Nash equilibria.

To find the non-cooperative equilibrium, we start the analysis at the last stage of the game. Imposing maximization conditions with respect to the choices of prices p_0 and p_1 by the network and the application, we find the network and application prices as respectively increasing and decreasing functions of the network access fee *s*.[32] In the first stage of the game, the network chooses fee *s* anticipating second stage equilibrium prices. The necessary condition for profit maximization is $\frac{d\pi_0}{ds} = \left(p_0 \frac{dq_0}{ds} + q_0 \frac{dp_0}{ds}\right) + \left(s \frac{dq_1}{ds} + q_1\right) = 0$. A marginal increase of s affects both profit streams of the network firm. The network's profit from users increases by $p_0 \frac{dq_0}{ds}$ and decreases by $q_0 \left| \frac{dp_0}{ds} \right|$. The profit from the application firm increases by q_1 and decreases by $s \left| \frac{dq_1}{ds} \right|$.[33] The network's choice of s maximizes the sum of the two profit streams. The effect of s on the network profit from users is $\frac{d\pi_{0u}(s)}{ds} = d \frac{d\left(a_1\left(2b_0b_1 + d^2\right) - 6b_1\left(b_0b_1 - d^2\right)s\right) - 2a_0b_1\left(2b_0b_1 + d^2\right)}{\left(4b_0b_1 - d^2\right)^2}$. The profit from users is decreasing at $s = 0$, since $\frac{d\pi_{0u}(0)}{ds} = d \frac{a_1 d\left(2b_0b_1 + d^2\right) - 2a_0b_1\left(2b_0b_1 + d^2\right)}{\left(4b_0b_1 - d^2\right)^2} < 0$.

---

32 Specifically, equilibrium prices are $p_0 = \frac{2a_0b_1 - da_1 - 3db_1s}{4b_0b_1 - d^2}$ and $p_1 = \frac{2a_1b_0 - da_0 + \left(2b_0b_1 + d^2\right)s}{4b_0b_1 - d^2}$.

Notice that $\frac{dp_1}{ds} > 0$ and $\frac{dp_0}{ds} < 0$, that is, as expected, the application price increases with the access fee s because the application firm faces a higher marginal cost, while the network price decreases as the application has a higher price. These two effects imply that sales of the access network (respectively application) increase (decrease) in the access fee *s*: $\frac{dq_0}{ds} = -b_0 \frac{dp_0}{ds} - d \frac{dp_1}{ds} > 0$ and $\frac{dq_1}{ds} = -b_1 \frac{dp_1}{ds} - d \frac{dp_0}{ds} < 0$.

33 Both profit streams of the network are concave in *s* and, therefore, the total network profit is concave is *s*.

Therefore, the fee s_u^* that would maximize only the access network profit from users is negative.

The effect of fee *s* on the access network profit from the application is $\frac{d\pi_{0a}(s)}{ds} = b_1 \frac{2a_1b_0 - a_0d - 4(b_0b_1 - d^2)s}{4b_0b_1 - d^2}$. This profit is increasing at $s = 0$, if $2a_1b_0 - a_0d > 0$. Then s_a^* is positive, and therefore s^* may be positive or negative $(s_u^* < s^* < s_a^*)$. The access fee s^* is positive when, at s = 0, the access profit from the application is increasing at a faster rate than the profit from users is decreasing. Figure 1 shows an example of that case. Figure 2 shows the relationship between the network's fee to the application, the network profit, the application's profit and the total industry surplus, which is the sum of the profits of the network, the profits of the application, and consumers' surplus.

The two-stage game has a unique sub-game perfect Nash equilibrium given by the following prices:

$$s^* = \frac{a_1(8b_0^2b_1^2 + d^4) - a_0b_1d(8b_0b_1 + d^2)}{2b_1(b_0b_1 - d^2)(8b_0b_1 + d^2)},\ p_0^* = \frac{a_0b_1(8b_0b_1 + d^2) - a_1d(10b_0b_1 - d^2)}{2(b_0b_1 - d^2)(8b_0b_1 + d^2)},$$

$$p_1^* = \frac{a_1(12b_0^2b_1^2 - 2b_0b_1d^2 - d^4) - a_0b_1d(8b_0b_1 + d^2)}{2b_1(b_0b_1 - d^2)(8b_0b_1 + d^2)}.$$

Figure 1. Network Profit Streams and Access Fee, *s**

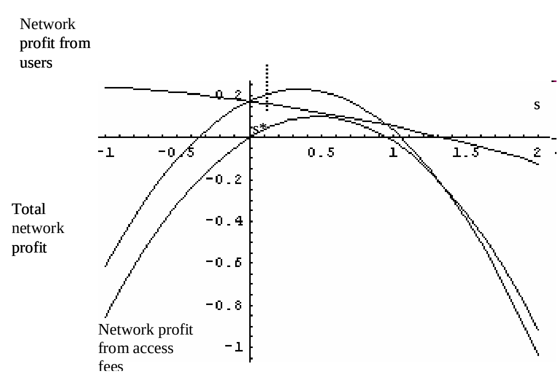

Figure 2. Network Profits, Application Profits, and Total Industry Surplus

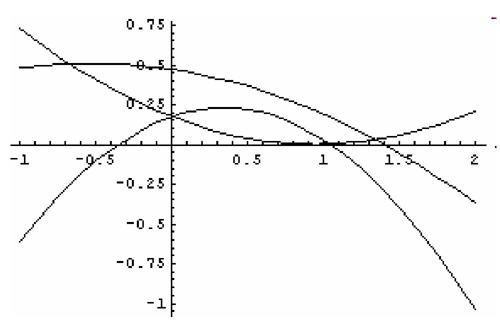

Thus, as Figures 1 and 2 demonstrate, total industry surplus is lower when the access network charges a positive fee to applications, even though a positive fee will typically be part of the equilibrium. Intuitively, this can be explained as follows: the fee acts as a marginal tax on the application and therefore increases its marginal cost and the price that the application charges to final consumers. Due to the complementarity between the application and the network, increasing the price of the application also hurts network sales. Thus, imposing a fee on the application would have a larger negative impact on total industry surplus than imposing the same fee on the consumers and no fee on the application. The same argument can be made in terms of network effects. There are network effects between the application and the network. Therefore, if the network imposes a fee on the application it will result in some negative effect on the network provider. For this reason, imposing a fee on applications reduces total industry surplus.[34]

2. PRICE DISCRIMINATING MONOPOLIST

The Internet, as it exists today, supports large numbers of applications and services. There is wide range in the willingnesses to pay for each type of service, and there is wide dispersion in its distribution. There is no simple index or measure of capacity or bandwidth use of an application that is closely correlated to the willingness to pay for that application. For example, bandwidth use is high for some highly valued services, such as video on demand, but bandwidth use is very low for information services, such as search or bidding in auctions in real time, which are also highly valuable.

In the absence of legally required non-discrimination, Internet broadband access providers may attempt to capture the consumer surplus that remains after uniform pricing. There are two reasons for this attempt. First, even in an unconstrained monopoly situation, price discrimination, based on differences in the elasticity of demand, increases profits. Second, uniform regional pricing, discussed above, constrains carriers' profits to duopoly levels, below the level that could be achieved through price discrimination. When selling to residential customers, a last mile monopolist carrier typically has the incentive to reduce the capacity of "plain" broadband Internet access service so

[34] Although the duopoly competition model for access with monopoly or duopoly applications had not yet been developed, there is no reason to believe that the main result on reduction of surplus by the imposition of fees on applications is going to be different.

that it can establish a "premium" service at a higher price as discussed below.

Suppose that information packets differ according to the willingness of end-users to pay for them. Let packet of type/function i be offered at price p_i and its demand be $D_i(p_i)$, $i = 1, \ldots, n$, under a price discrimination model. Alternatively, all packets could be sold at the same price p. Assuming that the cost of transmission is the same for all packets, in a price discriminating network the monopolist faces a cost, $C(\Sigma_i D_i(p_i))$, and its profits under discrimination (Π_d) are $\Pi_d = \Sigma_i p_i D_i(p_i) - C(\Sigma_i D_i(p_i))$. It is easy to show that maximization of the monopolist's profits implies $[p_i - C'(\Sigma_i D_i(p_i))]/p_i = 1/\varepsilon_i$, where ε_i is the elasticity of demand for packets of type i. Alternatively when all packets are sold at the same price, the monopolist maximizes profits under uniform pricing Π_u ("u" for uniform pricing) $\Pi_u = p[\Sigma_i D_i(p)] - C(\Sigma_i D_i(p))$. Maximization of uniform pricing profits implies $[p - C'(\Sigma_i D_i(p))]/p = [\Sigma_i D_i(p)]/[\Sigma_i D_i(p)\varepsilon_i]$, that is, in uniform pricing, the percentage of price to cost margin is a weighted average of the elasticities of demand for the various types of packages.

In general, the coordinated introduction of price discrimination schemes may reduce output. There is a general theorem in economics that price discrimination, which reduces total output, also reduces total surplus.[35] Thus, the first anti-competitive concern is that price discrimination may reduce output.

Two additional considerations reinforce this anti-competitive concern. First, most applications on the Internet exhibit network effects as described above. This means that the last transaction/sale/download is worth more to the consumer when sales of compatible applications are higher. For example, the Google search application is more valuable when Google has a larger audience. Using YouTube is more valuable when there are more subscribers to that web place. Additionally, more individual users decide to subscribe and to post on a web space when the web space has more subscribers. The existence of network effects implies that the efficient prices (total surplus maximizing prices) are below the perfectly competitive prices, that is, below marginal cost.[36] Broadband access

---

[35] This is contingent on serving all markets under uniform pricing, which holds here since I am starting with all markets served under "net neutrality." See Marius Schwartz, "Third-Degree Price Discrimination and Output: Generalizing a Welfare Result," *American Economic Review* 80 (1990): 1259–62.

[36] See Nicholas Economides, "The Economics of Networks," *International Journal of Industrial Organization* 14 (1996); 675–99, http://www.stern.nyu.edu/networks/Economides_Economics_of_Networks.pdf.

providers are charging, at best, duopoly prices, which are typically considerably higher than perfectly competitive prices. Thus, increasing present market prices as an effect of price discrimination will increase price divergence from efficient prices.

Second, the fact that application and content providers will be charged instead of subscribers is likely to mask the true cost of Internet service to residential subscribers and create additional price distortion and surplus loss.[37]

3. OLIGOPOLY CONCERNS

There is an additional concern in duopoly. Because broadband access competition is duopolistic in many areas, the creation of a "premium" service and the accompanying reduction in bandwidth capacity of plain service required to create it is likely to be coordinated among network access providers. The coordinated reduction of capacity in "plain" service is reminiscent of cartel behavior, such as two competing airlines deciding in a coordinated way to reduce their capacity in economy class. Therefore, the introduction of coordinated price discrimination may have anti-competitive consequences. In particular, if there is sufficient evidence that the markets for "plain" and "premium" services are sufficiently different, the cartelization of "plain" service is likely to be a Sherman Act Section 1 violation.

B. VERTICAL CONCERNS

There is also a variety of potentially anti-competitive vertical activity that could result in Sherman Act Section 2 violations as discussed below.

First, a carrier may favor its own content or application over that of independent providers. VOIP provided over broadband Internet by companies without a network infrastructure, such as Vonage or EarthLink, competes with traditional circuit-switched service provided by AT&T and Verizon and with VOIP provided by cable TV operators. Independent VOIP could be subject to discrimination. Additionally, both AT&T and Verizon are gearing to distribute video,[38]

---

[37] The generally more competitive market for large business customers will not shield them from the levies imposed by the access carriers.

[38] See Fred Dawson, "More Details on Verizon's Initial Video Launch," *xchange.com*, http://www.xchangemag.com/hotnews/59h231024228723.html (accessed April 10, 2008).

and could favor their video services over that of others. In the absence of non-discrimination rules, last mile carriers can leverage their market power in the Internet broadband access market to control/support their voice telecommunications market. This concern applies both to telecommunications companies who can degrade opponents VOIP service to protect their fixed line voice service and to cable companies who may degrade their opponents' VOIP service to protect their own VOIP service.

Similar concerns operate with regard to carriers' video services. It should be clear that, although active sabotage of a competitor's service is an obvious, and illegal, form of discrimination, network access providers do not need to use these tactics. To discriminate effectively against a VOIP competitor, it will be sufficient for the access provider to set a high fee for access to the "premium lane," which will effectively block profitable operation by the competitor whose operation in the "standard lane" has been degraded by the high allocation of bandwidth to the fast lane.[39]

Second, the anti-competitive concerns are hardly limited to the products and services currently provided by the firms with market power in the access market. Such carriers can also leverage market power in broadband access to the content or applications markets through contractual relationships. Two examples of this use of market power follow:

First, a carrier can contract with an Internet search engine (or other application, or video content provider) to put it in "premium" service, while searches using other search engines have considerable delays using "plain" service. In this setup, the "plain" service can be tweaked to be sufficiently slow so that consumers will choose to do almost all their searches with the search engine in "premium" service. By making a "take it or leave it" offer to the various search engines, the access carrier can extract a large part of the profits created by complementary goods, in this example, search engines. In effect, this type of strategy can determine who will be the successful search (or application, or content) company. It would give tremendous power to the network company without any obtrusiveness or the active sabotage of any individual company.

Second, in the same setup, a carrier can actively sabotage a search engine (or application, or content) company with similar results as above.

---

[39] *See* Nicholas Economides, "The Incentive for Non-Price Discrimination by an Input Monopolist," *International Journal of Industrial Organization* 16, no. 3: (1998): 271–84, http://www.stern.nyu.edu/networks/The_Incentive_for_Non-Price_Discrimination.pdf.

1. CALIBRATION OF POTENTIAL WELFARE LOSSES

There are no published estimates of the elasticity of demand for various Internet applications. Thus, it is very hard to estimate the exact effect of the proposed price discrimination scheme. However, Goolsbee, using early data, estimates the elasticity of demand for broadband Internet access to be approximately $\varepsilon = 3$, at a price of \$40 with marginal cost at \$25,[40] *i.e.*, at a 60% markup over cost.[41] We may assume, that a new price discrimination scheme would precipitate a moderate increase in average price of at least 20%. This would imply a deadweight loss ("DWL") of at least 6% of the annual total Internet broadband access bill, using the standard approximate calculation $DWL = (\Delta P)(\Delta Q)/2 = \varepsilon(QP)(\Delta P/P)^2/2$, where $\Delta P/P$ is the proposed percentage price increase, here 20%, and ε is the elasticity of demand, here $\varepsilon = 3$. OECD puts the number of broadband subscriptions in the United States at almost 60 million.[42] This brings the annual revenue to networks from broadband access to \$24 billion and the estimated direct welfare loss to residential consumers to roughly \$144 million annually. Currently, there is no good estimate of the additional welfare loss to business customers.

The above estimate is a moderate lower bound on the surplus losses that may be generated by price discrimination by the access networks. In addition to the direct losses to consumers, the proposed price discrimination scheme will decrease consumer surplus in a variety of ways:

1. It will decrease consumers' applications, and content providers' surplus because it will imply a further divergence from efficient pricing in the presence of network effects;

2. It will foreclose on the margin potential entrants in complementary applications and content markets;

---

[40] Austan Goolsbee, "The Value of Broadband and the Deadweight Loss of Taxing New Technology," *Contributions to Economic Analysis and Policy* 5, no.1 (2006): 13, http://journals.ohiolink.edu/ejc/pdf.cgi/Goolsbee_Austan.pdf?issn=15380645&issue=v05i0001&article=1505_tvobatdlotnt.

[41] Here marginal cost does not mean the cost of a single transmission. It rather means deployment of service to a customer.

[42] Organization for Economic Co-operation and Development, "OECD Broadband Statistics to December 2006," http://www.oecd.org/document/7/0,3343,en_2649_201185_38446855_1_1_1_1,00.htm.

3. It will decrease innovative activity of applications and content providers at the edge of the network; and

4. It will give the access providers the ability to choose which content and/or application will be successful removing the significant benefits of mix and match.

It is difficult to quantify the extent of these surplus losses. Noting, however, that the current cost of residential access is less than $24 billion, the profits of the complementary goods and services and applications plus consumers surplus from these are a large multiple of this amount.

2. POLICY IMPLICATIONS

The question posed to Congress is whether it should intervene now by imposing non-discrimination restrictions or if it should wait for antitrust suits to be filed and resolved. In my opinion, it is better to impose the non-discrimination restrictions by law because:

1. Suits take time and much damage can be done before they are resolved. The legal system is slow and lawsuits will not be resolved in "Internet time."

2. The abolition of "net neutrality" gives rise to a variety of antitrust concerns, while each suit would typically deal with one issue. Thus, delays may be compounded by the need for each type of suit to be adjudicated.

3. The Internet is a key essential network for growth of the U.S. economy. The United States is already lagging behind 14 countries in Internet penetration, as seen in Figures 3 and 4 below. Figure 4 shows that a number of countries with higher broadband Internet penetration than the United States have lower population densities, so U.S. population density does not explain the low penetration.[43] Since the Internet is a key factor for future growth, high penetration is desirable and adding price discrimination is unlikely to help.

---

[43] Iceland, Finland, Norway, Canada and Sweden have lower population densities than the United States, but have significantly higher broadband Internet penetration.

4. Increasing prices through two-sided pricing will not increase network traffic or contribute to network growth.

5. The abolition of "net neutrality" is likely to have significant negative consequences on innovation on the Internet, whether or not anti-trust violations occur in connection with the abolition of "net neutrality", and therefore it is in the public interest to prevent it by law.

Figure 3. Broadband Internet Penetration and per Capita Income[44]

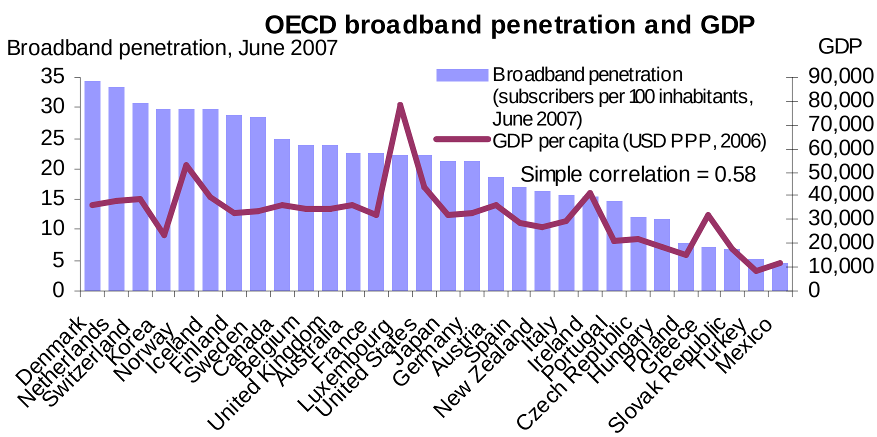

Source : OECD

[44] Organization for Economic Co-operation and Development, "OECD Broadband Statistics to December 2006," http://www.oecd.org/document/7/0,3343,en_2649_201185_38446855_1_1_1_1,00.html (accessed April 10, 2008).

Figure 4. Broadband Internet Penetration and Population Density[45]

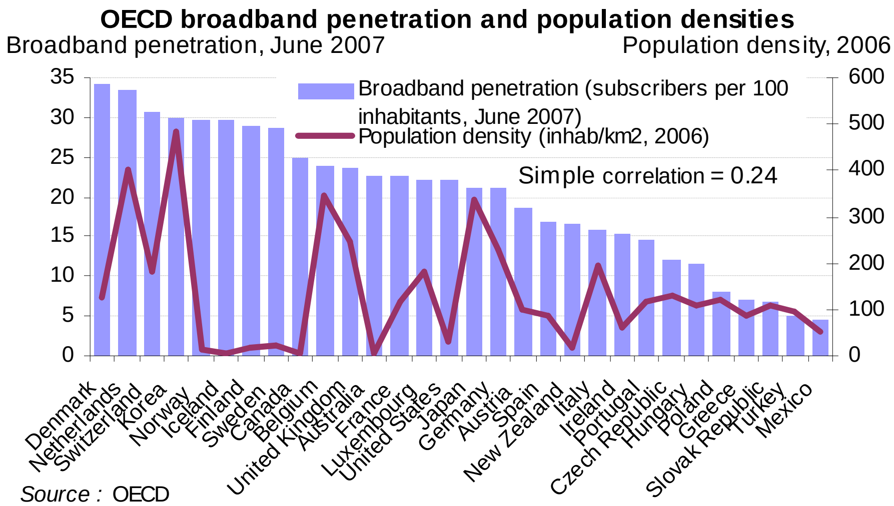

Source : OECD

IV. CONCLUSION

The Internet is the most important telecommunications network of the last fifty years. Enabled by public protocols and standards, and by significant advances in electronics, computers, fiberoptics, and laser technology, the Internet has been an engine for the growth of both the United States and world economies. Relying on public protocols, applications are developed to run across the Internet and content is disseminated on the Internet without the approval or consent of centralized Internet operators. Tremendous successes resulted such as the World Wide Web and all the applications that run on it, including big financial successes like Yahoo and Google, as well as big benefits of social interaction networks and great leaps in civil society through new discussion forums and formats.

The Internet, in its commercial form, is a relatively new network, with only a dozen or so years to date. Its tremendous acceptance and success has made it an essential part of both business and personal life. All previous electronic networks, including early successes, like AOL, have abandoned proprietary formats and folded into the Internet. The success of the Internet thus far has been based on openness and non-discrimination, which until recently, was

[45] Ibid.

guaranteed by U.S. telecommunications regulation. Recently, the abolition of this regulation has led to proposals by broadband Internet access providers that would radically change pricing on the Internet. This article shows that these changes are likely to hurt consumers and diminish innovative activities in complementary sectors such as computer applications and content dissemination. These pricing proposals, if implemented, are likely to raise a variety of significant anti-competitive concerns, outlined in detail in the article.

Among these concerns is the possibility that access providers will degrade and/or restrict capacity in traditional Internet access to force applications and content providers to use their new "premium" service. The possibility exists that this degradation and restriction of capacity will happen in a coordinated way, in a cartel-like fashion. This article demonstrates that, even in the absence of such discrimination, due to the existence of network effects, charging a fee to application and content providers is likely to both hurt consumers and to reduce the benefit that the Internet brings to society as a whole.

In addition, there are a large number of vertical anti-competitive concerns created by the absence of a non-discrimination policy. Access networks, if left unrestrained by non-discrimination rules, have incentives to favor their own services, applications, and content and to kill competing services, such as independent VOIP providers, which provide alternative telephone services over the Internet. Additionally, the access networks have incentives to leverage their access monopoly or duopoly market power in many other complementary markets by offering "take it or leave it" contracts. Thus, the access providers will be able to determine who will be the primary provider of search engines, content, and other applications and services. This would be highly detrimental to the consumers and industries that rely on the Internet.

The present question before Congress is whether to allow the Internet to be run without non-discrimination rules or whether to impose specific non-discrimination rules. A number of considerations favor imposing a specific rule supporting "net neutrality." First, litigation is very slow, and much damage can be done before the resolution of litigation establishes a clear rule. Second, there are a number of different antitrust concerns, and litigation will have to deal with each one at a time. Third, although the Internet is a crucial network supporting United States' economic growth, Internet penetration in the United States is low compared to many other countries with much lower per capita income. The imposition of discrimination is likely to amplify these problems. Fourth, because of network effects, the correct public policy is to subsidize the Internet, rather than increase its price. The price discrimination schemes discussed are likely to effectively increase the price consumers pay for Internet access. Finally, the innovation "at the edge" of the network that has flourished under the regime of "net neutrality" would be significantly threatened by discriminatory actions.

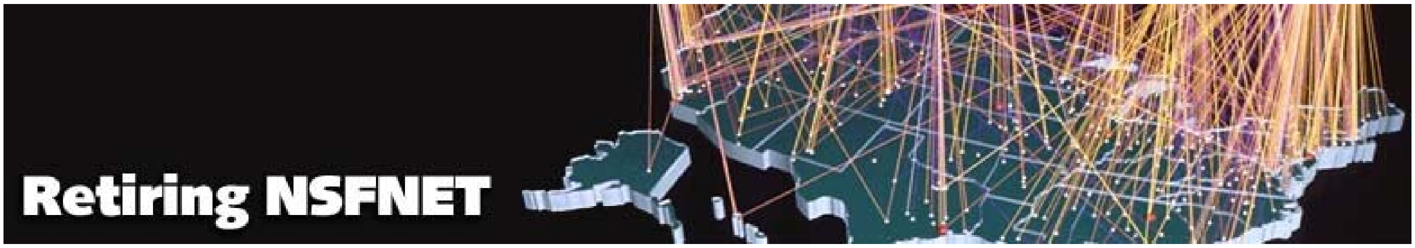

Print This Send This

Routing Analysis

- BGP Data Analysis
- VAST: Visualizing Autonomous System Topology

Traffic Analysis

- Flamingo

Middleware

- Shibboleth

Other Projects

- MI Streamnet
- RADb: The Routing Assets Database
- Virtual Dataset Repository

Papers & Presentations

Project History

- NSFNET
- Internet Performance Measurement and Analysis
- Merit GateD Consortium
- Route Server Next Generation
- Routing Arbiter

Network Research Home

Retiring the NSFNET Backbone Service: Chronicling the End of an Era

By Susan R. Harris, Ph.D., and Elise Gerich

Reprinted with permission from "ConneXions", Vol. 10, No. 4, April 1996.

April 30, 1996, marks the one-year anniversary of the final dismantling of the venerable NSFNET Backbone Service. After more than a year of planning, reconfiguration, shutdowns, and transitions, the U.S. Internet had completed its move to a new architecture composed of multiple backbones, linked at the new interexchange points.

The midnight NSFNET shutdown went remarkably smoothly, as did most of the events leading up to the final phaseout. This article looks back on the timelines, dependencies, delays, emergencies, and successes that marked the final year of the NSFNET. We begin by taking a brief look at the history of what was the world's largest and fastest network for research and education.

A Brief History of the NSFNET

The National Science Foundation inherited the responsibility for nurturing the U.S. Internet from the Advanced Research Projects Agency (ARPA). From its inception in 1985-1986, the NSFNET program laid the foundation of the U.S. Internet and was the main catalyst for the explosion in computer networking around the world that followed.[1] The first NSFNET, a 56Kbps backbone based on LSI-11 Fuzzball routers, went into production in 1985 and linked the six nationally funded supercomputer centers (the five NSF centers and the National Center for Atmospheric Research). Soon after the network's inception, the need for more advanced networking technology was indicated when rapid growth in traffic precipitated serious network congestion. In 1987, NSF issued a competitive solicitation for provision of a new, faster network service. The new service would provide a network backbone to link the six supercomputer centers and seven mid-level networks. The mid-level networks would in turn connect campuses and research organizations around the country, creating a three-tiered network architecture that remained in place until the end of the NSFNET backbone service.

In fall 1987, NSF selected Merit Network, Inc., and its partners MCI, IBM, and the State of Michigan to manage and re-engineer the new backbone service. Eight months after the NSF award, the NSFNET partnership delivered a new T1 backbone network that connected 13 sites: Merit, NCAR, BARRNet, MIDnet, Westnet, NorthWestNet, SESQUINET, SURAnet, and the NSF supercomputer centers. Two additional regional networks, NYSERNet and JVNCnet, were also served by the backbone, because each was collocated at a supercomputer center. Each of the 13 backbone nodes, known as Nodal Switching Subsystems, was composed of nine IBM RTs linked by two token rings with an Ethernet interface to attached networks. There were 14 T1s connecting the sites, on which a virtual topology was constructed. Each virtual path represented one-third T1 to the site.

In 1989 the backbone was reengineered, increasing the number of T1 circuits so that each site had redundant connections to the NSFNET backbone as well as increasing router capability to full T1 switching.[2] With this upgrade, the NSFNET's physical topology equalled its virtual topology. By then, the traffic load on the backbone had increased to just over 500 million packets per month, representing a 500% increase in only one year. Every seven months, traffic on the backbone doubled, and this exponential growth rate created enormous challenges for the NSFNET team.[3]

Upgrade to T3

To handle the increase in traffic, Merit and its partners introduced a plan to upgrade the backbone network service to T3. The NSF also wanted to add a number of new backbone nodes, and asked Merit to prepare proposals for the added cost of new nodes at T1 and T3 speeds, while the NSF issued a solicitation to the community for those interested in becoming new NSFNET sites. It was eventually decided by the NSF that the partners would increase the total number of backbone nodes on the NSFNET from 13 to 16, all running at 45 Mbps.[4] Additional sites served by the T3 NSFNET backbone service would include Cambridge MA (NEARNET), Chicago's Argonne National Lab, and Atlanta GA (SURAnet).

In late May 1990, Merit's cooperative agreement with NSF was modified to cover the additional work. By the end of the year, Merit, SDSC, and NCSA were connected to an early T3 service and began testing the new T3 routers with real traffic. In addition, a new T3 research and test network was implemented to parallel the existing T1 test facility.

Important architecture and equipment changes came with the new T3 network. The core backbone equipment was moved from the universities and supercomputer sites to MCI's points-of-presence (POPs), and the RTs were replaced with RS/6000s and a card-to-card forwarding architecture. Many of the techniques introduced in the T3 RS/6000 routers have since been adopted by commercial router vendors.

As the backbone network service was growing in complexity and was re- engineered, increasing focus and resources were needed to keep pace with more complex technical, business, and policy environments. To meet these organizational challenges, ANS was created and announced in September 1990. ANS began to provide service for NSFNET as a subcontractor to Merit, with IBM, MCI, and others continuing to infuse new technology to develop the infrastructure.

During 1991, a year of refining the new backbone technology, the T1 and T3 networks existed in parallel. Difficulties in tuning the new technology prevented the network from being moved to full production status until late in the year, when all sixteen backbone sites comprising the NSFNET service were connected to the new ANSnet national T3 infrastructure. With expansion work completed and improved performance validated, several sites began using the T3 for their primary traffic path by November 1991. A final round of testing in mid- December set the stage for moving the remaining NSFNET traffic to the new backbone service in early 1992. The network now exceeded the T1 structure in stability by a factor of ten, with fewer outages and errors in all categories.

The upgrade of the NSFNET backbone service to T3 was not only a technological and organizational challenge of the highest order. It also precipitated a greatly-needed, though contentious, community dialogue about the evolution and commercialization of the U.S. Internet. Internet Service Providers were springing up all over the country, from local dial-up providers to larger companies providing T1 and eventually T3 service, and there were now a growing number of vendors offering TCP/IP networking products and services.

During 1992, the National Science Board authorized an extension of Merit's cooperative agreement for eighteen months beyond the October 1992 expiration date in order for NSF to develop a follow-on solicitation for national networking, one that would accommodate the growing role of commercial providers and allow NSF to step back from actually operating a network to concentrate on supporting leading-edge research initiatives. NSF published a draft solicitation for community comment in 1992, and a new solicitation was issued in May 1993.

Early in 1994, awards for building the new architecture were given to Merit and USC's Information Science Institute for the Routing Arbiter service, to MCI for the vBNS, and to three providers for the Network Access Points: Sprint, MFS Datanet, and Bellcore, representing Ameritech and PacBell. NSF also awarded Merit a transition extension that began in May 1994 and lasted until April 1995, when the NSFNET backbone service would be retired and all connections would be switched to a new service.

Deadlines and Commitments

Moving the U.S. Internet to a new architecture in the months between the 1994 awards and the April 30, 1995 termination date was a frightening challenge for the regional networks, the ISPs, and the NSFNET partnership. Before the backbone could be decommissioned, four main tasks had to be accomplished by the networking community:

- Establish the Network Access Points (NAPs) and move them to production status.
- Attach to the NAPs the NSFNET and the ISPs that provided service to the regionals.
- Develop the RA Service by placing Route Servers at the NAPs and setting up a routing registry.
-

Move the regionals off the NSFNET and attach them to networks operated by ISPs.

According to NSF's ambitious transition schedule, the new NAPs would be available by August 15, 1994. The NSFNET backbone service would then attach to the NAPs, with all current attachments to the NSFNET remaining in place. The ISPs would then begin to attach to the NAPs, and regional networks that attached to NSFNET would begin to establish connections to the ISPs. By October 31, the regionals would cut over all traffic to the ISPs and disconnect their attachments to the NSFNET. Only the supercomputer centers would remain attached to the NSFNET. The vBNS would be deployed by January 1, 1995, and attached to the NAPs by February 1, 1995.

As it turned out, all of these actions were delayed, and revised deadlines established.

Establishing the NAPs

The first Network Access Point to go into production was the Washington, D.C. NAP (MAE-East, the Metropolitan Area Ethernet). MFS had been operating MAE-East since 1992, and MAE-East had served as a model for the NAPs as defined in NSF's solicitation. In fall 1994, MAE-East was upgraded from a 10Mbps Ethernet to FDDI; internetMCI and SprintLink, which had already attached to the MFS facility, upgraded their connections to FDDI, as did the NSFNET.

The Sprint NAP, a bridged FDDI/Ethernet hybrid, was up and running by the end of the summer; ANSnet/NSFNET, SprintLink, and internetMCI attached to it in September. The Sprint and Washington, D.C. NAPs began to carry much of the traffic for the U.S. Internet once networks began to move off NSFNET in November 1994, because the PacBell and Ameritech ATM NAPs were still being deployed and went into production several months later. Both facilities were physically in place by October 1994, but problems with ADSU performance and a concern with ATM switch buffer sizes led to a lack of confidence in the ability of the ATM NAPs to sustain the traffic load.

As a result, both PacBell and Ameritech decided to deploy interim configurations, and put FDDI LANs into production in March 1995. Some ISP routers on the FDDIs at these contingency NAPs were also connected to DS3 ports on the ATM switch, so they could pass traffic across the FDDI while still transmitting to ATM-connected peers. As of January 1996, this infrastructure is still in place at the PacBell and Ameritech NAPs.

Deploying the Route Servers

The Routing Arbiter service has two main components: the Route Servers, SPARC 20s deployed at the NAPs, and the Routing Arbiter Database, successor to the Policy Routing Database used to configure the NSFNET backbone service.

In November 1994, primary and backup Route Servers were shipped from ISI to each of the NAPs. Once the necessary data circuits, front-end systems, controllers, ATM switches, and FDDI bridges were installed and tested, addressing schemes worked out, security procedures implemented, and 24/7 network monitoring in place, the Routing Arbiter team began to set up peering sessions with customer routers at the NAPs. Out-of-band access -- a prerequisite for declaring the Route Servers fully in production -- became available several months later.

By April 1995, the Route Servers were peering with more than a dozen providers at the Sprint and Washington, D.C. NAPs. In July, production RA services were announced at the Sprint NAP, and announcements for the other NAPs soon followed. At each exchange point, the Route Servers began importing and exporting routes to numerous ISPs. The ISPs maintained sessions with other peers as well as the Route Servers, comparing the routing information from both sessions for consistency.

NACRs and the PRDB: the Long Goodbye

Merit originally planned a December 1994 retirement for the Policy Routing Database (PRDB), which had been used to configure the NSFNET's backbone routers since 1989. The PRDB would be replaced by the Routing Arbiter Database, which would then become part of the Internet Routing Registry (IRR) along with the RIPE NCC, MCI, ANS, and CA*net registries. The IRR would be an important global resource--a public repository of announced routes and routing policy in a common format, so that ISPs could use the information stored in any and all registries to configure their backbone routers, analyze routing policy, and build tools to help in these efforts.

The PRDB was established to maintain information about what were considered legitimate destination announcements from the various regionals. The primary goal of maintaining this information was to prevent routing loops. When BGP replaced EGP as the inter-domain routing protocol in 1994, suppression of routing loops no longer had to be so administratively controlled. The information in the PRDB was then mainly used to record routing policies such as path preferences and to generate the backbone configuration files.

NSF's follow-on solicitation for the new architecture specified a continuation of the function that the PRDB played in the T1/T3 NSFNET. The goal was to record global routing policy information based on each Autonomous System's policy. RIPE had pioneered this work in the European arena, and the data exchange format described in RIPE-181 (RFC 1786) was adopted as the "standard" for Internet Routing Registries.[5] The RADB adheres to this model.

The challenge was to establish and populate the RADB before the retirement of the NSFNET and the PRDB. By summer 1994, the RIPE NCC registry had been in production for two years, and CA*net and internetMCI were creating routing registries to support their customers. ANSnet would continue to use the PRDB until the RADB was established. But the dilemma was how to transition from NSFNET-centric information to the AS-specific information needed for the RADB, while continuing to provide a stable router configuration environment for the NSFNET service.

Merit's December target date for retiring the PRDB was based on the assumption that the regionals would be off the backbone by October 31. When it became clear that they weren't going to make that deadline and the PRDB would need to continue to support the NSFNET and its regionals well beyond the end of October, a plan was proposed to transition to the RADB to support the NSFNET in its last months.

The new situation presented several problems. First, the tools used to configure the NSFNET/ANSnet routers were based on PRDB attributes, not RIPE-181. Second, the RADB was not yet populated with data. And finally, the PRDB described AS690 policy with respect to its peer ASs on a per-prefix basis; in the RIPE syntax, the basis for describing routing policy was the Autonomous System where the route originated. With more than 40,000 prefix-based policies for the regionals, the PRDB was used to generate about 100 configuration files of around 250,000 total lines every two weeks, and those policies needed to be re-expressed in a RIPE- compatible format.

Continuing the Policy Routing Database for long-term support of ANSnet was inadvisable. If ANS continued to use the PRDB for AS690 routing after the transition, the PRDB's non-standard format would create a barrier to sharing global routing policies and building tools to aid with global routing. A solution had to be found that would provide stable routing through the transition, and, once the NSFNET was retired, allow the ANS registry to take its place alongside the other registries in the IRR.

To solve the problem, Merit proposed a modification to RIPE-181 -- a temporary attribute that would specify the peer or adjacent AS announcing the route to AS690. The community agreed to Merit's proposal, and the new expression came to be known as the advisory attribute. Merit now needed to quickly modify the PRDB configuration tools so they would generate the new attribute, populate the RADB with the data needed to generate AS690 configuration files, and make sure that the new configurations exactly matched those produced by the PRDB.

By December 1994, all the data in the PRDB had been converted to RIPE- 181-style expressions and entered in the RADB. By February, the RADB had been populated with RIPE-181-style Maintainer and AS Objects. The databases were running in parallel, with changes to the PRDB automatically reflected in the RADB. Other organizations whose routing information wasn't related to the NSFNET were also populating the RADB throughout the winter of 1994-95; this was another variable that had to be accommodated as the new database emerged.

Finally came the painstaking task of comparing the config files generated by each database. Merit's Dale Johnson went over the large, quarter-megabyte files line by line, adjusted the configuration tools to compensate for any differences in net lists, and repeated the process over and over until the configs matched perfectly. The RADB finally replaced the PRDB a week after the NSFNET was retired.

Moving the Regionals off NSFNET

NSF and Merit coordinated the process of moving the regionals to new Internet Service Providers, with Merit taking the lead in planning the transition. NSF's new Inter-Regional Connectivity program helped support new attachments not only for NSFNET peer networks -- regionals like SURAnet and NYSERNet that connected directly to the NSFNET backbone-- but also to downstream networks such as NevadaNet and MOREnet. Most of the regionals selected internetMCI or SprintLink as their ISP; CERFnet set up its own ATM connection to each NAP.

In mid-October 1994, NSFNET Program Director Priscilla Huston sent a letter to the regionals asking them to send a transition calendar and engineering overview to Elise Gerich of Merit and to her. Huston also asked the regionals to notify her if they weren't going to make the October 31 deadline for moving off the NSFNET.

As it turned out, none of the networks made it. MOREnet, one of the downstream regionals, missed by only a day; other networks slipped by as much as three or four months. The first NSFNET peer network to make the transition was CA*net, which faced a hard deadline from its link provider for terminating its connection to the NSFNET. The other cutovers were pushed back because of delays in provisioning the ISPs selected by the regionals, and because of reticence on the part of the regionals to move off the NSFNET backbone service.

On one or two occasions, networks that had made the transition had to pull back to full NSFNET connectivity because of deployment problems on the new ISP backbone. In general, though, once the regionals had selected an ISP and completed all the testing, re-routing, and reconfigurations necessary to make the switch, traffic flowed smoothly over the new infrastructure.

60-Day Notices: No Turning Back

Early in January, when SURAnet notified NSF and Merit that it was ready to move off the backbone, Merit sent ANS the first message to dismantle NSFNET backbone service -- a 60-day termination notice for ENSS 138 in Atlanta. The ENSSs (Exterior Nodal Switching Subsystems) were installed at regional networks attached to the NSFNET, and acted as end nodes for the backbone. This and subsequent termination notices were irrevocable; once sent, there would be no more NSFNET service through that node.

Later in January, NYSERNet and the Cornell Theory Center notified NSF and Merit that they were ready to terminate their NSFNET attachments. The other regionals and supercomputer centers followed suit, one by one, as the April deadline neared. On February 28, Gerich sent Jordan Becker of ANS the formal, 60-day notice for termination of the NSFNET backbone service at 19 locations:

```
ENSS 128  Palto Alto      April 30, 1995  midnight PST
ENSS 129  Champaign       April 30, 1995  midnight CST
ENSS 130  Argonne         April 30, 1995  midnight CST
ENSS 131  Ann Arbor       April 30, 1995  midnight EST
ENSS 132  Pittsburgh      April 30, 1995  midnight EST
ENSS 133  Ithaca          April 30, 1995  midnight EST
ENSS 134  Cambridge       April 30, 1995  midnight EST
ENSS 135  San Diego       April 30, 1995  midnight PST
ENSS 136  College Park    April 30, 1995  midnight EST
ENSS 137  Princeton       April 30, 1995  midnight EST
ENSS 139  Houston         April 30, 1995  midnight CST
ENSS 140  Lincoln         April 30, 1995  midnight CST
ENSS 141  Boulder         April 30, 1995  midnight MST
ENSS 142  Salt Lake City  April 30, 1995  midnight MST
ENSS 143  Seattle         April 30, 1995  midnight PST
ENSS 144  Moffett Field   April 30, 1995  midnight PST
ENSS 145  College Park    April 30, 1995  midnight EST
ENSS 146  DC              April 30, 1995  midnight EST
ENSS 147  MFS             April 30, 1995  midnight EST
```

The list included the NSFNET attachments at the NAPs, which were coexistent with ANSnet. ENSS 138 in Atlanta wasn't included, since a termination notice for that node had been issued earlier.

By March, backbone traffic had declined dramatically, but not quite as fast as NSF and Merit had expected. Gerich, concerned that the regionals and ISPs weren't moving fast enough, sent e-mail to the community noting that a significant amount of traffic was still traversing the NSFNET. Merit posted a histogram showing the top 10 originators of traffic into the backbone in February 1995, and reminded networks attached to nodes highlighted on the graph about the April 30 deadline.

Later that month, Merit discontinued the T1 safety net that had backed up the T3 infrastructure since 1992.

Black Friday and the Final Shutdown

By the middle of April, only seven regionals had completely severed their ties to the NSFNET backbone service. Other networks had cut over to a new service provider, but continued to peer with the NSFNET for backup purposes. As the final deadline neared, Merit and the NSFNET Executive Committee became concerned that these redundant connections would make it difficult to identify outstanding reachability issues before the April 30 cutoff.

To spot any pockets of unreachable destinations before it was too late, Merit on behalf of the

NSFNET Executive Committee notified the NSFNET community on April 14 that it would terminate peering sessions with all organizations still attached to the NSFNET Backbone service at 9:00 a.m. on Friday, April 21. On April 28, all sessions with the NSFNET service would be permanently terminated; ANS would terminate operation of the NSFNET Backbone service on April 30.

This announcement created quite a stir among the networks attached to the backbone. Several said that they'd lose their Internet connectivity completely if their NSFNET peering was shut down before April 30, and requested that their session stay up. One provider was still relying on his MAE-East NSFNET connection for all his East coast traffic; another requested clemency for a non-production peer router that was proving essential for network diagnostics. A midwest network's installation of a T3 circuit had been delayed; the operators weren't concerned about reachability if their NSFNET peering was shut down, but about capacityˆa large volume of traffic was still traversing the NSFNET, and cutting back to a T1 would lead to unacceptable response times. Merit made separate arrangements to accommodate each network, but held to the new deadline.

As it turned out, the test shutdown had to be postponed. In the early morning hours of April 21, Merit notified the community that it would have to delay the regular Friday backbone configuration run. The volume of routing configuration changes had increased so dramatically as networks switched to new providers that some of the files grew large enough to truncate during production, and produced corrupt configuration files. Merit wasn't confident that it would be able to produce complete and correct configurations in time for the normal 8:00 a.m. configuration window. Additional file space had to be allocated before the configs could be run, and Merit needed to work with several ASs to reduce the number of net lists in the config file. This meant postponing the Friday shutdown until Saturday, and delaying the NSFNET discontinuation until Tuesday, April 25. The test shutdown had indeed pinpointed at least one problem as a result of delayed transitions: the processing of several thousand simultaneous changes to router configurations was more than the PRDB could handle.

On Monday, one network jumped the gun, and surprised Merit and ANS by turning off its ENSS. The ANS staff noted that no harm had been done, but reminded the sysadmin that the plan was to manually turn off peering on the 25th, and shut off the ENSSs on the 30th. IBM was to physically remove the routers beginning in May.

On April 25, the peering sessions on 15 ENSSs were commented out of the configuration files and the NSFNET Backbone Service was, for all intents and purposes, terminated. The next Sunday evening at midnight, a dozen or so staff from Merit and ANS gathered in the University of Michigan NOC to turn off the ENSSs, one by one, at midnight in each respective time zone. One or two regional operations centers called the ANS NOC about unreachable ENSSs, but "mostly the NSFNET went away silently," as one ANS engineer remarked, "or rather, with only the sound of drives and fans spinning down in distant machine rooms."

On May 8, with Merit confident that the RADB was producing consistent configuration files for the ANSnet and ANS ready to take over configuration generation for AS690, the PRDB made a graceful exit. The new architecture was in place: internetMCI and SprintLink had absorbed the NSFNET regionals as their customers; the RADB and the databases maintained by the RIPE NCC, internetMCI, CA*net, and ANSnet had replaced the PRDB as a means of describing routing policy.

Farewell NSFNET! And congratulations to the hundreds of people who helped make the backbone such a great success.

Footnotes

1. Frazer, Karen D. "NSFNET: A Partnership for High-Speed Networking.", Merit Network, Inc., 1995.
2. Ibid.
3. Ibid.
4. Ibid.
5. Bates, T., Gerich, E. Joncheray, L., Jouanigot, J-M., Karrenberg, D. Terpstra, M. and Yu, J. "Representation of IP Routing Policies in a Routing Registry (ripe-81++)", October 1994. Subsequently published as RFC 1786, March 1995.

About the Authors

Susan Harris coordinated NANOG meetings, chaired the NANOG Program Committee, and moderated the NANOG email list. A Senior Science Writer at Merit, Susan provided user support for various operations and research projects and for Merit's regional backbone, MichNet. She was the author of three RFCs. Prior to joining Merit in 1992, Susan earned a Ph.D. in Near Eastern Studies from the University of Michigan.

Elise Gerich was a product manager at Juniper Networks. Prior to joining Juniper in 2001, she held various management and technical positions at Urban Media and Excite@Home, where she helped roll out the company's first broadband network. Prior to Excite@Home, Elise was Associate Director for National Networking at Merit Network, Inc. In this role, she worked closely with the U.S. regional networks to ensure a smooth transition from the NSFNET Backbone Service, served as Co-PI of the Routing Arbiter Project, and founded NANOG with colleague Mark Knopper. She is a former member of the Internet Architecture Board and a longtime participant in the IETF, NANOG, and RIPE.

Back to NSFNET Information.

Home » Network Research » Project History » NSFNET

About Merit | Services | Network | Resources & Support | Network Research
News | Events | Contact | Site Map | Merit Widget | Merit Network Home

1000 Oakbrook Drive, Suite 200, Ann Arbor, Michigan 48104-6794
Phone: (734) 527-5700 Fax: (734) 527-5790 E-mail: www@merit.edu

Search

Blog Home About

PolicyBlog

A Technology and Telecommunications Policy Blog

Monday, December 13, 2010

Twitter, Broadband and Innovation

Link Hoewing posted in PolicyBlog Broadband on June 08, 2009, 02:07 PM EST

Time magazine carries and intriguing story about the emergence of Twitter as a new communications tool. It lauds the service as providing more evidence of the innovativeness of the U. S. economy pointing to the emergence of a host of online services and applications such as Facebook, Google, Blogger and Wikipedia as evidence that the U. S. remains one of the most innovative countries in the world.

I think an untold part of this story however is role broadband deployment and connectivity played in the process. Recall AOL in its heyday. People were introduced to online services, email and Instant Messaging through AOL but the use of dial up access led to some real limitations that stifled the emergence of rapid innovation online. Why?

First, with dial up, users could not become a part of the Internet in any real sense. You could not leave a dial up connection "nailed up" as they used to say because the technology used up lots of capacity and caused congestion on voice networks. Voice remained the most important application for most people and dial up connections used the capacity of an entire copper line going to a home. People typically dialed in, checked their email, looked for some content checking a few popular content sites, and then disconnected.

Broadband technology (cable modem and ADSL initially) allowed people to connect to the Internet and stay connected all day if they wished. Computers became true "nodes" on the Internet and people became more integrated with the online world and with the content that was out there.

Second, capacity broadband capacity began to expand as I explained in an earlier blog post. Two-way communications of all kinds – voice, video, text – became common place on the Internet. Consumers became "prosumers" to use an Alvin Toffler phrase. They not only consumed content, they often created it.

Finally, the decline in the real price of PCs and the rapid deployment of broadband led to rapid growth in the number of consumers connected to the online world.

I was involved in the evolution and emergence of these broadband networks when I worked at Bell Atlantic. The company experimented with the use of ADSL to connect consumers to the much touted "Information Superhighway" which in those days was not really the Web but rather access to online video content. Bell Atlantic successfully trialed ADSL for on demand video but recognized as the trials went on that ADSL was better suited to accessing what became the true "Information Superhighway" – the World Wide Web.

As cable companies and telcos began to deploy broadband in the latter part of the 1990's, online applications began to grow and emerge. Early efforts to stimulate the use of web sites included companies like NBC which had one of the very first news sites that incorporated the use of video news clips. I remember this initiative and recall that after a year of offering the "iNBC' web news service, NBC withdrew it. One of the points it made in doing so was that without a "mass market" broadband market, interactive and video online services were not viable. This was in the late 1990's.

Just a few years later, broadband was used by a quarter of all homes on a rapid growth path upwards. As penetration rates grew, so did the emergence of online services like Facebook (2004) and the mass use of services like Instant Messaging. The "always on" character of broadband, its growing capacity in both directions and the rapid increase in uptake helped provide a platform for the emergence of more innovation in the form of new online services.

Broadband is not only a platform for innovation, it also represents innovation in many ways. The deployment of fiber by

Subsc

Categ

All C

Policy

Click

Recer

Verizon, for example, has resulted in the development of many new patents for everything from fiber connectors to "bendable" fiber lines. Speed itself – as with PCs and processors – is a form of innovation.

While we take for granted today the emergence of new innovations like Twitter, without broadband, the evolution of these new online services and applications would be far different.

By the way, I am an active user of Twitter too and would welcome feedback there if you'd like. I'm at "linkhoe" on Twitter.

Comments (1) | Permalink

SHARE

Tags: Twitter Broadband Innovation Facebook Google Blogger Wikipedia Deployment Fiber

Reader Comments

Is circuitous the new elegant?

S posted on 6/10/2009 4:02:24 PM Permalink

Post a Comment

Comments are moderated, and will not appear on this website until a moderator has approved them. The name you enter here will appear next to your comment. You must enter a valid email address to submit a comment.

Name * **Email** *

Website

Comments *

Please Add 8 and 2 and type the answer here *

Submit

All Co

Sele

Tag C

fcc
ne
int
ec
co
wir
bro
lte
fib
sei
int
tech
ope
inve
polic
pew
cybe
giga
smar
high
widg
broad
cable
Displ

About Policy Blog | Comment Policy | Disclaimer/Terms Of Use

Max Planck Institute for Software Systems

Glasnost: Results from tests for BitTorrent traffic blocking

Almost 100,000 users from locations around the world have used our tool, Glasnost, to test whether their BitTorrent traffic is being manipulated. On this page, we present preliminary results from these tests. The tests were conducted between March 18th, 2008 and January 27th, 2009.

We will update this page with more detailed results as we get more data from the tests. We also hope to uncover more cases of blocking as we refine our measurement tool and our analysis. So make sure to check back later. Alternately, you can stay up-to-date on our findings by subscribing to the glasnost-updates mailing list.

We have released the source code of our tool. You are welcome to download and inspect the code. Please contact us if you find any bugs or have questions, comments, or suggestions.

We published a paper on this work in the ACM Internet Measurement Conference 2008. It contains updated results up to July 25th, 2008 from more than 47,300 end users. You can download the paper in pdf format here.

Links to older versions of this page: 9.11.2008 and 25.07.2008.

Table of contents

1. HOW DO WE DETECT BITTORRENT BLOCKING BY ISPS?

At a high level, our test sets up a series of BitTorrent flows between an end user's host and our Glasnost test servers. We collect the packet trace for each flow on the server side, and we closely monitor both end points for any error conditions that might cause a flow to be aborted. If a flow is aborted by a control (RST) packet that was not sent by either of the end points, we report the flow as being blocked by some ISP along its path.

For more details on how we detect BitTorrent blocking, please click here.

2. ARE ISPS BLOCKING BITTORRENT TRAFFIC?

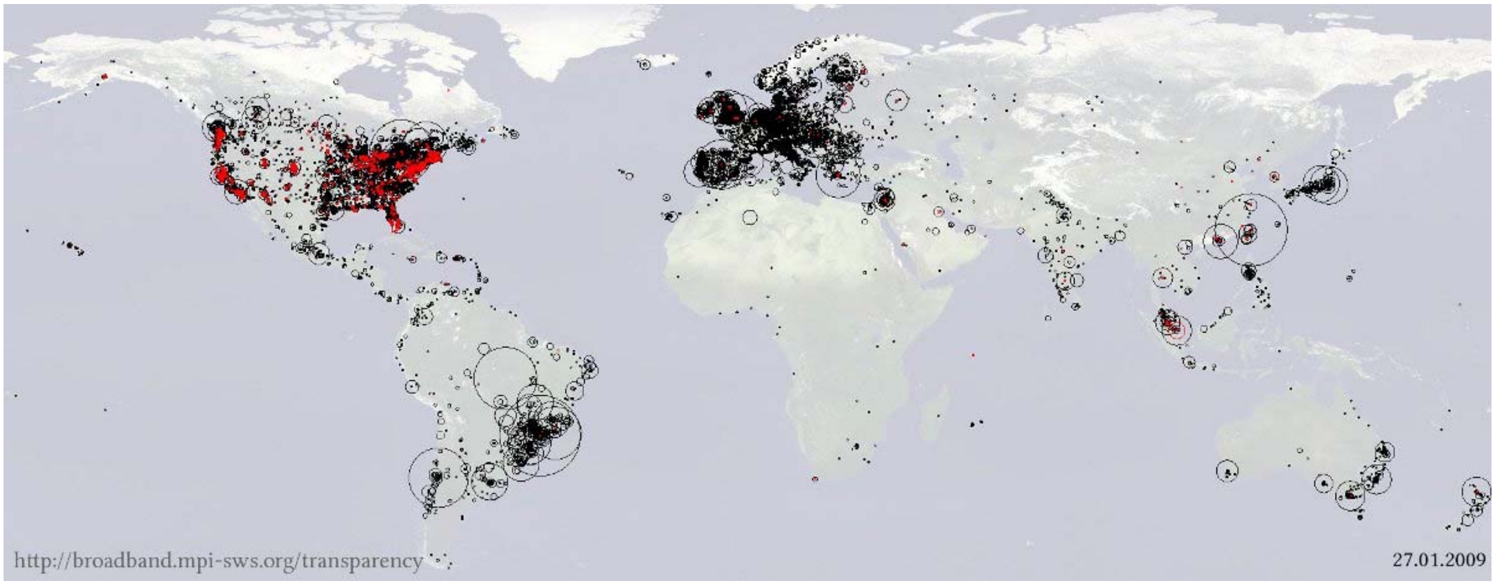

Note: This map was created using GeoLite data by MaxMind, available from http://www.maxmind.com/.

The map plots the geographic location of the 98,530 nodes that ran our BitTorrent tests. These hosts are distributed across 157 countries and 3,024 access ISPs. Hosts that found their BitTorrent transfers being blocked are marked in red. Circles represent multiple measurements from the same location; the bigger the circle the more the number of measurements from the same place.

Note: ISPs may throttle (rate-limit) BitTorrent traffic without blocking it. The results we present here are limited to hosts whose BitTorrent transfers to our servers are blocked, i.e., interrupted by RST packets generated by some ISP along the path. We are still actively investigating techniques to accurately detect throttling. So we do not report any results on rate-limiting BitTorrent traffic at this time and we do not mark such throttled hosts in red.

The table below shows for each country (a) the number of hosts that ran our test, (b) the number of hosts for which we detected BitTorrent blocking, (c) the number of distinct access ISPs from which our test was run, and (d) the number of these ISPs that contained one or more hosts for which we detected BitTorrent blocking.

Country	# measured hosts	# blocked hosts	# measured ISPs	# blocked ISPs
Afghanistan	2	0	1	0
Albania	12	0	3	0
Algeria	158	0	3	0
American Samoa	1	0	1	0
Andorra	5	0	1	0
Angola	1	0	1	0
Antigua and Barbuda	2	0	1	0
Argentina	961	0	36	0
Aruba	3	0	1	0
Australia	2047	1	76	1
Austria	410	0	34	0
Azerbaijan	2	0	1	0
Bahamas	3	0	2	0
Bahrain	14	0	4	0
Bangladesh	42	0	19	0
Barbados	17	0	1	0
Belarus	2	0	2	0
Belgium	479	2	21	1
Belize	3	0	1	0
Bermuda	1	0	1	0
Bhutan	1	0	1	0
Bolivia	18	0	7	0
Bosnia and Herzegovina	21	0	10	0
Botswana	1	0	1	0
Brazil	12727	4	128	2
Brunei Darussalam	7	0	1	0
Bulgaria	259	0	59	0
Cambodia	4	0	4	0
Canada	8636	23	161	10
Cayman Islands	3	0	1	0
Chile	1971	0	25	0
China	231	13	32	6
Colombia	216	0	22	0
Costa Rica	25	0	2	0
Cote D'Ivoire	1	0	1	0
Croatia	192	0	10	0
Cyprus	49	0	7	0
Czech Republic	196	0	33	0
Denmark	369	0	29	0
Dominica	3	0	1	0
Dominican Republic	47	0	2	0
Ecuador	22	0	8	0

Country	# measured hosts	# blocked hosts	# measured ISPs	# blocked ISPs
Latvia	56	0	19	0
Lebanon	15	0	6	0
Liechtenstein	4	0	2	0
Lithuania	198	1	27	1
Luxembourg	30	0	6	0
Macau	32	0	1	0
Macedonia	31	0	7	0
Madagascar	1	0	1	0
Malaysia	758	27	15	3
Maldives	9	0	2	0
Malta	42	0	7	0
Marshall Islands	1	0	1	0
Martinique	5	0	3	0
Mauritius	3	0	2	0
Mexico	644	0	24	0
Moldova, Republic of	10	0	2	0
Monaco	1	0	1	0
Mongolia	1	0	1	0
Montenegro	6	0	2	0
Morocco	10	0	2	0
Mozambique	2	0	2	0
Nepal	1	0	1	0
Netherlands	1961	0	40	0
Netherlands Antilles	16	1	6	1
New Caledonia	2	0	1	0
New Zealand	585	2	23	2
Nicaragua	6	0	3	0
Niger	1	0	1	0
Nigeria	3	0	2	0
Northern Mariana Islands	4	0	1	0
Norway	753	0	32	0
Oman	3	0	2	0
Pakistan	77	0	11	0
Palestinian Territory	19	0	2	0
Panama	47	0	5	0
Paraguay	30	0	4	0
Peru	88	0	4	0
Philippines	424	0	16	0
Poland	321	0	61	0
Portugal	1680	0	19	0
Puerto Rico	72	2	9	1
Reunion	10	0	5	0
Romania	598	1	64	1

Egypt	92	0	6	0
El Salvador	18	0	4	0
Equatorial Guinea	1	0	1	0
Estonia	81	2	8	2
Faroe Islands	19	0	4	0
Finland	828	4	34	1
France	969	0	27	0
French Guiana	2	0	1	0
French Polynesia	1	0	1	0
Gabon	1	0	1	0
Georgia	1	0	1	0
Germany	3619	1	84	1
Ghana	3	0	1	0
Gibraltar	2	0	1	0
Greece	1164	8	22	1
Guadeloupe	3	0	2	0
Guam	17	0	2	0
Guatemala	20	0	6	0
Guyana	3	0	1	0
Honduras	16	0	5	0
Hong Kong	560	0	17	0
Hungary	1440	1	47	1
Iceland	50	0	5	0
India	733	2	31	2
Indonesia	87	1	21	1
Iran, Islamic Republic of	53	0	14	0
Iraq	2	0	2	0
Ireland	544	6	27	3
Israel	971	3	17	3
Italy	2929	0	69	0
Jamaica	25	1	4	1
Japan	2227	0	119	0
Jordan	25	0	8	0
Kazakhstan	7	0	2	0
Kenya	1	0	1	0
Korea, Republic of	51	9	10	1
Kuwait	51	8	9	2
Russian Federation	249	1	85	1
Saint Lucia	5	0	2	0
Saint Vincent and the Grenadines	6	0	2	0
San Marino	3	0	1	0
Saudi Arabia	42	1	12	1
Senegal	2	0	1	0
Serbia	278	0	18	0
Seychelles	2	2	1	1
Singapore	448	103	16	2
Slovakia	111	1	20	1
Slovenia	158	0	15	0
South Africa	42	1	11	1
Spain	4160	1	46	1
Sri Lanka	24	0	4	0
Sudan	9	0	3	0
Sweden	674	1	40	1
Switzerland	408	0	33	0
Taiwan	2385	11	38	5
Tanzania, United Republic of	1	0	1	0
Thailand	169	1	13	1
Trinidad and Tobago	33	0	3	0
Tunisia	44	0	1	0
Turkey	66	0	7	0
Ukraine	40	0	25	0
United Arab Emirates	36	0	4	0
United Kingdom	6076	11	90	4
United States	28409	4318	871	43
Uruguay	40	0	4	0
Vanuatu	1	0	1	0
Venezuela	116	0	7	0
Vietnam	142	0	5	0
Virgin Islands, British	1	0	1	0
Virgin Islands, U.S.	7	0	4	0
Yemen	1	0	1	0
Zambia	1	0	1	0
Zimbabwe	1	0	1	0
Total	98530	4575	3024	110

3. DETAILS OF BLOCKED BITTORRENT TRANSFERS

1. All hosts which observed blocking did so in the upstream direction (i.e., when the client host attempted to upload data to one of our Glasnost servers). Only a handful of hosts observed blocking for downstream BitTorrent transfers.

2. We found widespread blocking of BitTorrent transfers only in the U.S. and Singapore. Interestingly, even within these countries, most of the hosts that observed blocking belonged to a few large ISPs.

3. Both in the U.S. and in Singapore, all hosts that suffered BitTorrent blocking are located in cable ISPs. We did not see any blocking of BitTorrent transfers from DSL hosts in these countries.

Most (4,054 of 4,318) U.S. hosts that observed blocking are located in Comcast and Cox networks. In Singapore, almost all blocked hosts are connected using the StarHub network. While we did observe blocking for hosts in 110 other ISPs (43 of which are in the U.S.), we did not see widespread blocking of BitTorrent traffic for hosts in those ISPs.

The table below shows the details of BitTorrent blocking for Comcast, Cox, and StarHub. For each ISP, we show (a) the number of distinct hosts we measured, and (b) the number of these hosts for which we detected BitTorrent blocking.

ISP	# measured hosts	# blocked hosts
Comcast	7719	3416
Cox	1940	638
StarHub	314	102

4. IS BITTORRENT BLOCKED ONLY DURING PERIODS OF NETWORK CONGESTION?

Recently it has been reported that Comcast defended its BitTorrent blocking before FCC as a necessary practice that is done only during periods of heavy network traffic. It is widely known that network traffic exhibits a strongly diurnal pattern. So we analyzed our data to see if hosts in Comcast and Cox networks see fewer of their upstream transfers blocked during early morning or weekends (when network load is generally low) than during other times of the day.

The left graph below shows (a) the number of measurements to Comcast hosts at different hours of the day and (b) the percentage of these measurements for which we observed BitTorrent blocking. The percentage of blocked tests remains high at all times of the day. Our data suggests that the BitTorrent blocking is independent of the time of the day.
The right graph below shows that the percentage of blocked BitTorrent connections remains fairly high even during the weekends for Comcast hosts.

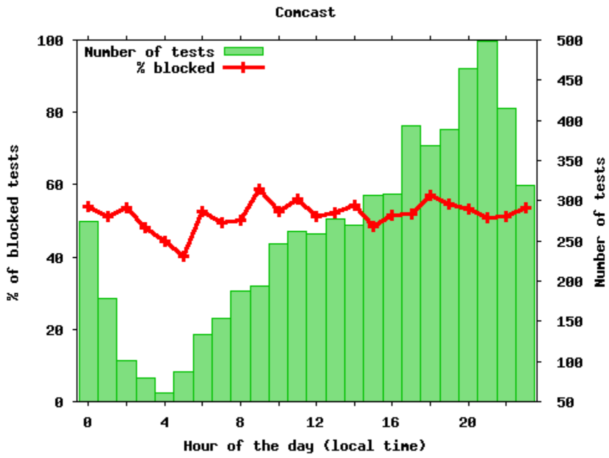

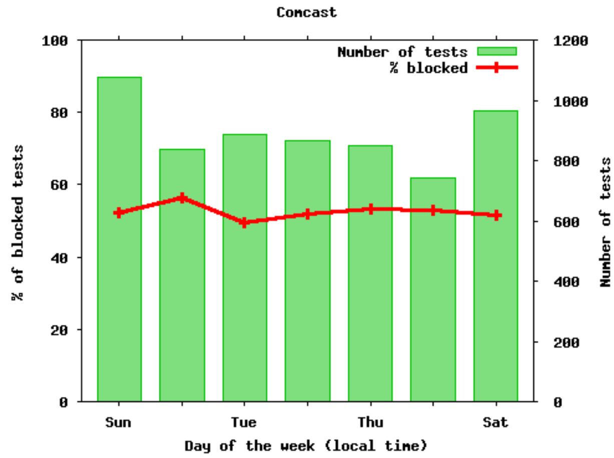

Similarly, the left graph below shows that Cox hosts suffer BitTorrent blocking at all times of the day. Note that the data for Cox is more noisy than Comcast, due to smaller number of measured hosts.
The right graph below shows that the percentage of blocked BitTorrent connections remains fairly high even during the weekends for Cox hosts.

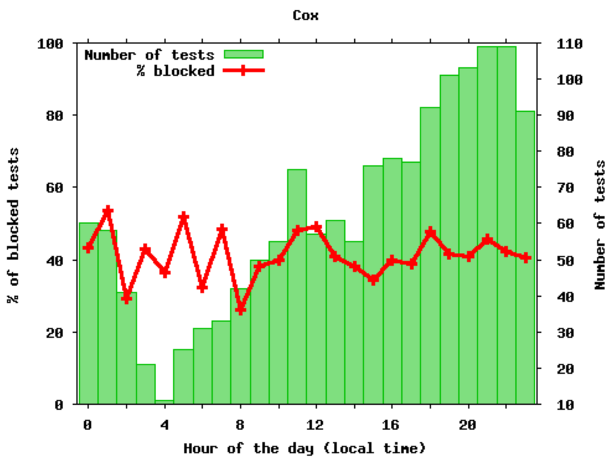

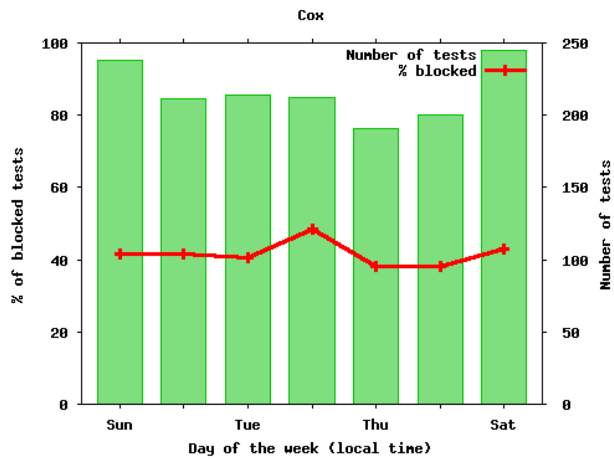

5. ARE ISPS CHANGING THEIR BITTORRENT BLOCKING POLICIES?

The graphs below show how the percentage of blocked BitTorrent connections changed during the last 10 months for Comcast, Cox, and StarHub customers. We notice a significant reduction in blocked BitTorrent tests for all three ISPs.

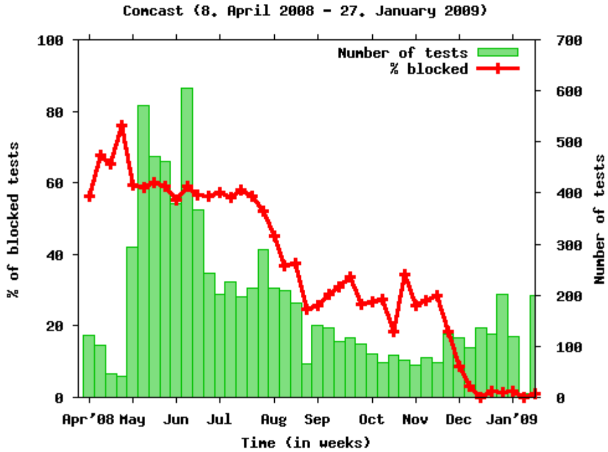

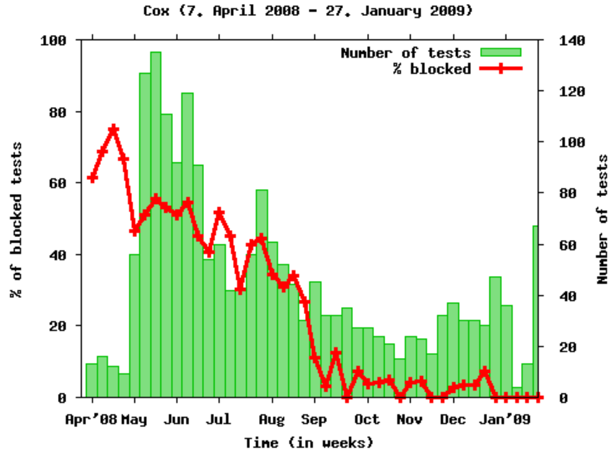

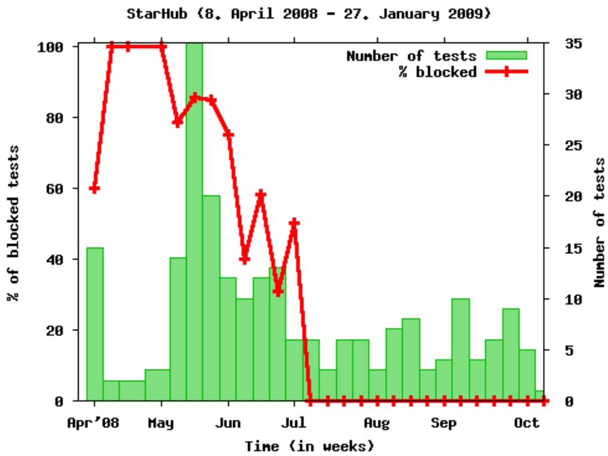

CONTACT

We are researchers at the Max Planck Institute for Software Systems. Our research focuses on characterizing residential broadband networks and understanding their implications for the designers of future protocols and applications. In case you have questions about this tool or our research, please visit our network transparency project webpage or contact us via e-mail: broadband@mpi-sws.org

Faculty

- Krishna P. Gummadi

Students

- Marcel Dischinger
- Andreas Haeberlen
- Alan Mislove

- Log In
- Forgot Password?
- Sign Up

-
- TV
- Movies
- Hulu Plus New

- Channels
- Most Popular
- Recently Added
- Collections
- Trailers
- Spotlight

search

Media Info

Get the latest info on Hulu's activites and executives.

- Media Info
- Company Timeline
- Press Quotes
- Awards
- Media FAQ
- Press Kit
- Hulu Screenshots
- Hulu Logos
- Hulu B-Roll
- Hulu Advertising B-Roll
- Hulu Plus
- Executive Team
- Content

- Distribution
- Benefits
- Products
- Widget Gallery
- Partners
- Advertising
- Ad Experience
- Case Studies
- Ad Specs
- User Feedback
- Product Tour

Company

Founded in March 2007, Hulu is operated independently by a dedicated management team with offices in Los Angeles, New York, Chicago and Beijing. NBC Universal, News Corp., The Walt Disney Company, Providence Equity Partners and the Hulu team share in the ownership stake of the company.

Mission

Hulu's mission is to help people find and enjoy the world's premium video content when, where and how they want it. As we pursue this mission, we aspire to create a service that users, advertisers, and content owners unabashedly love.

Overview

Hulu is an online video service that offers a selection of hit shows, clips, movies, and more at Hulu.com and numerous destination sites online and across four screens — PCs, TVs, mobile phones and tablets. Hulu's selection of premium programming is provided by more than 225 leading content companies, including FOX, NBC Universal, ABC, Lionsgate, MGM, National Geographic, Paramount, A&E Television Networks, PBS, and Warner Bros. Television Group. For more details on Hulu's service, check out the Hulu product tour and the Hulu Plus Guided Tour.

Content: Hulu brings together a large selection of videos from over 225 leading content companies, including FOX, NBC Universal, ABC, ABC Family, Biography, Lionsgate, Endemol, MGM, MTV Networks, National Geographic, Digital Rights Group, Paramount, PBS, Sony Pictures Television, Warner Bros. and more. Users can choose from more than 2,600 current primetime TV hits such as *The Simpsons*, *30 Rock*, *Lost*, *Glee* and *The Office* the morning after they air; classics like *Buffy the Vampire Slayer*, *The A-Team*, *Airwolf* and *Married...with Children*. Hulu also offers a wide array of more than a 1,000 movies and documentaries; and clips from Saturday Night Live; web originals like If I Can Dream, Fake It Til You Make It, The LXD and twentysixmiles; and other popular TV shows and movies.

User Experience

Hulu is *focused on quality and convenience* and strives to create the best possible online video experience.

- Hulu gives users the ability to customize their viewing experience online.
- Hulu allows users to watch favorites or discover new shows anytime — at home or on the road.
- Hulu's search feature helps users find any premium video online even if it is not

directly available on Hulu.com.

Hulu is *easy to use and share.* Simply go to www.hulu.com, and click on a video to watch right away.

- Hulu does not require a download of any software. Users only need a Flash 10.0.32 enabled computer and an Internet connection to enjoy.
- Hulu offers the freedom to share full-length episodes or clips via e-mail or embed on other Web sites, blogs and social networking pages.
- Hulu's clipping feature allows users to select a portion of the video they would like to share.

Hulu is *free* and *legal* through an advertising supported model.

- Videos are available for unlimited streaming; watch favorite shows and clips over and over, for free
- Videos contain fewer ads than on TV. Advertisements appear during normal commercial breaks
- Hulu acquires the rights to distribute its videos, making them available to users legally

Distribution: Hulu allows users to enjoy great videos on Hulu.com and on 40 other popular Web sites across the Web. Hulu videos are available on AOL, IMDb, MSN, MySpace, and Yahoo! in the U.S. as well as a growing network of personal blogs, fan sites, and other Web sites where users choose to embed the Hulu video player. Additionally, users can access Hulu content across four screen — PCs, TVs, mobile phones and tablets.

Advertising: Hulu gives advertisers an opportunity to associate their brands with premium online video content, connect with highly engaged consumers and extend their reach beyond Hulu.com to Hulu's distribution network. Additionally, Hulu offers and is committed to the continued development of innovative, new advertising experiences. Currently, Hulu partners with more than 400 advertisers including Johnson & Johnson, McDonald's, Visa, American Express, Best Buy, Chili's, DirectTV, GM, Intel, Nissan, State Farm, Unilever, Wal-Mart, Cisco, and Procter & Gamble.

- [Log In]
- [Forgot Password?]
- [Sign Up]

- hulu
- TV
- Movies
- Hulu Plus New

- Channels ▾
- Most Popular
- Recently Added
- Collections
- Trailers
- Spotlight

search

News and Information

All Current News Documentary and Biography Live Events and Specials Politics Science and Technology

Shows Videos

Search this channel Popular Searches: history channel fox news national geographic history fox

Filters: Captions Free

Sort by: Most Popular Today

What Would You Do?
subscribe

America's Most Wanted
subscribe

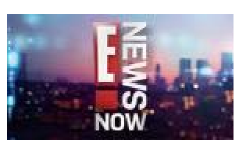

E! News Now
subscribe

ABC 20/20
subscribe

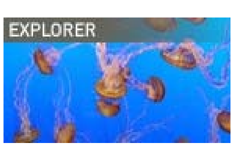

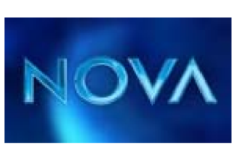

[Explorer](#)
[subscribe](#)

[NOVA](#)
[subscribe](#)

[Modern Marvels](#)
[subscribe](#)

[Biography](#)
[subscribe](#)

[Pawn Stars](#)
[subscribe](#)

[Cosmic Journeys](#)
[subscribe](#)

[SWAT](#)
[subscribe](#)

[ABC Nightline](#)
[subscribe](#)

[Paranormal State](#)
[subscribe](#)

[ABC News Specials](#)
[subscribe](#)

[Intervention](#)
[subscribe](#)

[Big Cat Week](#)
[subscribe](#)

[FORA.tv Science](#)
[subscribe](#)

[Cuba: Wild Island of the Caribbean](#)
[subscribe](#)

[The Strip Game](#)
[subscribe](#)

[CNBC Originals](#)
[subscribe](#)

-
-
- 1
- of
- 28
-

search Advanced Search | Search Plugin

Home TV Movies
Channels Labs Collections
Trailers Discussions Spotlight
Distribution
Advertising
RSS

Jobs Media Info Blog
Help Product Tour Contact Us

----- Original Message -----
From: Fred Humphries (LCA)
To: Josh Gottheimer
Sent: Wed Dec 01 13:58:51 2010
Subject: RE:

Below is our press statement:

We believe that today's net neutrality proposal from the FCC represents an important step forward. It's vital for net neutrality regulations to promote new investments by network operators to increase broadband speeds, while also ensuring that consumers easily can access content and services from both start-up and established companies with the necessary quality of service and at the highest speeds available. While we need to learn more about the proposal's all-important details, including as they relate to so-called 'specialized services,' we're encouraged that today's proposal attempts to address all of these needs."

Thanks

News & Analysis

16-Nov-2010 by John Moulding

TV Everywhere: more than one aggregation model

Convergence TV

One of the major challenges for Pay TV operators today is to develop TV Everywhere services that provide their subscribers with access to programming online and especially on-demand. Two models have developed for how the video is aggregated online: one where a service provider ingests, aggregates, stores and plays out the content provided by its content partners, and the other where a programmer manages its own online video portal and subscribers are effectively referred to that site from the Pay TV operator.

According to Ian Blaine, CEO of thePlatform, the online video publishing specialist and the company that manages the video backend for the TV Everywhere portals at North American cable operators Comcast and Rogers, neither approach is going to dominate. “Comcast puts the emphasis on its own destination site because they have a bunch of services in one place and a Comcast subscriber can get to everything. So their preferred approach is aggregation,” he says “But we work with companies who are putting content online at their own websites and we help them provide the connection to their operator partners.”

That ‘connection’ means the authentication that confirms that a consumer is a Pay TV subscriber, since the whole point of TV Everywhere is that online access is bundled with an appropriate channel subscription. So where a programmer is managing its own online video offer, visitors must prove they have a Pay TV channel subscription covering the content.

Who manages the content library will have an influence on the branding of the service, although the brands of both the Pay TV provider and the content owner will always be present, Blaine points out. “We have customers, like Rogers and Comcast, that ingest themselves and other operator customers that rely on content owners to ingest and manage content on the content owner website,” he adds. “There is a lot of experimentation and it is early days for both models. We think both models will win. Aggregators will continue to aggregate but content owners will also have a path for their customers who remain loyal to the brand.”

Having established its credentials for Pay TV operator TV Everywhere in North America, thePlatform has turned its attention to Europe as a potential market for its services. The company views the advent of multi-screen viewing as an opportunity and believes its experience with the Pay TV online authentication model means it is well placed to help service providers launch TV Everywhere services here.

For all Pay TV operators, TV Everywhere provides a way to defend their existing business and maintain the current pricing of their television by offering a better value service in the face of growing competition that includes attempts to bypass them by over-the-top services. Blaine views this as one of the main justifications but says multiplatform viewing can also be monetized through advertising against premium content online. And he points out: “This is the first time some of this content has been made available online legally.

“Consumer device options are starting to provide a pull [for multi-screen services] because there are so many video enabled devices that offer a good experience. So it is incumbent on service providers to meet that desire and I think there is more pressure to provide value on the subscription,” he explains.

Blaine says one of the benefits to a platform operator of taking responsibility for ingest, storage and playout themselves is the ability to ensure a consistent service quality for users. Other advantages include managing the standards across all platforms and the ability to advertise across content, which will help monetize the investment in multiplatform services. “It is actually not a great deal of work, since platform operators already do this for their own portal websites today,” he suggests.

In theory, service providers could offer ad insertion and targeting for use with their subscribers to boost the advertising model. “Nobody is targeting against consumer data today, but the information is important so they can advertise against it,” Blaine comments.

Most Read & Popular Content

News Blogs White Papers

These are the most read news stories of the past fortnight.

1. TF1 ready to expand connected TV offering
2. Telstra relying on Connected TV IPTV
3. Multi-screen viewing demands new audience measurement
4. The future of TV advertising, Part I
5. Consumers need to control ad targeting
6. Orange is introducing mobile-specific ads
7. Favourable results for consumer ad selection
8. YouTube will only charge for viewed ads

These are the most read blog items of the past month

1. TV advertising’s challenge to be liked

These are the most read white papers of the past 3 months

1. Over The Top TV (OTT TV) Delivery

One of the primary technology requirements for TV Everywhere is the authentication and rights management and thePlatform has published a mini-White Paper called 'TV Everywhere Certification Guide' that explains how the service provider and content owner hosting models can work in parallel. In particular, it explains how the system can work so that the Pay TV subscription model is not undermined when the video is managed at the content owner website.

The paper notes that the first challenge is knowing when an individual consumer subscribes to a particular Pay TV service. "The good news is that most TV service providers have built, or are in the process of developing, an authentication system in front of their billing system," it explains. "As a first step, the programming network's site can add an authentication widget to check against the TV service provider's authentication system. When a user enters a user name and password, the TV service provider's system will validate the person as an active customer and provide an associated list of channels to which they subscribe."

To overcome the fact that service providers have different technologies for authenticating customers, each with their own backend billing system and a variety of non-standardized web interfaces, thePlatform has introduced an 'Authentication Adaptor' service. This translates the user credentials into the format required by each authentication system. The website has to deal only with one unified authentication scheme, simplifying the process.

thePlatform adds: "To complement this authentication, TV service providers may want to establish policies that enable playback on specific PCs or devices. In these cases they can use thePlatform's Device Registry. As part of the sign-up process the customer registers their device with their account. With this information the service provider can set up policies for multiple devices, such as automatically authenticating the user if they are logging in from their broadband modem with a registered device. This is similar to the set-top box model today, where customers are not asked to log in to watch TV.

"When the customer takes their laptop on the road, additional policies can be applied. For example, if the user is off-network but logging in from a registered device, the system can just ask for their password. Apple iTunes employs this policy when customers are purchasing applications and content. As a last policy, if the customer is off-network and logging in from an unregistered device, the system can ask for user name and password."

The sharing of user credentials is an obvious worry for service providers and their content partners since it is a lot easier to share a laptop, user name and password with a friend than to share a set-top box. So the Device Registry can be used to set a limit on the number of approved devices or to require a device to log in via the broadband modem once per month.

Once a customer has been validated, the authentication system should return a list of channels to which they subscribe. The programming network can then personalise the site for the user based on their Pay TV subscription package, although there is, of course, the option to also show viewers content they do not subscribe to in order to encourage upgrades.

A further stage in the TV Everywhere process is to authorise the video to play. This goes beyond authentication and ensures the media rights associated with individual shows, like airdates, geographic restrictions and other business polices, are enforced. The authorisation process should not take longer than 250-500 milliseconds, according to thePlatform.

Related content:

Best ways to achieve TV Everywhere

Video: TV Everywhere ambitions present new security challenges

New products everywhere for TV Everywhere

Multiplatform TV report - March 2010

Anyscreen, Anywhere video report - July 2010

Related Posts

Telstra relying on Connected TV IPTV
Other news: Telenor and TiVo, gesture control for BSkyB...
Turner Broadcasting outlines Connected TV strategy
More top news from a busy week in convergence TV
Connected TV further extends Viasat reach
Other news: connected TV, 3D VOD, QoE, IPTV forecasts…
Other news: 3D, service provider tablets, connected TV…
Orange expects content rights to cover all screens
TF1 ready to expand connected TV offering
Welcome to the Connected TV Awards 2011
DISH Network uses connected TV to boost ARPU
Multi-screen TV: place-shifting has a role to play
Discovery: TVE defends pay model and benefits ads
Virgin Media will pioneer the Pay TV connected model
FetchTV reveals 2011 content and retail strategy
'Social-VOD' in a virtual world: another category of TV

About the author

2. Platforms Review
3. Enabling seamless digital media access in the home

   Home networks and virtual gatways
4. Eastern European Top 10 Multichannel Operators
5. Improving the relevance of TV search
6. 3DTV - Taking the TV experience to a new dimension
7. An Integrated approach to TV recommendations
8. Latin America Special Report

Videos Webcasts Industry Reports

These are the most viewed videos of the past month

1. The business rationale for changing cable modulation
2. Service providers ready to exploit online video
3. 30 second spot will be mainstay of free Connected TV
4. CDN analytics improves QoS and boosts content business
5. Content should be packaged for different 'media occasions'

These are the most viewed webcasts of the past 3 months

1. Next Generation User Interface
2. The Future of IPTV
3. The Business of Multi-platform TV
4. Adding Value to IPTV with Connected Home Services
5. Does television have what advertisers want?
6. Convergence
7. Advanced Advertising
8. Defining the 3DTV market opportunity
9. Impact of service assurance on IPTV profitability

John Moulding joined Videonet as editor at the start of 2010, having spent over 10 years writing about digital TV and the various technologies that have simultaneously disrupted and enriched the television business. With Videonet he is focused on the unstoppable march towards multiplatform, connected and personalized television. John was editor of Cable & Satellite International (now CSI) for six years before helping launch New Video Technology, and helped develop the IPTV World Series conference programmes from 2006-07. At home, he takes a Sky triple-play bundle, watches around one-third of content time-shifted, enjoys BBC iPlayer on television through the Wii, and eagerly awaits the arrival of YouTube on his own TV (the killer TV application for late on a Friday night). He is still loyal to channels - but can also remember when TV shut down after lunch.

These are the most read industry reports of the past 3 months

---

Want to add a comment?

Register at Videonet for free and you will be able to add comments on all our Blog & Analysis items.

Back to previous page

Contact Us

Please email press releases and news stories to the editor at: john@hubblemedia.co.uk.

If you'd like to advertise, sponsor a webcast, market an event or become a Videonet website contributor please contact:

Justin Lebbon - Publisher
+44 (0)7771890908
justin@hubblemedia.co.uk

Latest Tweet

VideonetNews: Multi-screen viewing demands new audience measurement #ftvads http://bit.ly/iaFR3s

Mon, 13 Dec 2010 | Follow us

RSS Feeds

Videonet @ LinkedIn

2009

MANAGING GROWTH AND PROFITS IN THE YOTTABYTE ERA

CHETAN SHARMA

Chetan Sharma Consulting

Passion for Innovation

Table of Contents

Executive Summary

In Q1 2009, the US market exceeded $10B in quarterly mobile data service revenues for the first time.[1] The subscription penetration in the US is well past 90% and the mobile data usage is on the rise. While the rate of new subscriptions has slowed, the pace of innovation is going very strong. It is quite apparent that the mobile industry is going through a significant transition from voice to data, from making calls to getting lost in applications and from voice communications to multimedia communications. Helped by the ever expanding wireless broadband networks, and release of hit devices every quarter, and consumer's insatiable appetite for information and content has brought us to the surge of a data tsunami that will shake the industry to its core.

As everything moves to digital, information repositories across the web are almost doubling every day moving rapidly to the yottabyte (YB) era. [2] The information and the desire and the capability to consume oodles of data is increasing exponentially. As a result the traffic – both Wireline and wireless is also increasing at a predictably fast rate.

In 2009, the *global yearly mobile data traffic* will reach a new milestone – 1 Exabyte (EB) or 1 Million Terabytes (TB). By 2016-17, the global yearly mobile data traffic is likely to exceed 1 Zettabyte (ZB) or 1000 Exabytes. By 2014, in the US alone, the total yearly mobile data traffic is likely to exceed 40 EB. How do you go about managing such growth in a profitable manner when the cost of supporting such traffic will increase exponentially despite the move to 4G?[3] Will the move to LTE offer some respite?

This paper discusses the analysis done by Chetan Sharma Consulting on the growth of mobile data traffic in the US market and how the ecosystem can apply some strategies to manage growth and profits. We built detailed models to estimate the rise of mobile data network traffic and discuss some solutions to handle such growth in this paper.

---

[1] Source: US Wireless Data Market update Q1 2009, Chetan Sharma Consulting. It was also the first time any nation exceeded the $10B mark in a quarter for mobile data revenues.

[2] For reference, 1 TB = 10^{12} bytes, 1 PB = 10^{15} bytes, 1 EB = 10^{18} bytes, 1 ZB = 10^{21} bytes, 1 YB = 10^{24} bytes

[3] For the purposes of this paper, we consider LTE as a 4G technology though it hasn't been officially designated as such. For more discussion on 4G, please see 4G: The State of the Union, Chetan Sharma, GigaOM, 2009

The Arrival of the Yottabyte Era

Mathematically, Yottabyte refers to a septillion or 1,000 billion terabytes. In 2006, IBM predicted that by 2010 the size of world's codified information base will be doubling every 11 hours.[4] That's beyond astronomical - paving the way to the yottabyte era.[5] When we were writing our Mobile Advertising book[6] in 2007, we started thinking about yottabytes and what it means to have all this data, how will this be accessed and understood.

What does that mean for information consumption and the resulting traffic? How do we extract intelligence out of this data without jeopardizing the network or the storage? How do we do that on a per user basis? How do we manage peaks and surges? How do we put the intelligence from the data back into the system to empower the feedback loop? There are significant opportunities in how yottabytes are managed and understood. Rather than being overwhelmed by the data tsunami, we should learn to tame the beast. As we wrote in the book, for analytics and intelligence, *data should be revered*. But for the network, *it needs to be managed*.

It is inevitable that the business models and financial metrics will also undergo a significant transition as a result of the growth in the mobile data segment. The pricing models, the customer acquisition strategies, the metrics that Wall Street pays attention to will change in due course. Instead of focusing exclusively on Average Revenue Per User (ARPU), we will start looking into Average Margin Per Subscription (AMPS), Average Connections Per User (ACPU), Customer and Family Lifetime Value that maximizes profits across all connections that a family or the user has. Instead of promoting family minutes, carriers will promote family terabyte (and someday yottabytes) plans where one can bundle usage across multiple devices used by a family.

Carriers will have to focus on how to both profit from the deluge of data as well as manage the surge in a cost-effective way that optimizes spectrum, financial, and the network resources. Companies who plan early and navigate this complex maze will benefit while the unprepared will be left to play catch-up. This also opens up the opportunity for the ecosystem to develop tools and technologies to help carriers manage growth as well as use the data generated intelligently to enhance the user experience.

[4] The toxic terabyte, IBM Global Technology Services, 2006, http://www-03.ibm.com/systems/resources/systems_storage_solutions_pdf_toxic_tb.pdf

[5] By some measures we are probably already in the yottabyte space as the growth since the prediction was originally made has been staggering

[6] Mobile Advertising: Supercharge your brand in the exploding wireless market, Chetan Sharma, Joe Herzog, and Victor Melfi, John Wiley, 2008. We wrote a section "Prepare for data everywhere – from gigabytes to yottabytes." Also see a paper written by Balasubramanium Venkitachalam of WPP titled "Your Guide to Profits: Insights through Business Analytics available at http://www.wpp.com/WPP/marketing/reportsstudies/yourguidetoprofits.htm

Mobile Media Evolution

Over the last 2-3 years, the consumption of digital media has evolved significantly. As content is becoming digital, as devices are becoming more powerful and able, and as the consumers are becoming dependent on mobile devices for their media needs, the mobile device is playing the central role in how digital media is consumed around the world. The digital rush has helped make mobile a $1.1 trillion (as of 2008) industry (Figure 1). As the demand for mobile content consumption increases, service providers are being rushed to enhance their infrastructure and keep up with the explosion of content and consumer interest.

The percentage of revenues coming from data services are increasingly becoming a big part of the overall service revenues of the operators with players like NTT DoCoMo of Japan relying on data services to provide in excess of 40% of their yearly revenues. Similarly, in US and Western Europe, data applications and services accounts for over 25% of their revenues. Globally, over 20% of the revenue now comes from data services (Figure 1).[7]

The main drivers for increased activity on the mobile devices are four-fold: better networks in the form of 3G (and future upgrades to 4G+), higher processing power devices being available for mass-market prices around the world, consumers becoming not only the consuming but also producing content at an exponential pace, and the opening up of the mobile ecosystem has attracted thousands of developers who are building compelling applications and services for various mobile platforms. As such from the early days of ringtones and graphics, the mobile ecosystem has evolved into more rich content experiences such as high-fidelity and multi-user mobile games, very high quality video in the form of multicast (though unicast is the one that is widespread), and social networking applications like Facebook and Twitter.

Additionally, the smartphone boom that followed the iPhone introduction in 2007 changed the dynamics of the market and how consumers view their mobile devices. It is interesting to note that on such integrated devices, consumers only spend less than 20% of their time on voice; rest is on other applications and services.

[7] Source: Chetan Sharma Consulting, 2009 (as of Q4 2008)

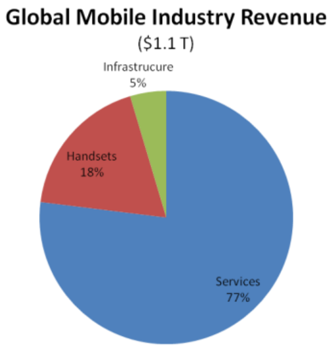

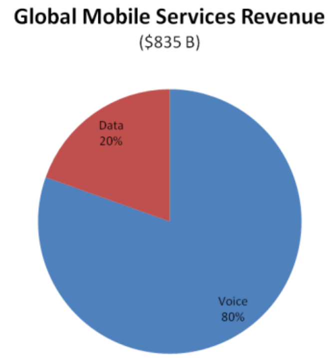

Figure 1. Global Mobile Industry Revenue Distribution[8]

Such a shift is also changing the service provider business models, how they run their business and plan for future growth. Mobile media and data services are the only driver for growth as voice revenues decline (worldwide). Significant mobile data usage is also putting strain on the operator's network and as such they are forced to come up with data expansion (3.5/4G) and alternate (Wi-Fi/Femtocell) strategies so that they can profitably stay ahead of the curve.

Also, the very definition of the mobile devices is changing. More and more consumer electronic devices are being launched with a wireless data connection (Wi-Fi or Cellular). Devices such as Kindle and Dash are introducing vertical devices that are changing the industry dynamics as well. Newer players are entering the marketplace and the competitive landscape is being impacted in many ways. Apple's appstore changed the way applications found their way onto consumer's handset. This made it easier and lucrative for a developer in a garage to launch new applications. The impact of open development and application platforms is rearranging the ecosystem in a profound way that is going to redistribute the revenues in the value chain.

While the opportunities to exploit mobile media remain strong, the ecosystem needs to worry about meeting the expectations of the consumers. They have to invest in infrastructure, developer ecosystem, and continuous flow of new and improved handsets to keep up with the growing interest. It is clear that as digital media consumption grows; mobile will be the front and center of this evolution.

[8] Source: Global Wireless Data Market update 2008, Chetan Sharma Consulting, 2009

Mobile Technology Evolution

At the end of 2008, the total mobile subscriptions crossed 4 billion.[9] Cell phone has clearly been the most adopted technology in the world. As it has moved from being a luxury to a utility to the necessity, both the cellular network technology and the devices have evolved to meet the growing demand and need for communications. From being used primarily for voice in the eighties and nineties, consumers are spending more time on their mobile devices on non-voice activities. Mobile devices have become the gadget of choice for both increasing the productivity to save time and for entertainment to kill time.

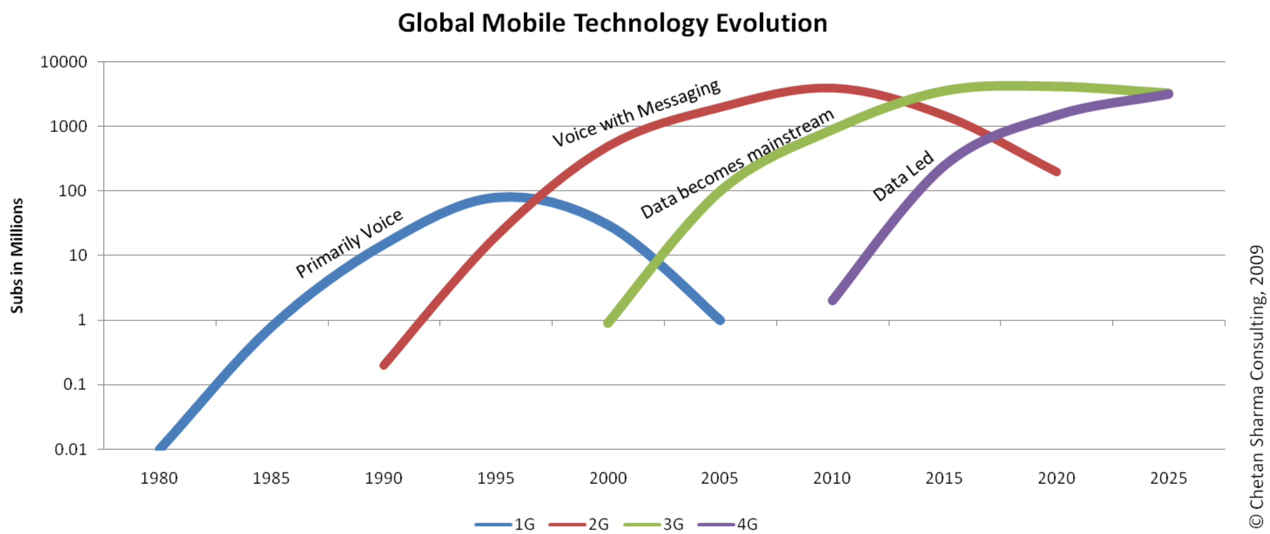

Figure 2. Global Mobile Technology Evolution[10]

If we look at the technology evolution (Figure 2) over the last 30 years, the first generation or 1G was primarily analog. With the arrival of the GSM Global Standard we moved into the digital or the 2G era but phones were still primarily voice communication devices with SMS based messaging starting to take off in various parts of the world by mid-nineties. Towards the tail-end of the last decade, NTT DoCoMo launched i-mode and defined mobile data. Soon, to meet capacity constraints and to exploit the promise of data services, 3G was launched with UMTS and EV-DO. Over time, most major carriers around the world (except in China, India, and Latin America)[11] have some form of 3G up and running by 2008. Starting next year, LTE deployments are expected to start and for the first time the motivation is solely data (initially).

[9] Source: Global Wireless Data Market update 2008, Chetan Sharma Consulting, 2009 http://chetansharma.com/globalmarketupdate2008.htm
[10] Source: Chetan Sharma Consulting, 2009
[11] China started their migration to 3G in Jan 2009 using TD-SCDMA, WCDMA and EV-DO paths.

It should also be noted that each generation of technology has a faster growth curve (compared to the previous generation). Put another way, the time it takes to amass the first 100M subscriptions shrinks with each evolution cycle. Another point that should be noted is that there is a significant overlap between technologies. Launch of 4G doesn't mean that 3G is going away anytime soon.

Each technology evolutionary step helps lower the cost of delivering the traffic. By using 4G, operators can lower their per MB costs and if they can aggressively provide coverage and move the heavy usage subscriptions to LTE, it will impact the bottom line.[12]

The Starting of the Mobile Data Tsunami

It has become a cliché to say that the iPhone has changed everything. Apple's iconic device has had a profound impact on the ecosystem on several fronts. First, Apple again taught the industry the power of simple user experience. Second, the appstore model has disrupted the traditional storefront model. Third, once exposed and addicted to mobile way of life, consumers are not looking back. Fourth, the flat data pricing is becoming the norm for the new smartphones especially in the US market.[13] Finally, all this means significant jump in network traffic. As the percentage of smartphone subscribers grow, so does the traffic – *by leaps and bound*.

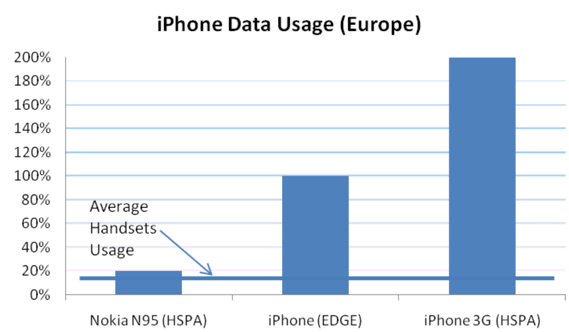

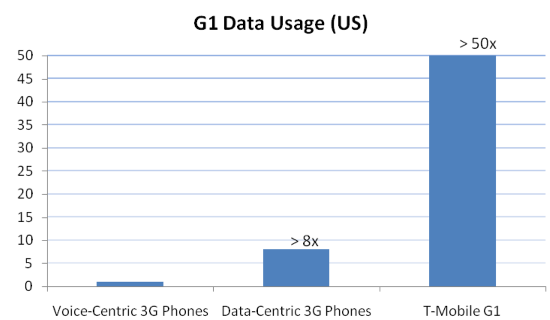

Figure 3. Data consumption driven by smartphones[14]

However, as we will show later in the paper, the main culprit of the network traffic growth is not the smartphones though they have a significant share, it's the data cards, USB dongles, and the netbooks (or smartbooks or notebooks or pick your favorite digital

[12] The move to LTE will mean faster connection which means more usage and hence more traffic but if this traffic is well managed as discussed later in the paper, the net of LTE deployment will be incremental cost savings. However, LTE might not be the most prudent strategy for every operator. The motivation and the need for 4G migration depends on a number of factors discussed on Page 16. LTE is also going to require substantial investment and will include a backhaul upgrade which is generally the biggest congestion point in the network.

[13] Not all operators in all countries are a fan of the "all-you-can-eat" data pricing. While they are offering similar devices and services, the pricing of data is based bandwidth consumption which is prudent in the long run.

[14] Source: T-Mobile, 2009

book). The data consumed by these devices is typically 5-10 times the average smartphone usage.

The usage data from the iPhone and the Android G1 (Figure 3) indicates that the data consumption increased 50-100 times compared to the previous generation of devices. As such the per user megabyte (MB) consumption is multiplying at an accelerated pace in almost every part of the world.

For example, in Hong Kong, the average MB/month consumption increased 266% (from 2007) to 40MB in 2008.[15] It is expected that the demand for data will continue its exponential trajectory. Cisco Systems estimates that the data consumption per user on a global basis will increase from 1 gigabyte (GB) in 2009 to over 14GB in 2015.[16] Operator Zain is reporting 12.5 TB/day traffic from 375,000 customers.[17] Verizon, US's largest operator reported over 3500 TB/month traffic in 2008 which could explode to 4,000,000 TB/month by 2014.[18] T-Mobile USA has indicated that their average mobile data consumption has shot-up to 3.7MB/user/month from a few KB/month just a couple years back.[19]

One should also appreciate the move towards providing wireless connectivity to all electronic devices from cameras to security alarms to energy meters. In fact, Wireless World Research Forum (WWRF) forecasts 7 trillion wireless devices serving 7 billion people by 2017[20] which translates into 1000 radios per person. Most will be in the form of sensors and tags that form the basis of forming ambient and context intelligence around us to truly make mobile device a remote control of our lives.[21] While this Always On Real-Time Access (AORTA) environment provides for instant access to information and intelligence, it also creates a mountain of data traffic that needs to be understood and managed.

The March Towards the Yottabyte Era

Consider average usage on a 3G iPhone: average session on NY Times - 0.5-1 MB, 10 min YouTube streaming - 10-15 MB, 15 min radio streaming: 12-16 MB, average application downloads: 2-10 MB, Average Facebook Sessions: 100-200 KB, etc. Eventually, they all add up to 0.5-1 GB of average data traffic/user/month. Similar trends can be found on devices like the Android G1 and the Palm Pre. One can well imagine that with increasing population of such devices and better bandwidths and

[15] Source: Key Statistics for Telecommunications in Hong Kong, Office of the Telecommunications Authority, Hong Kong, April 2009
[16] Source: Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, Jan 2009
[17] Source: Zain, 2009
[18] Source: Dick Lynch, Verizon Wireless, 2009
[19] Source: T-Mobile USA, 2008
[20] The WWRF Vision, Nigel Jefferies, WWRF, 2007
[21] Mobile Advertising: Supercharge your brand in the exploding wireless market, Chetan Sharma, Joe Herzog, Victor Melfi, 2008, John Wiley & Sons

devices, how the multiplier effect starts to take hold and move from Tera- to Peta- to Exa- to Zetta-to Yottabytes.

Methodology

To get a better understanding of the mobile data traffic growth in the US market, we built a bottom-up detailed model to estimate the current and future data traffic on the mobile networks in the US.[22]

For each of the major US carriers, we looked at the traffic generated by the data customers (excluding messaging traffic) across three distinct segments of devices – feature phones, smart phones, and data cards (including embedded devices). Then we analyzed the traffic generated by messaging and added to the non-messaging data traffic to calculate the overall mobile data traffic growth in the US market. We realized that only by looking at a very granular level can one fully grasp the various factors that impact network traffic growth.

The overall mobile data traffic in the US is expected to increase from 138 Petabytes (PB) in 2008 to 42 Exabytes (EB) in 2014 or by over 300 times (Figure 4). The main driver for this growth will be the data cards and the increasing population of the smartphones in the market.

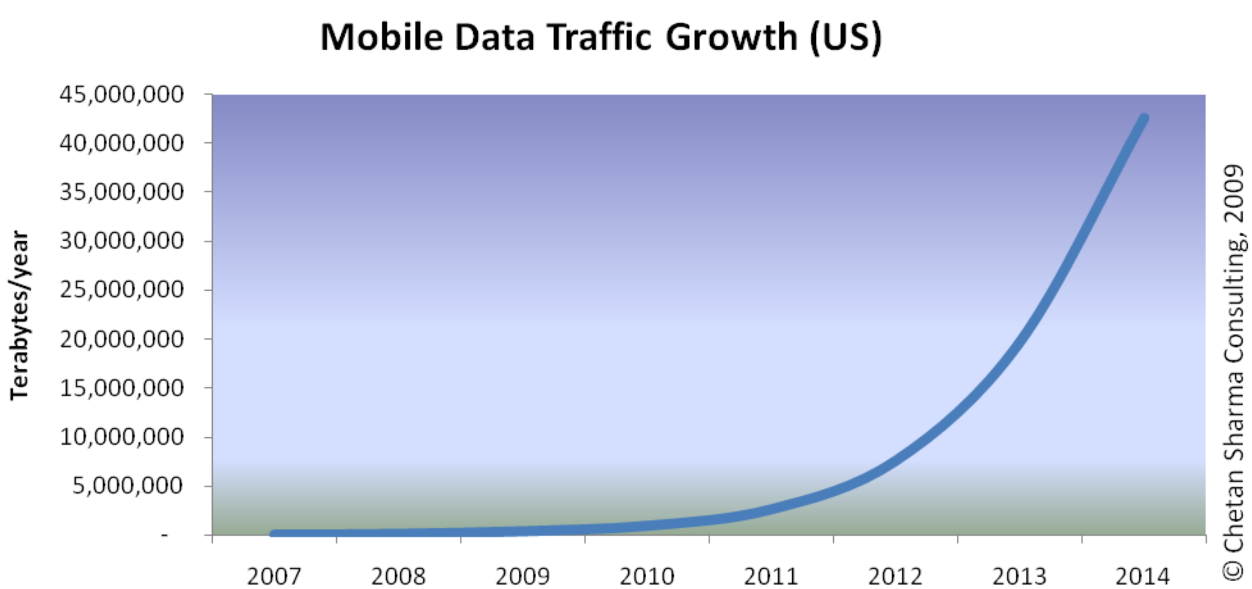

Figure 4. Mobile Data Traffic Growth in the US market[23]

[22] In this paper, we primarily explore the US market though the solutions discussed later in the paper apply to operators worldwide. Global mobile data traffic analysis is outside the scope of this paper. For reference, please see Cisco's global mobile data traffic study (note 16).

[23] Source: Chetan Sharma Consulting, 2009

This traffic can be managed and reduced by techniques discussed later in the paper like the introduction of the Femtocells and the Wi-Fi hotspots to offload traffic in fixed locations or by active congestion management and optimization. While LTE (or for that matter WiMAX) will help in reducing the cost of the traffic, it will also increase MB consumed compared to its predecessor technologies like EV-DO and HSDPA/HSPA, sometimes by 50-100%, thus neutralizing some the benefits.

The Distribution of the Mobile Data Traffic

As noted above, the data cards consume the most bandwidth on any major network in the western world. By mid-2009, data cards were accounting for over 73% of the data traffic in the US while smart phones were inching up with 24% (Figure 5). Though feature phones represented 75% of the device base, their contribution to the data traffic was only 4% (Figure 6).[24] Messaging - the biggest revenue generating category accounting for over 45% of the revenue had less than 1% of the traffic. What happens when the % share of the data cards increases? What if 25% of the subscriber base has a data card? What happens when phone-as-a-modem becomes more prominent or phone acts as a network conduit for projection screens? Network planning will need to take into account these scenarios.

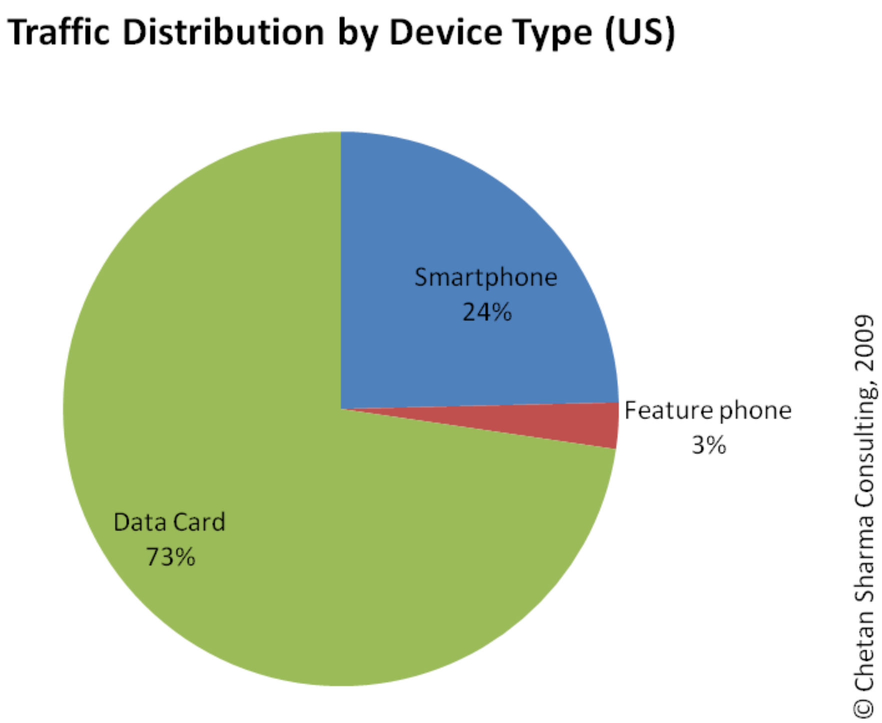

Figure 5. Mobile Data Traffic Distribution by Device Type[25]

[24] This profile will obviously look different for developing countries which have lower smartphone and data card penetration as of mid 2009

[25] Source: Chetan Sharma Consulting, 2009

Further, only a small percentage of the users are consuming a large share of data. According to Ericsson, 4% of the subscriber base account for 68% of the traffic and 96% of the subscriber base consume the rest i.e. 32% (Figure 7).

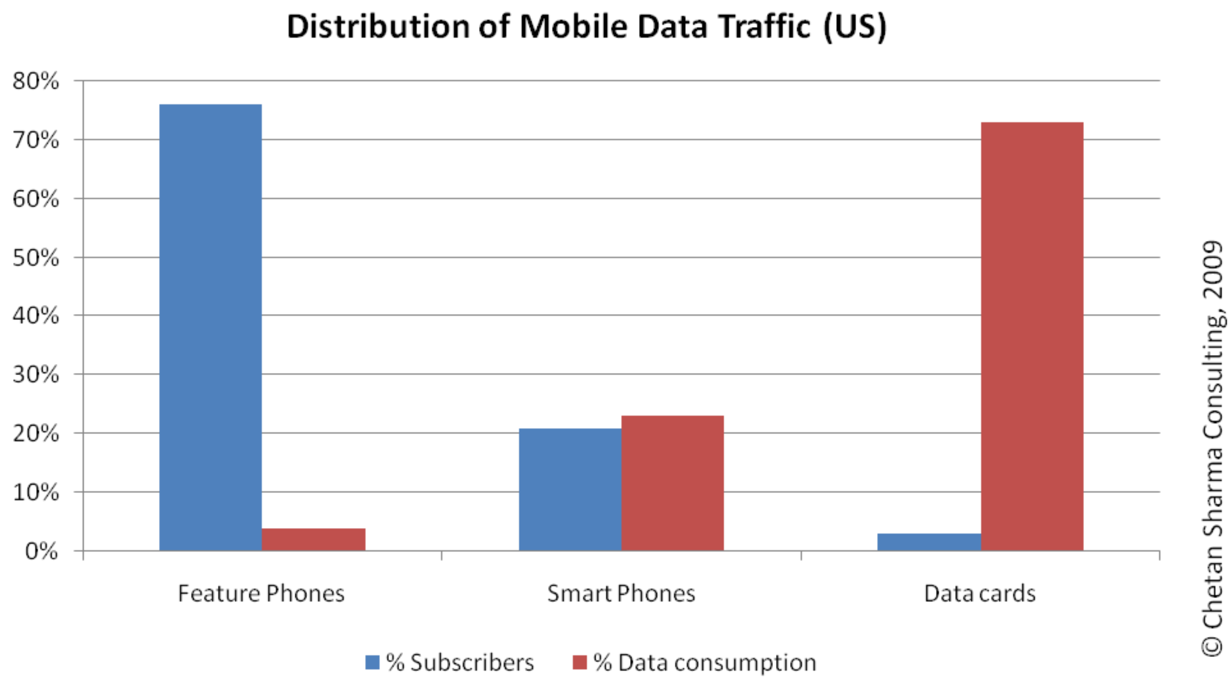

Figure 6. Mobile Data Traffic Distribution by Device Type[26]

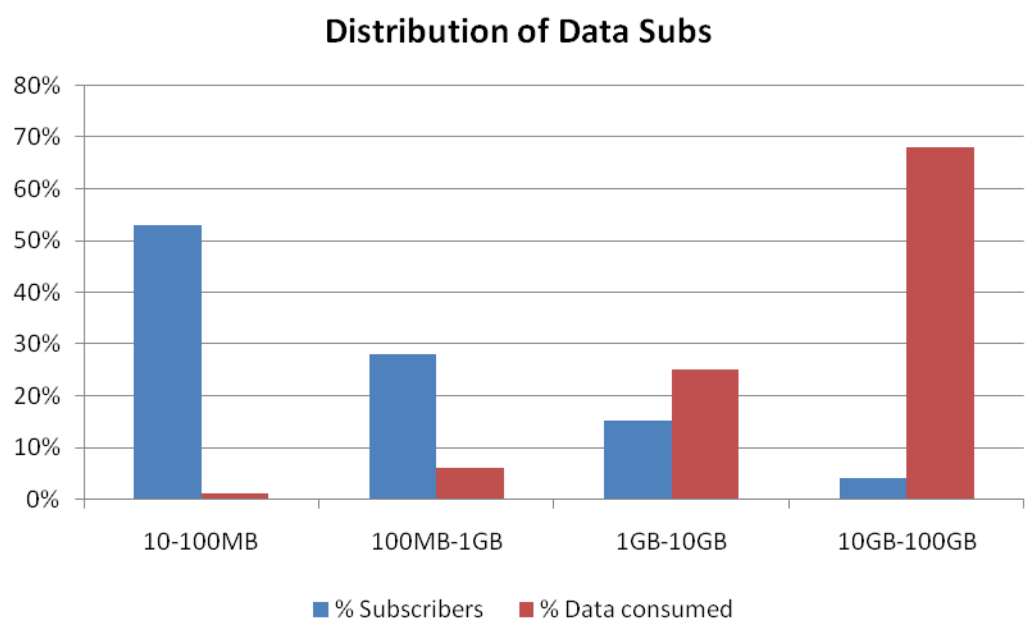

Figure 7. Distribution of Data subscribers by data usage[27]

[26] Source: Chetan Sharma Consulting, 2009
[27] Source: Ericsson, 2009

Comparing Mobile Data Growth to Wireline

To better gauge how the mobile data growth is going to progress, one should look at how data consumption grew in the Wireline world. If we look at the time-period 1996-2013 in the US,[28] one will notice that the mobile evolution has lagged Wireline growth by approximately 8-9 years (though with each passing year, the gap is closing).

As indicated in figure 8, the penetration of broadband grew in Wireline, so did the number of Petabytes consumed on a national basis for the US market. And the penetration grew with each enhancement of the technology from ISDN to FTTH. During the early years of mobile broadband in the US (2004-2008), the data consumption has followed similar patterns but since last year, with the advent of devices like the iPhone and with the increased penetration of the data-cards, the traffic is rising faster than it did in the Wireline world and we expect that by 2013, mobile data consumption will only be 5 years or less behind the Wireline data consumption.

While P2P (Peer-to-Peer) and video streaming have been the main reason Wireline consumption shot-up, in the mobile world it has been the browser usage and P2P on smartphones, netbooks, smartbooks, laptops, and similar devices that is contributing to the bulk of data traffic followed by streaming and application downloads. It is not that the video consumption won't grow on mobile; we think that the live broadcasts will be better handled by broadcast technologies which will free cellular spectrum from carrying the video load (see next section for more discussion).

Also, since peak usage significantly increases the network expenditure, the key statistics that need to be understood are around the applications that drive traffic during peak utilization, not necessary the applications that consume the most bandwidth across the entire traffic landscape. For example, P2P applications might consume a significant amount of traffic but their relative percentage during peak utilization might be smaller. As such, devoting too much time to P2P management might be that effective (though it needs to be tackled nevertheless).

So, as we try to understand mobile data consumption in the next decade, it will be worthwhile to keep an eye on how things are evolving in the Wireline space and what solutions are working in managing the traffic growth and can be adapted to the mobile environment.

[28] This profile will look different in advanced broadband nations such as Japan and Korea which are significantly ahead of the US in broadband penetration both in Wireline and wireless. Comparing broadband profiles of various nations is outside the scope of this paper.

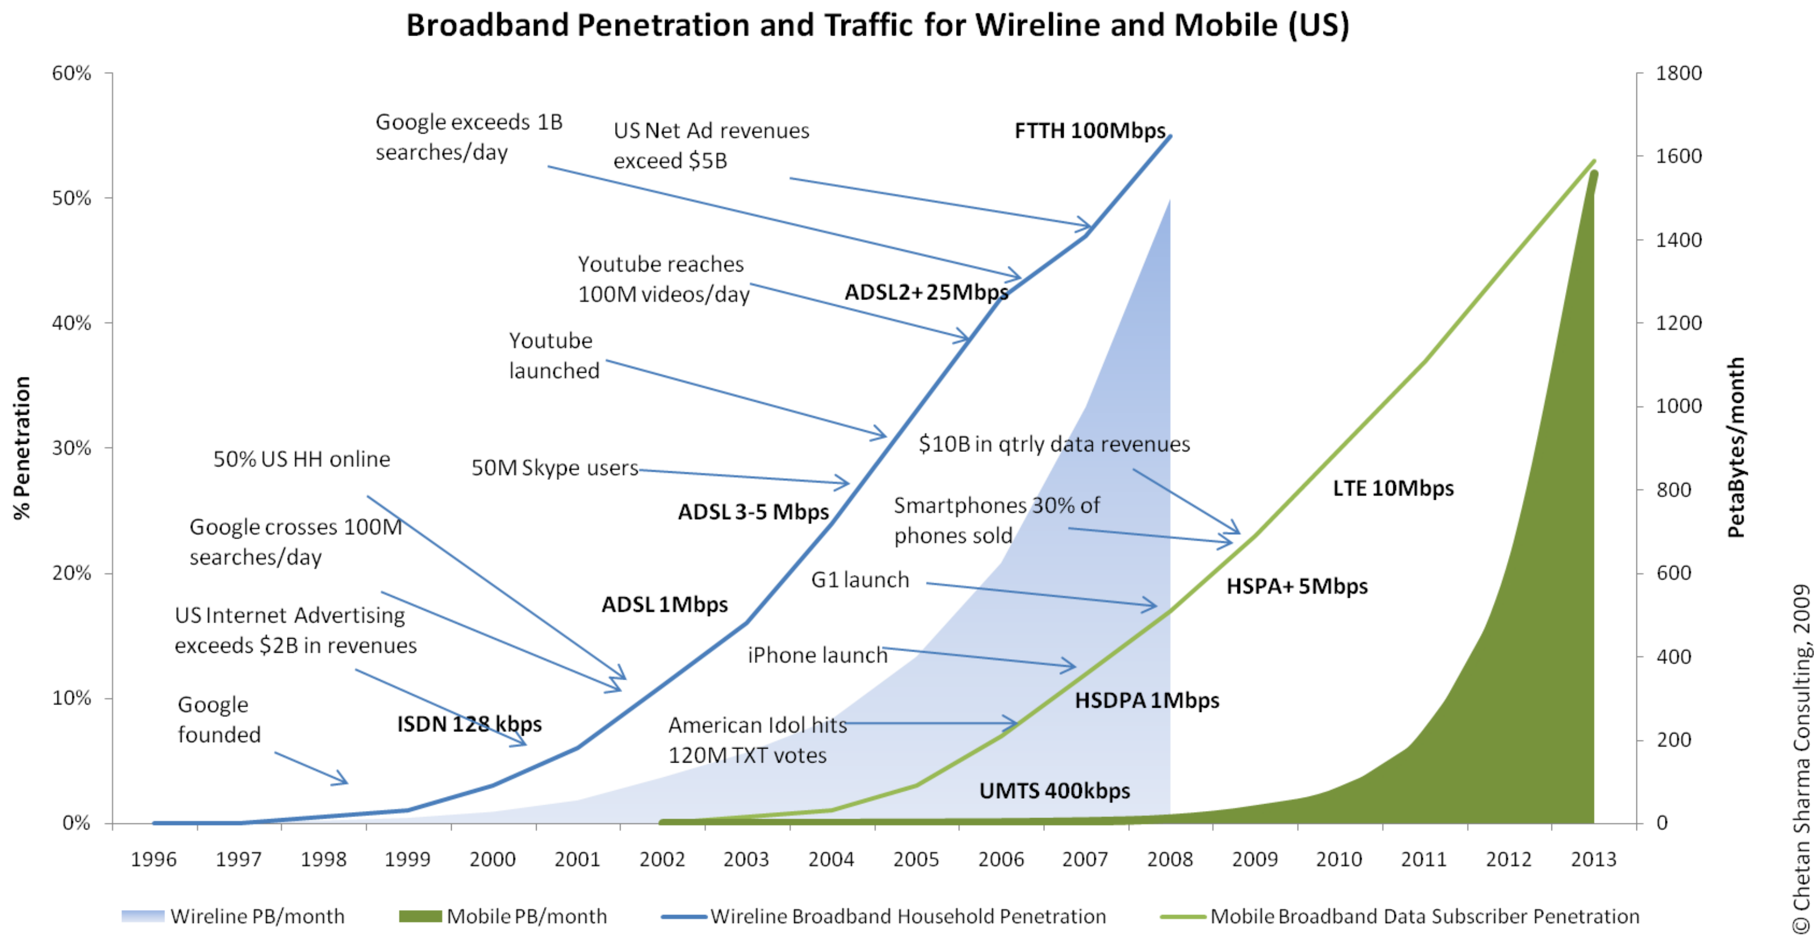

Figure 8. Broadband penetration and traffic for Wireline and Mobile data networks in the US (1996-2013)[29]

[29] Source: Chetan Sharma Consulting. Data Sources: Wireline Traffic Data from Minnesota Internet Traffic Studies (MINTS) - http://www.dtc.umn.edu/mints/igrowth.html. We took the mean of the yearly ranges. Wireline Broadband HouseHold penetration data from Pew Internet - http://www.pewinternet.org. Mobile Data traffic – Chetan Sharma Consulting analysis. Mobile Broadband Data Subscriber Penetration – Chetan Sharma Consulting analysis

Managing Growth and Profits in the Yottabyte Era

The growth in the network traffic can be quite injurious to the financial bottom line of the operator and to the industry's ability to maintain pace with the demand. Unless a long-term plan is put in place that addresses and manages the traffic at a very granular level, the cost incurred due to an explosive demand will become unsustainable by 2013. At that point the revenue being generated could fall below the cost of sustaining such traffic. However, if the operators attack the problem using several different strategies, the growth can be managed and brought in line with the technology evolution such that the industry can take advantage of the falling per MB costs.

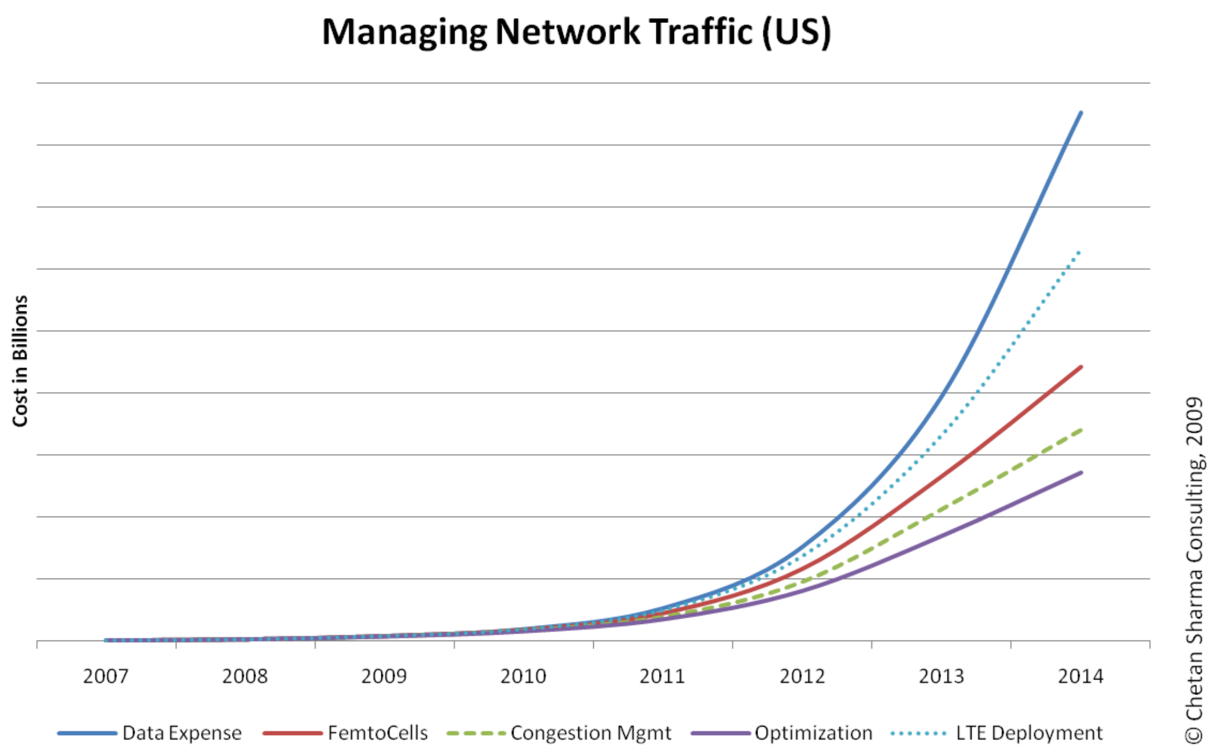

Figure 9. Managing Network Traffic Costs (US)[30]

Some of the strategies for managing network traffic growth are:

1. Faster LTE deployment
2. Femto-Cell deployment
3. Congestion Management
    a. Caching

[30] Source: Chetan Sharma Consulting, 2009

    b. Incentives to fill the troughs
4. Network Optimization
5. Adopt Broadcast Mobile Video
6. Redefine Smart phones

The eventual carrier strategy will depend on their particular situation which will differ on several fronts:

1. Spectrum position
2. 3G and 3.5G technology deployment lifecycle
3. Existing Broadband infrastructure investments
4. Network traffic and distribution of devices and traffic
5. Demographics and subscriber growth
6. Short-term and Long-term customer acquisition strategies in consumer and enterprise segments
7. Network Coverage and Pricing Plans
8. Competitive position
9. Financial resources
10. Multi-play position and strategy
11. Others

Depending on the carrier's broadband deployment and their position in the market, different set of solutions might be considered. For example, carriers who have just started to deploy 3G might want to hold-off on 4G while others who are 5-7 years into the lifecycle might be incented to move to 4G.

It is important to understand the importance of spectrum in the continued growth of mobile data services. In countries where sufficient spectrum is not allocated for 4G and the related services or the spectrum is not harmonized with rest of the world or the spectrum caps are imposed, these country will stand at a big disadvantage as the limited spectrum will but a technical barrier and the lack of harmonization a business one.[31] By 2010, the mobile broadband penetration will surpass fixed penetration globally. Countries that are behind the curve in spectrum allocation will lag behind as lack of spectrum will delay the launch of broadband services.[32]

However, one shouldn't just expect newer networks to take the load of the growing demand. One must consider a combination of strategies to lower the overall cost of managing the mobile data network traffic. Figure 9 shows the cumulative downward impact on network data traffic costs by different strategies.

[31] It will be difficult to get economies of scale that can help lower the price of equipment and services at a faster pace. Developing countries are the ones who will get most impacted by non-harmonized allocation of frequencies. Spectrum caps are also detrimental for mobile broadband growth as it limits the deployments. For example in South America, most nations have put spectrum caps (e.g. Argentina 50 MHz, Brazil 80 MHz, etc.).

[32] Full treatment of the impact of broadband on economies is beyond the scope of this paper. This subject is dealt with in detail in "Wireless Broadband: Conflict and Convergence," Vern Fotheringham and Chetan Sharma, John Wiley & IEEE Press, 2008 in chapter "Broadband and the Information Society."

Faster LTE deployment

The main driver for LTE is mobile data. As of mid-2009, several carriers have announced their LTE deployment plans. In the US, the two large operators Verizon and AT&T have indicated rollouts by the end of 2010 with more aggressive deployments in 2011 and beyond. If past experience is any indicator, it always takes more time to deploy and perfect new technologies esp. by the first-movers. LTE is expected to lower the cost of per MB delivered by 50-60% (compared to HSPA/EV-DO).[33]

However, these cost savings can only be delivered if there is sufficient coverage so that the data traffic can be off-loaded to the LTE network. Unless there is 80-90% coverage of the major metropolitan areas by 2012, we won't see LTE making a big dent in the network costs as the traffic will continue to accelerate through 2013-14. In fact, in our conservative model, LTE might only provide 1-3% cost savings.[34] If there is full deployment (80-90% POP coverage) and if a good majority of the data card users can be off-loaded to the LTE network, cost savings can be up to 25-30% by 2014.[35] This requires faster time-table of LTE deployment, device rollouts, and move towards an all-IP infrastructure.

Femto-Cell deployment

Significant majority of the data is consumed within the confines of a building, and most of the times, it is within homes or workplaces. This provides an excellent opportunity to off-load data traffic to the Wireline broadband connection and is especially important for the congested residential areas and for the heavy users. The consumption patterns in such areas and for such users need to be well understood to devise effective offloading strategies.

By encouraging users to deploy femtocells – both Cellular and Wi-Fi (along with Wi-Fi hotspots and corporate[36] Wi-Fi access points), carriers will benefit in multiple ways, namely, getting into the home media and communication management business and better indoor coverage (in addition to network cost savings). If majority of the data users were to deploy femtocells over the course of next 4-5 years, we can expect lowering of the network traffic cost by at least 25-35% depending on the extent of the Femtocell and Wi-Fi deployments.[37] Carriers who are laying down their off-loading strategy now will be better prepared to handle the traffic load that will start to accelerate further around 2011-12.

[33] Source: Chetan Sharma Consulting, 2009
[34] Source: Chetan Sharma Consulting, 2009
[35] Source: Chetan Sharma Consulting, 2009
[36] There could be interesting business models that can be developed that allow for significant offloading of traffic within the corporate premises.
[37] Source: Chetan Sharma Consulting, 2009

Congestion Management

The real cost of managing a network is in the planning for "peak traffic capacity." If this capacity can be managed through policy, quality of service, and congestion management, the peak time traffic costs can be lowered by 10-20% across the network. As shown in figure 7, only a small % of the users are typically consuming majority of the data. Additionally, only a fraction of the users typically go over the monthly capped limits (for example, 5GB/month in mid 2009). If their network usage can be managed, the cost savings will be meaningful.

Camiant a company that provides congestion management tools to the operators did some simulations to estimate cost savings. Using data from a real network, the model was initially setup to deliver 1Mbps/user 90% of the time. For the 10% of the time, the throughput was lower. By clamping down on the 2% of the users who go above the capped limit and lowering their performance to 100kbps during busy hours after they go over their allocated budget, the network savings (CAPEX and OPEX) were 18% over 3 years.[38]

Also, now is right time to start educating consumers about the "Quality of Service" and associated network traffic performances. The earlier we start doing that, the better off we will be. Wireless spectrum is a finite resource and it needs to be managed as such. It is reasonable to correlate Quality of Service with tiered service plans especially for peak traffic hours.

Additionally, congestion management should also fine-tune traffic by application for example; a streaming traffic is more sensitive to delays vs. an application download or a browsing session. A financial transaction is more sensitive than a twitter update. As such, the network should have the ability to handle traffic by priority and importance of the bits as assigned by the application provider, the carrier, and the consumer (for some twitter update might be more important than a bank transfer going through).

This also means that the regulators will have play a more constructive role in assessing how network growth might occur and what are the best ways to move the industry forward while preserving consumer interests at large.

Caching

Like Wireline, users often consume similar content on wireless connection as well. Whether it is going to the same websites or the blogs or downloading the same applications or P2P sharing or streaming the same clips, the pattern remains the same. Wireless networks can benefit greatly by using caching at various levels in the network to offload the core network from repetitious transfer of same data. There are differences in caching for Wireline vs. wireless for example, wireless sessions are usually shorter

[38] Source: Camiant. For more discussion, see whitepaper "Opportunity in the Air: Congestion Management and the Mobile Broadband Revolution" at www.camiant.com.

especially on smartphones and featurephones, the content freshness is more stringent requirement in mobile, caching capability of mobile devices is not as robust as Wireline, etc. As such, the caching technology will need to take into account the specific state of affairs in a given market but overall caching will help in managing the data traffic.

Incentives to fill the troughs

As we noted earlier, congestion management is about managing peaks to keep the incremental capacity needs to the minimum. If consumer behavior can be changed through incentives and education to use non-peak hours for bandwidth hungry applications and services, then the burden during peak hours can be lowered, sometimes significantly. Several types of incentive schemes can be designed to help shape consumer consumption patterns. This puts the emphasis back on Mbps (capacity) rather than MB consumed.

Network Optimization

Compression and transcoding have been around since the mobile web was envisioned some 10-12 years back. Whether it is WAP, full browser, applications, web services, and communication sessions – all can benefit from optimization on both ends, the device and the network. With full-browsers starting to become quite popular, the data traffic per web page delivered has started to rise. Similarly, application-based traffic is increasing but there are several opportunities for optimization from compressing and/or transcoding individual objects within a frame, page, or stream to the use of device cache or network-end points or content cache servers, data streams can be optimized. Developers should be required to adhere to strict traffic requirements so as not to make their applications overly chatty (unless they absolutely must).

Adopt Broadcast Mobile Video

Some 5 years ago, mobile video broke into the wireless landscape with a lot of promise but the uptake thus far has been less than stellar. There are two main reasons for the disappointing performance a) quality and b) pricing. While we have come a long ways from the 1-2fps video delivery services in the early days, there are still quality issues with mobile video on the cellular network especially if you are trying to do live video. As proven on Wireline, video is a huge driver for the traffic demand.

Over time, it is natural for this trend to follow on to the mobile world. However, doing live video over cellular doesn't make sense for cost and performance reasons. Operators will have to evaluate mobile broadcast technologies such as MediaFLO, DVB-H, iMB (Integrated Mobile Broadcast[39], etc.). They take the load off from the otherwise

[39] iMB is a 3GPP Release 8 technique that incorporates the broadcast mobile TV into the mobile network infrastructure but uses separate TDD spectrum.

burdened cellular networks. Obviously, one must build a compelling value proposition for broadcast mobile video for widespread adoption.

Redefine Smartphones

Smartphone has been an over-used term in our industry. While one can debate the basic elements that are needed to make a phone smart (instead of dumb), from a user's point of view it is all about the user experience and functionality. Whether you deliver it on a 300MHz or a 1GHz processor is irrelevant to them (though speed definitely helps).

As such, the industry should consider segmenting and deliver more functionality oriented phones to consumers. For example, if a certain segment of the population is enamored with social networking device, why not deliver them a device that is tightly integrated with the social networking apps so instead of using a browser or even a thick-client, one can get the same or better experience on a tightly integrated (sometimes embedded) application which lowers the bandwidth consumption, especially, as the frequency of use increases. INQMobile, a Hutchinson Whampoa company is doing just that by focusing on the functionality and applications and thus in the process redefining what a "smartphone" means.

New Business Models

It should be noted that the low-cost and bandwidth services like SMS (SMS already is) might end being the biggest revenue generator. However, the network must be planned to traffic in aggregate and have enough flexibility in the business models to help discover new high-revenue generating services.

It is quite apparent that the current business models[40] and pricing schemes will be inadequate to maintain the levels of current profitability. If the revenue equation stays flat with price pressure and the cost equation is only going up (at accelerated rates), at a certain point in the time graph, the cost of delivering data services overtakes the revenue being generated from them. So, new business models are needed that take the bandwidth consumption into account to manage the traffic especially during peak times.

Operators typically have great intelligence on voice usage but for data, the infrastructure and efforts are generally not on par. There is little understanding of what consumers are doing, which applications and services they are tuned to at any given instant, forecasting

[40] Some business models are only going to be new to certain countries. For example, the flat rate pricing phenomenon is not widespread. Many European and Australian operators charge by bandwidth consumption. Similarly, some operators are more adept at using mobile advertising as a means for continuous revenue flow than others, for example, Turkcell in Turkey. Countries who have moved aggressively into the flat rate pricing business models are the ones who are likely to shift their approach to pricing and consumption. The longer the delay in shifting, more problematic it will be in the future.

traffic spikes, etc. As cost of supporting data services exceeds the cost of managing voice services, and as the revenues from data services become more prominent than those from voice services, operators will have to pay much more attention to the specifics at a very granular level and design business models and pricing plans per the trends and forecasts.

There are practical limits to how much subscribers can talk. For example, in the US which is the most talkative country on the planet in terms of MOU (Minutes of Use),[41] voice traffic grew only 4% YOY in 2008.[42] Same can't be said of data traffic however, which grew over 150% during the same time-period.[43] The voice side of the equation can be set free with unlimited plans as the incremental cost of adding capacity is fairly low compared to the data side where the incremental cost to boost capacity is relatively quite high. As the mobile networks add not only the subscribers but also the data-enabled devices (including sensors and vertically focused devices) that can be connected to these networks, data side of the equation will dwarf voice traffic very shortly. This paradigm shift towards data will alter the economics of the industry and the markets which are in tune with the shifts of time will be able to respond better to the growing consumer expectations.

It should be noted that over the last 10 years there has been a gradual move from on-deck traffic to off-deck traffic with on-deck accounting for very little traffic in most developed markets. So, operators will have to rethink business models that are just based on selling bandwidth. They need to migrate to models that are more based on value to the end customer and the ecosystem. In fact, it will be wise to figure out the business models prior to the technology investments discussed in the previous section.

[41] The US has more than twice the MOUs (at 829) of the highest ranked European country and double the MOUs of ANY country with nearest being Canada at 444. Source: Merrill Lynch, "Global Wireless Matrix 4Q 2008."
[42] CTIA Year End Survey 2008, March 2009
[43] Source: Chetan Sharma Consulting, 2009

Conclusion

It is clearly a very exciting time in the mobile industry. The growth has been spectacular; its resilience a model for other industries, and its promise is something that keeps entrepreneurs on their toes. Within the next few months, for some of the leading operators, data revenues will overtake voice revenues for the first time.[44] Gradually, rest of the world will follow suit. However, this growth comes at with the cost of managing growth of data consumption from billions of devices and trillions of sensors around us.

In the coming years, there will be two types of opportunities that will be created, one that take advantage of the data being generated in a way that enhances the user experience and provides value and the other in technologies that help manage the traffic data that will continue to grow exponentially.

To be able to stay ahead of the demand, significant planning needs to go in to deal with the bits and bytes that are on the verge of exploding. New technical and business solutions will be needed to manage the growth and profit from the services. Relying on only 4G won't be an effective strategy to manage rising data demand. By introducing new business models and technology solutions such as femtocells, congestion management, optimization, broadcast video, new types of devices and others, carriers can manage the growth in the yottabyte era without negative impact on their profits.

[44] Technically, Operators in Philippines have crossed this threshold for sometime but they are more of an outlier in the global ecosystem. The most prominent operators which are nearing this milestone are NTT DoCoMo, KDDI, and Softbank.

Acknowledgements

The author acknowledges the assistance of Dr. Hugh Bradlow, Randy Fuller, Dr. Mallik Tatipamula, Sarla Sharma, Jon Linden, Sunil Jain, Chris Pearson, Jeff Giard, Ajit Jaokar, and others for their help in reviewing the paper.

Disclaimer: Some of the companies mentioned in this paper are clients of Chetan Sharma Consulting.

About Chetan Sharma Consulting

Chetan Sharma Consulting is a consulting and advisory firm started in 2001. We are focused on assisting companies in the mobile and voice communications sector with product management, technical due diligence, scenario planning, market and competitive research, patent & IP strategy, technology and business strategy. Our clients range from small startups with disruptive ideas to multinational conglomerates looking for an edge. We assist major brands formulate winning, profitable, and sustainable pervasive computing strategies.

Please visit us at www.chetansharma.com

About the Author

Chetan Sharma is President of Chetan Sharma Consulting and is one of the leading strategists in the mobile industry. Executives from wireless companies around the world seek his accurate predictions, independent insights, and actionable recommendations. He has served as an advisor to senior executive management of several Fortune 100 companies in the wireless space and is probably the only industry strategist who has advised each of the top 6 global mobile data carriers. Chetan has helped several global and local players in the ecosystem develop their mobile advertising strategies. Some of his clients include NTT DoCoMo, Disney, KTF, China Mobile, Toyota, Comcast, Motorola, FedEx, Sony, Samsung, Alcatel Lucent, KDDI, Virgin Mobile, Sprint Nextel, AT&T Wireless, Reuters, Qualcomm, Comverse, Motricity, Reliance Infocomm, SAP, Merrill Lynch, American Express, and Hewlett-Packard.

Chetan is the author or co-author of five best-selling books on wireless including *Mobile Advertising: Supercharge your brand in the exploding wireless market* and *Wireless Broadband: Conflict and Convergence.* His books have been adopted in several corporate training programs and university courses at NYU, Stanford, and Tokyo University. His research work is widely quoted in the industry. Chetan is interviewed frequently by leading international media publications such as *Time* magazine, *New York Times, Wall Street Journal, Business Week, Japan Media Review, Mobile Communications International,* and *GigaOM,* and has appeared on NPR, WBBN, and CNBC as a wireless data technology expert.

Chetan is an advisor to CEOs and CTOs of some of the leading wireless technology companies on product strategy and Intellectual Property (IP) development, and serves on the advisory board of several companies. He is also one of the most sought after IP strategist and expert witness in the wireless industry and has testified in some of the most important cases in the industry. Chetan is a senior member of IEEE, IEEE Communications Society, and IEEE Computers Society. Chetan has Master of Science degree in Electrical Engineering from Kansas State University and Bachelor of Science degree from the Indian Institute of Technology, Roorkee.

-----Original Message-----
From: Christopher Padilla
Sent: Wednesday, December 01, 2010 1:03 PM
To: Shomik Dutta; Edward Lazarus
Subject: IBM Statement

fyi... being shared w/reporters today.

"As the world becomes more instrumented and interconnected, vastly increased quantities of digital information are being transmitted every day over the nation's networks. Having fast and reliable access to that data enables smarter systems in areas such as energy, transportation, healthcare, retailing, and public safety. Today's announcement by the FCC represents a prudent and balanced approach to managing the growing volume of traffic on the Internet. IBM will continue to work with others in the IT industry as this process moves forward to support fair and reasonable rules that will encourage continued investment in smarter systems in the United States." -- Christopher Padilla, Vice President, Government Programs, IBM

Christopher Padilla
Vice President, Governmental Programs

Voxel Blogs

Peering Disputes: Comcast, Level 3, and You

December 2nd, 2010
Posted by Adam Rothschild

Coverage So Far

A number of folk have asked me why Level 3 and Comcast have been generating so much media attention over the past few days, and ultimately what this means for their hosting operations.

So, what's going on? Unfortunately, in a sea of rubbish, with authors ranging from Wall Street analysts to K Street lobbyists, it's difficult to find an accurate and objective accounting of what actually took place. For starters, I'd refer straight to the sources, which touch on why Level 3 and Comcast are unhappy with one another:

Level 3 Press Release: http://www.level3.com/index.cfm?pageID=491&PR=962

Level 3 Clarification: http://www.level3.com/index.cfm?pageID=491&PR=963

Comcast's rebuttals: http://blog.comcast.com/

Another good read is Dan Golding's GigaOM post (http://gigaom.com/2010/12/01/comcast-level-3-battle/). (While we disagree on several key points, it is nonetheless refreshing to read an analysis written by an author who's served his time in the trenches, and is a subject-matter expert on both the economics and technology in play.)

Game-Changers

In the absence of any real facts, one thing which *is* clear is that *both* Level 3 and Comcast grew their businesses in bold new directions, taking shortcuts and ignoring best-practices along the way.

In one corner, we have Comcast, freed from the shackles of AT&T as its sole provider, now aggressively attempting to establish itself as not merely an access provider, but a wholesale ISP which content hosters and smaller backbones might buy from. In doing so, Comcast took several key missteps, including treating "peering" as a profit center from practically day one. Where other access providers would be content with merely the "settlement free" exchange of traffic – an arrangement where both content originators and recipients exchange traffic at no cost, both avoiding having to pay a middleman to carry their bits – Comcast has made it clear it wants to collect money anywhere and everywhere possible. It would seem Comcast has again upped the ante, this time attempting the ultimate *chutzpah* of charging back its *vendors* for the "privilege" of servicing them. (Indeed, I'm a bit jealous I can't do this right now, though I'll certainly try with our next round of renewals.)

In the other corner, we have the financially-challenged Level 3, who's re-invented itself once again (as seems to happen once a year), shifting a lot of sales focus from wholesale IP transit and infrastructure to CDN. Though they arrived late to the game, they're coming on strong, targeting major content generators with cut-rate pricing. This is significant as Level 3 holds many large cable and DSL providers as its *customers*, and is effectively billing them for this new broadband-subscriber-bound traffic, as normal provider-customer relationships dictate. I can certainly sympathize with access providers trying to capacity plan around this influx of traffic. On a purely

Search

Subscribe RSS feed

Recent blog posts

Favorite Author?


Zachary Smith

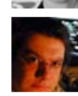
Adam Rothschild

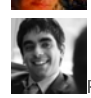
Raj Dutt


ngoulding


James Brinkerhoff

api Asia cloud computing events

contractual basis, I'm sure what Level 3 is doing is on the up-and-up. No evidence points to traffic being "stolen"; likewise "peering contracts" were not "broken", as they didn't exist in the first place. Nonetheless, they certainly co[illegible]ided the operational community some additional forewarning, or perhaps built their CDN as a ne[illegible]mous to the Level 3 backbone, maintaining an open peering policy and seeking out access ne[illegible]ange traffic with at no cost.

It's [illegible] realize that, with the exception of Comcast, every major US cable operator maintains *settlement-free* peering relations with regional service providers, CDNs, and large websites. (This is not to say they'll peer with any network off the street, however peering policies exist to set a fair bar – for instance to make sure that the party they're dealing with is professionally staffed and can route to them intelligently – not to discriminate against certain business models outright.)

CONTACT
Headquarters
Voxel Dot Net, Inc.
29 Broadway, 30th Floor
New York, NY 10006
Phone: 212.812.4190
Fax: 212.812.4195
Asia Headquarters
Voxel dot Net Pte Ltd
138 Cecil Street
Cecil Court #12-02
Singapore 069538

It's all about the "Benjamins"

What's ultimately under fire are simple economics, and how a jilted Comcast is trying to recover lost revenue. And who can blame them? A backbone provider sent me a diagram used to educate its costumers on the real issue:

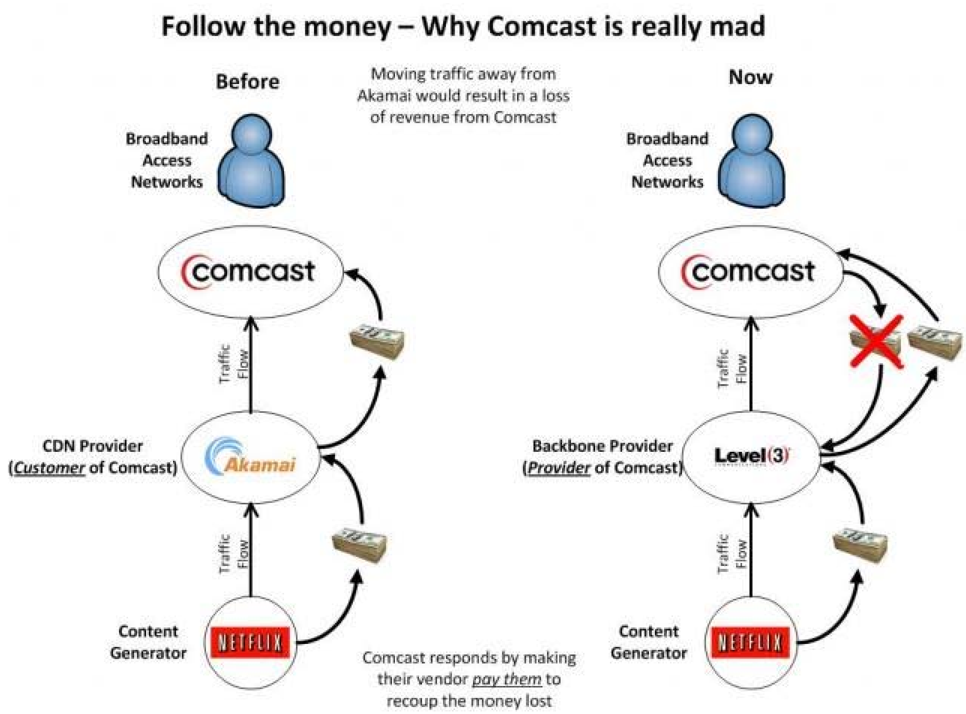

Fairness

In their filings, Comcast states repeatedly that major CDNs, specifically Limelight and Akamai, are paying Comcast for access to its customers. What they fail to state is that both companies lead the industry with their fair peering policies and massive exchange presences:

Akamai peering: http://www.akamai.com/peering/

LLNW peering: http://login.llnw.net/noauth/peering.cgi

Why, then, are they ponying up cash? Given a choice, I have little doubt both organizations would have entered into a settlement-free relationship with Comcast, as precedent dictates. I'm sure they don't consider it right that they pay Comcast to deliver traffic – rather, they agreed to it under protest, as it's the only viable way to serve their content at scale. Captive eyeballs and extortion or not (more on this point later!), CDN is a competitive market: Akamai and Limelight are paid by their customers to deliver traffic; if these bits get discarded, customers will move their business elsewhere. If they had the luxury of time, these organizations might have gone to the regulators; unfortunately content has a habit of moving rapidly between providers and contracts.

As additional facts are revealed, I'd be very curious to see the term sheets and sales orders companies like Akamai and Limelight signed with Comcast.

It's not about the ratios!

In an attempt to explain the issues, Comcast released a video of networking head honcho John Schanz discussing traffic ratios, disingenuously:

http://blog.comcast.com/2010/11/how-internet-peering-works.html

Traffic ratios date back to the days of "tier 1" telecom behemoths as major traffic sources, route miles, and hauling bits around the country to get between points "A" and "B". On a modern-day Internet , they are often cited as an excuse for denying peering where it might actually make business or technical sense. This is especially poignant in the frame of content-access negotiations – access providers are collocated at major carrier hotel locations; whether they absorb content from a free peer or paid-for transit connection, the routes and costs for hauling this traffic back to their broadband subscribers *remain the same*.

The Tata Problem

Amidst all the talk of foul play with Level 3, little attention has surrounded Comcast's relationship with Tata Communications, which I consider to be a far more egregious violation of their stated principles on Net Neutrality. As was the case with Level 3, Comcast purchases commodity IP transit service from Tata, as a means of reaching networks it doesn't maintain direct peering relationships with. Unlike Level 3 though, Comcast runs its ports to Tata at capacity, *deliberately*, as a means of degrading connectivity to networks which won't peer with them or pay them money.

Speaking off the record and respecting customer confidentiality, a Tata executive confirms, succinctly:

> *"[our] San Jose and New York links with Comcast are running full."*

One might explain the situation with Tata to their customers using a diagram like this:

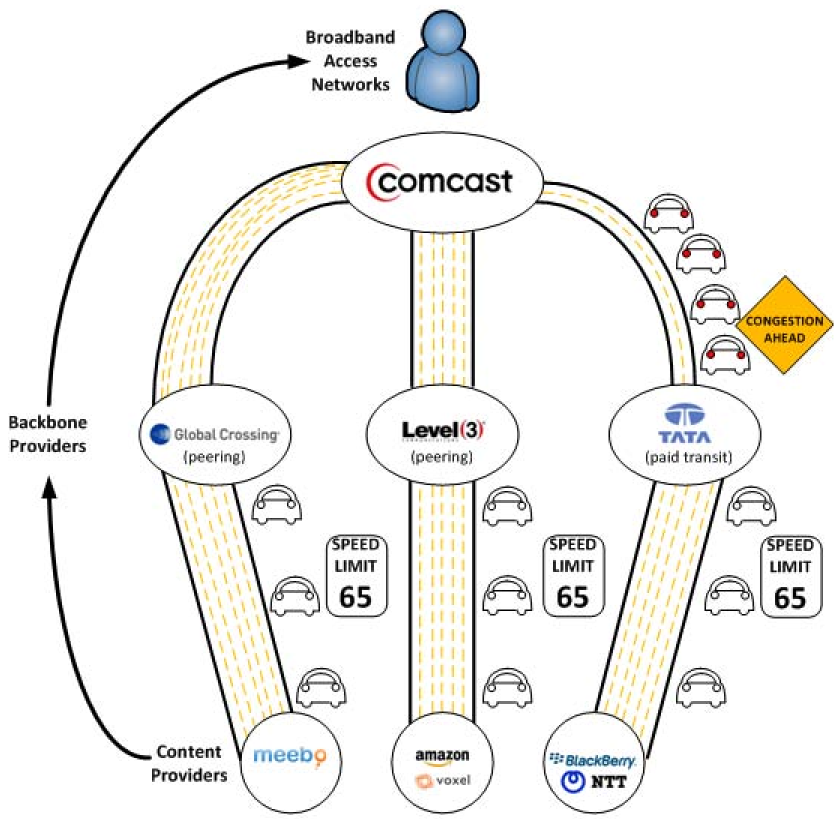

Indeed, testing reachability to Comcast eyeballs over Tata, we see loss statistics like this:

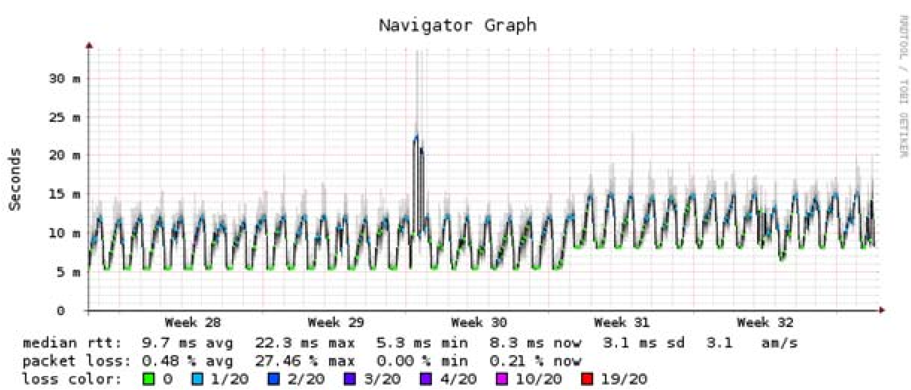

Given the high packet loss at certain times (likely corresponding to when these interfaces are saturated), serving even simple HTTP content is clearly not sustainable. In contrast, entering the Comcast network by way of Level 3, things look a lot better:

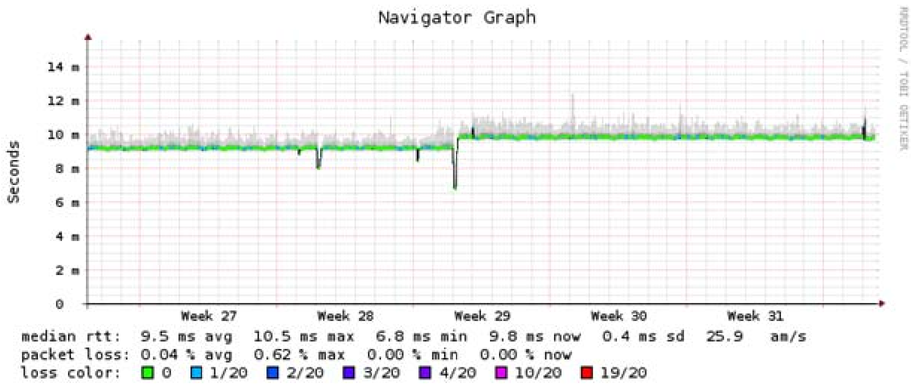

Perhaps Mr. Schanz could explain on the whiteboard why one need not worry, however a better course of action would be to simply purchase more capacity, through Tata and/or another provider. No stranger to Internet peering, Richard Steenbergen (CTO, nLayer) explains further:

> *"The true power of Comcast isn't in the size or scope of its network, it's in the captivity of its customer base.*
>
> *If Level 3 turned off Comcast for refusing to pay their contractually obligated transit bills, the traffic would be forced through massively congested Tata transit ports, and a huge number of Level 3's customers would take their business elsewhere as a result. If Comcast intentionally congests its transit providers and provides terrible service to its end users, which it has been doing for several months now, most of those users have no real alternatives to switch to.*
>
> *In other words, content is mobile, eyeballs are not. Comcast realizes that they can (ab)use their captive eyeballs to force content to pay them for access, without having to create a "100% down" partition like Cogent has done in the past. For a Netflix customer, 20% packet loss is effectively just as down as a hard partition.*
>
> *[…] while it should be every network's right to choose who they do and don't peer with, or buy transit from, things start to get murky when one network is abusing their franchise agreements and near monopoly or duopoly status in many markets. If users had an actual choice, and could get comparable broadband access elsewhere, then Comcast would be free to congest their network however they see fit. But that isn't the case, and this is where government involvement and Net Neutrality start to have legitimate grievances with Comcast's actions."*

Regulation and Disclosure

So, how should the FCC and other facets of the Federal Government intervene?

Perhaps at odds with my above criticisms, I believe in the power of the free market, and that networks should be left to negotiate (or not negotiate!) these issues without oversight. Telling network operators where they must or must not connect, and on what terms, is surely a recipe for disaster.

On the other hand, protecting the rights of broadband subscribers is of the utmost importance, given the scarcity of real competition on the last mile. If many Americans don't like the rates their cable providers charge, or the quality of service delivered, they're left to pound sand or downgrade to slow DSL. Just as the FCC is looking to codify "network management" practices on the last mile, I think it's fair to penalize monopoly/duopoly providers who fail to adequately manage their backbone and external capacity.

More immediately, I'm hoping for full disclosure of any commercial proposals and agreements between Level 3 and Comcast. Absent this data in its rawest form, it's impossible to form any intelligent opinions on the specific issues in play. For all we know, Comcast's commercial proposal might be fully reasonable, and limited to some basic cost-recovery of capex costs in router ports or Level3 transit install fees it must now incur on short notice. We just don't know!

Voxel Customer Impact

What does this all mean to your day-to-day hosting operations, as a customer of Voxel? Absolutely nothing, hopefully! The above commentary should serve only to shed light on a hot topic, not to present any grounds for immediate concern.

Voxel purchases a number of Ten Gigabit transit ports from Level 3, which provide a clear path for sending traffic to Comcast and other networks. When these circuits hit high "water marks" for utilization, we order more. Likewise, we carefully monitor for the quality of connectivity between our network and large broadband destinations; the above screenshots are only the tip of the iceberg for analytics. When we encounter a bad route, we engineer around it right away.

Holding our customers hostage, or deliberately saturating connections to prove a point, is at odds with how we do business, and is simply not in the cards.

(Disclaimer: Adam is Voxel's VP of Network Architecture; among his responsibilities are peering and transit strategy. No yellow journalists or Washington lobbyists were harmed in the making of this post.)

Comments

Mark Kamichoff

THU, 12/02/2010 - 19:48

Nice writeup.

The Tata problem reminds me of, among other things, the whole Cogent vs. ATDN fiasco a few years ago. Cogent, buying transit to ATDN through Verio, deliberately ran their links at capacity, to force ATDN into a private peering with them. The goal? I think it was just for marketing, so they could say they were tier 1. It sure gummed up things for us, who only had ATDN at the time. I think ATDN eventually caved, though.

I'd love to see an Internet without any involvement from the FCC, but if things like this continue to crop up, it might be a necessary evil.

- Mark

The Dude Dean

THU, 12/02/2010 - 21:17

This post is Comcastic! http://www.facebook.com/Comcastic

Rahul Tongia

FRI, 12/03/2010 - 02:24

Is the red "x" in the first diagram on the wrong arrow?

Rahul

Adam Rothschild

FRI, 12/03/2010 - 18:19

Rahul,

Wrong how? The 'X' shows the old flow of money from Comcast to Level 3; the other line shows how money is passing hands under the new arrangement, which Level 3 is upset about.

Gabriel

FRI, 12/03/2010 - 20:39

It is unclear to me if Comcast was initially a customer (who purchased transit) of Level 3 (as described in the article) or their connection was settlement-free (as suggested in the second diagram). If they were a customer and did buy transit, why would they not saturate the links with Level 3 as well in order to force the networks that refuse to peer with them to do so?

Voxel Guest

FRI, 12/03/2010 - 22:38

Doesn't it seem that LVLT is just blustering to improve its negotiating position prior to a private resolution with Comcast (perhaps with the FCC facilitating)?

Comcast seemed initially to have a superior negotiating position. However, when LVLT fans the flames of public sentiment "net-neutrality" its position becomes vastly improved. It seems unrealistic that the FCC will implement N-N following the recent google/verizon deal and the Sprnig 2010 federal court ruling noting the FCC's very limited authority in this area without further legislation.

From a purely negotiating standpoint, it seems wise for LVLT to push for N-N. Also, that would explain some of the disingenuous positions/claims by Comcast and LVLT.

Raul Martynek

SAT, 12/04/2010 - 09:25

Nice article Adam. My suspicion is that in the big scheme of things the money they are getting from Level 3 or lost from Akamai pales in comparison to what they are really trying to prevent: over the top video. The tens of millions of video RGUs (and the advertising they generate) are the real golden goose that they are trying to protect.

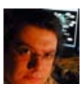

Adam Rothschild

SUN, 12/05/2010 - 13:54

Gabriel,

I've been trying to keep things terse and easy to parse, though since you asked... :-)

Historically, Comcast's relationship with Level 3 has been that of a paying transit customer. From the press releases both companies put out, it would seem they're looking to become a "paid peering" vendor ***to*** Level 3, which is what I illustrated in the second diagram.

Looking at technical details such as route propagation and BGP community tagging, it would appear as though Level 3 is currently providing Comcast with full transit. They provide Comcast with access to the Internet at large, and likewise, they re-advertise the routes they're learning from Comcast to their full base of customers and peers (with several exceptions; Comcast is utilizing outgoing BGP communities to request that Level 3 not advertise their routes to certain networks, for example TeliaSonera, hence forcing traffic down the congested Tata paths -- likely as a negotiating play).

I'm a bit curious myself as to how money passes hands at this very moment. It's possible the paid peering agreement hasn't yet kicked in. It's also possible it has, and that traffic accounting (ie sFlow/netflow, Juniper DCU, ...) is being used to allow Comcast to "charge back" just traffic from Level 3 customers/CDN destined towards them. Unfortunately one can't always surmise "layer 8+" details such as this from mere routing data! This is, in part, why I'm hoping for contractual disclosure by both parties.

Why their Level 3 interfaces show no obvious signs of saturation is an excellent question, and one I've yet to hear explained well.

JohnF

SUN, 12/05/2010 - 12:05

I would guess that there are probably two reasons; the new traffic will only switch over (fully?) in January and Comcast may need level3 a little more than they pretend they do. They are a major backbone provider. The control of eyeballs that Comcast has and the power it gives them is quite interesting and distressing.

[Editor: On the latter point, we agree.]

You said "it would appear as though Level 3 is currently providing Level 3 with full transit" in your comment above, I think that second L3 is supposed to be Comcast, unless you are referring to the CDN?

[Editor: Indeed, and thank you for the correction! Fixed.]

Charles Gucker

SUN, 12/05/2010 - 21:48

Adam,

It's pretty simple to see how the money used to (and partially still changes hands). Comcast purchased a 20 year fiber IRU's for a majority of their transport (and IP) backbone. This was not a "cheap" endeavor, costing Comcast (and putting against Level(3)'s debt) well north of 10's of millions of dollars. Also, Comcast buys a lot of voice services from Level(3), much like many MSO's (also worth millions, if not upwards of 10MM per year).

So, taking into account the dark fiber IRU's and voice services, it's not a far stretch that Level(3) wrote into their contract (like they have offered me in the past), zero dollar "ip transit" with 65000:0 (do not advertise to peers) appended to their routes. So, even tho the account is not handled by the peering group, I'm sure they weren't paying for the IP portion of their services.

[Editor: I've actually thought a lot about this possibility, though ultimately left it out of my original posting, absent hard facts and in the interest of brevity. I think it's possible Comcast's spend on Level 3 transport/voice/collocation-related services exceeds any IP-related spend (or credits?) by an order of magnitude, in a way making this whole struggle a non-issue. You'll remember a similar situation with AOL and Level 3 in the late 90s/early 00s...]

Now, since Comcast has invested 100's of millions of dollars into their infrastructure (after absorbing the old @home properties by way of AT&T, along with part of the Adelphia properties, MediaOne, et al), they expect to be paid to transit their network if your network does not meet their requirements for peering. In this case, due mainly as a result of the CDN business means Level(3*).*

[Editor: I disagree. Ratios are a relic, and are used by Comcast time after time to deny peering relationships which would result in actual benefits for both parties.]

Personally, I believe both sides need to save face in this dispute, so I believe that Level(3) will provide Comcast with concessions on their fiber IRU's and/or future voice services in lieu of a [Level(3)] paid peering arrangement [from Comcast]. This way Comcast wins by "saving" money, Level(3) wins by not having to "break" their transit-free status and retains their "customer". As both you and Dan pointed out, this is a "simple" business dispute with a finite amount of possible outcomes.

charles

P.S. As was pointed out, in your first diagram, the red X should be one link to the right. Comcast does not receive monies from L3, but L3 receives monies from Comcast.

*[Editor: Following the recent PRs, L3 ****will***** recieve money from Comcast, for on-net traffic.]*

Also, in your last response your statement "it would appear as though Level 3 is currently providing Level 3 with full transit." really should read "it would appear as though Level 3 is currently providing Comcast with full transit capabilities."

[Editor: Fixed, thanks!]

Lastly, Yahoo does not utilize transit from Tata to reach Comcast, they are behind Global Crossing. So, another content provider would need to be selected ;-)

[Editor: Comcast's best-path to Yahoo is through Tata. It was recently pointed out that Yahoo's return route is actually through Global Crossing, due to the Tata congestion issues, thus making Yahoo a bad example. Updates are forthcoming.]

1 2 next › last »

ADD NEW COMMENT

Deploying Premium Services Using Cisco Service Control Technology

Service providers are seeking to use an all-encompassing IP network for premium services, resulting in the delivery of higher-margin, higher-value service offerings. Cisco® Service Control technology enhances existing network infrastructure with subscriber and application awareness. Cisco Service Control is comprised of a programmable network element that creates a service layer for broadband networks, allowing operators to identify subscribers, classify applications, guarantee service-level performance, and charge for multiple IP services on a single IP network without costly equipment upgrades.

Challenge

The ability to offer voice, video, and data services on an all-encompassing IP network is the ultimate goal for many service providers. Yet "best effort" networks are not good enough to support these demanding service offerings. The first network operators to maximize network efficiencies for IP service delivery using granular analysis and reporting and real-time traffic policies combined with application-level quality of service (QoS) and subscriber-based metering will be well-positioned against the competition. Accompanying such objectives, however, is the need to make only modest incremental infrastructure investments and to control operating costs to secure healthy profits. Both cable and DSL operators are establishing "fat pipes" for delivering advanced IP services into users' homes. However, the existing transport networks are constrained by an inability to easily and cost-effectively identify and meter individual subscriber usage by application. The all-encompassing IP network must be capable of carrying virtually any "triple play" service—voice, video, and data—yet infrastructure continues to fall short of these goals and prevents providers from profitably maximizing network investment, limiting their ability to create new business models or easily customize services to individual subscriber preferences.

Meanwhile, the number of Internet subscribers continues to grow along with the number of intelligent portable devices. The market for both broadband and mobile operators is positioned to accept new premium service offerings such as voice over IP (VoIP), online gaming, music downloads, video on demand (VoD), and streaming television. Such services offer the potential to dramatically increase average revenue per user (ARPU) for service providers, which further increases the overall value of their network assets.

The lack of intelligence in the network's data plane causes significant issues for operators in their attempt to deploy premium services. The network must be intelligent enough to understand "who" is doing "what" during any subscriber session. Subscriber and application awareness must be efficiently integrated into the network to meet the essential needs of profitable premium service delivery.

Subscriber and Application Awareness Facilitates New Service Exchanges

Subscriber acceptance is predicated on perceived value. New business models must be tested as providers seek to develop the right formulas to value various offerings. As more and more digitized content becomes available, a plethora of new business opportunities emerges for content-based services. However, to bring a diversity of content to the subscriber, network providers and third-party content suppliers will need to cooperate with each other to use their value proposition. However, for service providers to pursue partnerships with content suppliers in exchange for access to their pipelines or revenue shares, the network requires service intelligence. Content usage must be tracked by subscriber, protected from piracy, metered, and appropriately valued. Operators have an opportunity to take advantage of QoS and other network capabilities for a share of third-party content revenue. The concept of a “service exchange” can facilitate new IP deliverables based on partnerships between operators and content providers.

Services such as VoIP or VoD have more rigorous performance requirements that allow for zero latency. Those provider networks that are capable of isolating traffic flows and applying application-level QoS to VoIP traffic or VoD will be more attractive to users. A network is needed that can classify applications easily or identify third-party providers of VoIP. Moreover, by identifying services that might be riding an operator’s network for free, a provider can truly differentiate its own “branded” VoIP service traffic from best-effort traffic or extend QoS guarantees to that third party for a share of the profits. Such arrangements further demonstrate the opportunity for network “service exchange” agreements.

The Importance of Service Control for Emerging Multimedia Services on Mobile Networks

In the mobile sector, with the introduction of high-speed access and video-capable handsets, multimedia services are rapidly gaining popularity among mobile subscribers. The ability to make video and audio recordings virtually anywhere subscribers go and instantly share them with friends and family is extremely powerful and will drastically influence the way subscribers communicate with each other. With propagation of third-generation (3G) and fourth-generation (4G) networks, mobile providers will face the same usage analysis, traffic optimization, and security issues that confront broadband providers as more advanced applications migrate to mobile from wire-line networks.

Multimedia Message Service (MMS) messaging is the first application to enable a near real-time sharing experience of pictures and short video or audio clips of “captured moments.” However, as mobile subscribers master this new technology, the demands of multimedia sharing on the network become enormous and will be followed by powerful real-time multimedia applications such as VoD, video broadcasts, and video phone conversations or even peer-to-peer (P2P) communications. These applications use the powerful Real Time Streaming Protocol (RTSP) to enable multimedia-streaming applications permitting multiple audio and video channels to be simultaneously streamed while delivering a robust multimedia experience to the end user.

The opportunities for advanced IP service delivery on broadband and mobile networks are virtually endless; however, there are gaps that remain in the network infrastructure that must be filled before vision can truly catch up with reality. Cisco Service Control is specifically crafted to fill these infrastructure gaps.

Solution

The Cisco Service Control Platform is comprised of a programmable network element that creates an intelligent overlay, enabling network operators to identify subscribers, classify application-level traffic, guarantee service performance, and charge for content-based services. The solution allows providers to address the gaps in premium service deployment and to customize solutions for individual subscribers while effectively charging for new service offers.

Cisco Service Control technology is comprised of both hardware and software integrated into a state-of-the-art, dedicated network device, providing detection and control capabilities. Typically, the Cisco Service Control Engine resides “in traffic” behind an IP aggregation point and can be configured redundantly to meet high-availability requirements. Using the platform’s Layer 7 stateful deep packet-inspection capability, the solution can accurately identify application use by individual subscriber. The Cisco Service Control solution has a set of characteristics and architectural attributes built to perform real-time traffic classification, accounting, and control. To undertake stateful deep packet inspection at multigigabit speeds, a specific hardware architecture is required that is capable of maintaining the state of each network conversation, while implementing deep and detailed inspection of every data packet through the application or Layer 7 network layer. The result is a solution that can detect specific protocol signatures and classify all traffic for a given network session.

Using Cisco Service Control, operators can:

- Granularly analyze how subscribers are using network resources and more effectively assess trends or evaluate the acceptance of new services or business models (refer to Enhancing Usage Analysis Using Cisco Service Control at: http://www.cisco.com/go/servicecontrol).
- Prioritize and guarantee performance for premium services such as VoIP, interactive gaming, or VoD.
- Transparently identify application flows requiring preferential treatment and signal other network elements to set up QoS for packet transport.
- Track service usage to create revenue-sharing opportunities between providers and content partners.
- Mitigate against security threats to the provider network and redirect infected subscribers to technical centers for resolution (refer to Providing Service Security with Cisco Service Control at: http://www.cisco.com/go/servicecontrol).
- Monitor quality and service delivery of premium services such as track voice or streaming traffic’s jitter and packet-loss parameters.
- Charge for individual services or suites of services all individually metered and running on common transport network.
- Monitor and control VoIP traffic traversing an IP network, whether originating from the operator’s own voice service or that of a nonfacility-based operator. This provides the means to not only track service quality, but also help ensure that adequate resources are available for these services.
- Monitor usage and quality of nonfacility services used by the subscriber base from off-net destinations.

Application-Level QoS: Performance Guarantees

The application-aware and subscriber-aware service overlay created by Cisco Service Control technology provides multigigabit analysis and is powered by Cisco Service Control engines, while the Cisco Service Control Application for Broadband helps ensure that providers can now move beyond best-effort services. Service providers can offer guaranteed performance and QoS for sensitive IP applications such as VoIP and VoD, while taking advantage of existing investments in the common transport network (refer to Using Cisco Service Control for Traffic Optimization at http://www.cisco.com/go/servicecontrol).

This capability helps ensure the appropriate priority is granted to application traffic throughout all network segments, from the first and last mile of the network to the network edge and core. Cisco Service Control adequately monitors network-resource availability so that appropriate actions are taken in the event of oversubscribed network resources. Providers can:

- Detect application usage directly from the data stream, control bandwidth, and mark traffic for preferred treatment as it flows through the network.
- Set up QoS for specific application streams without requiring network integration with application servers. This greatly reduces implementation and operational costs.
- Perform stateful traffic classification to treat multiflow, multimedia application traffic in a single context, helping ensure suitable QoS for the entire application session.
- Go beyond classification and truly understand application semantics to make suitable network resources available for the needs of each application session and take appropriate action if not.

The stateful deep packet-inspection capabilities of Cisco Service Control greatly simplify signaling requirements, and Cisco Service Control engines can simultaneously track hundreds of thousands of flows and maintain an application-layer state for every one.

Accurate Analysis and Control of New Multimedia Traffic

The Cisco Service Control solution is capable of performing application-layer (Layer 7) stateful deep packet inspection of RTSP traffic going over its control channel and associating it with the traffic for all of its audio and video channels. This capability helps enable Cisco service applications to accurately analyze and control multimedia traffic regardless of its network origin and correctly charge for traffic. Mobile service providers who have already deployed the Cisco Service Control solution can deliver and charge for multimedia services without compromising profits.

With Cisco Service Control technology, premium IP services can be managed, controlled, and delivered on a converged IP network capable of meeting the needs of individual service providers to astutely analyze usage by subscriber, classify and guarantee application-level performance, and meter and charge for any number of IP applications. The resulting profitability closes the gap between vision and reality.

Business Benefits

Cisco Service Control technology allows operators to:

- Increase ARPU by customizing services to meet individual subscriber needs
- Effectively converge network deployments and operations

- Make a small incremental investment for existing networks and use overall network investment for new premium service delivery
- Reduce operating costs and capital investment by using a single network for multiple IP services
- Test and trial new offerings without additional investment, to better amortize infrastructure costs across a multitude of service offerings
- Meet the demands of latency- and jitter-sensitive applications such as VoIP and VoD
- Increase overall customer satisfaction and reduce turnover by offering truly differentiated or customized service offerings

Architecture

Figure 1 indicates where Cisco Service Control resides in the network and how it effectively uses existing infrastructure investment.

Figure 1. Cisco Service Control in the Network

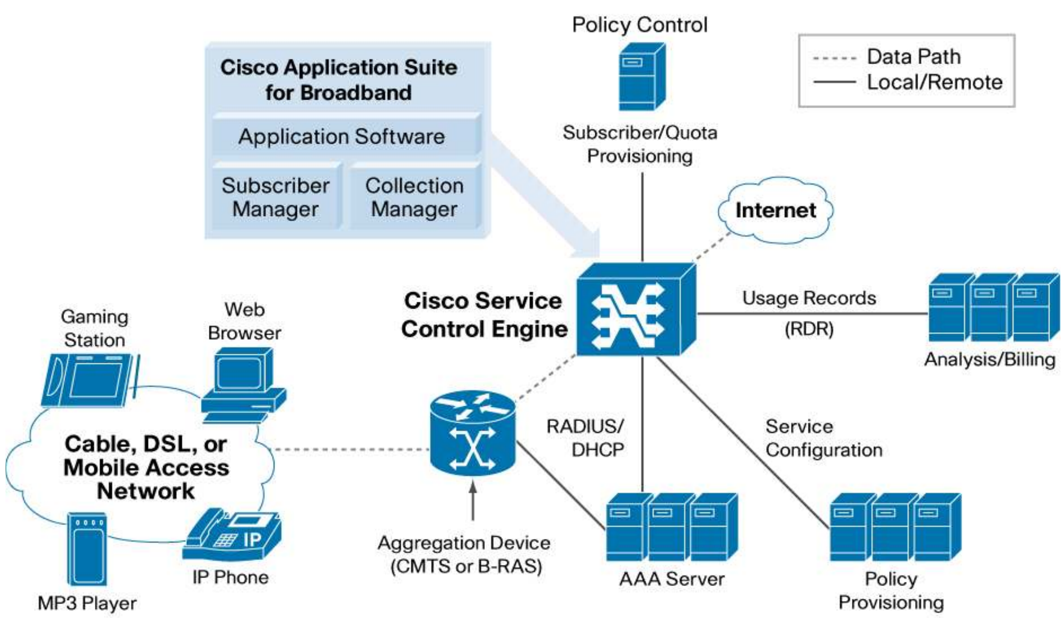

Product Offering

- Cisco SCE 1000 Series Service Control Engine
- Cisco SCE 2000 Series Service Control Engine
- Cisco SCE 8000 Series Service Control Engine
- Cisco Service Control Application Suite for Broadband
- Cisco Service Control Collection Manager
- Cisco Service Control Subscriber Manager

Why Cisco

Cisco offers industry-leading service control solutions, delivering multigigabit performance and stateful deep packet inspection as well as worldwide technical assistance and support. Cisco is speeding the evolution of networks from generic transport to platforms offering higher-value, higher-margin services. Programmable, scalable, and designed for the communications sector, Cisco Service Control technology accelerates network delivery of advanced IP services. The Cisco Service Control platform adds intelligence, stateful deep packet inspection, and multigigabit analysis to existing network infrastructure, so carriers can identify and charge for dissimilar content applications while simultaneously managing different applications' performance requirements. The Cisco Service Control solution is deployed in more than 450 companies worldwide.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C22-479388-00 07/08

Search

Search Keyword:
Tag: PLR

Total 7 results found.
Ordering Newest First

All Sections
content: Products (7)

PLR

The PacketLogic Real-Time Enforcement platforms (PLR) utilize multiple hardware platforms that run the same operating software. Each of the platforms offer consistent feature richness enabled through the different PacketLogic software modules: LiveView, Filtering, Traffic Shaping, and Statistics. The PacketLogic hardware platforms offer a range of configurations from the entry-level 4 Mbps PL5600 through through 2 Gbps PL7720. The mid-range PL8720 is a 2RU unit with up to 10Gbps throughput. At the top of the line is the high-end PL10000 series with capacity up to 80Gbps and 5M subscribers per system. It consists of a modular AdvancedTCA (ATCA) chassis solution in two sizes – PL10005 5RU and PL10014 12RU.

Hardware Platforms: PL10000 Datasheet PL10000 Datasheet (A4) PL8720 Datasheet PL8720 Datasheet (A4) PL7720 Datasheet PL7720 Datasheet (A4) PL5600 Datasheet PL5600 Datasheet (A4) PL1200 Datasheet PL1200 Datasheet (A4) PL1420 Datasheet PL1420 Datasheet (A4) Procera Bypass Switch Datasheet Procera Bypass Switch Datasheet (A4)

Tags: PLR | PL8720 | Products | PL10000 | PL7720 | pl5600 | Procera Bypass Switch

PL8720

PL8720 extends Procera Network's midrange PacketLogic™ product solutions. It is a conveniently packaged 2RU rack-mounted appliance that that supports configurations of up to eight (8) channels of Gigabit Ethernet (GE) and four (4) channels of 10 Gigabit Ethernet. This makes the PL8720 suitable for access/edge deployments in broadband networks, or the WAN connection of educational campuses and enterprises. The network interfaces, hard drives (HDD) and power supplies are easily accessible, which enables field replacement in case of failure. It also enables a mix of copper (RJ-45) and fiber (SX/LX/SR/LR) channels. All PacketLogic platforms utilize the same firmware which gives PL8720 the feature-richness PacketLogic is renowned for and that makes it the preferred DPI product for in-line deployments. It also enables PL8720 to co-exist in networks with PL5600 , PL7720 and the high-end PL10000. PL8720 runs all PacketLogic DPI software modules – LiveView, Filtering and Traffic Shaping. Statistics typically runs on a dedicated hardware, but can run internally on the PL8720 for small deployments. PL8720 Datasheet PL8720 Datasheet (A4)

Tags: PLR | PL8720 | Products

Procera Bypass Switch

Procera Bypass Switch (PBS) is an active external bypass switch that adds an extra level of resilience to PacketLogic™ deployments that require exceptional availability. This enables even the most cautious network manager to deploy PacketLogic inline. The use of configurable heartbeat pulses allows PBS not only to detect a link failure but also system failures where the heartbeat packet fails to pass. Heartbeat frequency as low as 10 ms allows instant fail-over in case of a failure. With PBS you can also manually switch to bypass mode for systems maintenance. PBS is configured and administered via a serial communication console port. An intuitive web GUI (graphical user interface) makes setup and administration quick and easy. The PBS can also be administrated via CLI or SNMP and supports SNMP traps for failure notification. PBS consists of a 1RU host system that can hold up to four PBS bypass modules. This enables one single PBS host system to manage the redundancy for one to four PacketLogic units. The PBS comes with redundant power supplies and can be ordered with either 48V DC or 90-240V AC. The stand-alone and external architecture of PBS enables it to manage installations with any PacketLogic system, i.e. from the entry-level PL5600, through PL7720, PL8720 all the way up to the PL10000 Series.

Procera Bypass Switch Datasheet Procera Bypass Switch Datasheet (A4)

Tags: PLR | Products | Procera Bypass Switch

PL7720

PL7720 is Procera Network's midrange member of the PacketLogic™ product family. It is a conveniently packaged 2RU rack-mounted appliance that enables one or two channels of gigabit Ethernet (GE). This makes PL7720 suitable for access/edge deployments in broadband networks, or the WAN connection of campuses and enterprises. The network interfaces, hard drives (HDD) and power supplies, are front-mounted modules, which enables field replacement in case of failure. It also enables a mix of copper and fiber channels. All PacketLogic platforms utilize the same firmware which gives PL7720 the feature-richness PacketLogic is renowned for and that makes it the preferred DPI product for in-line deployments. It also enables PL7720 to co-exist in networks with PL5600 and the high-end PL10000. PL7720 runs all PacketLogic DPI software modules – LiveView, Filtering and Traffic Shaping. Statistics typically runs on a dedicated hardware, but can run internally on the PL7720 for small deployments. PL7720 Datasheet PL7720 Datasheet (A4)

Tags: PLR | Products | PL7720

Software Modules

Tag Cloud

DRDL Technology PIC PL10000 PL1200 PL1420 pl5600 PL7720 PL8720 PLR Procera Bypass Switch Products PSM

The PacketLogic firmware consists of four software modules – LiveView, Traffic Shaping, Filtering and Statistics. All software modules rely on Procera's state-of-the-art identification and classification engine DRDL. The modules are fully integrated which means that everything is managed in the same user interface. It also means that common objects are used to define groups of users (local hosts), applications, properties etc. A change to an object will affect all occasions where that object is used. This gives an easy overview of the network traffic as well as the PacketLogic administration.

The PacketLogic software is managed and monitored through the graphical (GUI) admin client that is available for Windows, Linux and Mac. Other interfaces include SNMP, Syslog, CLI (command line interface) for systems configuration, and a complete Python API for automation of PacketLogic administration. PacketLogic's centralized management manages common rule-sets and firmware editions across multiple PacketLogic systems.

LiveView

The traffic is presented in real-time in the PacketLogic module LiveView. This is the first view that meets the PacketLogic administrator. LiveView shows all traffic going through the PacketLogic system, from overview down to a specific connection. The real-time capability enables you to support your users on the fly and resolve issues when they actually occur.

Traffic Shaping and Filtering

Policies are enforced in the Traffic Shaping and Filtering PacketLogic modules. Traffic Shaping can force certain traffic to a defined level. It can also be used to secure that it does not exceed a set threshold. This way capacity can be designated to critical and sensitive traffic. The Traffic Shaping module also contains prioritization where traffic based on labeling is forwarded in a certain order in case of congestion. PacketLogic uses a pure non-disruptive queuing mechanism.

The Filtering module is a capable Layer 7 firewall. Traffic can be filtered or allowed based on all granular properties provided by DRDL. This enables forensic control of unwanted and hazardous traffic. The Filtering module also provides rewrite functionality that can redirect traffic, change the QoS label DSCP (DiffServ) and perform automated DoS/DDoS protection.

Statistics

The data presented in real-time in the LiveView module is aggregated in the PacketLogic Statistics module. Statistics validates that the policies provide the intended results. This is valuable information for Marketing and Product Management to explore what the users do online, for Abuse Management to track malicious behavior, and for Company Management to get weekly reports.

Software Modules Datasheet Software Modules Datasheet (A4)

Tags: PLR | Products

PL10000

The PL10000 series is a modular AdvancedTCA (ATCA) chassis solution in two sizes – 5RU (rack units) and 12RU. All are 19" rack-mounted systems that fit the depth of a telecom rack. It delivers capacity of 40 Gbps Full Duplex, i.e. 80 Gbps throughput, and 5M subscribers per system. This makes it the unchallenged leader in DPI (Deep Packet Inspection). The high capacity makes PL10000 a future-proof investment that will grow with your customer base and changing conditions. The PL10000 series is based on the AdvancedTCA industry standard which guarantees carrier-grade performance and full resilience on all vital components like power supply and cooling. PacketLogic PL10000 is purpose-built with Tier-1 broadband network deployments (both wired and wireless) in mind. Service providers now have a platform that will support millions of subscribers while giving them the business intelligence, service creation, network visibility and control required to successfully roll out new revenue-generating services and optimize network performance. Designed as a carrier-grade platform, the PL10000 gives service providers the scalability, reliability and flexibility they need to meet network requirements for today and tomorrow. It has four times the capacity of its nearest competitor, enabling a single deployment to handle millions of broadband subscribers. This increased capacity, in combination with pertinent asymmetric traffic support and the "5 9s" reliability that carriers demand, enables deployment of PacketLogic anywhere in the network. Functionality has not been compromised in favor of capacity. DRDL™ (Datastream Recognition Definition Language) is at the core of what makes PacketLogic stand out from the competition. This state-of-the-art traffic identification and classification engine drives all PacketLogic platforms.

PL10005 (5RU) PL10005 is the 10G solution in the PL10000 series. Based on the industry standard AdvancedTCA (ATCA), it delivers 20 Gbps processing capacity, i.e. 10 Gbps Full Duplex, in a single chassis with reliable performance that guarantees high availability. Every PL10005 handles up to 3M subscribers per system. This makes PL10005 the preferred single-channel 10GE solution for broadband networks. PL10014 (12RU) PL10014 is the flagship multi-10GE member of the PL10000 series. Based on the industry standard AdvancedTCA (ATCA), it delivers up to 80 Gbps processing capacity, i.e. 40 Gbps Full Duplex, in a single chassis with reliable performance that guarantees high availability. Every PL10014 handles up to 5M subscribers per system. This capacity and performance makes PL10014 the unchallenged capacity leader in carrier-grade DPI (deep packet inspection).

PL10000 Datasheet PL10000 Datasheet (A4)

Tags: PLR | Products | PL10000

PL5600

PL5600 is Procera Networks' entry-level member of the PacketLogic product family. The intuitive licensing is available in five bandwidth configurations – up to 4, 10, 24, 45 or 100 Mbps. The PL5600 comes pre-packaged with all PacketLogic modules – LiveView, Filtering, Traffic Shaping, and Statistics.

The PL5600 is an easy to deploy 1RU 19” rack-mounted hardware platform with six gigabit Ethernet (GE) Base-T copper interfaces (RJ-45), i.e. three channels. The use of a common firmware across all PacketLogic platforms enables PL5600 to reside in environments with a mix of PacketLogic hardware platforms.

DRDL™ (Datastream Recognition Definition Language) is at the core of what makes PacketLogic stand out from the competition. This state-of-the-art traffic identification and classification engine drives all PacketLogic platforms from the PL10000 to the PL5600.

PL5600 Datasheet PL5600 Datasheet (A4)

Tags: PLR | Products | pl5600

[back to press releases]

Sandvine Internet Report: Average is Not Typical

Sandvine report highlights the diversity of subscriber experiences in fixed and mobile networks

Waterloo, Canada; October 20, 2010 – Sandvine, (TSX:SVC; AIM:SAND) a leading provider of intelligent broadband network solutions for cable, DSL, FTTx, fixed wireless and mobile operators, today announced that it has launched an Internet traffic trends report, entitled “Fall 2010 Global Internet Phenomena”, based on data from cable, DSL and mobile service provider networks, spanning four regions worldwide including, Asia-Pacific, Europe, Latin America and North America.

This is the eighth report in an ongoing series of Internet phenomena & traffic analysis Sandvine has published since 2002. Sandvine’s global view, which includes over 200 service provider customers spanning more than 80 countries, makes the report the most comprehensive of its kind in the industry.

Sandvine’s major findings reveal the subtle yet substantial differences between behavior patterns of consumers in various regions when connecting to the Internet. By contrasting Internet usage with previous Sandvine reports, analysis shows that even within regions traffic trends have changed over the past six-to-twelve months.

“The Internet has a unique way of bridging international gaps and bringing people together,” said Dave Caputo, president and CEO, Sandvine. “Yet, interestingly, as we all plug into this international network to satisfy our social, professional and entertainment needs, our access patterns and online behaviors have become as unique as we are. The Internet is one single source that satisfies 500 million people, so it is no wonder that an average user does not exist.”

Regional Internet findings:

- In the United States, Netflix represents more than 20 percent of downstream traffic during peak times and is heaviest between 8-10 p.m.
- The Asia-Pacific region ramps up their Internet usage at 5 a.m. and their median monthly data consumption is close to 12 gigabytes per household compared to 4 gigabytes in North America
- In Europe, zSHARE has become the dominant leader for storage and back-up services. It accounts for 3 percent of downstream traffic during peak periods
- Behaviorally, some subscribers in Latin America use the Internet the same regardless of a fixed or wireless connection. For example, close to 1/3 of traffic on wireless and fixed networks is real-time entertainment such as YouTube or PPStream
- Overall there is a wide variation between the amount of time Internet connections are active. For example, in North America the average time a fixed connection is active is 3 hours, whereas in Asia-Pacific it's closer to 5.5 hours

Another major driver affecting worldwide Internet behavior is the increased availability of 3G and 4G networks. Internet mobility has become as accessible as fixed line in many regions and subscribers are taking full advantage of the flexibility that converged networks offer.

“This is the first report where we compared the behaviors of fixed and mobile users,” said Mr. Caputo. “For a subscriber the Internet is the Internet, regardless of when, where or how they connect to the network and that is consistent in our findings. Usage plans and personalized services that appeal to the broadband-individual, rather than the broadband-household have become the Internet of today.”

Fall 2010 Global Internet Phenomena report - methodology

This study was based on a representative cross-section of the world’s leading fixed and mobile data providers serving spanning four regions worldwide including: Asia-Pacific, Europe, Latin America and North America. Data was gathered over a two-month period in August and September 2010 and captured the bits-per-second, per protocol and the number of active hosts per protocol on the network.

The data gathered in Sandvine’s global Internet traffic report is completely subscriber-anonymous. No identifiable information of any kind, including IP addresses were collected during this study. Sandvine’s network equipment analyzes data from an application utilization level and is not content aware.

With service provider customers in over 80 countries serving hundreds of millions of fixed and mobile subscribers, Sandvine is enhancing the Internet experience worldwide. For more information on Sandvine's platform, products, partners and applications please visit www.sandvine.com.

Visit http://www.sandvine.com/news/global_broadband_trends.asp for the executive summary of the Fall 2010 Global Internet Phenomena report.

-30-

ABOUT SANDVINE

Sandvine's network policy control solutions focus on protecting and improving the quality of experience on the Internet.

Our award-winning network equipment and software helps DSL, FTTx, cable, fixed wireless and mobile operators better understand network traffic, manage network congestion, create new services and revenues, mitigate traffic that is malicious or undesirable to subscribers, deliver QoS-prioritized multimedia services and increase subscriber satisfaction. With service provider customers in more than 80 countries serving hundreds of millions of broadband and mobile data subscribers, Sandvine is enhancing the Internet experience worldwide.

MEDIA CONTACT
Jennifer Ross
Sandvine
+1 519 880 2232
jross@sandvine.com

INVESTOR RELATIONS CONTACT
Rick Wadsworth
Sandvine
+1 519 880 2400 ext. 3503
rwadsworth@sandvine.com

AIM NOMAD
Andrew Chubb, Simon Bridges
Canaccord Adams Limited
+44 0207 050 6500

FORWARD-LOOKING STATEMENTS

Certain statements in this release constitute forward-looking statements or forward-looking information within the meaning of applicable securities laws and are made pursuant to the "safe harbor" provisions of such laws. Statements related to potential benefits of, and demand for, Sandvine's products or services including statements with respect to the features and benefits that may be achieved through the use of Sandvine's products or services and the relative position of these products vis-à-vis competitive offerings in the industry are forward-looking statements which are subject to certain assumptions, risks and uncertainties. These risks and uncertainties include such factors as rapid technological changes, changes in customer architecture and equipment deployment requirements, the introduction of competing technologies, the risks and uncertainties of new product introductions, dependence on key supplies and other similar factors that may cause the actual results, performance or achievements of Sandvine to differ materially from the results, performance, achievements or developments expressed or implied by such forward-looking statements. Readers are cautioned not to place undue reliance on such statements. Sandvine assumes no obligation to update or revise any forward-looking statements, whether as a result of new information, future events or otherwise.

December 13, 2010
Login | Register Now

Advertising Age®

Search Advertising Age
Advanced Search

MEDIAWORKS

Fox to Use Hulu Inventory for Advertiser 'Make-Goods'

Network's Maneuver Pushes TV, Online Audiences Closer

by Brian Steinberg

Published: November 23, 2010

NEW YORK (AdAge.com) -- Fox has secured agreements with about a dozen advertisers to supply them with inventory from online-video site Hulu to make up for ratings shortfalls on its broadcast network, according to the News Corp. network's top ad-sales executive. The move is the latest signal that marketers are growing more comfortable with the idea that consumers who watch TV via the web are comparable to a more traditional TV audience.

Fox routinely purchases inventory from Hulu -- owned by its parent, News Corp., as well as NBC Universal and Walt Disney Co. -- to sell as part of broader sponsorship packages or in the normal course of sales, said Jon Nesvig, Fox Broadcasting's president-sales, in an interview. This season, he said, the network's need to provide additional inventory to marketers due to lower-than-guaranteed ratings on its fall schedule -- a practice known in the industry as providing "make goods" -- prompted a different solution. Mr. Nesvig declined to identify the advertisers that had agreed to accept Hulu viewers as make-goods for traditional TV viewers.

Fox's effort follows that of the CW, which this season started selling packages of ad inventory that encompassed its TV

RELATED CONTENT

Toby Byrne Named Head of Fox Ad Sales
Fox Takes a Break for Blu-Ray, and Paramount Without Partners?
Tuning In: NBC's Odd Nod to Fox's 'Glee'; World Series Notes
Jon Nesvig, TV-Sales Dean and Fox Vet, to Retire
More...

network as well as its streaming-video website, CWTV.com. The networks' desire to now sell these packages, rather than keep advertisers primarily buying TV inventory, shows media companies rushing to adapt as new technology erodes the typical audience for prime-time TV, the priciest part of their schedules. And Fox's ability to use Hulu inventory to make up for broadcast shortfalls shows that advertisers are coming along.

"This is what television is going to be," said Michael Bologna, director-emerging communications, at WPP's Group M. He predicts the emergence of an "aggregation model," with networks cobbling together audiences from any number of viewing opportunities, whether they hew close to the couch-potato method of watching the boob-tube or hail from new behaviors, such as watching online.

Group M and other big media buyers are encouraging TV networks to offer and negotiate packages that include both standard TV viewers and viewers watching TV programs online, according to Mr. Bologna. "From what I'm hearing, this move is supported by a large part of the television community," he said.

Group M had not yet made a deal to buy Hulu inventory but is in the midst of discussing the idea, Mr. Bologna added.

In years past, Fox would have tried to keep its clients buying as much TV inventory as possible, while marketers might have blanched at the notion of using Hulu, a video service in its relative infancy, to replace viewers not watching mainstream TV. These days, however, audiences have dispersed among a plethora of new devices and viewing behaviors, ranging from playback on digital video recorders to downloads from iTunes to video on demand.

To be sure, Fox has had challenges this season. Its much-anticipated drama, "Lone Star," fell flat right out of the gate and its airing of the 2010 World Series, which featured teams not based in New York, Chicago or Los Angeles, lasted only five games and did not draw as big an audience as past broadcasts, particularly last year's matchup between the New York Yankees and the Philadelphia Phillies.

Reaching an agreement to mix online inventory with TV-ad purchases isn't the easiest thing to accomplish. The CW uses online-impressions data from DoubleClick as well as Nielsen VideoCensus data to give advertisers a sense of how its shows are watched online. Nielsen has been working on a plan to provide commercial ratings for shows watched on TV or online, but only so long as the ads streamed online match those that aired on TV -- which will only prove useful to the networks if people who watch TV programs via the web eventually see the same amount of ads as they might on TV, not fewer, as has been the norm in recent years.

Marketers need to determine the number of unique viewers an ad streamed in online video will reach, Mr. Bologna suggested, as well as the number of commercials seen during a break. Hulu is known for running fewer ads in a commercial break than TV networks do, meaning that Fox may be supplying its clients with inventory that would help them stand out better in certain cases than a traditional appearance on television.

Ad Network + Exchange Guide
Custom Programs
Tech Marketing Center Channel
Search Marketing Channel
Video Advertising Channel
The PR Factor
Mobile Marketing Guide

SERVICES

Subscriptions
Media Kit
Help Center
List Rental
Share & Save
RSS (What is RSS?)

Staff Contacts

Browse Archive

Download the Print Edition (Coverleaf/Texterity)

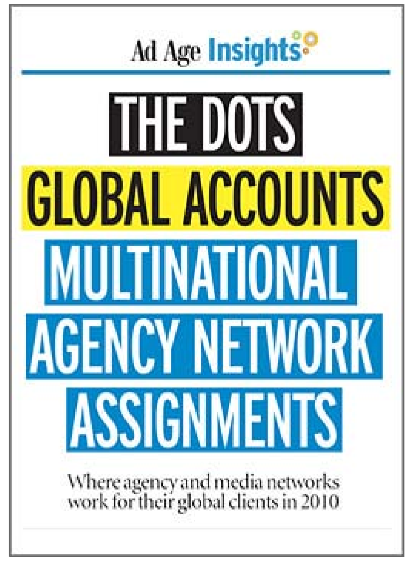

Which agency network represents which global clients in which countries? This report will tell you.

CONFERENCE COVERAGE

ME* Conference 2010

News, thoughts and insights from the brightest minds in marketing, media and publishing from Advertising Age's first ME* Conference.

SPECIAL REPORT

The Media Vanguard Awards

The first MVAs, created to honor transformative, next-generation media, drew hundreds of entries from major media companies (broadcasters, newspapers, magazines and digital natives) and marketers. See profiles on all 55 winners.

SPECIAL REPORT

Magazine A-List 2010

Propelled by smart audience strategies and diversified revenue streams, this year's A-List -- topped by Magazine of the Year, People StyleWatch -- shows what economic recovery looks like

4 Comments

Sort by **Date** | Popularity

By Jayter | jackson hts, NY November 23, 2010 07:06:14 pm:

The reason one buys TV is broad reach, fast reach, concurrent reach and impact. The demand for this drives the pricing. While I'm sure the networks would try and sell streaming makegoods as a good thing - "adding continuity", that is generally not why one buys TV. One generally buys a spot for concurrency, not continuity on TV. I guess what I'm saying is that if the butcher shorts me on the filet mignon I buy, he had better be prepared to make it up with many, many times the weight if he is going to make me good with hamburger. The oldest game in the book is oversetimating performance when making upfront network guarantees and then making up the shortfall (which they KNEW would happen due to their overestimates) with extremely low rated audience deficiency units (ADU's) - the type of units buyers would rather not buy. If the new ADU's are going to be microscopically rated streams that lack the concurrency, then advertisers would be out of their minds to let the agency buyers accept these. If the nets don't have the on-air ADU's, there is always the dreaded "Cash Back".

If advertisers actually let their agency buyers accept these impressions on a one for one basis as opposed to perhaps a 5:1 basis, then the lunatics truly have the keys to the asylum.

By cliffc | STL, MO November 24, 2010 06:11:17 am:

Give FOX credit for creativity in finding a way out of its current predicament. However, I would question whether an "aggregation model" with audiences cobbled together from different media will be acceptable to advertisers? The industry is still working on figuring out the cross-media audience delivery riddle. How can publishers, broadcasters, agencies and advertisers possibly identify and value these types of resource allocation trade-off decisions?

By justallie | New York, NY November 24, 2010 02:00:28 pm:

Well done, FOX. It's time to start thinking outside of the TV box... my full thoughts here:
http://justallie.wordpress.com/2010/11/24/blurring-the-lines/

By ayerpro69 | QUOGUE, NY November 29, 2010 09:28:19 am:

Fox should use Hulu type eyeballs as part of the overall sales package. Advertisers understand and embrace all this as long as you sell it upfront rather than as a consolation prize.
MJ Ayer, NY, NY

To submit comments you must be registered. Please Login

SPECIAL REPORT

Media Mavens 2010

These 18 Agency and Marketer Executives Faced the Great Recession Head on and Found Success Stretching Media Budgets and Bolstering Their Companies' Bottom Lines

100 Leading Media Companies

Ad Age's 30th annual ranking of the 100 Leading Media Companies, by net U.S. media revenue.

Premium Content for DataCenter subscribers: Media Family Trees 2010 database, including profiles, holdings, revenue and articles links for the Media 100.

SPECIAL REPORT

Your Media Questions Answered

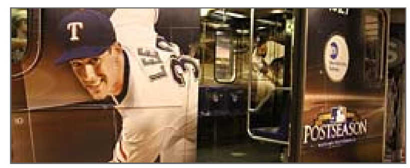

Complex Media Landscape Has You Confused? Here's Some Help.

MOST READ MOST E-MAILED

1. Book of Tens: Biggest Stories of the Year
2. Book of Tens: Social-Media Campaigns We Liked
3. Book of Tens: Campaigns Creativity Loved
4. Book of Tens: Advertising Icons of the Past Decade
5. Book of Tens: Best Marketer Apps

or Register Now

Username:

Password:

Comment:

remorse. Preizinh

Type the two words:

Post Comment

Note: Comments submitted to AdAge.com are posted automatically and will include the user name with which you registered. Ad Age reserves the right to delete comments that are insulting or personal in nature. Comments may be used in the print edition at editorial discretion. **Comments are restricted to 500 words or less.**

Share & Save (?)

Google
netvibes
Bloglines
newsgator
widgetbox
Windows Live

RSS Feed
Save on Del.icio.us
Digg this
Reddit
Bookmarks
LinkedIn

RECENT MEDIAWORKS

Time Inc. Names IAB CEO Randall Rothenberg Its First Chief Digital Officer
Even Oprah Needs A Sales Pitch
Answering Wikileaks, FarmVille's Hard Soil and Wasting Twitter on the Young
People Magazine Names Karen Kovacs Publisher

AD AGE WHITEPAPERS

The Dots: Global Accounts, Multinational Agency Network Assignments
Global Media Habits 2010
2010 Creativity's Awards Report
Building Brands Online
World's Hottest Brands: BRIC (Brazil, India, China)

More...

›› MORE MEDIAWORKS ARTICLES

Search Advertising Age

Advanced Search

Advertising Age: Your Online Source for Marketing and Media News

NEWS Latest Marketing and Advertising News | Marketing and Advertising Agency News | Best Ads From Creativity | CMO Strategy | DataCenter | Digital Marketing and Advertising News | Global Marketing and Advertising News | Hispanic Marketing | Madison+Vine | MediaWorks - Media and Marketing News | People & Players | This Week's Issue | Your Opinion

Netalyzr: Illuminating Edge Network Neutrality, Security, and Performance

Christian Kreibich[†], Nicholas Weaver[†],
Boris Nechaev*, and Vern Paxson[†]

TR-10-006

May 2010

Abstract

In this paper we present Netalyzr, a network measurement and debugging service that evaluates the functionality provided by people's Internet connectivity. The design aims to prove both comprehensive in terms of the properties we measure and easy to employ and understand for users with little technical background. We structure Netalyzr as a signed Java applet (which users access via their Web browser) that communicates with a suite of measurement-specific servers. Traffic between the two then probes for a diverse set of network properties, including outbound port filtering, hidden in-network HTTP caches, DNS manipulations, NAT behavior, path MTU issues, IPv6 support, and access-modem buffer capacity. In addition to reporting results to the user, Netalyzr also forms the foundation for an extensive measurement of edge-network properties. To this end, along with describing Netalyzr's architecture and system implementation, we present a detailed study of 112,000 measurement sessions that the service has recorded since we made it publicly available in June 2009.

†International Computer Science Institute, 1947 Center St., Ste. 600, Berkeley, CA 94704
* Helsinki Institute for Information Technology, PO Box 19800, 00076 Aalto, Finland

This work was partially supported by funding provided to ICSI through National Science Foundation grants CNS: 0722035 ("Architectural Support for Network Trouble-Shooting"), NSF: 0433702 ("CCIED: Center for Internet Epidemiology and Defenses"), and CNS: 0905631 ("Invigorating Empirical Network Research via Mediated Trace Analysis"), and also by funding provided by Amazon.com. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors or originators and do not necessarily reflect the views of the National Science Foundation. We are deeply grateful to the Netalyzr users for enabling this study and to our beta testers for the insightful comments and feedback. We would particularly like to thank Mark Allman, Paul Barford, Scott Bradner, John Brzozowski, Randy Bush, Niels Bakker, Richard Clayton, Chris Cowart, Keith Dawson, Adrian Dimcev, Holger Dreger, Brandon Enright, Kevin Fall, Carrie Gates, Andrei Gurtov, Mark Handley, Theodore Hong, Kelly Kane, Matthew Kogan, Keith Medcalf, Thomas Narten, Michael Ross, Wooter Wijngaards, and Richard Woundy.

1. INTRODUCTION

For most Internet users, their network experience—perceived service availability, connectivity constraints, responsiveness, and reliability—is largely determined by the configuration and management of their *edge network*, i.e., the specifics of what their Internet Service Provider (ISP) gives them in terms of Internet access. While conceptually we often think of users receiving a straight-forward "bit pipe" service that transports traffic transparently, in reality a myriad of factors affect the fate of their traffic.

It then comes as no surprise that this proliferation of complexity constantly leads to troubleshooting headaches for novice users and technical experts alike, leaving providers of web-based services uncertain regarding what caliber of connectivity their clients possess. Only a few tools exist to analyze even specific facets of these problems, and fewer still that people with limited technical understanding of the Internet will find usable. Similarly, the lack of such tools has resulted in the literature containing few measurement studies that characterize in a comprehensive fashion the prevalence and nature of such problems in the Internet.

In this work we seek to close this gap. We present the design, implementation, and evaluation of *Netalyzr*,[1] a publicly available service that lets any Internet user obtain a detailed analysis of the operational envelope of their Internet connectivity, serving both as a source of information for the curious as well as an extensive troubleshooting diagnostic should users find anything amiss with their network experience. *Netalyzr* tests a wide array of properties of users' Internet connections, starting at the network layer, including IP address use and translation, IPv6 support, DNS resolver fidelity and security, TCP and UDP service reachability, proxying and firewalling, anti-virus intervention, content-based download restrictions, content manipulation, HTTP caching prevalence and correctness, latencies, and access-link buffering.

We believe the breadth and depth of analysis *Netalyzr* provides is unique among tools available for such measurement. In addition, as of this writing we have recorded 112,000 runs of the system from 86,000 different public IP addresses, allowing us to construct a large-scale picture of many facets of Internet edge behavior. The measurements have found a wide range of behavior, on occasion even revealing traffic manipulation that the network operators themselves did not know about. More broadly, we find chronic over-buffering of links, a significant inability to handle fragmentation, numerous incorrectly operating HTTP caches, common NXDOMAIN wildcarding, impediments to DNSSEC deployment, poor DNS performance, and deliberate manipulation of DNS results.

We begin by presenting *Netalyzr*'s architecture and implementation (§ 2) and the specifics of the different types of measurements it conducts (§ 3). We have been operating *Netalyzr* publicly and continuously since June 2009, and in § 4 report on the resulting data collection, including flash crowds, their resulting measurement biases, and our extensive calibration tests to assess the correct operation of *Netalyzr*'s test suite. In § 5 we present a detailed analysis of the resulting dataset and some consequences of our findings. We defer our main discussion of related work to § 6 in order to have the context of the details of our measurement analysis to compare against. Finally, we summarize in § 7.

2. SYSTEM DESIGN

When designing *Netalyzr* we had to strike a balance between a tool with sufficient flexibility to conduct a wide range of measurement tests, yet with a simple enough interface that unsophisticated users would run it—giving us access to a much larger (and less biased towards "techies") end-system population than possible if the measurements required the user to install privileged software. To this end, we decided to base our approach on using a Java applet to drive the bulk of the tests, since (*i*) Java applets run automatically within most major web browsers, (*ii*) applets can engage in raw TCP and UDP flows to arbitrary ports (though not with altered IP headers), and, if the user approves trusting the applet, contact hosts outside the same-origin policy, (*iii*) Java applets come with intrinsic security guarantees for users (e.g., no host-level file system access allowed by default runtime policies), and (*iv*) Java's fine-grained permissions model allows us to adapt gracefully if a user declines to fully trust our applet.

The resulting system includes about 5,000 lines of Java for the applet (as well as some JavaScript to implement the client side of some test connections) and 12,000 lines of Python for the different servers. Figure 1 shows the conceptual *Netalyzr* architecture, whose components we now discuss in turn.

Application Flow. Users initiate a test session by visiting the *Netalyzr* website and clicking **Start Analysis** on the webpage with the embedded Java test applet. Once loaded, the applet conducts a large set of measurements probes, indicating test progress to the user. When testing completes, the applet redirects to a summary page that shows the results of the tests in detail and with explanations (Figure 2). The users can later revisit a session's results via a permanent link associated with each session. We also save the session state (and server-side packet traces) for subsequent analysis.

Front- and Back-end Hosts. The *Netalyzr* system involves three distinct locations: (*i*) the user's machine running the test applet in a browser, (*ii*) the *front-end* machine responsible for dispatching users and providing DNS service, and (*iii*) multiple *back-end* machines that each host both a copy of the applet and a full set of test servers. All back-end machines run identical configurations and *Netalyzr* conducts all tests in a given client's session using the same back-end machine.

The front-end machine runs Linux 2.6 on a 2.5 GHz Intel

[1] `http://netalyzr.icsi.berkeley.edu`

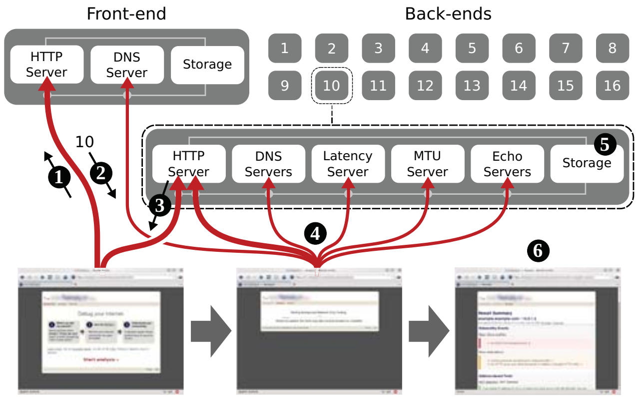

Figure 1: *Netalyzr*'s conceptual architecture. ❶ The user visits the *Netalyzr* website. ❷ When starting the test, the front-end redirects the session to a randomly selected back-end node. ❸ The browser downloads and executes the applet. ❹ The applet conducts test connections to various *Netalyzr* servers on the back-end, as well as DNS requests which are eventually received by the main *Netalyzr* DNS server on the front-end. ❺ We store the test results and raw network traffic for later analysis. ❻ *Netalyzr* presents a summary of the test results to the user.

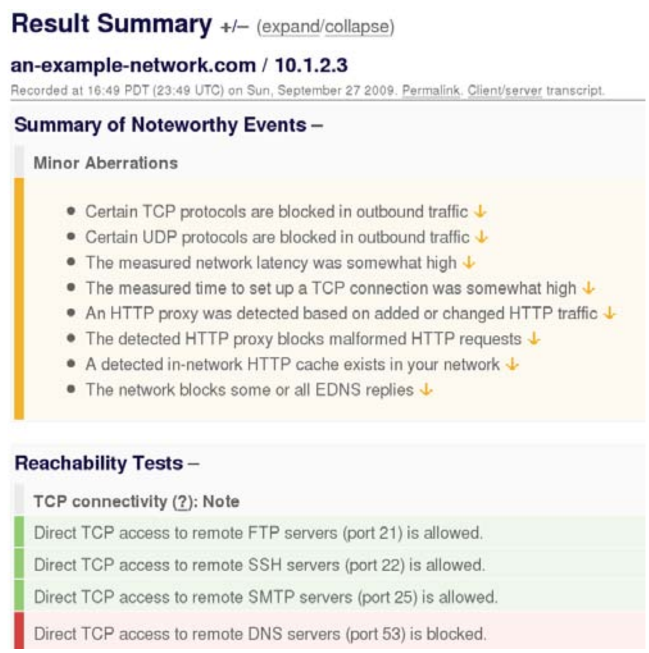

Figure 2: A partial screen capture of *Netalyzr*'s results page as seen by the user upon completion of all tests. The full report is 4–10 times this size, depending on whether the user expands the different sections.

Xeon machine with 8 GB of memory, physically located at our institute. We manage the back-end machines using Amazon's EC2 service [1] to facilitate scalability. These hosts are virtual 2.6 GHz AMD Opteron machines with 1.8 GB of memory and run Linux 2.6. At peak load times we employ 20 back-end hosts.

Front-end Web Server. Running on the front-end machine, this server provides the main website, including a landing/dispatch page, documentation, FAQs, an example report, and access to reports from previous sessions. The server employs a pre-forked pool of multithreaded child processes. The front page also includes a Java dispatch applet that ensures that the user has Java installed and then directs the user to a randomly selected back-end server to load-balance the actual testing process. Finally, the front page rate-limits visitors to a fixed number of measurements per minute per back-end server.

Back-end Web Servers. The back-end web servers host the actual measurement applet (so that its probe connections to the server accord with the same-origin policy) and perform HTTP testing and overall session management. When sending the measurement applet, the server includes a set of configuration parameters, including a globally unique session ID.

Measurement Applet. The Java applet implements 38 types of tests, some with a number of subtests. We describe them in detail in Section 3. The applet conducts the test cases sequentially, but also employs multithreading to ensure that test sessions cannot stall the entire process, and to speed up some parallelizable tasks. As tests complete, the applet transmits detailed test results to the back-end server; it also sends a continuously recorded client-side transcript of the session. Finally, we sign our applet with a certificate from a trusted authority so that browsers indicate a valid signature.

DNS Servers. An instance of this server runs on the front-end as well as the back-end machines. On the front-end, it acts as the authoritative resolver for two subdomains, `.`*n*`.na.edu` and `.`*n*`.na.org`, while on the back-ends it receives DNS test queries generated directly from the applet rather than through the user's DNS resolver library. The server interprets queries for specific names as commands, generating replies that encode values in A and CNAME records. For example, requesting `has_edns.`*n*`.na.edu` will return an A record reflecting whether the query message indicated EDNS support. The server also accepts names with arbitrary interior padding to act as a cache-busting nonce, ensuring that queries reach our server.

Echo Servers. An array of simple TCP and UDP echo servers allow us to test service-level reachability and content modification of traffic on various ports. The servers mostly run on well-known ports but do not implement the associated application protocol. Rather, they use their own simple payload schema to convey timing, sequencing, and the requester's IP address and source port back to the client. An additional server can direct a DNS request to the user's public address to check if the user's NAT or gateway acts as a proxy for external DNS requests.

Bandwidth Measurement Servers. To assess bandwidth, latency, buffer sizing, and packet dynamics (loss, reordering, duplication) we employ dedicated UDP-based measurement servers. Like the echo servers, these use a custom payload schema that includes timing information, sequence numbers, instructions regarding future sending, and aggregate counters.

Path MTU Measurement Server. To measure directional path MTUs, we use a server that can capture and transmit raw packets, giving us full access to and control over all packet headers.

Storage. To maintain a complete record of server-side session activity, we record all relevant network traffic on the front- and back-end machines, except for the relatively high-volume bandwidth tests. Since Java applets do not have the ability to record packets, we cannot record such traces on the client side.

Session Management. The back-end web servers establish and maintain session state as test sessions progress, identifying sessions via RFC 4122 UUIDs. We serialize completed session state to disk on the back-end hosts and periodically archive it on the front-end. When viewing a session summary, the front-end web server redirects the request to the appropriate back-end (encoded in the session ID) if it does not have the state locally, and the back-end web server does the opposite, with cycle detection to avoid looping.

3. MEASUREMENTS CONDUCTED

We now describe the types of measurements *Netalyzr* conducts and the particular methodology used. We begin with layer 3 measurements (addressing, fragmentation, MTU, raw performance, IPv6 support) and then progress to higher layers (general service reachability, DNS, HTTP), finishing with a discussion of user feedback and tests we chose to omit.

3.1 Network-layer Information

Addressing. We obtain the client's local IP address via the Java API, and use a set of raw TCP connections and UDP flows to our echo servers to learn the client's public address. From this set of connections we can identify the presence of NAT, and if so how it renumbers addresses and ports. If across multiple flows we observe more than one public address, then we assess whether the address flipped from one to another—indicating the client changed networks while the test was in progress—or alternates back and forth. This latter implies either the use of load-balancing, or that the NAT does not attempt to associate local systems with a single consistent public address but simply assigns new flows out of a public address block as convenient. (Only 1% of sessions included an address change from any source.)

IP Fragmentation. We test for proper support of IP fragmentation (and also for MTU measurement; see below) by sending UDP payloads to our test servers. We first check for the ability to send and receive fragmented UDP datagrams. In the applet → server direction, we send a 2 KB datagram which, if received, generates a small confirmation response. Due to the prevalence of Ethernet framing, we would expect most clients to send this packet in fragments, but it will always be fragmented by the time it reaches the server. We likewise test the server → applet direction by our server transmitting (in response to a small query from the client) a 2 KB message to the client. This direction will definitely fragment, as the back-end nodes have an interface MTU of 1500 bytes.

If either of the directional tests fails, the applet performs binary search to find the maximum packet size that it can successfully send/receive unfragmented.

The applet also tries to send and receive packets with 1471 bytes of UDP payload (normally yielding a 1499-byte IP packet) which would maximize the payload on an Ethernet network without fragmentation. This checks for the existence of an "MTU hole", where packets can be sent unfragmented by the endpoint but cannot be refragmented properly when passing through a path MTU bottleneck, either because the bottleneck is functioning incorrectly or the host sent the packet with `DF` set.

Path MTU. A related set of tests conducts path MTU probing. The back-end server for this test supports two modes, one for each direction. In the applet → server direction, the applet sends a large UDP datagram, resulting in fragmentation. The server monitors arriving packets and reports the IP datagram size of the entire original message (if received unfragmented) or of the original message's initial resulting fragment. This represents a lower bound on MTU in the applet → server direction, since the first fragment's size is not necessarily the full path MTU. (Such "runts" occurred in only a handful of sessions).

In the server → applet direction, the applet conducts a binary search beginning with a request for 1500 bytes. The server responds by sending datagrams of the requested size with `DF` set. In each iteration one of three cases occurs. First, if the applet receives the `DF`-enabled response, its size is no more than the path MTU. Second, if the response exceeds the path MTU, the server processes any resulting ICMP "fragmentation required" messages and sends to the applet the attempted message size, the offending location's IP address, and the next-hop MTU conveyed in the ICMP message. Finally, if no messages arrive at the client, the applet infers that the ICMP "fragmentation required" message was not generated or did not reach the server, and thus a path MTU problem exists.

Latency, Bandwidth, and Buffering. We measure packet delivery performance in terms of round-trip latencies, directional bandwidth limits, and buffer sizing. With these, our primary goal is not to measure capacity itself (which numerous test sites already address [31]), but as a means to measure the sizing of bottleneck buffers, which can significantly affect user-perceived latency. We do so by measuring the increase in latency between quiescence and that experienced during the bandwidth test, which in most cases will briefly saturate the path capacity in one direction and thus fill the buffer at the bottleneck.

Netalyzr conducts these measurements in two basic ways. First, early in the measurement process it starts sending in the background small packets at a rate of 5 Hz. We use this test to detect transient outages, such as those due to a poor wireless signal.

Second, it conducts an explicit latency and bandwidth test. The test begins with a 10 Hz train of 200 small UDP packets, for which the back-end's responses provide the baseline mean latency used when estimating buffer sizing effects. The test next sends a train of small UDP packets that elicit 1000-byte replies, with exponentially ramping up (over 10 seconds) the size in slow-start fashion: for each packet received, the applet sends two more. In the second half of the interval, the applet measures the sustained rate at which it receives packets, as well as the average latency. (It also notes duplicated and reordered packets over the entire run.) After waiting 5 seconds for queues to drain, it repeats with sizes reversed, sending large packets to the server that trigger small responses. Note that most Java implementations will throttle sending rates to ≤ 20 Mbps, imposing an upper bound on the speed we can measure.

IPv6 Adoption. To measure IPv6 connectivity we have to rely on an approximation because neither our institution nor Amazon EC2 supports IPv6. However, on JavaScript-enabled hosts the analysis page requests a small logo from `ipv6.google.com`, reachable only over IPv6. We report the outcome of this request to our HTTP server. Since we cannot prevent this test from possibly fetching a cached image, we could overcount IPv6 connectivity if the user's system earlier requested the same resource (perhaps due to a previous *Netalyzr* run from an IPv6-enabled network).

3.2 Service Reachability

To assess any restrictions the user's connectivity may impose on the types of services they can access, we attempt to connect to 25 well-known services along with a few additional ports on the back-end. For `80/tcp` and `53/udp` connectivity, the applet speaks proper HTTP and DNS, respectively. We test all other services using our echo server protocol as described in Section 2.

In addition to detecting static blocking, these probes also allow us to measure the prevalence of proxying. In the absence of a proxy, our traffic will flow unaltered and the response will include our public IP address as expected. On the other hand, protocol-specific proxies will often transform this non-protocol-compliant response into an error, or simply abort the connection. Such proxies can reside on the end host (e.g., as part of an AV system) or in the network, with additional protocol information such as banners or headers often suggesting the source.

3.3 DNS Measurements

Netalyzr performs extensive measurements of DNS behavior, since DNS manipulations and subtle errors can have a major impact on a user's network experience. We implement two levels of measurement, *restricted* and *unrestricted*. Restricted measurements comply with Java's default same-origin policy, which for most JVMs allows the lookup of arbitrary names but only ever returns the IP address of the origin server, or throws an exception if the result is not the origin server's address. If however the user trusts the applet, then we can look up arbitrary names through the system's DNS resolver unrestrictedly, allowing us to conduct substantially more comprehensive testing. We refer to names corresponding to *Netalyzr*'s actual domain as *internal*, and any others as *external*. We can only look up the latter if unrestricted.

As mentioned earlier, our DNS authority server interprets requests for specific names as commands telling it what sort of response to generate. We encode Boolean results by returning the IP address of the back-end service for *true* and the address of an unrelated host in our institution for *false*. For results that return names, we indicate failure with the hostname `return_false`.

In our discussion, we abbreviate the fully qualified hostname of the back-end node as follows. First, *n*`.na.edu` stands for *node*`.netalyzr.icsi.berkeley.edu` (likewise *n*`.na.org` stands for *node*`.netalyzr.icir.org`). Second, if we give only a hostname `name`, it stands for `name.`*node*`.netalyzr.icsi.berkeley.edu`. Finally, we indicate the presence of a pseudo-random nonce value (to ensure cache penetration) using "*nonce*" in the name.

Glue Policy. One important but subtle aspect of the DNS resolution process concerns the acceptance and promotion of response data in the Authoritative or Additional records of a response, commonly referred to as *"glue"* records. Acceptance of such records can boost performance by avoiding future lookups, but also risk cache poisoning attacks [6]. Assessing the acceptance of these records is commonly referred to as "bailiwick checking," but the guidelines on the procedure allow latitude in how to conduct it [11]. *Netalyzr* leverages glue acceptance to enable tests of the DNS resolver itself.

We first check acceptance of arbitrary A records in the Additional section by sending lookups of special names (made distinct with nonces) that return particular additional A records. We then look up those additional names directly to see whether the resolver issues new queries for the names (which would return *false* when those names are queried directly) or answers them from its cache (returning *true*), indicating that the resolver accepted the glue. We then likewise check for caching of Authority A records. Finally, we check whether the server will automatically follow CNAME aliases. In this test, the response provides an Answer of a CNAME for `return_false`, with an Additional record encoding `return_false` as *true*. Thus, the query evaluates as *true* only if the resolver accepts the A record associated with the CNAME.

DNS Server Identification and Properties. We next probe more general DNS properties, including resolver identity, IPv6 support, `0x20` support [8], respect for short TTLs, port randomization for DNS requests, and whether the user's NAT, if present, acts as a DNS proxy on its external IP address.

When able to conduct unrestricted DNS measurements, we identify the resolver's IP address (as seen by our server) by returning it in an A record in response to a query for `server.`*nonce*`.`*n*`.na.edu`. This represents the address of the final server sending the request, not necessarily the one the client uses to generate the request. During our beta-testing we changed the applet code to conduct this query multiple times because we observed that some hosts will shift between DNS resolvers, and some DNS resolvers actually operate as clusters.

We test IPv6 AAAA support by resolving `ipv6_set.`*nonce*. We expect the resolver to request at least an A record for this name, and if it supports IPv6 then also a AAAA record. We discard the server's reply for the A record and then then resolve `ipv6_check.`*nonce*. When the A record request for this name arrives, the server checks whether it saw a AAAA request for the previous name (which might have arrived after the original A request, and thus could not have been reported initially), which it indicates by whether it returns *true* for the second A request. By proceeding in this fashion, we can assess resolver support for IPv6 even if the client itself does not support it.

Queries for the name `0x20` return *true* if the capitalization in a mix-cased request retains the original mix of casing. This detects non-`0x20`-compliant resolvers that change the capitalization of requested names.

If the DNS resolver accepts glue records for nameservers (NS responses in Authority or Additional), we leverage this to check whether the resolver respects short TTLs. Responses to the name `ttl0` or `ttl1` place a glue record for `return_false` in the Authoritative section with a TTL of 0 or 1 seconds, respectively. A subsequent fetch of `return_false` reveals whether the short TTLs were respected. (We can't simply use A records for this test because both the browser and end host may cache these records independently.)

We also use lookups of `glue_ns.`*nonce* to measure request latency. If the DNS resolver accepts glue records, it then also looks up `return_false.`*nonce* to check the latency for a cached lookup. We repeat this process ten times and report the mean value to the server, and also validate that `return_false.`*nonce* was fetched from the resolver's cache rather than generating a new lookup.

Finally, we test DNS port randomization. For unrestricted measurements, we perform queries for `port.`*nonce*, which the server answers by encoding in an A record the source port of the UDP datagram that delivered the request. For restricted measurements, the applet sends several queries for `dns_rand_set` and then checks the result by a query for `dns_rand_check`, with the latter resolving as *true* if the ports seen by our DNS server appeared non-monotone.

EDNS, DNSSEC, and actual DNS MTU. DNS resolvers can advertise the ability to receive large responses using EDNS [29], though they might not actually be capable of doing so. For example, some firewalls will not pass IP fragments, creating a de-facto DNS MTU of 1478 bytes for Ethernet framing. Other firewall devices may block all DNS replies greater than 512 bytes under the out-of-date assumption that DNS replies cannot be larger. While today small replies predominate, a lack of support for large replies poses a significant concern for DNSSEC deployment, as it will result in unpredictable performance degradation when DNS replies exceed unstated and hidden limits.

We measure the prevalence of this limitation by issuing lookups (*i*) to determine whether requests arrive indicating EDNS support, (*ii*) to measure the DNS MTU (for unrestricted measurements), and (*iii*) to check whether the resolver requests DNSSEC records. For the first, we look up `has_edns`, which returns *true* if the request contained an EDNS `OPT` pseudo-record. Responses for `edns_mtu` encode the advertised EDNS MTU in the lower 16 bits of an A record, and `wants_dnssec` returns *true* if the `DO` ("use DNSSEC") flag is set in an EDNS pseudo-record.

That a DNS resolver advertises (via EDNS) the ability to receive large responses does not guarantee that it actually can. We test its ability by requesting names `edns_medium` and `edns_large`, padded to 1300 and 1700 bytes, respectively. (We pad the replies to those sizes by adding Additional CNAME records.) Their arrival at the client indicates the resolver an indeed receive larger DNS replies.

During beta-testing we made this test more precise: the server answers requests for `ednspadding_`*X* with a response padded to exactly *X* bytes of DNS payload. We use this mechanism and binary search to determine the actual maximum supported by the resolver (whether or not it advertises EDNS).

NXDOMAIN Wildcarding. Some DNS operators configure their resolvers to perform "NXDOMAIN wildcarding", where they rewrite hostname lookups that fail with a "no such domain" error to instead return an A record for the IP address of a web server. The presumption of such blanket rewriting is that the original lookup reflected web surfing, and therefore returning the impostor address will lead to the subsequent HTTP traffic coming to the operator's web server, which then typically offers suggestions related to the presumed intended name. Such rewriting—often motivated by selling advertisements on the landing page—corrupts the web browsers' URL auto-complete features, and, worse, breaks protocol semantics for any non-HTTP application looking a hostname.

If unrestricted, the applet checks for this behavior by querying for a series of names in our own domain namespace, and which do not exist. We first look up www.*nonce*.com. If this yields an IP address, we have detected NXDOMAIN wildcarding, and proceed to probe the behavior in more detail, including simple transpositions (www.yahoo.cmo), other top-level domains (www.*nonce*.org), non-web domains (fubar.*nonce*.com), and domain internal to our site (nxdomain.*n*.na.edu). The applet also attempts to contact the host returned for www.*nonce*.com on 80/tcp to obtain the imposed web content, which we log.

DNS proxies, NATs, and Firewalls. Another set of DNS problems arise not due to ISP interference but misconfigured or misguided NATs and firewalls. If the applet operates unrestricted, it conducts the following tests to probe for these behaviors. First, it measures DNS awareness and proxying. Our servers answer requests for entropy.*n*.na.edu with a CNAME encoding the response's parameters, including the public address, UDP port, DNS transaction ID, and presence of 0x20 encoding. The applet sends such DNS requests directly to the back-end server, bypassing the configured resolver. If it observes any change in the response (e.g., a different transaction ID or public address), then we have found in-path DNS proxying. The applet makes another request directly to the back-end server, now with deliberately invalid format, to which our server generates a similarly broken reply. If blocked, we have detected a DNS-aware middlebox that prohibits non-DNS traffic on 53/udp. The applet then issues direct queries for the names edns_large and edns_medium (discussed above), and now also edns_small (a 400-byte response with EDNS), to check whether the NAT or firewall has problems handling either EDNS replies or large DNS responses.

During beta-testing we added a series of tests for the presence of DNS proxies in NAT devices. NATs often include such a proxy, returning via DHCP its local address to clients as the DNS resolver location if the NAT has not yet itself acquired an external DNS resolver.[2] Upon detecting the presence of a NAT, the applet assumes the gateway's local address is the *a.b.c.*1 address in the same /24 as the local IP address[3] and sends it a query for entropy.*n*.na.edu. Any reply indicates with high probability that the NAT implements a DNS proxy. In addition, we can observe to where it forwards the request based on the client IP address seen by our server.

During our beta-testing we became aware of the possibility that some in-gateway DNS resolvers act as open relays *for the outside* (i.e., for queries coming from external sources), enabling amplification attacks [22] and other mischief. We thus added a test in which the the applet instructs the back-end DNS server to send a UDP datagram containing a DNS request for entropy.*n*.na.edu to the public IP address of the client to see if it elicits a resulting response at our DNS server.

Name Lookup Test. Finally, if unrestricted the applet looks up a list of 70+ common names, including major search engines, advertisement providers, financial institutions, email providers, and e-commerce sites. It uploads the results to our server, which then performs reverse lookups to test the forward lookups for consistency. This testing unearthed numerous aberrations, as discussed below.

3.4 HTTP Proxying and Caching

For analyzing HTTP behavior, the applet employs two different methods: using Java's *high-level API*, or its *low-level TCP sockets* (for which we implement our own HTTP logic). The first allows us to assess behavior imposed on the user by their browser (such as proxy settings), while the latter reflects behavior imposed by their access connectivity. (For the latter we take care to achieve the same HTTP "personality" as the browser by having our server mirror the browser's HTTP request headers to the applet so it can emulate them in subsequent low-level requests.) In general, the applet co-ordinates measurement tasks with the server using URL-encoded commands that instruct the server to deliver specific kinds of content (such as cache-sensitive images), report on properties of the request (e.g., specific header values), and establish and store session state.

Proxy Detection. We detect proxy configuration settings by monitoring request and result headers, as well as the server-perceived client address of a test connection. Differences when using the high-level API versus the socket API indicate the presence of a configured proxy. We first send a low level message with specific headers to the web server. The server mirrors the headers back to the applet, allowing the applet to conduct a comparison. Added, deleted, or modified headers flag the presence of an in-path proxy. To improve the detectability of such proxies, we use eccentric capitalization of header names (e.g. User-AgEnt) and observe whether these arrive with the same casing. A second test relies on sending an invalid request method (as opposed to GET or POST). This can confuse proxies and cause them to terminate the connection. A final test sets the Host request header to www.google.com instead of *Netalyzr*'s domain. Some proxies use this header's value to direct the outgoing connection [13]. The applet monitors for unexpected content—either Google's HTML banner, or a 302 redirect to a country-specific Google page. If seen, this represents a significant security vulnerability, as such proxies will allow Java and Flash to violate same-origin policies arbitrarily. However, we saw only a handful of instances of such behavior.

Caching policies, Content Transcoding, and File-type Blocking. We next test for in-network HTTP caching. For this testing, our server provides two test images of identical size (67 KB) and dimensions (512·512 pixels), but each the color-inverse of the other. Consecutive requests for the image result in alternating images returned to the applet. We

[2]Once the NAT obtains its external DHCP lease, it then forwards all DNS requests to the remote resolver.

[3] We assume this is the address, rather than probe for it, to avoid creating any apparent scanning activity.

can thus reliably infer when the applet receives a cached image based on the unchanged contents (or an HTTP 304 status code, "Not Modified"). We conduct four such request pairs, varying the cacheability of the images via various request and response headers, and including a unique identifier in each request URL to ensure each session starts uncached.

The applet can also identify image transcoding or blocking by comparing the received image's size to the expected one. In the post-beta codebase, the applet uploads any changed content for off-line analysis.

Finally, we test for content-based filtering. The applet downloads (*i*) an innocuous Windows PE executable (notepad.exe), (*ii*) a small MP3 file, (*iii*) a bencoded BitTorrent download file (for a Linux distribution's DVD image), and (*iv*) the EICAR test "virus",[4] a benign file that AV vendors recognize as malicious for testing purposes.

3.5 User Feedback

Because we cannot readily measure the physical context in which the user runs *Netalyzr*, we include a small, optional questionnaire in the results page. Some 19% of the users provided feedback. Of those, 57% reported using a wired rather than a wireless network; 17% reported running *Netalyzr* at work, 79% from home, 2% on public networks, and 2% on "other" networks.

3.6 Intentional Omissions

We considered several tests for inclusion but decided not to do so for one of two reasons. First, some tests can result in potentially destructive or abusive effects, particularly if run frequently or by multiple users. In this regard we decided against tests to measure the NAT's connection table size (which could disrupt unrelated network connections purged from the table), fingerprint NATs by connecting to its internal web-administration interface (which might expose sensitive information), general scanning either locally or remotely, and sustained high-bandwidth tests (such as BitTorrent throttling, for which alternative, bandwidth-intensive tests exist [10]). Another reason to omit a test concerns potential long-term side-effects for the users themselves. These could occur for technical reasons (e.g., we contribute towards possible upload/download volume caps) or legal/political ones (e.g., tests that attempt to determine whether access to certain sites suffers from censorship). Finally, we do not store tracking cookies in the user's browsers, since we do not aim to collect mobility profiles and can manage sessions using state on our servers.

4. DATA COLLECTION

We began running *Netalyzr* publicly in June 2009 and have kept it available continuously. We initially offered the service as a "beta" release (termed BETA), and for the most part did not change the operational codebase until January 2010, when we rolled out a substantial set of adjustments and additional tests (RELEASE). These comprise about 68% and 32% of the measurements, respectively. Unless otherwise specified, discussion refers to the combination of both datsets.

Website Operation. To date we have collected 112,239 sessions from 86,252 public IP addresses. The peak rate of data acquisition occurred during the June roll-out, with a maximum of 1,452 sessions in one hour. This spike resulted from mention of our service on several web sites. A similar but smaller spike occurred during the January relaunch, resulting in a peak load of 373 sessions in one hour.

Calibration. We undertook extensive calibration of the measurement results to build up confidence in the coherence and meaningfulness of our data. A particular challenge in realizing *Netalyzr* has been that it must operate correctly in the presence of a wide range of failure modes. While we put extensive effort into anticipating these problems during development, subsequent calibration served as a key technique to validate our assumptions and learn how the tests actually work on a large scale. In addition, it proved highly beneficial to employ someone for this task who was not involved in developing the tests, as doing so avoided incorporating numerous assumptions implicitly present in the code. Finally, we emphasize the importance of capturing subtle flaws in the data and uncovering inconsistencies that would otherwise skew the analysis results or deflate the scientific value of the data.

We based our calibration efforts on the BETA dataset, using it to identify and remedy sources of errors before beginning the RELEASE data collection. To do so, we assessed data consistency individually for each of the tests mentioned in § 3. We emphasized finding missing or ambiguous values in test results, checking value ranges, investigating outliers, confirming that each test's set of result variables exhibited consistency (e.g., examining that mutual exclusiveness was honored, or that fractions added up to a correct total), ensuring that particular variable values complied with corresponding preconditions (e.g., availability of raw UDP capability reliably enabling certain DNS tests), and searching for systematic errors in the data.

To our relief, this process did not uncover any major flaws in the codebase or the data. The most common problems we uncovered were ambiguity (for example, in distinguishing silent test failures from cases when a test was not executed at all) and inaccuracies in the process of importing the data into our session database. The RELEASE version of the codebase only differs from BETA in the presence of more unambiguous and extensive result reporting (along with the addition of new tests).

Identified Measurement Biases. A disadvantage of website-driven data collection is vulnerability to sudden referral surges from specific websites—in particular if these entail a technologically biased user population that can skew our dataset. In addition, our Java runtime requirement could

[4] `http://www.eicar.org/anti_virus_test_file.htm`

discourage non-technical users whose systems do not have the runtime installed by default. It also precludes the use of *Netalyzr* on many smartphone platforms. We now analyze the extent to which our dataset contains such bias.

The five sites referring the most users to *Netalyzr* are: stumbleupon.com (25%), lifehacker.com (14%), slashdot.org (13%), google.com (7%), and heise.de (7%). The context of these referrals affects the number of sessions we record for various ISPs. For example, most users arriving from slashdot.org did so in the context of an article on alleged misbehavior by Comcast's DNS servers, likely contributing to making their customers the biggest share of our users (10.9% of our sessions originate from Comcast's IP address ranges). Coverage in Germany via heise.de likely drove visits from customers of Deutsche Telekom, accounting for 2.6% of the sessions. We show a summary of the dominant ISPs in our dataset in Table 3 below.

The technical nature of our service introduced a "geek bias" in our dataset, which we can partially assess by using the `User-Agent` HTTP request headers of our users to infer browser type and operating system. Here we compare against published "typical" numbers [33, 34], which we give in parentheses. 39.8% (90%) of our users ran Windows, 8.1% (1.0%) used Linux, and 14.3% (5.9%) used MacOS. We find Firefox over-represented with 60.9% (28.3%) of sessions, followed by 18.8% (59.2%) for Internet Explorer, 15.6% (4.5%) for Safari, and 2.9% (1.7%) for Opera. This bias also extends to the choice of DNS resolver, with 12% of users selecting OpenDNS as their DNS provider.

While such bias is undesirable, it can be difficult to avoid in a study that requires user participation. We can at least ameliorate distortions from it because we can identify its presence. Its primary effect concerns our characterizations across ISPs, where we endeavor to normalize accordingly, as discussed below. We also note that technically savvy users may be more likely to select ISPs with fewer connectivity deficiencies, which would mean the prevalence of problems we observe may reflect underestimates.

5. DATA ANALYSIS

We now turn to an assessment of the data gathered from *Netalyzr* measurements to date. In our discussion we follow the presentation of the different types of tests above, beginning with layer 3 measurements and then progressing to general service reachability and specifics regarding DNS and HTTP behavior.

5.1 ISP and Geographic Diversity

We estimate the ISP and location of *Netalyzr* users by inspecting reverse (PTR) lookups of their public IP address, if available; or else the final Start-of-Authority record in the DNS when attempting the PTR lookup. We found these results available for 97% of our sessions.

To extract a meaningful organizational name, we started with a database of "effective TLDs," i.e., domains for which the parent is a broad, undifferentiated domain such as `gouv.fr` [19], to identify the relevant name preceding these TLDs. Given this approach, our dataset consists of sessions from 6,868 organizations (see Table 3 below for the 15 most frequent) across 182 countries, as shown in Figure 3. Activity however was dominated by users in the USA (48.2%), the EU (31.3%, with Germany accounting for 9.7% and Great Britain for 7.9%), and Canada (5.4%). 10 countries contributed sessions from more than 1,000 addresses, 46 from more than 100, and 97 from more than 10.

Figure 3: Global locations of *Netalyzr* runs.

5.2 Network-Layer Information

Network Address Translation. Unsurprisingly, we find NATs very prevalent among *Netalyzr* users (90% of all sessions). 79% of these sessions used the `192.168/16` address range, 16% used `10/8`, and 4% used `172.16/12`. 2% of the address-translated sessions employed some form of non-private address (either public or not allocated for private use). We did not discern any particular pattern in these sessions or their addresses; some were quite bizarre.

Port sequencing behavior. For more recent *Netalyzr* runs we have tracked potential NAT port renumbering explicitly, recording port numbers as seen by both the client and the server for a batch of 10 TCP connections. Of 19,510 sessions, 33% exhibit port renumbering. Of these, 8.9% appear random,[5] while 89.0% renumber in a strictly monotone-increasing fashion. We find a median "spread" for this sequence (range from smallest port to largest, inclusively) of 10, indicating renumbering that exactly reflects the tests we generated. A number of sessions have much higher spread, however (with a mean of 102). For these we have ruled out little-endian increments (i.e., by 256 rather than by 1) for other than a handful of sessions, but have not at this point assessed whether sessions with higher means contain significant forward jumps. Such jumps could occur due to effects other than the NAT concurrently processing additional connections separate from our measurements. Identifying and removing these would then enable us to estimate the level of multiplexing apparently present in the user's access link.

[5] We use a Wald-Wolfowitz test with sequence threshold 4 to measure randomness.

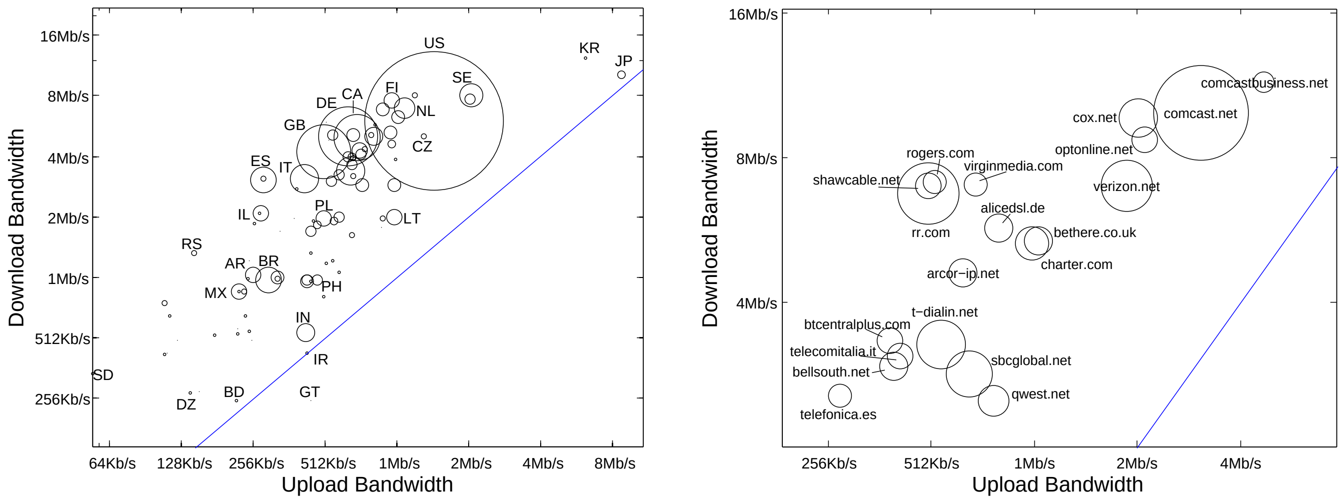

Figure 4: Average up/downstream bandwidths for countries with $\geq$ 10 sessions (left) and the 20 most prevalent ISPs (right). Circle areas are proportional to prevalence in the dataset, and diagonal lines mark symmetric upload and download capacity.

IPv6. We found IPv6 support to be rare but non-negligible: 4.5% of sessions fetched the logo from `ipv6.google.com`. As discussed above, this represents an upper bound due to possible caching effects.

Fragmentation. Overall, we find that fragmentation is not as reliable as desired [16, 26]. The RELEASE included a significant evaluation of UDP fragmentation behavior, for which we found 8% of the sessions unable to send 2 KB UDP packets, and likewise 8% unable to receive them.

We also found that 3% of the sessions which *could* send 2 KB packets could *not* send 1500 B packets. We find that 88% of these sessions come from Linux systems, strongly suggesting the likely cause to be Linux's arguably incorrect application of Path MTU discovery to UDP traffic, sending unfragmented UDP with `DF` set unless the system previously received an ICMP "fragmentation required" message from the recipient's path. Java, likewise, does not appear to retransmit in the face of such ICMP feedback, instead raising an exception which *Netalyzr* reports as a failure.

Regarding the path MTU from our server to the client, 80% of the sessions exhibited a path MTU of 1500 B, followed by 1492 B (15%) which suggests a prevalence of PPP over Ethernet (PPPoE). We also observe small clusters at 1480 B, 1476 B, 1460 B, and 1458 B, but these are rare. Only 1% reported an MTU less than 1450 bytes.

For sessions with an MTU < 1500 B, only 58% had a path that successfully sent a proper "fragmentation required" ICMP message back to our server. This finding reinforces that systems should avoid PMTU for UDP, and for TCP should provide robustness in the presence of MTU black holes [18].

Latency and Bandwidth. Figure 4 illustrates the balance of upstream vs. downstream capacities for countries and ISPs. Figure 5 shows the distribution of download bandwidths for particularly prominent ISPs. Two years after the study by Dischinger et al. [9] our results still partially match theirs, particularly for RoadRunner.

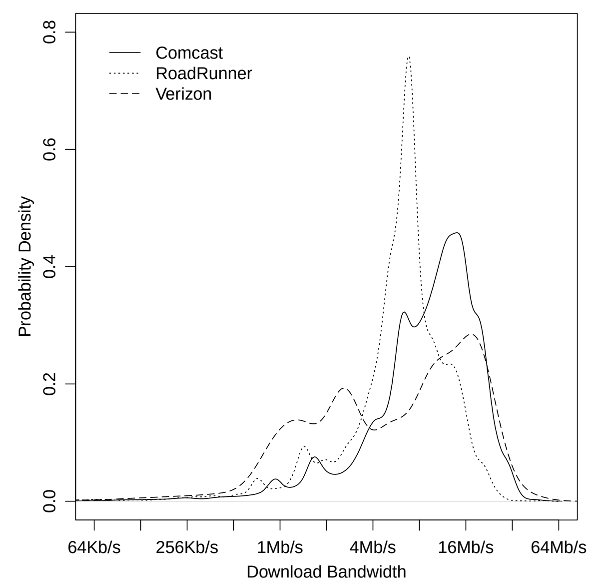

Figure 5: PDF of download bandwidths for the three most prominent ISPs in our dataset.

From the most aggregated perspective, we observed an average download bandwidth of 6.7 Mbps and for upload 2.7 Mbps. We find far more symmetric bandwidths for sessions that users self-reported as at work (10 Mbps/8.2 Mbps), and reported home connections exhibited far more asymmetry and lower bandwidth (6.3 Mbps/1.6 Mbps). Public networks exhibited less download bandwidth but more symmetry (3.4 Mbps/2.3 Mbps).

We saw less variation in the aggregate perspective for quiescent latency. Sessions reported as run at work had an average latency of 100 ms, while home networks experienced 120 ms and public networks 180 ms of latency.

Network Uplink Buffering. A known problem [9] confirmed by *Netalyzr* concerns the substantial over-buffering present in the network, especially in end-user access devices. *Netalyzr* attempts to measure this by recording the amount of delay induced by the high-bandwidth burst of traffic once it exceeds the actual bandwidth obtained. We then infer the buffer capacity as equal to the sustained sending rate multiplied by the additional delay induced by this test. Since the test uses UDP, no back-off comes into play to keep the buffer from completely filling, though we note that *Netalyzr* cannot determine whether the buffer did indeed actually fill to capacity.

When plotting measured upload bandwidth vs. inferred upload buffer capacity (Figure 6), several features stand out. First, we note that because we keep the test short in order to not induce excessive load on the user's link, sometimes *Netalyzr* cannot completely fill the buffer, leading to noise, which also occurs when the bandwidth is quite small (so we do not have a good "quiescence" baseline). Next, horizontal banding in the figure reflects commonly provided levels of service.

Most strikingly, we observe frequent instances of very large buffers. Vertical bands reflect common buffer sizes, which we find fall into powers of two, with many sessions exhibiting buffers of 128 KB or 256 KB in size. Even with a relatively fast 8 Mbps uplink, such buffers can easily induce 250 ms of additional latency during file transfers. For a not atypical 1 Mbps uplink, such buffers translate into well over 1 sec queueing delays.

We can leverage the biases in our data to partially validate these results. By examining only Comcast customers, we would naturally expect only one or two buffer sizes to predominate, due to more homogeneous hardware deployments—and indeed the Ruthann figure for just Comcast manifests sizes mainly at 128 KB and 256 KB. In this figure, another more subtle feature stands out with the small cluster that lies along a diagonal. Its presence suggests that a small number of customer have access modems that size their buffers directly in terms of time, rather than memory.

In both plots, the scattered values above 256 KB that lack any particular power-of-two alignment suggest the possible existence of other buffering processes in effect for large UDP transfers. For example, we have observed that some of our notebook wireless connections occasionally experience larger delays during this test apparently because the notebook buffers packets at the wireless interface (perhaps due to use of ARQ) to recover from wireless congestion.

Yet even given noise introduced by other sources, the conclusion is inescapable: over-buffering is endemic in access devices, and they would significantly benefit from dynamically sized buffers that introduce only a fixed delay before dropping packets.

Packet Duplication, Reordering, Outages, and Corruption. The bandwidth tests deliberately stress the network, not only to test the buffer capacity but to induce duplication or reordering. For these tests, the bottleneck point receives 1000 B packets at up to 2x the maximum rate of the bottleneck. Only 1% of the uplink tests exhibited packet duplication, while 16% included some reordering. For downlink tests, 2% exhibited duplication and 33% included reordering. The prevalence of reordering qualitatively matches considerably older results [2]; more direct comparisons are difficult because the inter-packet spacing in our tests varies, and reordering rates fundamentally depend on this spacing.

In addition, the RELEASE data includes the background monitoring process that enables us to check for transient outages. We define an outage as a period with a loss of $\geq$ 3 background test packets (sent at 5 Hz) in a row. We find fairly frequent outages, with 9% of sessions experiencing one or more such events (45% of these reflect single loss bursts, while 28% included ≥ 5 bursts). These burst are generally short, with 48% of sessions with losses having outages $\leq$ 1 sec.

We also find a significant correlation between such bursts and whether the user reported use of a wireless vs. wired network, with 10% of the former sessions exhibiting at least one outage, versus only 5% of the wired sessions.

Finally, analysis of the server-side packet traces finds no instances of TCP or IP checksum errors. We do see UDP checksum errors at an overall rate of about $1.6 \cdot 10^{-5}$, but these are heavily dominated by bursts experienced by just a few systems. The presence of UDP errors but not TCP might suggest use of selective link-layer checksum schemes such as UDP Lite.

5.3 Service Reachability

Table 1 summarizes the prevalence of service reachability for the application ports *Netalyzr* measures. As explained above, for TCP services we can distinguish between blocking (no successful connection), application-aware connectivity (established connection terminated when our server's reply violates the protocol), and proxying (we observe altered requests/responses). For UDP services we cannot in general distinguish the second case due to the lack of explicit connection establishment.

The first four entries likely reflect ISP security policies in terms of limiting exposure to services well-known for vulnerabilities and not significantly used across the wide-area (first three) or to thwart some forms of email spam (SMTP). For this latter, the fraction of blocking in fact appears lower than expected, suggests that many ISPs may employ dynamic blocking for SMTP or other methods to fight bot infections, rather than wholesale blocking of all SMTP.

The prevalence of blocking and termination ("BLOCKED") for FTP, however, likely arises as an artifact of NAT usage: because FTP uses a separate data

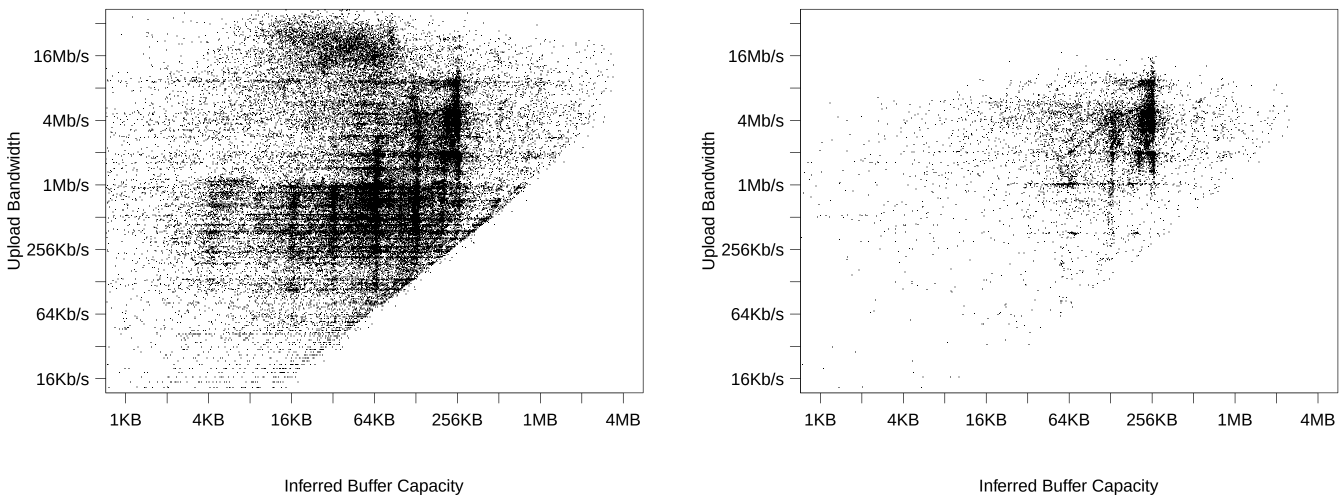

Figure 6: Inferred upload packet-buffer capacity (x-axis) vs. bandwidth (y-axis), for all sessions (left) and Comcast (right).

channel, many NATs implement FTP proxies, which presumably terminate our FTP probing when observing a protocol violation in the response from our server.

Somewhat surprising is the prevalence of blocking for `1434/udp`, used by the Slammer worm of 2003. Likely these blocks reflect legacy countermeasures that have remained in place for years even though Slammer no longer poses a significant threat.

The large fraction of terminated or proxied POP3 connections appears due to in-host anti-virus software that attempts to relay all email requests. In particular, we can identify almost all of the proxying as due to AVG anti-virus because it alters the banner in the POP3 dialog. We expect that the large number of terminated IMAP connections has a similar explanation.

We found the prevalence of terminated SIP connections surprising. Apparently numerous NATs and Firewalls are SIP-aware and take umbrage out our echo server's protocol violation. We learned that this blocking can even occur without the knowledge of the network administrators—a *Netalyzr* run at a large university flagged the blockage, which came as a surprise to the operators, who removed the restriction once we reported it.

Finally, services over TLS (particularly HTTPS, `443/tcp`) are generally unmolested in the network, as expected given the end-to-end security properties that TLS provides. Thus clearly if one wishes to construct a network service resistant to network disruption, tunneling it over over HTTPS should prove effective.

5.4 DNS Measurements

Selected DNS Server Properties. We measured several DNS server properties of interest, including glue policy, IPv6 queries, EDNS, and MTU. Regarding the first, most resolvers behave conservatively, with only 22% accepting any glue records present in the Additional field, and those only doing so for records for subdomains of the authoritative server. Similarly, only 25% accept A records corresponding to CNAMEs contained in the reply. On the other hand, resolvers much more readily (63%) accept glue records when the glue records refer to authoritative nameservers.

We find `0x20` usage scarce amongst resolvers (1.8% of sessions). However, only 4% removed capitalizations from requests, which bodes well for `0x20`'s deployability. Similarly, only a minuscule number of sessions incorrectly cached a 0-TTL record, and none cached a 1 sec TTL record for two seconds.

We quite commonly observe requests for AAAA (IPv6) records (12% of sessions), perhaps largely due to a common Linux default rather than a resolver property, as 38% of sessions with a Linux-related User-Agent requested AAAA records.

The prevalence of EDNS and DNSSEC in requests is significant but not universal, due to BIND's default behavior of requesting DNSSEC data in replies even in the absence of a configured root of trust. 51% of sessions used EDNS-aware DNS resolvers, with 48% of sessions DNSSEC-enabled. Most cases where we observe an advertised MTU show the BIND default of 4096 B (94%), but some other MTUs also occur, notably 512 B (3.4%), 2048 B (1.7%) and 1280 B (0.3%)

The prevalence of DNSSEC-enabled resolvers does not mean transition to broad use of DNSSEC will prove painless, however. For EDNS sessions with an advertised MTU of $\geq$ 1800 B, 14% failed to fetch the large EDNS-enabled reply and 2.1% for the medium-sized one. This finding suggests a common failure where the DNS resolver is connected through a network that either won't carry fragmented UDP traffic or assumes that DNS replies never exceed 1500 B (since `edns_medium` is unlikely to be fragmented). Since

		Interference (%)		
Service	Port	Blocked	Closed	Proxied
NetBIOS	139 T	51.0	1.0	
SMB	445 T	50.3	1.0	
RPC	135 T	46.2	1.2	
SMTP	25 T	25.7	8.2	1.0
FTP	21 T	20.0	3.7	0.1
MSSQL	1434 U	11.1		
SNMP	161 T	7.4	0.2	
BitTorrent	6881 T	6.7	0.5	
AuthSMTP	587 T	6.6	0.2	0.7
SecureIMAP	585 T	6.2	0.2	
Netalyzr Echo	1947 T	6.1		
SIP	5060 T	5.8	4.9	
SecureSMTP	465 T	5.7	0.3	<0.1
PPTP Control	1723 T	5.5	5.1	<0.1
OpenVPN	1194 T	5.3	0.2	
DNS	53 T	5.2	0.8	
IMAP/SSL	993 T	5.1	0.2	<0.1
TOR	9001 T	5.0	0.2	
POP3/SSL	995 T	5.0	0.3	<0.1
IMAP	143 T	5.0	6.7	0.2
POP3	110 T	4.0	7.4	6.1
SSH	22 T	3.6	0.1	<0.1
HTTPS	443 T	2.3	0.4	<0.1
HTTP	80 T		3.8	5.3

Table 1: Reachability for services examined by *Netalyzr*, for all attempted connections. "Blocked" reflects failure to connect to the servers, "Closed" are cases where an in-path proxy or firewall terminated the established connection after the request was sent, likely due to a protocol violation. "Proxied" indicates cases where a proxy revealed its presence through its response, excluding the "closed" cases. Omitted values reflect zero occurrences.

DNSSEC replies will likely exceed 1500 B, the prevalence of this problem suggests a potentially serious deployment issue that will require changes to the resolver logic.

The Release data includes a full validation of DNS MTU up to 4 KB. We find that despite not advertising a large MTU, almost all sessions (95%) used a resolver capable of receiving messages over 512 B. However, a significant number of sessions (16%) exhibited a measured DNS MTU of 1472 B, suggesting an inability to receive fragmented traffic. This even occurred for 11% of sessions that explicitly advertised an explicit EDNS MTU > 1472 B.

A similar problem exists in the clients themselves, but often due to a different cause. When the client directly requests `edns_large`, `edns_medium`, and `edns_small` from the server, 14.5%/4.5%/1.3% failed, respectively. This suggest two additional difficulties: network devices assuming DNS replies do not exceed 512 B (both `edns_large` and `edns_medium` fail) or networks that do not handle EDNS at all (all three fail).[6] We find this high failure rate quite problematic, as the experiences with NXDOMAIN wildcarding and DNS lookups (§ 5.4) clearly demonstrate that DNS resolvers can behave in an adversarial manner. Thus, sound DNSSEC validation requires implementation on the end host's stub resolver to achieve end-to-end security, which requires that end hosts can receive large, EDNS-enabled DNS messages.

Another concern comes from the continued lack of DNS port randomization [6]. This widely publicized vulnerability was over a year old when we first released *Netalyzr*, but 5% of sessions used monotone or fixed ports in DNS requests, Manual examination suggests that these cases mostly reflect small resolvers run by individuals or institutions— no major ISP showed significant problems with this test.

In terms of DNS performance, it appears that DNS resolvers may constitute a bottleneck for many users. 9% of the sessions required 300 ms more time to look up an name within our domain versus the base round-trip time to our server, and 4.6% required more than 600 ms. (We can account for likely at most 100 ms of the increase due to our DNS server residing at a different location than the back-end servers.)

When the user's resolver accepted glue records (53% of sessions), we could directly measure the performance of DNS requests answered from the resolver's cache. Surprisingly, 10% of such sessions required over 200 ms to look up *cached* items, and 3.7% required over 500 ms. Such high latency suggests a considerable distance between the client and the resolver, and for example we found 15% of sessions that used OpenDNS required over 200 ms for cached answers compared to 9% for non-OpenDNS sessions.

Finally, we note that numerous resolvers reflect BIND implementations: 32% of the sessions used resolvers that match a BIND fingerprint in terms of glue policy, CNAME processing, and request options.

NXDOMAIN Wildcarding. We find NXDOMAIN wildcarding quite prevalent among *Netalyzr* users. 28% performing this test found NXDOMAIN wildcarding for `www.`*nonce*`.com`. Even excluding users of both OpenDNS (which wildcards by default) and Comcast (which started wildcarding during the course of our measurements), 21% show NXDOMAIN wildcarding. This wildcarding will disrupt features such as Firefox's address bar, which prepends `www.` onto failed DNS lookups before defaulting to a Google search.

Of further concern is the number of users affected by NXDOMAIN wildcarding that causes broader collateral damage. Excluding Comcast and OpenDNS users[7], 44% of sessions with NXDOMAIN wildcarding also showed wildcarding for non-`www` names. Wildcarding all addresses mistakenly assumes that only web browsers will generate name lookups.

DNS Proxies, NATs, and Firewalls. Many NATs and

[6] We note that the failures we observe could instead be due to heavy packet loss. However, each failure would require five consecutive losses just after a successful non-EDNS query. Furthermore, such failures should not particularly favor one type of query over another, yet we observe only 0.09% of sessions for which `edns_medium` succeeded while `edns_small` failed.

[7] Comcast only wildcards names beginning with `www.`, while the default OpenDNS behavior wildcards all invalid names.

Domain	All Lookups (%) Failed	Blocked	OpenDNS (%) Failed	Changed
www.nationwide.co.uk	2.3	<0.01	1.6	0.01
ad.doubleclick.net	1.5	1.99	1.6	1.27
www.citibank.com	1.3	<0.01	0.9	0.03
windowsupdate.microsoft.com	0.7	0.02	0.5	0.01
www.microsoft.com	0.7	<0.01	0.4	0.01
mail.yahoo.com	0.6	0.02	0.4	0.17
mail.google.com	0.4	0.02	0.3	0.13
www.paypal.com	0.4	0.04	0.1	0.03
www.google.com	0.3	0.01	0.2	76.71
www.meebo.com	0.3	0.03	0.2	0.79

Table 2: Reliability of DNS lookups for 10 selected names (reflecting 107,000 sessions, 11,000 of which used OpenDNS).

firewalls are DNS-aware and may act as DNS proxies. Although we find 99% able to perform direct DNS queries, 11% of these sessions show evidence of a DNS-aware network device, where a non-DNS test message destined for `53/udp` failed (but proper DNS messages succeeded). Far fewer networks contain mandatory DNS proxies, with only 1.2% of DNS-capable sessions indicating such in the form of changed DNS transaction ID.

Although most NATs don't automatically proxy DNS, most contain DNS proxies. We found 67% of the NATs would forward a DNS request to the server (with this measurement restricted to the cases where *Netalyzr* correctly guessed the gateway IP address). Of these, only 1.8% of the sessions contained their own recursive resolver, rather than forwarding the request to a different recursive resolver. Finally, although rare the number of NATs providing open DNS resolution *externally* accessible is still significant. When queried by our server, 4.4% of the NATed sessions forwarded the query to our DNS servers. Such systems can be used both for DNS amplification attacks and to probe the ISP's resolver.

DNS Reliability of Important Names. DNS lookups can fail for a variety of reasons, including an unreliable local network, problems in the DNS resolver infrastructure, and failures in the DNS authorities or paths between the resolver and authority. Table 2 characterizes some failure modes for 10 common domain names. For general lookups, "failure" reflects a negative result or an exception returned to the applet by `InetAddress.getByName()`, or a 20 sec timeout expiring. "Blocked" denotes the return of an obviously invalid address (such as a loopback address).

We explored reliability for OpenDNS users in more detail. OpenDNS not only performs NXDOMAIN wildcarding, but also wildcards SERVFAIL (for the latter, returning the IP address of `hit-servfail.opendns.com`). Thus for queries generated by OpenDNS users we can distinguish between failures which occur between the client and OpenDNS, and server failures due to problems between OpenDNS and the DNS authority for the domain. OpenDNS also includes powerful features to change other names. For `www.google.com`, OpenDNS will act as a proxy by default, redirecting users transparently through an OpenDNS server. For other domains, OpenDNS allows users or domain administrators to block "undesirable" names, with OpenDNS instead returning the address of various blocking servers.

Some behavior immediately stands out. First, regardless of resolver, we observe significant unreliability of DNS to the client, due to packet loss and other issues. Caching also helps, as highly popular names have a failure rate substantially less than that for less common names. For example, compare the failure rate of `www.nationwide.co.uk` to that of `mail.google.com`, for which we presume resolvers will have the latter cached significantly more often.

Second, we observe high reliability for the DNS authorities of the names we tested. Only 14 sessions had OpenDNS returning the SERVFAIL wildcard in response to a legitimate query. (One such session showed many names failing to resolve, obviously due to a problem with OpenDNS's resolver rather than the authority servers.)

Third, we can see the acceptance of DNS as a tool for network management and control. All but the `www.google.com` case for OpenDNS represent user or site-admin configured redirections. For domains like `mail.yahoo`, the common change is to return a private Internet address, most likely configured in the institution's DNS server, while blocking of `ad.doubleclick` commonly uses nonsense addresses (such as `0.0.0.0`), which may reflect resolution directly from the user's hosts file (as suggested on some forum discussions on blocking `ad.doubleclick`).

The DNS results also included two strains of maliciousness. The first concerns an ISP (Wide Open West) that commonly returned their own proxy as an answer for www.google.com or search.yahoo.com (but not sites such as mail.google.com or www.yahoo.com). Deliberately invalid requests to these proxies return a reference to "phishing-warning-site.com", a domain parked with GoDaddy. We also observed similar behavior for customers of sigecom.net, cavtel.net, rcn.net, fuse.net, and o1.com.

Second, in a few dozen sessions we observed malicious DNS resolvers due to malcode having reconfigured an infected user's system settings. These servers exhibit two signatures: (*i*) malicious resolution for windowsupdate.microsoft.com, which instead returns an arbitrary Google server to disable Windows Update, and (*ii*) sometimes a malicious result for ad.doubleclick.net. In these latter (less frequent) instances, these ad servers insert malicious advertisements that replace the normal ads a user sees with ones for items like "ViMax Male Enhancement" [12].

5.5 HTTP Proxying and Caching

8.6% of all sessions show evidence of HTTP proxying. Of these, 32.4% had the browser explicitly configured to use an

HTTP proxy, as the server recorded a different client-side IP address only for HTTP connections made via Java's HTTP API. More interestingly, 90.8% of proxied sessions showed evidence of a mandatory in-path proxy for all HTTP traffic. (These are not mutually exclusive—the overlap is explained by users that are double-proxied.) We detect such proxies by several mechanisms, including changes to headers or expected content, requests from a different IP address, or in-network caching. A proxy may announce its location through `Via` or `X-Cache-Lookup` response headers. The applet follows such clues by attempting a direct connection to such potential proxies with instructions to connect to our back-end server, which succeeded in 11.0% of proxied sessions. The reported names can be net-local hostnames (such as "`CLT-PRXY-04`" or "`Bastion`") or fully qualified domain names. Of the announced proxies, 25.2% used a domain matching that of the client's PTR record.

We rarely observed caching of our 67 KB image (5.3% of sessions cached at least one version of it). Manual examination reveals that such caching most commonly occurred in wireless hotspots and corporate networks. Two South African ISPs used in-path caching throughout, presumably to reduce bandwidth costs and improve latency.

The infrequency of such caches perhaps represents a blessing in disguise, as they often get it wrong. A minor instance concerns the 55.8% of caches that cached the image we specified it as weakly uncacheable (no cache-specific HTTP headers). More problematic are the 37.8% that cached the image despite strong uncacheability (use of headers such as `Cache-control: no-cache,no-store`, a fresh `Last-Modified` timestamp expiring immediately). Finally, 5.3% of these broken caches failed to cache a highly cacheable version of the image (those with `Last-Modified` well in the past and `Expires` well into the future, or with an `ETag` identifier). Considering that 41.5% of all HTTP-proxied connections did not gain the benefits of caching legitimately cacheable content, we identify considerable unrealized savings.

Network proxies seldom transcode the raw images during this test, but it does occur. 0.05% of the sessions showed transcoding of one or more of the fetched images, detected as a returned result smaller than the expected length but $>$ 10 KB. Manual examination of a few cases verified that the applet received a proper HTTP response for the image with a reduced `Content-Length` header, and thus the network did indeed change the image rather than merely truncate.

In-path processes also only rarely interrupt file transfers. Only 0.7% of all sessions failed to correctly fetch the `.mp3` file and 1.0% for the `.exe`. Slightly more, 1.3%, failed to fetch the `.torrent` file, suggesting that some networks filter on file type. However, 10% filtered the EICAR test "virus", suggesting significant deployment of either in-network or host-based AV. As only 0.36% failed to fetch all four test-files, these results do not reflect proxies that block all of the tests.

5.6 ISP Profiles

Table 3 illustrates some of the policies that *Netalyzr* observed for the 15 most common ISPs. As mentioned above, the relative lack of SMTP blocking amongst several major ISPs could reflect that some IPS perform dynamic response to block spam-bots in their network. Likewise, a few ISPs do not appear to filter Windows traffic outbound from customer connections. They might however block these ports inbound, which we cannot determine since *Netalyzr* does not perform inbound scanning. cannot determine if these ports are unblocked on inbound traffic.

Another characteristic we see reflects early design decisions still in place today. Although DSL always offered the ability to provide direct Ethernet connections, many DSL providers initially offered PPPoE connections rather than IP over Ethernet [32]. DOCSIS-based cable-modems, however, always used IP-over-Ethernet. We can see the effects of this transition for Verizon customers, as only 9% of Verizon customers whose reverse name suggests they are FiOS (fiber to the home) customers manifest the PPPoE MTU, while 69% of the others do.

A final trend concerns the growth of NXDOMAIN wildcarding, especially ISPs wildcarding all names rather than just `www` names. During *Netalyzr*'s initial release, Comcast had yet to implement NXDOMAIN wildcarding, but began wildcarding during Fall 2009.

We also confirmed that the observed policies for Comcast match their stated policies. Comcast has publicly stated that they will block outbound traffic on the Windows ports, and may block outbound SMTP with dynamic techniques [7]. When they began widespread deployment of their wildcarding, they also stated that they would only wildcard `www` addresses, but we did observe the results of an early test deployment that wildcarded all addresses for a short period of time.

6. RELATED WORK

There is a substantial existing body of work on approaches for measuring various aspects of the Internet. Here we focus on those related to our study in the nature of the measurements conducted or how data collection occurred.

Network performance. Dischinger et al. studied network-level performance characteristics, including link capacities, latencies, jitter, loss, and packet queue management [9]. They used measurement packet trains similar to ours, but picked the client machines by scanning ISP address ranges for responding hosts, subsequently probing 1,894 such hosts autonomously. In 2002 Saroiu et al. studied similar access link properties as well as P2P-specific aspects of 17,000 Napster file-sharing nodes and 7,000 Gnutella peers [24]. They identified probe targets by crawling the P2P overlays, and identified a large diversity in bandwidth (only 35% of hosts exceeded an upload bandwidth of 100Kb/s,

ISP	Sessions	Country	Blocked (%) Win	Blocked (%) SMTP	Blocked (%) MSSQL	DNS Wildcarding Type	DNS Wildcarding %	PPPoE (%)	Medium
Comcast	13,403	US	99	8		www	33		Cable
RoadRunner	5,544	US				www	63		Cable
Verizon	3,854	US	7	14		www	83	33	DSL/Fiber
SBC	2,938	US	51	73					DSL
Deutsche Telekom	2,523	DE	76			all	48	56	DSL
Cox Cable	2,187	US	92	77	88				Cable
Charter Comm.	1,665	US	95	23	32	all	62		Cable
Qwest	1,334	US	18	6		all	51	70	DSL
BE Un Limited	1,276	UK		49					DSL
Arcor	1,139	DE	33					6	DSL
BellSouth	1,080	US	62	69	96			16	DSL
Alice DSL	1,032	DE	30			www	62	69	DSL
Shaw Cable	1,018	US	6	61					Cable
telecomitalia.it	918		7		14	all	64	65	
Optimum Online	866	US	97	78		www	77		Cable

Table 3: Policies detected for the top 15 ISPs. We indicate blocking when $>$ 5% of sessions manifested outbound filtering, particularly for Windows services (TCP 135/139/445). We infer PPPoE from path MTUs of 1492 B.

8% exceeded 10Mbps, between 8% and 25% used dial-up modems, and at least 30% had more than 3Mb/s downstream bandwidth) and latency (the fastest 20% of hosts exhibited latencies under 70ms, the slowest 20% exceeded 280ms). Maier et al. analyzed residential broadband traffic of a major European ISP [17], finding that round-trip latencies between users and the ISP's border gateway often exceed that between the gateway and the remote destination (due to DSL interleaving), and that most of the observed DSL lines used only a small fraction of the available bandwidth. Various techniques have been developed for measuring bandwidth or latency directly, including Sting [25], IGI [14], YAZ [27], and Iperf [20]. In addition, numerous websites offer throughput tests aimed at home users [31]. We could in principle incorporate the techniques underlying some of these tools into our measurements, but have not at this point in order to keep our main focus on ways in which users have their connectivity restricted or shaped, rather than end-to-end performance.

Network neutrality. Several studies have looked at the degree to which network operators provide different service to different types of traffic. Dischinger et al. provided a downloadable tool enabling users detect whether their ISP imposes restrictions on BitTorrent traffic. They studied 47,000 sessions conducted using the tool, finding that around 8% of the users experienced BitTorrent blocking [10]. Bin Tariq et al. devised NANO, a distributed measurement platform, to detect statistically and policy-agnostically whether a given ISP intentionally or accidentally causes degraded performance for specific classes of service [28]. They evaluate their system in Emulab, using Click configurations to synthesize "ISP" discrimination, and source synthetic traffic from PlanetLab nodes. Beverly et al. leveraged the "referral" feature of Gnutella to conduct TCP port reachability tests from 72,000 unique Gnutella clients, finding that Microsoft's network filesharing ports are frequently blocked, and that email-related ports are more than twice as likely to be blocked as other ports [3]. Reis et al. used JavaScript-based "web tripwires" to detect modifications to HTTP-borne HTML documents [23]. Of the 50,000 unique IP addresses from which users visited their test website, approximately 0.1% experienced content modifications. Weaver et al. examined properties of TCP RST packets observed in the network traffic of four sites [30]. They identified the operational specifics and apparent policy goals underlying a set of reset injection products, including filtering implemented by the "Great Firewall of China". Nottingham provided a cache fidelity test for XMLHttpRequest implementations [21], analyzing a large variety of caching properties including HTTP header values, content validation and freshness, caching freshness, and variant treatment. NetPolice [35] measured traffic differentiation in 18 large ISPs for several popular services in terms of packet loss, using multiple end points inside a given ISP to transmit application-layer traffic to destinations using the same ISP egress points. They found clear indications of preferential treatments for different kinds of service. Finally, subsequent to *Netalyzr*'s release, Huang et al. released a network tester for smartphones to detect hidden proxies and service blocks using methodology inspired by *Netalyzr* [15].

Address fidelity. Casado and Freeman investigated the reliability of using a client's IP address—as seen by a public server—in order to identify the client [5]. Their basic methodology somewhat resembles ours in that they used active web content to record and measure various connection properties, but also differs significantly with regard to the process of users running the measurements. They instrumented several websites to serve an `iframe` "web bug", leading to narrow data collection—users had to coincidentally visit those sites, and remained oblivious to the fact that measurement was occurring opportunistically. They found that 60% of the observed clients reside behind NATs, which typically translated no more than seven clients, while 15% of the clients arrived via HTTP proxies, often originating from a diverse geographical region. Maier et al. [17] found that DHCP-based address reuse is frequent, with 50% of all addresses being assigned at least twice per day. Finally, Beverly and Bauer's Spoofer Project [4] employed a downloadable measurement client to measure the extent to which end systems can spoof IP source addresses. They analyzed an extensive longitudinal dataset, finding that through the period of study a significant minority of clients could perform arbitrary spoofing.

7. SUMMARY

The *Netalyzr* system demonstrates the possibility of developing a browser-based tool that provides detailed diagnostics, discovery, and debugging for end-user network connectivity. Visitors who ran the *Netalyzr* applet conducted 112,000 measurement sessions from 86,000 public IP Addresses. *Netalyzr* both reveals specific problems to individual users and forms the foundation for a broad survey of edge-network behavior. Some systemic problems revealed include difficulties with fragmentation, the unreliability of path MTU discovery, restrictions on DNSSEC deployment, legacy network blocks, frequent over-buffering of access devices, poor DNS performance for many clients, and deliberate manipulations of DNS results. The tool remains in active use and we aim to support it indefinitely as an ongoing service for illuminating edge network neutrality, security, and performance.

8. ACKNOWLEDGEMENTS

We are deeply grateful to the Netalyzr users for enabling this study and to our beta testers for the insightful comments and feedback. We would particularly like to thank Mark Allman, Paul Barford, Scott Bradner, John Brzozowski, Randy Bush, Niels Bakker, Richard Clayton, Chris Cowart, Keith Dawson, Adrian Dimcev, Holger Dreger, Brandon Enright, Kevin Fall, Carrie Gates, Andrei Gurtov, Mark Handley, Theodore Hong, Kelly Kane, Matthew Kogan, Keith Medcalf, Thomas Narten, Michael Ross, Wooter Wijngaards, and Richard Woundy. We thank Amazon.com for supporting our EC2 deployment. This work was supported by the National Science Foundation under grants NSF CNS-0722035, NSF-0433702, and CNS-0905631.

9. REFERENCES

[1] Amazon.com. Amazon Elastic Compute Cloud (Amazon EC2). `http://aws.amazon.com/ec2/`, August 2006.

[2] J. Bennett, C. Partridge, and N. Shectman. Packet reordering is not pathological network behavior. *IEEE/ACM Transactions on Networking (TON)*, 7:789–798, 1999.

[3] R. Beverly, S. Bauer, and A. Berger. The Internet's Not a Big Truck: Toward Quantifying Network Neutrality. In *Proc. PAM*, 2007.

[4] R. Beverly, A. Berger, Y. Hyun, and k claffy. Understanding the efficacy of deployed internet source address validation filtering. In *Proc. IMC*, 2009.

[5] M. Casado and M. Freedman. Peering through the Shroud: The Effect of Edge Opacity on IP-based Client Identification. In *Proc. NSDI*, 2007.

[6] Chad R. Dougherty. CERT Vulnerability Note VU 800113: Multiple DNS implementations vulnerable to cache poisoning, July 2008.

[7] What ports are blocked by Comcast High-Speed Internet? `http://lite.help.comcast.net/content/faq/What-ports-are-blocked-by-Comcast-High-Speed-Internet`.

[8] D. Dagon, M. Antonakakis, P. Vixie, T. Jinmei, and W. Lee. Increased DNS Forgery Resistance Through 0x20-bit Encoding. In *Proc. CCS*, 2008.

[9] M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu. Characterizing Residential Broadband Networks. In *Proc. IMC*, 2007.

[10] M. Dischinger, A. Mislove, A. Haeberlen, and K. Gummadi. Detecting BitTorrent Blocking. In *Proc. IMC*.

[11] R. Erzs and R. Bush. Clarifications to the DNS Specification. RFC 2181, IETF, July 1997.

[12] M. Fauenfelder. How to get rid of Vimax ads. `http://boingboing.net/2009/01/16/how-to-get-rid-of-vi.html`, January 2009.

[13] R. Giobbi. CERT Vulnerability Note VU 435052: Intercepting proxy servers may incorrectly rely on HTTP headers to make connections, February 2009.

[14] N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE Journal on Selected Areas in Communications*, 21(6):879–894, 2003.

[15] J. Huang, Q. Xu, B. Tiwana, and M. Mao. The UMich Smartphone 3G Test. `http://www.eecs.umich.edu/3gtest/`.

[16] C. Kent and J. Mogul. Fragmentation considered harmful. *ACM SIGCOMM Computer Communication Review*, 25(1):87, 1995.

[17] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On dominant characteristics of residential broadband internet traffic. In *Proc. IMC*, 2009.

[18] M. Mathis and J. Heffner. Packetization Layer Path MTU Discovery. RFC 4821, IETF, March 2007.

[19] Mozilla. Effective TLD names. `http://mxr.mozilla.org/mozilla-central/source/netwerk/dns/src/effective_tld_names.dat`.

[20] NLANR/DAST. Iperf. `http://iperf.sourceforge.net/`, March 2008.

[21] M. Nottingham. XMLHttpRequest Caching Tests. `http://www.mnot.net/javascript/xmlhttprequest/cache.html`, December 2008.

[22] V. Paxson. An analysis of using reflectors for distributed denial-of-service attacks. *ACM SIGCOMM*

Computer Communication Review, 31(3):38–47, 2001.
[23] C. Reis, S. Gribble, T. Kohno, and N. Weaver. Detecting In-Flight Page Changes with Web Tripwires. In *Proc. NSDI*, 2008.
[24] S. Saroiu, P. Gummadi, S. Gribble, et al. A measurement study of peer-to-peer file sharing systems. In *Proceedings of Multimedia Computing and Networking*, volume 2002, page 152, 2002.
[25] S. Savage. Sting: a TCP-based network measurement tool. In *Proc. USITS*, 1999.
[26] P. Savola. MTU and Fragmentation Issues with In-the-Network Tunneling. RFC 4459, 2006.
[27] J. Sommers, P. Barford, and W. Willinger. Laboratory-based calibration of available bandwidth estimation tools. *Microprocessors and Microsystems*, 31(4):222–235, 2007.
[28] M. Tariq, M. Motiwala, N. Feamster, and M. Ammar. Detecting network neutrality violations with causal inference. In *Proc. Emerging Networking Experiments and Technologies*, 2009.
[29] P. Vixie. Extension Mechanisms for DNS (EDNS0). RFC 2671, IETF, August 1999.
[30] N. Weaver, R. Sommer, and V. Paxson. Detecting Forged TCP Reset Packets. In *Proc. NDSS*, 2009.
[31] Web resources. Bandwidth/latency measurement sites. `http://www.speedtest.net`, `http://helpme.att.net/dsl/speedtest/`, `http://www2.verizon.net/micro/speedtest/hsi/`, `http://www.dslreports.com/stest`.
[32] Wikipedia. Point-to-point protocol over ethernet. `http://en.wikipedia.org/wiki/Point-to-Point_Protocol_over_Ethernet`, January 2010.
[33] Wikipedia. Usage share of operating systems. `http://en.wikipedia.org/wiki/Usage_share_of_operating_systems`, January 2010.
[34] Wikipedia. Usage share of web browsers. `http://en.wikipedia.org/wiki/Usage_share_of_web_browsers`, January 2010.
[35] Y. Zhang, Z. M. Mao, and M. Zhang. Detecting traffic differentiation in backbone ISPs with NetPolice. In *Proc. IMC*, 2009.

BIA Kelsey

Digital Strategies for Broadcasting

Roku's Channel Store Brings OTT Option to Local TV

By: Jed Williams 1 December 2010

Roku Channel Store

With cable subscriptions declining and a host of nonlinear, IP-connected devices emerging to offer new, incremental revenue possibilities, the time has come for local TV to get serious about an OTT strategy. To capitalize on this opportunity at a market level, local operators should be asking, “what does an effective OTT strategy entail?” and “what steps can be taken today?”

Roku, an OTT set-top box provider with an open application environment and channel store for content developers to build on, is probing these same questions and thinks it may have the answer. Partnerships, curated content, revised workflow and advertising are all part of the solution.

Roku’s channel store already features a growing swath of audio and video content, including Netflix, Hulu Plus, MLB.TV, the NHL, Pandora and Roku Newscaster, which culls together a cross-section of news from national outlets. While sports and entertainment/drama are particularly appealing, Jim Funk, Roku’s VP of business development, told BIA/Kelsey that customers of the channel store have demonstrated “lots of interest in local stations.”

CEO Anthony Wood cautions that these users “don’t watch linear programming,” creating a new challenge for stations to leverage their existing brand strengths while packaging their content to appeal to a new breed of consumer.

Because of Roku’s open network, developers can create content channels that are compatible across other similar platforms (OTT boxes, gaming consoles, Blu-ray players). But participating local stations must curate their content to meet the needs of a digital video audience, which has shown the tendency to snack on smaller samples (e.g., YouTube vignettes) instead of longer-form programming (partially explaining the demise of FLO TV). There is an e-commerce billing service, with a revenue split, if a developer chooses to charge for content.

Once content decisions are made and internal workflow is adjusted to serve the OTT channel, broadcasters must partner with an online video platform (Brightcove, thePlatform and New Lion, to name a few) to upload, encode and deliver their video assets to the to the Roku store. These companies also target mobile and PC screens, so OTT can be part of a multi-screen relationship.

Roku currently draws most of its revenue from box sales, and channel development is free. Wood envisions building ad network and e-commerce capabilities in the near future. As part of the ad network, participating local stations would cede some of their inventory to Roku to monetize. Local advertising options could emerge, too.

MHz Networks provide an early model for locals to follow. In February, the Washington, D.C.-based non-commercial broadcaster introduced its premium international programming in the on-demand Roku channel store at $9.95 per month.

Rights issues must still be solved to make the migration to Roku and other OTT platforms more seamless. If broadcast rights and OTT rights are separate entities, then networks and syndicators may be reticent to approve their content for delivery in local OTT offerings.

BOOKMARK

Tags: local tv, MHz Networks, OTT, Roku

No Comment

Leave your response!

Add your comment below, or trackback from your own site. You can also subscribe to these comments via RSS.

Be nice. Keep it clean. Stay on topic. No spam.

Name (required)

Mail (will not be published) (required)

Website (optional)

You can use these tags:
`<a href="" title=""> <abbr title=""> <acronym title=""> <b> <blockquote cite=""> <cite> <code> <del datetime=""> <em> <i> <q cite=""> <strike> <strong>`

This is a Gravatar-enabled weblog. To get your own globally-recognized-avatar, please register at Gravatar.

Submit Comment

Subscribe to our blog

Enter your e-mail address to receive the Digital Strategies for Broadcasting feed.

SUBSCRIBE

BIA/Kelsey Blogs

- Digital Strategies for Broadcasting
- Global Yellow Pages
- Local Media Watch
- Mobile Local Media

Connect with BIA/Kelsey

Latest Blog Posts

- ILM:10: Pandora … Coming to a Car Near You
- ILM:10: CBS' Ezra Kucharz Envisions 'the Next Decade in Local'
- FM on Cell — Is It Worth the Effort?
- TBD Shelves Community Ad Network After Marginal Results
- Roku's Channel Store Brings OTT Option to Local TV
- Wheelin' and Dealin': Belo Jumps Into Daily Deals

Most Popular Categories

- Digital Strategies for Broadcasting (RSS) (64)
- Featured (RSS) (59)
- television (RSS) (40)
- Radio (RSS) (38)
- Advertising (RSS) (30)
- broadcasters (RSS) (28)
- Digital Media (RSS) (26)
- Local Media (RSS) (21)
- Winning Media Strategies (RSS) (19)
- Internet (RSS) (18)
- The New Radio Model (RSS) (18)
- Mobile/Handheld (RSS) (17)
- General (RSS) (16)
- Headline (RSS) (14)
- video (RSS) (10)

Authors

- BIA Financial Network (21)
- Ben Manilla (2)
- Dan Halyburton (2)
- Ed Czarnecki (1)
- Georgina Santilli (5)
- Gerardo Tabio (1)
- Jed Williams (26)
- Jim Hopes (2)
- John Kelsey (1)
- John Stevens (2)
- John Sutton (1)
- Lori Dillé (2)
- Mark Fratrik (22)
- Mark O'Brien (1)
- Matthew Booth (1)
- Michael Boland (7)
- Michael J. Hackmer (16)
- Michael Taylor (3)
- Mike Henry (2)
- Mackenzie Lovings (7)
- Neal Polachek (1)
- Omar Wardak (1)
- Paul Marszalek (1)
- Patrick Buono (1)
- Peter Krasilovsky (23)
- Richard Warner (1)
- Rick Ducey (58)
- Skip Pizzi (2)
- Steve Marshall (4)
- Steve Passwaiter (8)
- Thomas Buono (1)
- William Redpath (1)

Recent Comments

- Jed Williams on Google TV Ads Extends Reach With DirecTV Partnership — Will SMBs Follow?
- Rick Ducey on Over The Top (OTT) Homes Top Cable by 2013
- Dave on Google TV Ads Extends Reach With DirecTV Partnership — Will SMBs Follow?
- Bruce Roberson on Local TV and the FCC's National Broadband Plan

- Groundswell on Over The Top (OTT) Homes Top Cable by 2013

Most Commented

- Over The Top (OTT) Homes Top Cable by 2013
- Google TV Ads Extends Reach With DirecTV Partnership -- Will SMBs Follow?
- Spectrum Values in the 2.5 GHz Band
- Auction 73 Results - 700MHz Spectrum
- An Insider's View of NAB 2008

Popular Terms

FEMA Microsoft Steve Passwaiter Georgina Santilli digital media local advertising Digital Media Mark Fratrik Facebook Arbitron ActiveAccess TV FCC MEDIA Access Pro DTV Kelsey Group mobile Advertising ATSC BIA Google online Internet NAB Rick Ducey sales AOL iPhone multiplatform NAB Radio Show Common Alerting Protocol CBS Radio Thomas Buono BIA Financial Network CAP BIA Advisory Services broadcasters newspapers Digital TV television HD radio Twitter EAS Pandora OTT NBC EBS OMVC

Contact Us | Log in | Digital Strategies for Broadcasting (RSS) | Comments (RSS) | Subscribe

- Home
- New Shows
- Genres

- Posted by Sonya in on 09 9th, 2010 | No Comments
  Celeste Bright: Revelations (S01:E07)
  PLAY VIDEO
- Posted by Sonya in on 09 2nd, 2010 | No Comments
  Celeste Bright: S.O.L. (S01:E06)
  PLAY VIDEO
- Posted by Sonya in on 08 31st, 2010 | No Comments
  Celeste Bright: Fear No Evil (S01:E05)
  PLAY VIDEO
- Posted by Sonya in on 08 24th, 2010 | No Comments
  Celeste Bright: True Colors (S01:E04)
  PLAY VIDEO
- Posted by DominionSeries in on 08 19th, 2010 | No Comments
  Dominion: Stranger Danger (S01:M01)
  PLAY VIDEO
- Posted by Sonya in on 08 17th, 2010 | No Comments
  Celeste Bright: The Company She Keeps (S01:E03)
  PLAY VIDEO
- Posted by Sonya in on 08 3rd, 2010 | No Comments
  Celeste Bright: T.M.I. (S01:E02)
  PLAY VIDEO

- About Rowdy Orbit

- About Rowdy Orbit
- |
- Terms of Service
- Register
- Log in
- |
-
- SHARE

- Home
- In The News
-
- About
- Actors
- Contact
- Crew
- Gallery
- Music
- Press Kit
- The Map

Share |

- Find Webisodes
  - Featured
  - In The News
    - Buzz
    - Press
  - Uncategorized
  - Video
    - Behind the Scenes
    - Season 1
    - Season 2
      - with Subtitles/con Subtítulos

About

YLSE

“A Web Series for our Generación”

YLSE is a not-so politically correct dramedy about a modern Americana and the quirky characters that make up her life. We follow ambitious, single, thirty-something Ylse as she juggles career, a not-so-successful love life and a family who doesn’t understand her progressive American ways…all this, on her quest to achieve her version of the American dream; becoming just like her idol Oprah. Think: Bridget Jones with a bicultural twist.

ABOUT **YLSE**:

-We have **Two successful seasons** already under our belt.

-Recipient of **TWO NEW MEDIA AWARDS**:

-2010 IMAGEN AWARDS “Best Internet Program” Winner.

-2009 RASQUACHE Film festival “Audience Favorite” Winner.

-We are on the ground floor of New Media content creation that targets the extremely desirable, under-explored and growing Hispanic market.

-SEASON 3 is in-development

WHAT our FANS ARE SAYING:

“…As a Latina it just makes me so proud to know that something like Ylse not only exist but is possible. You’ve inspired me to write my own story about finding and using my voice”. – Icess Fernandez A fan/ a blogger

“ I hugely enjoyed Ylse yesterday, then read an interview with you, all the while marveling at how it is possible to come from opposite sides of the world and yet to end up so alike. I am a Russian-Jewish immigrant born in the former USSR (Ukraine). My family arrived in the US when I was eight (with nothing, of course, as was your case)….I now live in France with my family.“ –Julia Kogan

“…watch webisodes of *Ylse*, about a smart, passionate Mexican-American journalist who prays to Oprah ….Then you go pray to whatever virgencita you believe in (or Oprah) that the shows continue or end up on some premium cable channel”. -Bilingual in the Boonies, a fan/a blogger (found Ylse on Latino Future Magazine)

WHO is our CORE TARGET MARKET?

*Source MarketResearch.com/emarketer July 1, 2008 **Source: Public Opinion Study of CA Latina electorate/HOPE 2009

*By 2012, nearly 30 million US Hispanics will be online. (More than 50% of the US Hispanic population)

*Many of them will be native-born Hispanics rather than immigrants.

*This group of young Latinos will be bilingual and proud of its Hispanic origins, and marketers cannot assume that mass-market, English-language campaigns will reach them effectively.

*Multiple executions and strategies may be necessities for marketers trying to reach this growing population.

*Bilingual Powerhouse…more fierce competition by bilingual sites, they only stand to gain by offering more bilingual and bicultural content.

**English dominant Latinas are much more likely to be connected to “new media” than Spanish dominant Latinas (the worldwide reach of our show maintains the importance of our Spanish language component)

** There is an important opportunity for corporations, government institutions, political parties and candidates to increase their market share or their electoral appeal by directly communicating through “new media” with those segments of the Latina population that are more culturally assimilated and with those higher up on the socio-economic scale.

** Other factors that impact access to the Internet are age, educational level, and presence of children at home. A substantial majority of Latina voters between the ages of 18-34, those with a college or graduate degree, and mothers with children living at home are connected to “new media

Ylse Fans

YLSE on Facebook

Like

150 people like YLSE

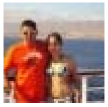 Elizabeth
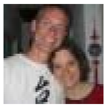 Diana
 Antonio
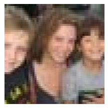 Brenda
 D En

 Pechan
 Patty
 Manny
 Heiko
 Alan

 Carlos
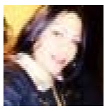 Cynthia
 Blanca
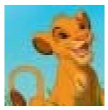 Sep


Syed Muham

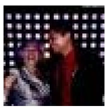 Jorge-Luis
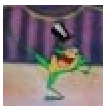 Richard
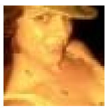 Claudia
 Elena
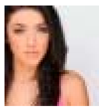 Yvette

 Blanca Elodia
 MAynor
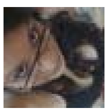 Nathaly Karir
 Jose
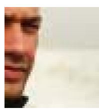 David

 Jose
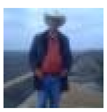 Juan
 Javier
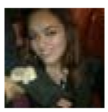 Karla
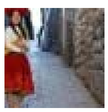 Luz

Facebook social plugin

FEDERAL COMMUNICATIONS COMMISSION

OPEN INTERNET WORKSHOPS

SPEECH, DEMOCRATIC ENGAGEMENT, AND THE OPEN INTERNET

Washington, D.C.

Tuesday, December 15, 2009

ANDERSON COURT REPORTING
706 Duke Street, Suite 100
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190
www.andersonreporting.net

PARTICIPANTS:

Introduction of Workshop and Speakers:

STUART BENJAMIN
Distinguished Scholar in Residence, FCC Office of Strategic Planning

Panelist Presentations:

MICHELE COMBS
Vice President of Communications, Christian Coalition of America

GLENN REYNOLDS
Blogger, Instapundit; Beauchamp Brogan Distinguished Professor, University of Tennessee College of Law

JONATHAN MOORE
Founder & CEO, Rowdy Orbit IPTV, LLC

RUTH LIVIER
Actress, Writer and Producer, YLSE.

GARLIN GILCHRIST
Director of New Media, Center for Community Change.

BOB CORN-REVERE
Partner, Davis Wright Tremaine LLP

JACK BALKIN
Knight Professor of Constitutional Law and the First Amendment, Yale Law School

ANDREW SCHWARTZMAN
President and CEO, Media Access Project

* * * * *

P R O C E E D I N G S

(1:05 p.m.)

MR. BENJAMIN: Welcome to the first of these workshops on speech democracy -- sorry, speech, democratic engagement, and the open Internet.

I'm Stuart Benjamin. This is actually my first day here as the distinguished scholar-in-residence. And this is the first of three workshops that are going to be examining the significance of the Internet's openness for different interests relevant to the open Internet proceeding. So, this one is obviously on speech and democratic engagement. There will also be one on innovation and investment. And then finally, there will be one on consumers and transparency.

After that, the Commission will then use information gathered from these workshops and other sources to then have three more workshops that will look more specifically at how rules might be drafted. So, this is the beginning of a fairly long process.

But the true beginning of this event is statements. We have, actually, three Commissioners here: Commissioners Capps, McDowell, and Clyburn. So, we'll begin with statements from them in that order.

So, Commissioner -- and none of these Commissioners need any introduction, I assume. So, Commissioner Copps, please take it away.

MR. COPPS: Thank you and welcome to Stuart and your first day at the FCC. This is actually my first minute at the FCC today because Mignon and I just got back from Memphis a couple of minutes ago. And we were out there for a hearing on digital inclusion last night, which went very well and gave us a lot of the good kind of information we want as we set about doing a broadband plan.

I am very happy to see everybody here and the workshop for speech, democracy, and the open Internet. This dialogue is part of a really open and transparent process that Chairman Genachowski has put in place, probably the most

open, transparent, and extensive process that I have seen in my going on nine years at the FCC now.

And I think the Commission is on the move. I supported Chairman Genachowski's decision to move forward with a notice of proposed rulemaking on open Internet principles and to add two more to the four Internet principles that we adopted in 2005, and to make all six of them rules of the Commission.

Those actions taken together, 2005 and 2009, I think are sober, clear-eyed recognition that the Internet must never be about powerful gatekeepers, never be about walled gardens. It must always be about the smoothest flow of communications among people. Such speech should not be stifled.

We have confronted similar tensions before between the concerns of network operators and the interests of citizens in communicating freely. For as long as this Commission has existed, entrenched powerful network operators have argued that harm will inevitably result from pro- consumer decision.

In the 1950s and '60s, the government was told that the entire phone network would be crashing down and compromised if innovations like Hush-a-Phone and Cart-a-Phone were attached to the end of the telephone line. In the early 1980s, the Department of Justice was told that breaking up Ma Bell would leave the United States literally unable to respond to a nuclear threat.

In recent years we were told that forcing telecom carriers to accept enforceable network neutrality rules would jeopardize their financial future, as they were consolidating. And in 2007, we were told that wireless carriers couldn't make an open access model work until these very same carriers then changed their mind and came out in favor of such a model.

I recount this brief history just to remind you that we need to proceed thoughtfully and with a healthy dose of skepticism.

As we work to deliver a worthy national

broadband plan, I'm excited about the potential of broadband for the citizens of this country. Broadband intersects with just about every great challenge confronting our nation, and that's really the interesting part. I think that this exercise that Julius has started here is finally starting to permeate the country.

People understand we're not just a bunch of broadband geeks talking technical language at the FCC, but that this goes to every significant challenge facing the United States of America today, whether it's the creation of jobs, expanding educational opportunity, health care, energy dependence, the degradation of the environment, opening the doors of equal opportunity. None of those challenges can be solved unless a component of the solution is broadband.

In other words, broadband is pertinent to the resolution of every single national challenge facing the United States of America and every one of our citizens must have access to this enabling technology to participate fully, meaningfully, productively in American life.

The genius of the Internet is its openness and its dynamism and its availability to one and all. And that's why I have believed for a long time that the FCC needs to play a proactive role in preserving the Internet as a vibrant place for democratic values, innovation, and economic growth.

A solid democracy in the future's going to depend on broad pipes, good private sector vision, and thoughtful public policy to make sure that everyone has access to the information infrastructure they need to exercise their citizen rights and responsibilities. History teaches us that if a special interest has a technical capacity and the financial incentive to interfere, there're going to be some bad apples who are going to try. That's why we are moving to network neutrality. Given what's at stake, we need hard and fast rules, not just idyllic principles and an honesty system arrangement to keep them from doing so.

There are founded and unfounded fears in this debate, to be sure. And we need to have all the facts while considering the future ramifications for this powerful tool. I don't believe that the importance of our open Internet proceeding can be overstated. It's about safeguarding America's broadband users, whether -- wherever they are and however they choose to access the Internet.

So, with that in mind, those are my thoughts as we go into thinking about network neutrality.

One more thought before I give the platform to my colleagues. There's another dimension -- it's related to mention, but just -- we need to build, I think, a little bit on where we are right now, and that's what I want to emphasize today.

Broadband, we've already seen, has so many effects on various aspects of our life. This other dimension we haven't failed to -- we have failed to talk enough about is really what broadband can do, not just for our economy, and our competitiveness, but for citizen engagement, for our democratic dialogue, for the civic give-and-take we need to underpin our democracy. So, broadband is more than just a technology, again, more than just something that you can quantify in terms of jobs and other things like that. It goes to this intangible heart of what democracy is all about.

So, this is what I want to see us emphasize in our broadband plan. And I don't think we can come up with a roadmap for what's the media of the 21st century going to look like in the next 64 days. But I want our broadband plan to evidence an understanding that civic engagement and these intangible values that I talked about, go to the heart and soul of that plan. And if it's going to be a viable plan, it has to address them. And I want that broadband plan to demonstrate that we have a commitment here at the FCC to do that, and I think we do.

I'm delighted with the folks the Chairman has brought to the Commission. I know Steve Wallman's going to be working on media, and Stuart, a lot of other people here. So, this is good. But it's also something that can't wait.

You know, this broadband and the implications of this infrastructure for our democracy goes back to -- what we're trying to do goes back to the beginning of the country. I'm reminded of Thomas Jefferson. You all remember that famous quote that he had about whether he would prefer a government without newspapers or newspapers without government, and we all know that he ended up saying, well, newspapers without government.

But our folks at -- friends at Free Press, with their usual diligence, dug up the rest of Thomas Jefferson's quote. It's a fascinating thing, because right after he said what I just said, he said: But I should mean that every man should receive those papers and be capable of reading them.

But I should mean that every man should receive those papers and be capable of reading them. Isn't that something? He's talking about deployment, getting newspapers -- which was the information infrastructure of his time -- out to every American. Out, ubiquitously. And he's talking about adoption. People knowing how to read, people having access to those newspapers through enlightened public policy, like postal subsidies.

So, you know, the technology may be new, some of the jargon we're going to be using in this debate is new, but the challenge is as old as our country and it's a democratic challenge and it's always with us. It's with us in Jefferson's time, it's with us now, and 50 years. And we've got some new whiz-bang technology that anyone can even think of will be trying to think, how do you fit that in to the heart and soul, to the fabric of our country.

So, I am grateful that you are all here. I'm grateful you're talking about this, and I'm looking forward to your comments.

MR. BENJAMIN: Mr. McDowell.

MR. McDOWELL: Thank you, Stuart and Commissioner Copps. It's great to hear you quoting Thomas Jefferson. I want to hear more of that. It's great. Great Virginian, and pretty soon, not only that --

SPEAKER: (inaudible)

MR. McDOWELL: Not only that, but I'm going to try to make you a Blue Devils fan, too.

So, on that note, Stuart, it's great to have you here. I've gotten to know Stuart Benjamin a little bit over the years. He comes to us from my alma mater, Duke University. And so it's terrific to continue this tradition of Duke lending us some of their distinguished professors.

Of course, we've had two chief economists in the past couple of years, as well. One from the Fuqua School and then the other from the Undergrad School as well; Michele Connelly, who came back as chief economist.

But, in any case, welcome, and you've

already started making controversies, which is fantastic. You weren't even on the job yet when you became controversial here at the FCC, so that's the sign of great progress. And keep pushing the envelope and stimulating the debate and I look forward to working with you going forward.

And I'm also pleased to be joining this distinguished panel of diverse views that we have here today, and not the least of whom are my distinguished colleagues, Mike Copps and Mignon Clyburn.

But I really want to thank the panelists in particular for taking the time to prepare their thoughtful comments about their experiences and their perspectives. And we'll be hearing today about many innovative uses of the Internet to broadly disseminate information, share viewpoints, and organize citizens to participate in the political process -- whether at the local, regional, or national level -- and from any and all political philosophies.

And I actually had this written down. This was a -- already in my remarks, but you just repeated it as well. But, as you said last week, and just five seconds ago, you hailed the genius of the Internet for its openness, its dynamism, its availability to one and all. And I certainly agree with that sentiment. In fact, I think it's safe to say that we all agree that the freedom of expression made possible by the Internet is stirring, even a bit stunning, if you're old enough to remember life before the Net, as I am.

Nowhere is this truer than in the context of political speech, the fundamental lifeblood of our democracy. Now, political speech takes many different forms, from candidates giving speeches and running broadcast ads to concerned citizens coalescing around a cause and organizing themselves through the Internet.

The number of voices and perspectives accessible on the Net are literally uncountable. At no time in American history have our fellow citizens had more communications power at their fingertips. Yet, the political component of the Internet success story rests, to some degree, on choices made by communications policy makers over the 15 years since the Net was privatized. At that historic moment, policymakers decided on a bipartisan basis to allow the Internet to grow largely unfettered by government mandates. They chose to rely instead on a cooperative model of decision making and self-governance by a large number of interested parties -- private and public -- working together in collaborative forums, associations, and other non- governmental bodies. That model gave birth to the most innovative and adaptive communications tool that the world has known, and the fastest-penetrating technology human beings have ever created.

So, while I'm pleased that we will gain a deeper understanding today of the Internet's effect on our national civic life, it is also appropriate to discuss the legal and policy implications of any potential expansion of the government's role in regulating this medium of expression. It probably comes as no surprise that I have some serious concerns from both a Constitutional and policy standpoint about placing more regulatory mandates on the Internet.

My legal review begins, as always, with the wording of the Constitutional or statutory provisions at issue. The plain language of the First Amendment, of course, is a check on government power to burden speech and not on a private party's choice to speak or stay silent. Efforts to advance First Amendment values through additional government regulation risks turning over 200 yeas of First Amendment jurisprudence on its head. As a result, new rules likely will receive heightened scrutiny from reviewing courts and, based on past precedent, would seem particularly vulnerable when the regulated entities operate in a competitive environment that leave open other avenues for disseminating speech.

In addition to Constitutional considerations, however, I have some plain old practical worries, one of which is the real prospect of the government getting new regulation wrong, no matter how well-intentioned its actions might be. The risks of getting it wrong are high in a fast- evolving arena like the realm of the Internet, where consumer demands and business plans evolve on a near daily basis, if not hourly. And the consequences of getting it wrong are much more significant when it is the government that misjudges a changing marketplace rather than a private sector player competing against other and possibly savvier business rivals.

When a company guesses wrong and builds a business model that fails, the failure hurts the company's employees and investors, but rarely does that failure impair a competitive marketplace because consumers have the power to chose among other options. When the government guesses wrong, however, and imposes a regulatory regime based on unsubstantiated fears about the future, it can distort the development and deployment of new services by all providers for years to come. In short, what is frequently considered to be market failure is really the result of regulatory failure. Until such harmful policies become reversed, consumers lose out.

Concerns about well-meaning government misjudgment are not limited to my side of the political aisle. I recently ran across a quip from the economist John Kenneth Galbraith, a man who wasn't shy about advocating for federal intervention on occasion. Yet frustration with one government agency prompted even Galbraith to say, "You will find that the State Department is the kind of organization which, though it does big things badly, does small things badly, too."

As proud as I am of the FCC and its people, I don't think we are smarter than the collective wisdom of the millions of Americans or billions of people worldwide who comprise the private sector.

Producing ineffective regulation is one thing, but issuing harmful regulation would be, of course, irresponsible. The latter would be especially troubling when the service or product at issue concerns speech, including efforts that would burden one party's ability to speak in order to advance the speech of others. Sometimes the effect of the government's past failed efforts in this area is mostly an economic one. Consider the time and resources wasted in the 1990s on regulatory constructs such as the FCC's video dial tone rulemaking and its statutory successor known as open video systems, or OVS. Both were designed to bring new competing multichannel video services into being on telephone networks. Video dial tone would have inserted one independent multichannel provider onto the telco platform, but barred the network operator from having any control over the content.

OVS, on the other hand, would allow telcos to offer their own and slightly deregulated, multichannel programming service, but only if they reserved two-thirds of their system capacity for one or more independent multichannel providers. Neither regulatory scheme delivered any appreciable new service options to consumers.

For those few companies who tried to make the concepts work, how many resources were diverted away from productive uses that would have responded to actual consumer demand?

Other failed government efforts have imposed a more insidious burden by potentially depriving consumers of the ability to receive speech that they might highly value. One example in this regard is the Commission's old prime time access rule, or PTAR. That restraint, first imposed in 1970, effectively barred the 3 TV broadcast networks -- and yes, for the young folks listening, at one point there were only 3 TV channels -- but it barred them from supplying program to affiliated stations for the first 30 minutes of the most heavily viewed evening hours. The idea was that the freed up time would give rise to a wealth of new and better programming, more specifically targeted at local viewers' needs and desires. And what new and improved programming did the government's rule produce? A lot of syndicated game shows. The law of

unintended consequences prevailed, once again.

The FCC finally gave up on that rule in the mid- 1990s when it became clear that the broadcast network bottleneck, as it was known back then, that argument was no longer credible in light of competition provided by cable and other sources.

So, some say that it's too early to judge the Constitutional and policy implications of any Net neutrality regulations on speech. I think the better approach is to start considering the potential ramifications now, in an effort to avoid repeating old mistakes. But in the spirit of the First Amendment, let the debate begin.

So, thank you to all of you for coming here today and for this chance to speak with you. I look forward to reviewing the byproduct of this panel today.

Thank you.

MR. BENJAMIN: Commissioner Clyburn --

SPEAKER: I was going to go next.

MR. BENJAMIN: Oh, I'm sorry. Go.

SPEAKER: It is my pleasure to introduce the right honorable gentle lady from South Carolina, the former Chair of the Public Service Commission of the Palmetto State, The Honorable Mignon Clyburn.

Ms. CLYBURN: Thank you, Commissioner. He's just getting me back for something I -- yeah, something I did in the past.

So, I'd like to thank my colleagues. Welcome, Stuart, to the FCC. It's going to be a lot of fun for you. It's already starting out that way, I hear.

Good afternoon, everyone. I believe there is no better way to begin the Commission's workshop for our open Internet proceeding than with a focus on an Internet's role in enabling speech, journalism, culture, and democracy.

The Internet has had an incredible impact on each of these elements, and our role is to preserve and enhance those things that enable greater democratic participation, and a better informed citizenry. One of the core reasons we've begun this proceeding is to ensure that the Internet remains an unbiased platform where all speakers, including historically underrepresented voices, can reach all audiences on the same footing and with extremely low barriers to entry. This is an incredible development.

As many of you know, for years I owned and operated a weekly newspaper in Charleston, South Carolina, where I learned firsthand the challenges that face individuals and small businesses trying to compete in the traditional media marketplace. These challenges were particularly acute in light of the difficulty of competing against media outlets with established distribution networks. Today, thanks to an open Internet, a community newspaper like the Costal Times or a budding journalist essentially has the same distribution network as the Washington Post or, dare I say it, the Memphis Commercial Appeal. I just got back this morning.

Just think about what this opportunity affords all Americans. I guess I should say, all connected Americans. But that is another topic for another day.

The Internet's openness is also particularly important for those minority voices, which have traditionally encountered a whole host of barriers to reaching audiences through traditional media.

How long have we wrangled with the problem of minority media ownership? This struggle is, in large part, due to the difficulty it would be for minority media owners to have meaningful access to their target audiences. And while that struggle is important and ongoing, and is likely to continue for some time, when it comes to the Internet, the opportunity is here and now. That is, as long as the Internet remains an open platform.

The open Internet is fundamentally different than how broadcast TV or radio and cable television have developed, historically. There are no gatekeepers on the open Internet that determines who gets to speak, and what they get to speak about. If anything, it is a unique opportunity for the First Amendment to truly flourish.

We have here with us today an impressive and diverse group of panelists who have come together today from across the country, including, I see, from West Virginia, California, Tennessee, and -- yes -- South Carolina, to share their personal experiences with the open Internet. I am looking forward to our panelists' presentations on how the Internet's openness has affected their ability to do the important work they do, whether in journalism, political organizing, or cultural creation.

Our open Internet proceeding has just begun, and we will begin discussing these issues -- we will be discussing these issues for some time. But today's workshop will be an essential part of the foundation for our future work.

Thanks to the staff for their terrific work, and to the panelists for traveling here today, and for those of you joining us here, and online, for taking the time to be here today.

Thank you very much.

MR. BENJAMIN: Thanks to Commissioners Copps, McDowell, and Clyburn.

Just as a reminder, again, this is the first of three workshops that are going to be focusing on issues related to the open Internet. And, just by way of background, the ideas is that these workshops are going to be part of the information gathering process for the Commission in devising the approach that it has toward the open Internet. And Chairman Genachowski and the Commissioners are really committed to this as a data-driven and fact-based process.

So, what do we mean by that? In this, as in any debate, you can start with preconceived notions and let those guide you. But the Commissioners think that that's a mistake.

Of course, people have different perspectives, but that's no excuse for blinders. Of course, the questions relevant to the open Internet have different policy implications, and your policy orientation will influence that. But you should start with the data, you should start with the facts and have the facts and the data really drive the process.

That's the idea behind this workshop and the other workshops and other information gathering that the Commission will be engaged in.

So, this is the beginning of a process. It's one that will develop over time where we're going to be trying to get information into the Commission the best information we can, and have that really drive things.

Now, the question for today is, what are the effects of the Internet's openness, on speech, civic participation, democratic engagement, journalism, and cultural creativity. So, the panelists are going to talk about the relationship between the Internet's openness and democratic speech and participation, including its impact on things like citizen journalism as well as blogging, political organizing, and cultural expression, discussing the costs and the benefits of the Internet's historic openness.

The panelists are also going to discuss how the Internet differs from other communications media as a platform for their uses. Each panelist is going to talk for five minutes. Our clock, I think, is working. So, I won't have to give you the hook, but if I give you this sign, that means that you've not been paying attention and the time has run out.

Those of you in the audience, if you have questions, Zach right there will have index cards. Please write down questions on index cards, and you can pass them forward. Those of you on the web, you can send -- you can e-mail any questions to new media at FCC.gov. Those of you in the audience, you could also e-mail your questions -- if your handwriting is really terrible -- and then I will aggregate those questions and I will put them to the panelists as time permits. The hope is, with only five minutes per panelist, we'll actually have some time for some back and forth.

So, with that, let me introduce the first speaker to my immediate left, Michele -- I'm sorry -- Michele Combs?

MS. COMBS: Combs.

MR. BENJAMIN: Combs. I have already met her and I've already forgotten. So, she began her role in politics as an early age when she became the State Chair of the South Carolina Young Republicans. And under the guidance of the late Lee Atwater, a fellow South Carolinian -- Commissioner Clyburn will want to know -- who was the campaign manager for George H.W. Bush and former RNC Chair, she managed to take the chapter to one of the best of the country and was elected as Young Republican of the Year in 1989.

She was then chosen to head up America 2000, Educational Service Corp. in the early 1990s. She brought together children -- corporations and children from all over the state, and raised thousands of dollars to better the education system in the state. She also produced corporate and political events all over the country. And then, in 2002, she became the director of special projects for the Christian Coalition of America. And in 2003, she became the director of communication, and over the years she's worked for Republicans on the state and -- local, state, and national levels.

And so, without any further ado, Michele.

MS. COMBS: Thank you. Good afternoon, Commissioners and staff members of the FCC. Thank you for inviting my organization to participate at this important workshop.

Use of the Internet has allowed the Christian Coalition to amplify the voices of millions of hardworking, pro-family Americans in a way that has revolutionized their ability to be heard and to engage in the political process. The Christian Coalition is excited about its ability to use the Internet to reach millions of our members worldwide. Today our members can follow the Christian Coalition through Twitter tweets, Facebook updates, YouTube messages, frequent daily blog posts, and our e-mail blasts. And users can use Digg to bookmark our content to friends.

Last year during the presidential election, thousands of churches were able to download our presidential and congressional voter guides. And in the last two cycles, we've distributed over 100 million voter guides. The Internet has made these distributions easy and inexpensive. As you can see, increasingly our members are using the Internet for social media tools to connect with the Christian Coalition.

In fact, right now, our homepage features a link for our Twitter users to engage in the health care debate in Congress. On this page, we encourage our members to do five things that only the Internet allows us to do: Sign an online petition, tweet a link to our action center, tweet a message to their members of Congress, tweet messages to their health care friends about this debate, and tweet our Christian Coalition home page to their friends.

Our web page also links to our YouTube,

where the Christian Coalition maintains a Christian Coalition channel, where our members can watch and listen to discussions about policy debates we are involved in. And, increasingly, our site incorporates embedded videos which makes our site more informative and user-friendly.

The mobile world enables the Christian Coalition to reach our members directly without a big media filter on real-time basis through text messaging and tweets. This is extremely important to help our members engage in the political process, especially when there is a time-sensitive vote that is happening in Congress and state legislature.

For example, in 2005, an amendment was introduced in the Bankruptcy Reform Act that would deny bankruptcy protection for those that protest at abortion clinics. Once we became aware of this amendment, we immediately sent out hundreds of thousands of e-mails to our online members, describing the amendment and encouraging them to contact their representatives immediately.

Through the open Internet, we were able to quickly and effectively communicate information pertinent to our 3 million members and they in turn were able to swiftly relay their opinions to those who represent them in Congress. Thanks in part to our active members, we were able to successfully defeat this amendment.

The Christian Coalition's use of the Internet is to help us and our members engage in the political process is what we believe that preserving and protecting an open Internet is so important. We believe that organizations such as the Christian Coalition should be able to continue to use the Internet to communicate with our members worldwide without a phone or cable company snooping in our communication and deciding whether to allow particular communication to proceed, slow it down, block it, or offer to speed it up only if the author pays to be on the fast lane. Unfortunately, in the last couple years, we have seen network operators block political speech, block content, and block the most popular applications on the Internet. As you know, one cable company was discovered to have blocked consumers' ability to the King James Bible.

Unfortunately, however, vendors of deep-pocket inspection technology are actively marketing to allow network operators to inspect the content of my communications and decide whether to slow down or block the speech, or perhaps put my speech on the fast lane if my organization pays the broadband provider an additional tax to have access to their fast lane that connects to the Internet. I am concerned about this level invasion, and to the privacy of the members of my organization, and its burden on our freedom in association and speech. If you care about grassroots democracy, I respectively suggest it ought to concern you as well.

A nonprofit family organization like the Christian Coalition depends on an open Internet to allow us to compete against big media companies and deep-pocketed political organizations to reach our members and allow our members to contact their representatives. The open Internet levels the playing field for organizations like ours, and we have seen how the Internet has allowed the Christian Coalition to amplify its voice to be heard on important policy debates.

The Christian Coalition does not seek burdensome regulations. We generally believe that less government is better than more government. However, we support a free market of ideas and commerce on the Internet.

However, this time, we need very limited rules of the road to protect the new public cyber square. Any threat to the ability of organizations and groups to reach the American public at a very low cost -- without permission -- is simply unacceptable and strikes at the heart of citizen participation and well functioned democracy in the 21st century. Consequently, we support the Chairman's proposal to codify some basic lightweight rules that would preserve the qualities that have made the Internet the most market-driven tool for communication, speech, and innovation ever invented.

Increasingly, faith-based groups are turning to the Internet to promote their political rights and to engage in what Ronald Regan called the hard work of freedom. The on ramps to activism must be kept open and accessible for citizen activists across the political spectrum so we can continue to do that work. And our time.

MR. BENJAMIN: Thank you very much. I think we have next Glenn Reynolds. Glenn, are you there?

While we're booting up Glenn, he is the excellently named Beauchamp Brogan, distinguished professor of law at the University of Tennessee College of Law, where he teaches constitutional law, administrative law, and Internet law. He's a graduate of Yale Law School, where he was a founder of the Yale Law and Policy -- the Yale Law and Technology Association, and he writes regularly for Popular Mechanics, the Wall Street Journal, the Washington Examiner, Forbes, and many other publications. He's probably best-known for his blog at instapundit.com. It also has a regular show on PJTV.com. He previously blogged for FOX News and MSNBC.

Glenn, I can -- we can see you, so take it away. Uh-oh. Wait, hold on. Hold that thought, we're not getting your audio. Can you -- try talking now. Hmm. We're -- seem to be experiencing technical difficulties. We will go on. Glenn, with your indulgence, we will go on to Jonathan Moore and come back to you.

All right. So, Jonathan Moore, our next speaker, is the -- yes, no jokes here about the fact that this is a Federal Communications Commission. Or maybe you can just insert your own joke here. Too easy.

All right. Jonathan Moore is the founder and CEO of Rowdy Orbit IPTV, LLC, located in Martinsburg, West Virginia. Rowdy Orbit is an online platform and destination featuring professionally produced web series for a highly unrepresented consumer base: African-American, Hispanic, Latino, and Native American viewing audiences.

Before starting Rowdy Orbit, Jonathan Moore was a 16-year seasoned advertising creative executive who developed successful general and multicultural executions for Fortune 500 clients, such as the dearly departed Pontiac, McDonald's, the Prudential, Novartis, U.S. Army, Seagrams, Partnership for a Drug Free America, Sony Music, Mazda, Missouri Division of Tourism, and Burger King.

And before I hand it over, I should tell you all -- because I'm not doing this well myself -- you should talk right into the mike so that they get the audio very well.

MR. MOORE: Okay. I want to thank the FCC for inviting me to speak today about network neutrality and free speech. It's truly an honor.

My name is Jonathan Moore, my company is called Rowdy Orbit. It's a company I started out of frustration. Frustration over the lack of representation of people of color in traditional media.

But the failure of media companies to provide and produce programming for and about people of color made me realize there was an opportunity to create a successful platform to truly showcase culturally relevant web shows. But how? I can't afford a broadcasting license and I couldn't own a cable channel. So, the only way to create and deliver a solid platform was the Internet. And that's why I'm such a strong supporter of network neutrality.

So, my story is, a few years ago I moved to West Virginia to launch my own business, Rowdy Orbit. What is Rowdy Orbit? Rowdy Orbit is an online platform featuring professionally-produced original programming, specifically to and for the African-American, Hispanic, Latino, and Native American viewing audiences. Featuring, hosting, and developing great storytelling that would be shelved by the traditional process.

Some examples of -- we have a show called Check, and it's all about an African-American female superhero trying to find this secret academy. We also have Hulce, should we go into that. We also have Kendrick. It's three African-American women who talk about love, drama, and -- with their whole entire family in the backdrop of New York City. What makes them unique? They have one pretty skinny woman, and the other two women are plus-size women, and we feel as though that would have never made it to network television. And there's also Chef Melvin, who's also a cook who didn't make it to the Food Network channel because of his exuberant personality.

Rowdy Orbit is a clear alternative platform, as you can see. a conduit to prove, your passion and idea has merit, an unobstructed direct line to an underserved viewing audience, a launching pad springboard for quality multicultural programming. Every -- each and every web series on Rowdy Orbit has a goal: Either stay independent or gain the necessary buzz, momentum, and business metrics for a company to acquire their intellectual property. This represents a perfect platform for web series creators to create value, build an audience, grow their business, and some cases, live out the American dream.

We simply enable web series creators to take direct action, grow without added weighted restrictions, and express with freedom to delivery quality creative product without subjective roadblocks or process barriers that are not touched by traditional media.

At Rowdy Orbit, we're opening doors, connecting creators with opportunities, creating a viable space for great ideas and content to devise silos and stereotypes.

This is all possible because the Internet provides the access to inexpensive online tools to streamline operations, cost controls, control costs, all while connecting with like-minded individuals across the country.

Most importantly, the Internet has allowed me to innovate despite running a bootstrap operations. It has allowed me to develop a competitive vantage by delivering quality and relevant programming to a visibly under-represented audience. People of color are able to go online and use the Internet to tell the stories without first getting the approval from an executive that's culturally removed and doesn't value or understand the needs of our community.

Rowdy O. wouldn't be possible without an open Internet. Without network neutrality, for example, since the start of developing the business plan and as of today, December 15th, my initial out of pocket investment has been $526. That's correct. $526. The modern example of the basic business story.

When an immigrant comes to the country -- comes to the United States, they only have a few dollars in their pocket. You're looking at one of the prime examples, and Rowdy is a direct result. Low-cost opportunity, plus minority ownership, plus hard work, total a thriving alternative platform that celebrates the freedom of expression via a quality creative product.

MR. BENJAMIN: Thanks, very much. Now I think we do have Glenn Reynolds. Glenn, are you there? Let's try again for the -- one more time --

MR. REYNOLDS: I'm here.

MR. BENJAMIN: Great. Glenn, take it away.

MR. REYNOLDS: Hi, thanks for having me. I'm glad this is finally working -- still a few bugs in the system, I guess. Well, someone earlier said we're not just a bunch of Internet geeks here, and I don't know about the rest of you, but I actually am a broadband geek and my involvement in this stuff has been pretty much from the level of a hobbyist.

Before I started blogging, I ran a record label -- it's still sort of alive -- and produced a number of bands and marketed them mostly through the late, lamented MP3.com, and ran a series of Internet radio stations, and did a lot of stuff like that -- also, with a zero to low initial investment. And had a lot of fun, sold some records, met people around the world, some of whom I'm still in touch with, and found it a great creative outlet for me and it was for quite a lot of other people.

I started blogging in 2001. I teach Internet law and I also try to do various hands-on things on the Internet so that I can stay grounded and, you know, with a little real world cred.

I went to blogger.com, I followed the instructions to set up my blog, which took me about 15 minutes, I started blogging on August 8th of 2001. And I thought, well, you know, if it goes well I might have a few hundred high-quality readers of academics and journalists I know, and stuff like that. And actually, I remember at the end of August, I had gotten up to 600 pages a day and I really thought I was getting somewhere.

Now, I've got between 4- and 500,000 page views a day. It's all done over the Internet at effectively no cost for me, and it has reached people all around the world and gets me e-mail and introduces me to all kinds of smart people who do all kinds of interesting things, many of them not traditional smart people jobs. And I felt it -- tremendous.

It's also served as a platform for a lot of people doing journalism. My blog is called Instapundit, and most of what I do is basically links and punditry. I occasionally do hard news coverage, if I happen to be somewhere where it's happening and I have readers who are correspondents in the classic sense. They correspond with me.

I've had people who report from Afghanistan and from Iraq and other places via my blog from time to time. Other people have really turned blogging into a much more serious journalistic model. Michael Yon, for example, who's currently in Afghanistan as an embedded journalist has probably spent more time in Iraq and Afghanistan than any other journalist as an embed, and funded entirely by donations from his readers who read his blog at michaelyon-online.com.

He writes like a 21st century Ernie Pile. He takes great photos. He's a special forces veteran himself, so he's got a lot of experience. And he also actually makes more money this way than he would probably as a correspondent for a major news outlet, assuming major news outlets could afford to have correspondents these days.

Michael Totten is another one -- michaeltotten.com -- who does a similar sort of reporting from a lot of the world's hot spots: Lebanon, Iraq, all over. And there are quite a few other bloggers doing this sort of thing.

Probably disproportionately, they've been military bloggers because it turns out, there are a fair number of people with military experience or at least experience getting along in the world's hotspots who have interesting things to say and felt that it was under covered for the last few years.

Beyond that, we have people who report news just because they happen to be there when it happens, and this happens via blogs -- it happens particularly now via Twitter. And now that you have video cameras and digital cameras, and cell phones and all integrated, and you can stream video to the web from your iPhone, and things like that, we get much more on the spot reportage as news happens. And even though traditional news organizations are really strained for resources, there's one group of people who are always going to be on the scene when news happens, and that's the people who are already on the scene when news happens.

And traditional news organizations are beginning to incorporate them to some degree as well as a way of just increasing their reach and getting first hand fast response coverage of things, take advantage of what's already out there.

Beyond this, we're also seeing huge numbers of people who are expressing themselves through political satire, and such, on things like YouTube -- some of those people go on to have bigger TV careers. Alfonzo Rachel, who started out on YouTube, is now on Pajamas TV and sometimes finds himself in other places. Ditto a guy named Steve Crowder. And all of this is made possible simply because the Internet's cheap.

If you want to produce video and send it around the world on the Internet, you need an inexpensive camera and an inexpensive computer and an inexpensive Internet connection, and some place to host it, like YouTube or Motion Box, or whatever.

In the old days, you had to have satellite trucks and stuff like that. In fact, I covered the Knoxville Tea Party protest for Pajamas TV. I used a MacBook Pro with a pretty good video camera connected to it by FireWire and my Verizon EVDO broadband card to send stuff live to the Pajamas TV studio in L.A. over the Internet. And as I walked past all the million-dollar satellite trucks from the broadcast people, one of the guys looked at me and said you're really making us obsolete, aren't you? And

I was like, maybe a little bit, yeah.

Why can you do this? Because the barriers are low, the equipment is cheap, the service is cheap. Once you're on the Internet, everybody is equal. And that, I guess, really brings us to the point of today's presentation.

I mean, my first experience of that -- and probably really my inspiration for blogging -- was when Mickey Kautz started his blog and I followed a link from Slate.com, which was sort of one of the early big deal online magazines. And I followed the link to Mickey Kautz's site, which was something run by one guy in his own little server, and he looked just as good. And I realized that from the standpoint of the reader, it was seamless. The Internet's all just one big place.

There aren't the prestige addresses in the same sense there are on the television dial, or the cable selector, or anything like that. And it really opens things up to let people communicate in all sorts of ways, and to draw audiences based on their ability to appeal to people. And I think that's a wonderful thing, on a couple of levels.

I mean, for guys like me, the Internet's just a big playground. I'm a geek. I have fun with it. I have a lot of hobbies and it gives me a way to do things.

For the society at large, it brings a lot of voices into the conversation that have previously been excluded for a whole lot of different reasons. And I think that's terrific, too. And I certainly am in favor of efforts to keep the Internet a level playing field. Though, as I mentioned in my prepared statement, I have certain concerns regarding the ability of regulators to really keep up with technological and economic change.

MR. BENJAMIN: Glenn --

MR. REYNOLDS: And at that, I will bring my segment to an end.

Thank you so much for having me on this somewhat buggy connection, and I hope it's worked out.

MR. BENJAMIN: Great, thanks very much. You actually just ended at the point I was going to ask you. So that was telekinesis.

All right. Our next person is Ruth Livier, an actress, writer, producer who was recently featured on the cover of Written By -- which is the Writer's Guild of America magazine -- as the first person to join the Writer's Union for her work in new media.

She has an award-winning dramedy -- is it, YLSE?

MS. LIVIER: YLSE.

MR. BENJAMIN: YLSE, now in it's second season. And she also has a web series, Osiris, in development. She's had a long career in many other roles going everywhere from Playhouse's Disney to Arlen, Texas -- for those of you who are King of the Hill fans.

Without any further ado, Ruth.

MS. LIVIER: Thank you. Online video is in its infancy. It was only in 2005 -- a mere four years ago -- that a then unknown site called YouTube was launched, effectively beginning the online video revolution that has transformed the Internet from an information and communications tool to an entertainment destination.

The Internet is quickly becoming a new avenue for independent content, diverse content, and minority small business ownership. By nearly all accounts, online video has one of the fastest consumer adoption rates among any form of media distribution. For example, in March 2009, Nielsen reported 9.6 billion videos streamed by 130 million web users in the month.

As this chart illustrates, in the course of one year, growth of online video has continued to rise dramatically. During this period, Nielsen estimates that online video usage in the United States has grown by 24 percent.

Advertisers have also begun to flock to online video, spending $587 million in online video advertising in 2008. While this figure is tiny compared to the billions spent on television advertising, online video remains one area of advertising growth even in this economic downturn.

While the major media companies have taken to the Internet by allowing users to legally watch content ranging from classics such as The Dick Van Dyke Show to the latest episodes of your favorite shows, what is largely unavailable through their extensive libraries is a wider range of diverse, culturally-relevant programming keeping pace with the changing face of America.

The U.S. Census Bureau's report of 2007 affirmed that the minority population in the United States climbed to 100.7 million people, or 33.6 percent of the total population, meaning that 1 in 3 persons in the United States is a minority. And yet, the most recent NAACP report found that minority communities are still not adequately represented on broadcast television relative to their numbers.

Now, this lack of balanced and equitable representation translates into fewer job opportunities for minorities, fewer opportunities to make a living as a writer, director, or as an executive in our business, and consequently less access to lend our voices to entertainment content that is created.

For minority content creators, the low barriers to entry on the Internet have allowed us to take initiative, create, produce, and distribute our stories and develop a financially viable business. But my ability and the ability of many other new media content creators is dependent upon my ability to reach an audience unimpeded and unencumbered by gatekeepers and filters.

This ability to prove our market has very strong socioeconomic repercussions for our communities. As an American and as a Latina, I got tired of seeing a disproportionate amount of negative stereotypes in traditional media. As a creative, new media empowered me. I saw in the web an opportunity in whatever small way to partake in redefining those unfair and unbalanced perceptions. The only requirement was than an audience be able to find these little six-minute webisodes that I was writing.

I wanted to tell a story that was more in line with the vast majority of the 47 million Latinos who are contributing positively to our communities, rather than the oftentimes biased portrayals of mainstream media.

Just last week, the Pew Research Centers Project for Excellence in Journalism issued a report chastising the media's portrayal of Latino Americans. My show, YLSE, is a bicultural dramedy about a modern American Latina, someone with big dreams fighting through other people's low expectations, juggling, career, a not-so-successful love life, and a family who sometimes doesn't understand her progressive American ways.

There is no way I could have gotten my show done in traditional media. I was quickly dissuaded from trying, many times. With no precedent for such programming and a custom of producing content that can be resold in other, most notably European countries -- European markets -- I wasn't surprised to hear the skepticism I encountered.

We started production two years ago on the show and we are effectively now finding an audience, our fan base. As you can see by the fine quotes, it reaches beyond a target niche market. It is even broader than I could have imagined.

Our numbers have multiplied. Season one we had a few thousand. Season two--now we are closing in on half a million halfway through season two. And we achieved all of this with zero marketing dollars.

My startup production company has been able to tell a relevant American story because the Internet allows us the unique opportunity to compete for the hearts, eyes, and minds with billion dollar media conglomerates of today. Our success is based on the popularity, the strength of our content, our audience's ability to find us, and the sweat equity of an amazing, talented team of creative people.

The Writer's Guild of America West represents 8,000 writers of television, film, and new media. Being the first person to join the WGA because of my work on the web shows -- that the web can nurture new voices -- gives us an opportunity to perfect our craft and be added to the mix of professional writers and producers effectively diversifying the talent pool.

As a small business owner, I hope to create jobs and open opportunities for others. Allowing ISPs to construct barriers to entry or to offer uncompetitive advantages to those with the deepest pockets will only serve to derail the innovation and entrepreneurial spirit that have driven the Internet since the beginning.

It was the same freedom that allowed innovators to create the web in the first place. As the inventor of the World Wide Web, Tim Berners-Lee stated, when I invented the web I didn't have to ask anyone's permission. I didn't have to ask anyone's permission to create YLSE. I was able to work hard, pull together the talent, and launch a successful web series.

An open Internet drives innovation and encourages the independent American entrepreneurial spirit. Failure to maintain the open Internet will mean we are repeating the mistakes of our own past.

The financial interest and syndication rules were designed to protect the ability for independent voices to be heard. Repealing them has had detrimental effects. It opened the door to media consolidation, and traditional media -- which allowed a mere handful of companies to dominate the entertainment industry and shape the news and information that gets to the American public and the world.

Until the Internet, a viable distribution alternative to traditional media has never existed. A neutral, non-discriminatory Internet is a market-driven, equal playing field for all Americans regardless of ethnicity or socioeconomic standing. It allows viewers to find the information and entertainment they want without filters or gatekeepers. It is critical to the vitality of the American public discourse and our democracy.

A neutral Internet is our best opportunity for diverse voices to partake in the molding of the American perception and perspective.

Thank you.

MR. BENJAMIN: Thanks. Our next speaker is Garlin -- sorry -- Garlin Gilchrist, who is the director of new media at the Center for Community Change, which advocates for public policies in the interest of low-income people -- particularly low-income people of color. He was -- in 2005, he co-founded and contributes to the SuperSpade: Black Thought at the Highest Level, one of the nation's leading black political blogs. He also co-founded Black Net Action, a strategic collective of black bloggers and online activists who coordinate online campaigns across the country.

Take it away.

MR. GILCHRIST: Thank you, Stuart. Thank you, FCC commissions and staff for inviting me to testify here at this very important hearing.

I'd like to tell a personal story about my own personal evolution as an activist, as an organizer, and how that connects to the open Internet and the things that it enables Americans to do, and enables -- how it enables people to participate in the democratic process.

I'm an engineer by training, and as a young person growing up in the city of Detroit -- and later in its suburbs -- I witnessed firsthand the fallibility, the challenges, and the promises of a city, and of a state, that is really dependent on a single flavor of manufacturing. The challenges of living in this place really opened my eyes to what's visible, what's possible, and what's promising about open networks and about the open Internet.

So, I always wanted to pursue my own passion for technology that developed when I was a young person. As a pathway for empowerment, and a way to inspire others to take collective action, to build stronger communities, and to finally realize the American dream as it is most relevant to them.

So, as I came of age I sought for ways to connect that passion for technology to how to improve the lives of others. When I worked at General Motors as an intern, I found ways to do this professionally. I helped contribute, to create, and then manage two websites that I think really set the table for my own activism.

One was called Webhands.org, which was an initiative that was created as an Internet portal to connect volunteers and donors to local nonprofits in their own geographic locations. This was a way for sort of the small advocacy organizations to really reach out to donors that may have been inaccessible to them or may have been inaccessible or invisible to them. And this is the type of thing that the open Internet enables because it allows people to reach audiences that they never even knew existed.

Another was called Patentdonors.org, which enabled small businesses and institutions of higher learning -- big and small -- to research and donate their unused patents and intellectual property to small businesses and investors who were interested in taking advantage of that technology.

So, these were very important ways for these institutions to take advantage of the web, take advantage of its low barriers to entry, take advantage of its environment to increase competition, and increase participation in order to really be competitive.

So, these projects really opened my eyes to the use of open networks and the tool for encouraging participation, encouraging empowerment, and social change. And building on this foundation as a student, community organizer, and political organizer, I've been armed with the technology, infused with passion, and inspired by metaphor of open gates that I think applies very well to the open Internet.

Openness is a defining characteristic of democracy. Openness means available, visible, and flexible, sort of like an open gate. If a gate is open, anyone can walk through it and enjoy what's on the other side.

Now, I'd like to make a reference to Commissioner Clyburn, who also agrees that the question of who can access these open networks is an important but separate question. And it's not the focus of today's discussion, but I do want to, at least, highlight that as an important thing that I hope that the Commission will seek to address.

So, in the context of political engagement, open gates really is a metaphor that has presented itself repeatedly throughout history. Walled gardens and gated cities have characterized the kingdoms of ancient history in the same way that they characterize our current media and technical landscapes. We need to make sure that we design regulations thoughtfully, deliberately, and appropriately to ensure that openness is always protected, no matter what the situation is.

So, the characteristics of openness, the characteristics of flexibility, transparency, and disruptive potential are what's made the Internet so empowering. Individuals can communicate with strangers, friends, and family without hindrance. This flexibility leads to the fear of retribution for those communications, and that freedom inevitably leads to political change.

So I can illustrate this point, using a few examples from my own experience, that involve video, authentic storytelling, as well as how marginalized people are organizing online to take advantage of open networks to impact the political process.

One I'd like to talk about is with who's known as the Tea Party rapper. So, this year there's been a movement building among disenchanted people with the current President, the current Congress, and their policies, and they've organized marches and demonstrations in various parts of the country. Well, one individual wanted to really tell his story about what his opinions were about how policy was being designed through video. Well, through a site like YouTube he produced a music video independently, without money out of his own pocket, with a very small investment, with no marketing budget. He produced a rap video that was viewed by hundreds of thousands of individuals that really incited not only a level of frustration, but a level of inspiration for people that agreed with his movement.

He was able to reach people all across the country and even all across the world with his message. This is something that would have never been possible throughout -- before the advent of the Internet, before the advent of a communications platform that was so open, so accessible, and so inexpensive.

Another example is voter mobilization. During the final weeks of the 2008 presidential campaign, advocates really turned their attention to making sure that as many people got to the polls to vote as possible. Well, through the open Internet, they were able to find ways to make their advocacy stronger backed up by academic research.

So there was a Yale political science research paper, called "Social Pressure and Voter Turnout," that was published in February 2008. It's authors included Alan S. Gerber and Donald P. Greene, and they discovered that direct mailings informing people that their neighbors, friends, and family could find out whether or not they voted actually had an impact of increasing voter turnout by at least 5 percent.

So, the liberal group MoveOn.org decided to take this research that was available to them via the open platform of the Internet, via universities that could share their knowledge and understanding with the rest of the world for no cost. They were able to use this to influence their activism. So, they created a video series where they actually had an interactive video where people could insert another person's name into a video that said, hey, it was a fake newscast. It said your candidate lost the election because you didn't vote, and it would have their name in it and everything. Very interactive.

That video was shared and viewed over 21 million times, and that video -- backed up by the strong science and the strong evidence that was available to these folks because of the open Internet -- really impacted voter turnout in a very positive way.

So, the web is -- finally, the web is uniquely suited to tell authentic stories. I'd like to tell one last example of this from my organization, the Center for Community Change.

This past summer, as we've been debating health care reform legislation in this country, our group thought it was important to make sure all American stories were heard. And so we took to -- went out to different state, county, and country fairs in rural states across America --

Nebraska, Tennessee, Maine, and Missouri -- and we asked people about their experience with their current health care coverage -- their current health care system.

We connected these people online to be able to tell their stories visually and in written form and, finally, we shared those stories with people all across the country, including members of Congress and their staff. We hosted a briefing at -- on the Hill where we allowed these folks to come and tell their stories. And these real stories that were viewed hundreds of thousands of times really had an impact and helped make sure that the voices of rural America were heard in this debate.

That's something that, again, never would have been possible without open networks. So I want to just be clear in saying that I am in firm support of open networks being protected, and being supported, and being encouraged.

Thank you.

MR. BENJAMIN: Thanks. Next we have

Robert Corn. Bob is a partner at the law firm Davis, Wright, Tremaine and has extensive experience in First Amendment law and communications, media, and information technology law, counseling clients writing in the area, serving as litigation counsel. Speaking -- he also successfully petitioned George Pataki for a posthumous pardon for Lenny Bruce, which I think is unique; I could be wrong.

And his clients include the A&E television networks, American Association of Advertising Agencies, Association of National Advertisers, CBS, Motion Picture Association of America, the National Association of Broadcasters, Viacom, and Playboy Enterprises.

Without any further ado, Bob.

MR. CORN-REVERE: Well, Stuart, thanks for the introduction, and I'd like to thank the Commission for inviting me to participate. It's -- this is important conversation on the Internet and free speech.

At the outset I just need to make clear

that the views I'm expressing today are mine alone, they're not that of any client. I'm not here to represent anybody. And my observations are based on my experience as a student of the First Amendment, as a practitioner in the field of Constitutional law, and as a former FCC staff member.

I think all you have to do is listen to a panel like this and you get a sense of how the Internet has transformed communications in America. These are really inspiring stories of how individuals have gained a greater ability to communicate and connect than any time ever before.

Even the transmission of Glenn Reynolds from, apparently, the Apollo 13 spacecraft indicates how barriers have been lowered and more things are possible and more connections can be made. But these revolutionary developments have also transformed the law and, specifically, the way in which American courts view new technologies under the First Amendment.

This has always been a rather uneasy

relationship. From cinema to broadcasting, and from cable television to satellites, courts were slow to recognize the application of traditional First Amendment principles to new media. Legislators and regulatory agencies treated different categories -- and created different categories and classifications for communications technologies as one emerged. And courts established different levels of Constitutional protection based on those different categories. Some of us have always believed that this jurisprudential approach never made sense, but to whatever extent it once did, it is entirely untenable in the age of media convergence.

In some cases, we've already been able to make some progress in extending First Amendment protections to new media while in other cases, as some of you may know, we're still working at it. But the Internet was revolutionary because it was the first new communications technology the courts found to be fully protected by the First Amendment, from the outset.

In Reno v. ACLU, the Supreme Court could find no basis for qualifying the level of First Amendment protection -- First Amendment scrutiny -- that should be applied to online communications, because it found the new medium to be as diverse as human thought. As Judge Dalzell wrote in the District Court opinion that led to the Supreme Court's ruling, the Internet is a never-ending worldwide conversation and the most participatory form of mass speech yet developed.

On the strength of those findings, courts viewed content regulation as a threat to the open Internet, and they uniformly invalidated both the CDA and its successor statute, the Child Online Protection Act, as well as a number of state versions of this -- essentially the same law.

Now, ironically, the same Congress that adopted the CDA also declared that it was a national policy for the United States to preserve the vibrant and competitive free market that presently exists for Internet and other interactive computer services, unfettered by federal and state regulation. Now, that was a different part of the same statute, but one that recognized that these technologies had evolved because of a minimal regulatory view. And that those technologies grew up without regulation.

Now, as technology is evolving at or after the increasing pace, the central question we face is whether an open Internet will better be preserved by the hands-off policies that led to its creation, or by new regulations?

The FCC and its network neutrality notice of proposed rulemaking cites a number of its prior policies that help to serve as building blocks for competitive telecommunications, including the Carter phone line of decisions, and the Commission's computer inquiries. These decided as precedent for -- to support the idea that the government can act to keep the network open. But the FCC didn't invent the Internet any more than Al Gore did.

These pre-World Wide Web decisions --

and in some cases, pre-Internet decision -- may have contributed to an increasingly competitive telecommunications sector, but they were not adopted with the Internet in mind.

The difficulty, too, in looking to prior Commission policy is that it's not difficult to find counter policies that went the opposite way. For many years, the Cable Act and the MFJ prohibited AT&T and the telephone companies from getting into video services -- something that wasn't changed until the 1996 Telecommunications Act.

As Commissioner McDowell noted, the FCC's effort to create open platforms has been -- how shall we call this -- an incomplete success. Both video dial tone and open video systems were regulatory constructs crafted by legislators and regulators, with the goal of ensuring an open platform --

(Interruption)

MR. BENJAMIN: This is not a referendum on what you're saying. Does anybody have any idea what's going on?

All right. No, the record will reflect -- we will add 30 seconds to his clock. Bob -- no, keep going.

MR. CORN-REVERE: Okay. Ultimately, as we debate the question of how to preserve an open Internet, we face two potentially competing risks.

One, the risk that's been outlined by some suggesting that allowing the unregulated Internet to continue will risk adverse actions by various companies. The other, the risk that we will impose a level of government regulation on the Internet that court decisions, so far, have freed us from.

The choice that people will make on those -- in evaluating those two questions -- will be informed by whether they consider government to be a greater risk, as it has in the past, or whether or not they consider the potential for adverse action by corporations to be a greater risk.

I've outlined my take on that, and the

written comments that I've provided that go into some detail as to the reasons I've reached the conclusions that I do. And my conclusion is that there is a greater risk from inserting more government regulation over the Internet and bringing us back to the failed experiment with media regulation that we had in the 20th century.

MR. BENJAMIN: Great, thanks. Next is Jack Balkin, the Knight Professor of Constitutional Law and the First Amendment at Yale Law School, the founder and director of Yale's Information Society Project, and the director of the Knight Law and Media Program at Yale.

He's -- has a blog, Balkanization; also has written several books, including Processes of Constitutional Decision-making, and Cultural Software: A Theory of Ideology. His most recent -- latest book co-edited with Riva Segal is The Constitution in 2020. Jack?

MR. BALKIN: Thanks very much, Stuart. And I'd like to thank the Commission for inviting me to speak.

I was asked to come and speak in my capacity as a First Amendment scholar to talk about why the open Internet is crucial both to freedom of speech and to democracy.

So, I want to start by asking about really what the gifts are that the Internet gives us. And start by talking about what digital networks make possible.

First of all, they allow people to become active speakers and creators, instead of merely passive consumers of information and entertainment. Secondly, they decentralize innovation. They give people abundant opportunities to create and use new applications for communication and creativity. And third, they allow people to form new kinds of social relationships, groups, and communities.

In short, digital networks allow people to participate in culture, society, and politics in all kinds of new ways. They can do it individually or in groups, they can do it locally, nationwide, or around the world.

The ability to participate is the Internet's great gift to mankind. Participation is also central to the First Amendment. Some scholars tell you that the point of the First Amendment is liberty, other people say that it's democracy. I tend to combine the two. For me, the point of the First Amendment is to foster a democratic culture. That's a culture in which ordinary people have a say about the forces that shape them and make them who they are.

A culture is democratic, not because people vote on it. But because they get to participate in making it. A participatory and democratic culture requires more than just protecting political speech. It requires a vibrant public sphere that makes self government possible.

Participation, however, means very little if we need permission to participate. And an open Internet means that we can speak, organize, and innovate without getting anybody's prior permission. This idea, too, has deep roots in the values underlying the First Amendment.

One of the earliest conceptions of freedom of speech was freedom from prior government restraints. And I'm sure you all know the saying that the real freedom of the press belongs to the person who owns one.

Until recently, nobody could gain access to mass communications unless they had the permission of a large media company, like a newspaper or a television station. And even then, they only got access on the broadcaster's terms. They were often heavily edited, and good luck getting access if you said something a little oddball or unpopular.

An open Internet changes all this. People can reach audiences that only large media corporations could reach before. They can route around traditional media gatekeepers who often functioned like private censors. They can create new tools and applications for speaking, communicating, and organizing, all without having to get anybody's prior permission.

Imagine a world in which you had to get permission from Internet service providers before you could create a platform like Typepad or YouTube, or upload content onto Flickr or Facebook. Free speech and democracy thrive precisely because we don't have to ask somebody's permission before we speak, before we engage in politics, before we upload files, or before we create a new social media application.

An open Internet is an Internet that is open to new content and new applications. It's an Internet where your ISP doesn't try to block you or shut you down for daring to compete with its favored content providers. The Internet lets us route around the old gatekeepers. But the new gatekeepers are the broadband companies who own and operate the conduits through which everybody speaks.

Broadband companies are private organizations. Their private interest is in making money and pleasing their shareholders, but their business is affected with a crucial public interest: Promoting free expression and democratic participation, and giving ordinary people the opportunity to create, to innovate, to spread innovation, and organize politically.

Broadband companies further the public interest when they operate as open, non-discriminatory platforms for other people's innovation and as open non-discriminatory conduits for other people's speech. But their private interests inevitably lead them to play favorites. Their incentives mean that even in broadband companies have no plans to censor unpopular speech, they won't really want or enforce a level playing field for private speech and innovation. That's why there is inevitably a conflict between the public interest and private interests, and that is the reason why regulation is necessary.

To preserve the great promise of the Internet, we must confront this conflict of interest head on. When companies act as primary conduits for other people's speech, they should not be able to discriminate in content or applications -- and they must be transparent about how they maintain and manage their networks to promote efficiency. That's the simple principle.

Now, seeing that regulation is on the way, broadband companies have begun to argue that they have a Constitutional right to block applications and discriminate against content, and that any attempt to keep them from maximizing their profits in this way violates the First Amendment.

Nothing could be further from the truth, and I'll be happy to talk about this more during Q&A. But for right now, let me just explain that under the First Amendment, Congress can make both telephone and cable companies into common carriers who must take on all traffic.

Congress can certainly require a much milder, non- discrimination requirement, like network neutrality. The thing to remember is this, the First Amendment protects speech. It does not protect business models. The FCC's job is to make sure that communications companies serve the public interest as well as their own private interest. This idea has been the basis of telephone and cable regulation for decades.

The public interest demands that we secure the benefits of an open and participatory Internet for this country. Network neutrality rules are a good place to start.

MR. BENJAMIN: Thanks. Finally, we have Andrew Schwartzman, president and CEO of the Media Access Project. He has directed the organization since June 1978. This is a public interest telecommunications law firm that represents the public interest in promoting First Amendment rights to speak and hear. He has written and spoken extensively, is well known, I think, to everybody in the FCC and everybody in Washington.

Andrew?

MR. SCHWARTZMAN: Thank you very much. Much of my written statement is similar to and elaborates on some of the points that Professor Balkin made and is responsive to some of the things that my friend Bob Corn-Revere had to say.

I would simply emphasize that much of what the debate is about here is whose First Amendment is it? And I would argue that Internet service providers are not functioning as speakers when they are functioning as conduits of speech of others.

Knowing that we're going to have all of these creators on this panel, I chose to focus on what is, candidly, a more negative kind of presentation today.

In the latter half of the 19th century, the telegraph came into being, quickly became a monopoly under financier Jay Gould, leveraged a deal with the Associated Press to keep other telegraphy companies from being able to function. And, indeed, some people think that this monopoly manipulated the outcome of the 1876 presidential election in favor of Rutherford B. Hayes.

Now, that was a long time ago. Analogies are limited, but there are some similarities -- and one is that the best way to protect against such abuses involves erecting safeguards against them.

The problem is not speculative. Here's what I know. Four years ago, Telus, one of Canada's largest Internet service providers surreptitiously blocked customer access to a website operated by a union with which it was engaged in a labor management dispute. Seven hundred other websites were collateral damage as a result of this blockage. They stopped it, but they continued to maintain they have the legal right to do so.

A few months later, the communications carrier named Madison River adopted a practice of blocking voice -- ended telephone calls placed by customers using its own DSL Internet facilities -- presumably as a competitive factor. They entered into a consent agree with the FCC, again without admitting that there was any legal problem.

Two years ago, Verizon Wireless initially blocked text messages from NARAL. Verizon Wireless said that it does not accept issue-oriented abortion war, et cetera, programs

-- only basic, general politician-related campaigns. It said that for now, Verizon Wireless will not accept any programs that are issue-oriented from any organization that seeks to promote an agenda or distribute content that, in its discretion, may be seen as controversial or unsavory to any of our users. They said general information on campaigns were acceptable to the extent that the content involved is in Verizon's sole discretion, not issue-oriented or controversial in nature. Now, Verizon backed down, but it never issued written guidelines until August 2008, when it commendably issued very good guidelines.

More recently, we had the Comcast bit torrent controversy, in which the Commission found that Comcast was blocking -- surreptitiously blocking peer-to-peer protocols. The complaint demonstrated, among other content, Comcast blocked transmission of a file containing the text to the King James Bible.

Initially, Comcast denied that it was

doing so. And after this was proven to be untrue, Comcast lied to the Commission, saying that this blocking only took place at times of congestion. Months later, it was established that actually they were blocking it at all times and had been doing so for several years.

Now, that's what we know. More importantly, what we don't know is how many other instances of blockage or degradation of service there might have been and how many are taking place right now. After all, Comcast initially denied it was blocking bit torrent. Only after it was confronted with conclusive evidence discovered by the accident -- that one of its customers was a highly experienced computer engineer who was puzzled by his inability to upload barbershop quartet music that was in the public domain -- only then did it come to light. Determining when such misbehavior takes place will only become more difficult.

Advancing technologies, including deep packet inspection, are much more sophisticated and much less easily detected than Comcast's blatant ham-handed exploits. The great exchanges we face in maintaining free expression on the Internet arise because of what we do not know, and because of what new techniques may have been, or may soon be, developed.

And it's in that context that I address the question in my written testimony of how to address the First Amendment to the network neutrality debate.

Thank you.

MR. BENJAMIN: So, thanks very much to everyone. I have a couple of questions that I would like to throw out and we can hopefully have some discussion among us. And again, people should feel free to give those questions to Zach, or new media at FCC.gov, is that -- I think that's the right e-mail address.

If I -- let me just start with a question for Michele, Garlin, and Glenn, if he still exists. Which is -- so there's -- Andrew just talked about this issue of blocking. And I'm

-- you've all said about how the Internet has opened up avenues for you that cable doesn't open up, that broadcast doesn't have, there's this huge communicatory capacity of it.

How and why is it realistic that they would actually try to block any of your messages? What's your fear of what would happen if the companies can freely block? What's the story you would tell as to how and why this would happen and how it would impact you?

I don't know if you want to start, Michele, then Garlin, then we can throw it to Glenn.

MS. COMBS: I think Andrew mentioned it. The incident that happened with NARAL, even though we're on totally opposite ends of the political spectrum, we actually did an op-ed with Nancy Keegan on this issue.

And we're very action-oriented. So, once we get an issue out, we have to have it out immediately with -- I gave an example of how we stopped an amendment on a bill. So, because we're so action oriented, and if they stopped and blocked our content, then it would significantly hinder the way we communicate through our supporters.

MR. GILCHRIST: So, to speak to that very specifically, I think that the challenge is it is not clear where the certain gatekeepers, what their biases would be and what their opinions would be on different issues. And so, as you have such a diverse group of interests, organizing around all sorts of issues, it's really unclear as to how much they can or cannot trust those gatekeepers.

And the best way to ensure that the discourse is as vibrant and as diverse as it can be is to eliminate those gatekeepers or prevent them from blocking that sort of content.

The case for Comcast blocking bit torrent is particularly interesting because activists as we learn more about the Internet and learn more about these tools and learn more about how to take advantage of the open Internet will want to take advantage of it in all sorts of ways. And that might mean sharing content that is -- that very large content -- which is really the issue in this bit torrent case where the argument was that congestion on the network was happening as a result of this sharing.

Well, the reality is, as speech and as content and as the things that we create become more advanced, become more substantive and become, frankly, just larger, we're going to need ways and pathways and avenues to share those. And so, this sharing should not be hindered or should not be blocked or should not be encumbered by interests who organizers, activists -- who citizens do not understand their biases and cannot trust.

MR. BENJAMIN: Glenn, if you're there, do you want to say anything? Or have we lost you entirely?

MR. REYNOLDS: Yeah. No, I'm here. I am absolutely concerned with the notion of gatekeepers blocking content that they regard as controversial. This often plays out as ridiculous. Some of you may have heard, Apple was blocking an iPhone app that turned public figures into bobble heads because they said they wouldn't allow anything that ridicules public figures.

Corporations in general tend to be very risk- averse, very control-oriented and, basically, kind of dumb about stuff like this. And it's probably easiest to give them a hard and fast rule that says, you don't get to do that.

I want to, by the way, weigh in and say that I am inclined to agree with Jack on the First Amendment issue. You know, you may have a First Amendment right to speak in terms of what you put on your Comcast or your Verizon website, but I don't think you have a First Amendment right to function as an editor with regard to the speech of your customers.

Customers of an Internet service provider are not akin to reporters at a newspaper or something like that. And I just don't think that that First Amendment holds any water at all.

MR. BENJAMIN: Okay, that's a nice

segue. I now wanted to ask Bob, Jack, and Andrew a related question on this point about blocking.

The -- Bob made a point in his opening statement, yes there were instances of blocking and they were reversed. So, in fact, the NARAL text message -- I've forgotten -- within 48 hours they had reversed their position and allowed it. So, if it -- given that in the past, Madison River -- well, I guess the FCC smacked them down, but maybe there would have been a competitive pressure on them as well.

So, if we think -- let's just -- if we think that by and large, competitive pressures will at least diminish this danger. Then, why do we need open Internet rules, if we think that the market will respond? And I will let you -- we can go down the panel -- we can start with Bob and then Jack and then Andrew.

MR. CORN-REVERE: Okay, let me start with that. And I'll start with some of the examples that have been raised. First of all, you know, I think listening to the stories of the other panelists and what they're capable of doing on the web provides sort of a powerful response to this. I mean, it shows what people are actually doing and their story is replicated millions and billions of times around the country and around the globe.

What you hear are a few isolated examples of where people have suggested there has been bad behavior. The drive for network neutrality rules accrues if you believe that the behavior will become a significant concern among a significant number of network operators -- that it won't be disciplined by competition or by the adverse reaction of customers.

And if -- where there's market power, more general laws like antitrust can't be brought to bear. Where it has occurred, it's been mainly through consumer pressure that has corrected the misbehavior right way. I think the real question that we face is the one that Commission McDowell identified at the outset. And that is, if one company makes a mistake or engages in bad behavior, that's one thing, but when the government does that or makes a mistake, it affects everybody.

And as in the case of keeping the telcos out of electronic publishing, it takes decades to undo it. That raises a couple of related First Amendment issues. And not the ones that Jack started by talking about, whether or not people have -- are free not to be a common carriers.

That actually is an interesting question, because there is some case law in that, where restrictions were placed on telephone carriers and whether or not they could get into the electronic publishing. And the development of that case law was cut short by the passing of the 1996 Telecommunications Act. But I'd love to hear what precedence you have in mind when you say that anyone can be forced to a common carrier?

But there's a broader First Amendment issue. And that is, to what extent do you insert government regulation of broadband networks that is inevitably going to extend to other areas?

Keep in mind, the government's agenda isn't just to keep an open Internet. Even in the network neutrality proceeding, the FCC describes its mission as to preserve an open, safe, and secure Internet. And once regulatory jurisdiction is established, it then extends to other areas.

The FCC is conducting an inquiry right now on whether or not it can regulate content and exert regulatory jurisdiction generally over all platforms, including the Internet, in the name of protecting children. And that would apply to all types of different content.

There are other areas in which the FCC is looking to exert regulations. And those present First Amendment issues of a different order, and they're the kinds of things that I thought we were emerging from.

MR. BENJAMIN: Jack?

MR. BALKIN: There are several different points -- I hope I can remember them all.

First, you would get no consumer pushback if you didn't discover that there was something going on. And one of the things that these open Internet rules are trying to do is promote transparency.

So, if you believe that -- if you believe in competition and consumer pushback as a way to keep broadband providers honest -- you shouldn't be opposed to transparency requirements.

The second thing is that, generally -- the example of Comcast with bit torrent is a wonderful example of what's going on. It's a kind of combination -- I'm going to take up what Glenn said -- it's a kind of combination of not invented here, plus you might eat into our profits in the way we're imagining our business model.

That is, the problem wasn't that Comcast was opposed to barbershop quartets or the King James Bible. Comcast thought, oh my god, here's a new application, we didn't invent it. We don't know how it works. It might, in fact, interfere with the ways we plan to make money through our favorite content providers. So, let's cut it down.

That is much more likely to be the way in which you get bad behavior in the future. And since with an open Internet you always getting new kinds of applications, nobody every knew about bit torrent, you know, five years ago, right?

The -- you're going to get the same pattern repeated over and over again. And if you don't have transparency, then you won't be able to know what's happening until much later. So, it's really important to create a baseline for fair competition now. Saying that you have to adapt to new applications, not just shut them down, that's the real problem.

The final point, I guess, that Bob asked me is, what's the source of the idea that you can be made into a common carrier? Well, Bob and I can't be made common carriers, we wouldn't be very good at it. On the other hand, that if you're a conduit for other people's speech, yes, you can. And a classic example of that is in Turner. Every single justice of the Supreme Court agreed that cable companies, which were traditionally treated like broadcasters, could have been turned into common carriers if Congress had chosen.

Now, it's true. Sometimes there are 5-4 decisions. But when nine justices take the same view of this, it seems to me this is pretty settled First Amendment law.

MR. BENJAMIN: Andrew.

MR. SCHWARTZMAN: Well, in my prepared testimony I went into more detail on what I alluded to in my oral statement, which is much of the problem here is what we don't know and what we can't find out about. And the incidents that we become aware of are probably only the tip of the iceberg. Transparency, as Professor Balkin says, is a critical element in this, but standing alone -- even that's not enough.

Nor are marketplace forces sufficient to restrain the kinds of problems that we're talking about. For most Americans, they have limited choices in broadband providers. And the switching costs through early termination fees, through equipment charges, and bundling with other kinds of services, promotional rates that lock them into contract periods. The switching costs are very high. So, it is difficult to get a true consumer response in that kind of environment.

The last point I would make -- again, this is discussed in greater detail in my written statement -- is, the carriers can't have it both ways. If they want to benefit from the protection that comes from being a dumb terminal that is provided by the 1996 Telecommunications Act, Section 230-C, so that they are immune from tort liability for other people's postings and other people's content over the Internet. And if they wish to take advantage of the protection -- the Digital Millennium Copyright Act from secondary liability, for copyright infringement -- as long as they take it down when asked, then they cannot turn around and claim to be First Amendment speakers at the same time. They benefit from not being speakers, not being editors, from just passing on content, because that's what they're really doing.

MR. BENJAMIN: Let me ask a question now of Ruth and Jonathan. I'm just -- because you all are providing video over the web, have you had any difficulty on either -- with mobile services or home services ever not having your -- not being able to -- people not being able to access your video programs? Is that a realistic fear going forward?

MS. LIVIER: Well, I'm not the techie in my company. But I do -- I have received e-mails from people that, yeah, they haven't been able to download it quickly enough, sometimes, depending on their carrier.

And the big fear with that, really -- and then we look at it and try to fix it and the show is also available through a mobile capability. But the fear in that is that I won't be able to reach my audience and prove that there is a market for my kind of content.

So, it's really important for me to be able to have a direct access to an audience and for their ability to be able to download the show.

I mean, you know, they are paying -- the consumer is already paying for a certain level of speed to download. So, you know, ISPs are already getting paid in that sense.

But my audience's ability to find me is just important to my business. Because that way I can prove my market.

MR. MOORE: For us, it's the fear of competition and being completely locked out. Because the telecom industry, they own Hulu and they own Fancast. Now, we're talking about prioritizing content. Who should -- whose content should be viewed quicker versus someone else's? Which competes -- which is an unfair competitive advantage.

Now, when we start to talk to individuals about quality of content, we have to talk about, okay, if you're on Comcast please forgive us, which puts us in a great disadvantage and completely locks us out of the marketplace. And now the competitive advantage becomes they aren't as good as we are. Now, we're locking out individuals, we can't compete, we're talking about loss of advertising revenue, we're talking about loss of viewerships. We're talking about basically going out of business.

MR. BENJAMIN: So, if you'll forgive me, the academic in me -- let me push back on that a little bit.

MR. MOORE: Okay.

MR. BENJAMIN: And this can be for, I guess -- for all four of you, plus Glenn, who are out there actually producing, you know, Michele, Garlin, Ruth, Jonathan, and Glenn.

So, these companies have spent a lot of money to develop these very fast systems, right? They've got a heavy investment, and they say, well, look, we just want to be able to charge -- we'll charge a non-discriminatory rate. It'll just be a non-discriminatory, significant fee if you want to get high quality of service. And then anybody who wants to can pay that fee and then get access to the higher quality of service. So, I -- on its face, that seems like that's a fair, neutral position.

MR. MOORE: Yeah, but now you're -- it's a gatekeeper mentality. Now you're deciding who and what becomes -- who and what's quality. Now you're separating the individual and the access. And it's all based on capital.

Now, we're talking about discriminatory access. And you're creating an unfair advantage. Yes, I understand a tiered system, but a tiered system is nothing more than a barrier. It's a barrier between access and freedom of expression.

MR. BENJAMIN: I think Garlin and Michele want to jump in on -- Garlin then Michele?

MR. GILCHRIST: Yes, I think that's problematic because the historical analogy -- I'm not a lawyer, as my distinguished panelists are. But that's a poll tax. It's essentially -- it's introducing a barrier that it is thought to be or described as non-discriminatory, but it is discriminatory because it prevents folks who do no have access to resources to pay that fee.

The ability to create content. And if

the purpose, or if one of the plan goals of an open Internet is to facilitate mass participation -- or as close to full participation as possible -- in the democratic process, in democratic and civil engagement and discourse. We need to remove any barriers. Whether they are designed and thought to be discriminatory, or in effect are. Which is what this sort of poll tax or this notion of a tiered system would end up being.

MR. BENJAMIN: Michele?

MS. COMBS: Yeah, basically since we're a grassroots organization, a lot of our state chairmen would have to pay. You know, everybody has their own website. We help them with their website, but they wouldn't be able to afford this and it would really leave out so many grassroots organizations like ours. So, I think it -- I totally agree with all three of you. It would just really hinder our ability to survive as a grassroots organization.

And a lot of our state chairman, too, are small business people -- which all have websites -- which would really hinder their abilities to have their business online that so many people survive on.

MS. LIVIER: We're already paying, I mean that's the thing. I mean, as startups, we don't have a lot of funds, you know. That's in our definition. So we're already paying to put our video up for a certain amount of bandwidth and then the consumer's already paying. So we're already paying for that freeway of communication that's been built. So, I don't know what more they're asking us to pay.

MR. BENJAMIN: Glenn, I can't see you. Is there something you wanted to add in here?

Hearing nothing. All right, let me jump to a question from the Internet which, actually, is best put to Glenn and to Andrew Schwartzman. Many advocates of the regulation proposed in the notice of proposed rulemaking claim that it's necessary to regulate Internet service providers so as to preserve democratic engagement and free speech. However, the proposed regulations, unlike the original four principles, do not cover content providers, even though such, as search engines, which have effectively become gatekeepers.

Should the regulations include such providers? Glenn, are you there? You can take that first, if you want. Hearing nothing, Andrew?

MR. SCHWARTZMAN: The short answer is, no. They are not functioning as carriers of other people's content, there is complete open entry. In an unregulated market, they have not benefited from access to public rights of way, or other kinds of benefits that come as being carriers. So, people on the edge are not properly subject to the kind of regulation we're talking about here. This is for conduits.

MR. BENJAMIN: Bob -- it looks like Bob and Jack want to jump in. Bob?

MR. CORN-REVERE: Yeah, if I could just respond to that?

Andy, I am glad to hear you draw the line somewhere. You know, you might be willing to say that neutrality regulation shouldn't extend to search engines, but you're not in agreement with everybody. I mean, one academic testified before Congress that we should have a federal search commission -- much like the Federal Communications Commission -- because of the importance and the dominance of some search engines out there.

You know, it's the kind of regulatory mission creep that I think is likely to occur if you adopt network neutrality regulations. And also the kind that we've seen in the past with other efforts to maintain and to create from a regulatory standpoint. Some kind of public participation -- I mean, consider the example of public educational and governmental access channels on cable television. They were created as -- and mandated as part of the 1984 Cable Act, and then expanded in the '92 act, and then in 1996, there was a requirement that if cable operators were willing to allow indecent communication on access channels, they had to block and segregate that speech to another part of the -- to their channels.

Again, it's an example of where something was created for one purpose from a regulatory standpoint to serve democratic principles, but then was used for regulatory principles in -- for regulatory purpose in another way. And it's the kind of thing that we've seen again and again.

Even an organization like ICANN, which simply comes up with the -- or manages the domain name system, is susceptible to this as well. In 2005, it was on the cusp of approving a triple-X domain to be used as another voluntary measure for people to be able to filter adult content from the web if they wanted to. But before the contract could be finalized with ICANN, the White House, barraged with e-mails that this would put a stamp of approval on adult material, put pressure on the Commerce Department who, in fact, even threatened ICANN with withholding access to the server if dot-XXX were approved.

And so again, the neutral principles of the domain name system were compromised because once you exert political control over these mechanisms, you distort what you are trying to protect.

MR. BENJAMIN: Jack?

MR. BALKIN: Well, I'm sympathetic to making a distinction between businesses that aren't conduits and businesses that are conduits. Between businesses whose major job is traffic for people, and businesses who basically are producing their own content and shipping it out.

I'm a little puzzled, though, at the argument Bob seems to be making. The argument that Bob seems to be making is that if you have any regulation, it will turn into regulation on an unrelated area.

I don't think that's true. Now, I'm going to be with him on the idea of fighting the expansion and mission creep of regulation that's designed to serve a valuable public interest and then is used for other purposes. So, he and I will be on the same side of that.

But it doesn't strike me that the

examples he gave are examples of how a network neutrality rule would be misused. They're examples of additional regulations that he and I might both be opposed to that I think, you know, is -- presents a different problem.

I think, for example, the search engine problem, I think you and I would both agree that the problem that search engines pose -- to the extent that they pose any problem is very, very different than the problem posed by conduits.

MR. CORN-REVERE: Let me just address that last point, because I think we've already seen examples of that very kind of mission creep. That affects even arguments over broadband and the Internet.

Andy, you cited Section 230 as an example of where if Internet service providers are going to be immunized from private litigation, then they ought to accept the obligations of neutral principles. And yet, if you'll recall, Section 230 was adopted -- or at least, it was proposed originally -- as an alternative to the

Communications Decency Act. And Congress being Congress decided, why not adopt both? And as it turned out, Section 230 has been a valuable insulating factor that has promoted the growth of the Internet by allowing people to make choices whether or not to host content or not. And to have a measure of editorial discretion as well as, you know, to operate as speakers or as conduits.

So, you know, I think it's not hard -- you don't have to look very far to find examples of where even things that were adopted to protect the open Internet are now being used as justifications for extending regulation.

And Jack, I do disagree with you. I think inevitably, when you exert the jurisdiction, the other regulations inevitably follow.

MR. BALKIN: But you're not opposed to Section 230, are you?

MR. CORN-REVERE: No, not at all.

MR. BALKIN: But that was regulation passed by Congress. Why aren't you opposed to it?

MR. CORN-REVERE: That was a measure

that over -- well, it did two things: It overturned a judicial decision that imposed liability on someone who was not exerting editorial control, and it also allowed Internet service providers to provide a service to their customers of discriminating between content that they consider to be objectionable and not.

That is -- not all law is regulation. The network neutrality regulations you're talking about are prescriptive rules to give government oversight over how networks are used, and in some cases, designed. And the impact, I think, is quite different.

MR. BALKIN: So, you're telling me Section 230 is not regulation? I'm surprised to hear that.

MR. BENJAMIN: And on that tantalizing note, I apologize, but I want to respect our audience's patience. It's now 3:01. So, with apologies for the other questions that didn't get asked, I want to thank everybody on the panel for a great workshop.

(Whereupon, at 3:01 p.m., the PROCEEDINGS were adjourned.)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

/s/Carleton J. Anderson, III

Notary Public # 351998

in and for the Commonwealth of Virginia

My Commission Expires: November 30, 2012

FEDERAL COMMUNICATIONS COMMISSION

OPEN INTERNET WORKSHOPS

CONSUMERS, TRANSPARENCY, AND THE OPEN INTERNET

Washington, D.C.

Tuesday, January 19, 2010

ANDERSON COURT REPORTING
706 Duke Street, Suite 100
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190
www.andersonreporting.net

PARTICIPANTS:

Introduction of Workshop and Speakers:

JOEL GURIN
Chief, Consumer & Governmental Affairs Bureau, FCC;

JULIUS KNAPP
Chief, Office of Engineering and Technology, FCC

Opening Remarks:

JON LEIBOWITZ
Chairman, Federal Trade Commission

HON. KONRAD W. VON FINCKENSTEIN
Chairman, Canadian Radio-telephone and Telecommunications Commission

Panelist Presentations:

PARUL P. DESAI
Vice President, Media Access Project

RON DICKLIN
Co-Founder and Chief Technical Officer, Root Wireless

GERALD R. FAULHABER
Professor Emeritus, Wharton School, University of Pennsylvania

JOEL KELSEY
Policy Analyst, Consumers Union

FERNANDO LAGUARDA
Vice President, External Affairs and Policy Counsel, Time Warner Cable

SASCHA MEINRATH
Director, Open Technology Initiative, New America Foundation

PARTICIPANTS (CONT'D):

JAY MONAHAN
Vice President and General Counsel, Vuze, Inc.

NICHOLAS WEAVER
International Computer Science Institute,
University of California, Berkeley (Netalyzr)

DAVID E. YOUNG
Vice President, Federal Regulatory Affairs,
Verizon

* * * * *

P R O C E E D I N G S

(1:32 p.m.)

CHAIRMAN GENACHOWSKI: Let me welcome everyone to this Workshop on Consumers, Transparency, and the Open Internet.

I appreciate the turnout. I appreciate Commissioner Copps, Commissioner Baker being here. I'll mention some of the other people who are here in a minute.

Let me start by thanking Joel Gurin and Julie Knapp for working so hard to put this together. Joel, of course, is the new -- I think I can still say new -- chief of our Consumer Bureau. I think you have another week of newness.

And Julie Knapp -- I don't want to say the opposite of new, but Julie has been an institution at FCC for so long, as we know; runs our Office of Engineering and Technology.

And, as you all know, this workshop really brings together two of the core themes that have been at the heart of the work at the Commission: Consumers and technology. And so to have Joel and Julie work together in organizing this hearing is just terrific.

This workshop is one of a series of workshops on the Open Internet Proceeding that the Commission launched a few months ago, a proceeding designed to develop rules to preserve a free and open Internet.

At its core, the open Internet Proceeding is about protecting and empowering consumers: Preserving users' control over the Internet and their Internet experience. When we talk about consumers here, by the way, we mean -- I mean -- hopefully, this will be something that gets talked about on the panel today -- both ordinary consumers who are interacting with the Internet at home, but also a kind of class of consumers very interested in the topic we're talking about today -- engineers, CTOs, others who want to innovate on the web and are consumers of the Internet in that sense.

The Open Internet Proceeding is about preserving consumers' freedom to access content

and apps of their choosing on the Internet, to produce and distribute content -- the freedom to innovate without permission.

Now all the principles that we've proposed and that had been the subject of discussion at the various workshops that we've had and that we will have have a common purpose -- to benefit consumers. But there's one principle in particular that I'm sure will get a lot of attention today and that is the Sixth Principle that we've proposed, the principle of transparency, a principle that's particularly important for consumer protection and empowerment.

And it's particularly important for at least a couple of reasons in my view. One is making sure that all consumers of the Internet have the best information possible help make markets work most effectively; will help make sure that it's consumers who are picking winners and losers on the Internet. So one, the focus on transparency has that in mind, making the market work, empowering consumers to pick winners and losers.

And second, the transparency principle has the opportunity to minimize government involvement in disputes by increasing the chances that, with full information, problematic behaviors will be less likely to develop and that if anything develops, they'll be known about early and participants in the ecosystem in good faith will have more of a chance to work them out, minimizing government involvement.

And so when we think about consumers and transparency in this area and others, those are important frames for how I look at it. And I think it's -- you know, something -- the comments just came in -- the first round of comments just came into the Proceeding, but I think that this point is something that was embraced pretty broadly by many of the filers. I wouldn't pretend that I've read all 120,000 people and organizations that have filed in this proceeding already, but we do know that a large and diverse group of commenters have said that transparency can play a critical role in preserving the Internet's openness.

We saw many, many different people stand up to make that point. We had a constructive joint filing form Verizon and Google. We've seen in this the growing common ground that is developing and that I think can develop on this issue of such great importance for the country.

Consumers and transparency, I should say, since everyone is assembled, it's essential to the Open Internet Proceeding. It's also something that other members of the staff and the Commission are focusing on in other areas of our work. There's been a lot of discussion, and there will continue to be discussion around consumers and transparency as part of our broadband plan.

Several months ago, the Commission released a Notice of Inquiry, looking at transparency, information provision across all telecom services.

One of the reasons for this is the opportunity that's presented in this area by new information technologies. And there is much less of an excuse now for making sure that consumers have information than there was before, because there are so many different ways that timely, fresh information can be in the hands of consumers and all participants in the marketplace.

We have just a terrific group of panelists here today, thanks to Zach Katz and Joel and Julie for organizing this. Let me note a couple of distinguished participants we have here today and thank them for coming over to the Federal Communications Commission so they can participate.

Jon Leibowitz, the Chairman of the Federal Trade Commission, is here and will speak and participate. Jon is an institution in Washington, extremely knowledgeable about consumer issues, transparency issues, the Internet; and has made empowering and protecting consumers the core mission of the Federal Trade Commission. We're very pleased to have them here. Very much look forward to his comments.

We also have here today the Honorable Conrad von Finckenstein, who is chairman of the Canadian Radio Television and Telecommunications Commission, the Canadian CRTC, which has been struggling with these issues as well. At all three of our agencies, there's been a lot of work over several years thinking about the consumer and marketplace issues raised by this extraordinary new technology.

We're very pleased to have Chairman Finckenstein here participating in this. We have consumer advocates, broadband service providers, content creators, apps providers, developers of transparency and information tools for the Internet. I couldn't be more pleased by this panel that we have put together here. I thank you all for participating, and we're all looking forward to a vibrant, transparent discussion.

Let me acknowledge -- now we have three of my colleagues on the Commission here -- Commissioner Copps and Commissioner Clyburn and Commissioner Baker. And we don't have anything scripted, but, Commissioner Copps, if you'd like to make a few remarks before we start, we'd certainly appreciate that.

And so, as Commissioner Copps comes up, let me thank you all for participating, and, again, we look forward to a vibrant discussion today. Commissioner Copps, please.

(Applause)

COMMISSIONER COPPS: Thank you very much and good afternoon to everybody. I am delighted to be here with our Chairman in welcoming all of you to this very important workshop, and I want to acknowledge at the outset the great work Chairman Genachowski has set in motion here and his leadership in returning this agency to what it's supposed to be, and that is a consumer protection agency. And in that regard, I am really happy to welcome Mr. Leibowitz, Chairman Leibowitz, here who is hard at work on the same job of turning his agency back into what it was always intended to be, a consumer protection agency.

I want to recognize the presence of our

good friend from Canada here, the Chairman of the Commission there -- Conrad von Finckenstein.

It was my great pleasure to be in Ottawa about 18 months ago and visit with him there, and learn of the many good things our friends and neighbors to the north are undertaking.

So thank you all. It's a very distinguished panel, and I'm not going to delay your deliberations further, except to highlight the importance of what it is that you're about -- the principles that we fought for in the Internet Policy Statement four years ago, focused squarely on consumer rights. Again, this is a consumer agency, and while just everybody gains from the availability of an Open Internet, nobody gains more than consumers.

And I have been advocating long and hard for the Commission to establish a mechanism to ensure that consumers have continued access to a vibrant and open Internet.

You know, much as the Internet has transformed us, we haven't yet really begun to realize, I think, the potential of this technology to remake the world we live in. And the opportunities to throw obstacles into the openness and the vitality of the Internet are legion as to such things as consolidating content control or gaming the architecture of the net or creating an environment where people cannot only speak but also be heard.

The tremendous challenges are all ahead of us, and they're all central to this theme of transparency that you're going to be talking about today. So without further ado, I will let you get on with that, but we appreciate the expertise and the dedication that you bring here. And I see my colleague, Commissioner Clyburn here, and would like to welcome her and see if she would like to make a few words of introduction.

COMMISSIONER CLYBURN: Thank you, Commissioner Copps, Mr. Chairman, and all of you incredible guests. It's good to put faces on names that I've been reading about over the -- especially the last several months as I decided to migrate to the North, so to speak.

It's no secret that the primary area in which this current Commission has sought to improve is transparency. That concept applies to both industries we regulate and to ourselves.

With respect to industry, we're in the process of conducting an inquiry into whether or not communications service providers are supplying consumers with the information they require to make informed decisions about the services and plans most beneficial to them.

Internally, the way in which we have conducted our efforts towards producing a national broadband plan, holding a record number of public workshops, issuing numerous public notices, and utilizing every conceivable form of social media known to humankind makes clear that we are committed to a transparent FCC.

But I am not convinced, however, that consumers and content and application developers have all the information they need. Consider the unexplained fees some wireless service providers include in the consumers' monthly bills. Consider the actual speeds of wire line broadband Internet connections appear to be only a fraction of the advertised speeds. And consider at least one known example of undisclosed blocking of consumers' lawful Internet traffic by a major cable provider in 2008.

We have a long way to go. I am pleased to see major broadband service providers participating in today's workshop. Broadband providers play an essential role in controlling the entry and exit points that consumers and content and application providers depend on to access the Internet.

I am eager to hear about how these providers see their responsibilities with respect to adequately disclosing their service terms and how their networks are being managed.

Transparency is a fairly easy word to say, but I sincerely hope that it does not become a mere buzzword. It is essential that consumers and developers have all the information necessary

for them to make informed decisions. I am all for a thriving industry, but not when it occurs on the basis of consumer confusion or misperception.

When all the players have fully and clearly disclosed their services, plans, and management of their networks, only then can we have a market that flourishes on the basis of the quality of services provided and can -- an American consumers can truly reap the benefits of the industry with perhaps the greatest innovative potential among all.

Thank you all for allowing me to share these few thoughts with you this afternoon, and I look forward to working with you to develop a meaningful and robust Sixth Principle. Thank you very much.

COMMISSIONER CLYBURN: Commissioner Baker is going to take (inaudible).

COMMISSIONER BAKER: I will be very brief, because I am just really grateful for our great panel that have come from so far and thrones beyond to inform us. This is a workshop, so we are here to learn.

I want to thank Chairman Genachowski for the terrific processes he's put -- he's put a transparent process together so that we can talk about transparency.

I really do encourage broad participation. I want to also thank Chairman Leibowitz and Chairman von Finckenstein for being here. Your leadership is very important, and we are grateful for your participation.

I was privileged last week to actually go to the Innovation and Investment Workshop. It was in Cambridge, so we had a little less familiar crowd. But hello to those online as well.

These are really turning out to be very, very helpful to us as we look forward to the careful and thoughtful consideration that we're giving to these rules.

There we heard from some industry representatives and some policy experts about the various forces and the considerations that shape the way that Internet entrepreneurs and the network operators interact in their usage of the Internet today, and how they'd like to do it tomorrow.

While there was broad agreement from all the parties on the need to keep the Internet open, there was certainly a divergence of views on the best way to do it.

I think that transparency about Internet practices and policies will be an important factor in taking that discussion forward and may be a path to more common ground. To me, transparency and the open Internet go hand-in-hand. The more that is known about the policies and the practices of a particular Internet service, the better members of the public at large are served.

Whether a casual web user or a hard-core Internaut, a kid with a cool idea for a single app or a software developer who makes her living from launching crazy -- that was the term from last week -- crazy web-based ideas, everyone can benefit from more, better, and clearer information about what happens on a particular Internet connection.

So again, I want to thank Chairman Genachowski for putting this together. We're going to develop an ample record, even more than we already have. So as we move forward, if we make any decisions, it will be based on a record.

I really do think that these issues are critical to the future of the Internet. I think it's critical to the successful deployment of broadband to all Americans, and I think it's really critical to our economy. So thanks for being here.

CHAIRMAN GENACHOWSKI: So let me again thank each of my colleagues. I think one of the strands that you hear from all of the comments, from all the Commissioners, are the importance of transparency, the importance of this process to air ideas, to roll up our sleeves, to really understand how we can make sure that this vital medium serves the country well.

So thank you, Commissioner Copps, Commissioner Clyburn, Commissioner Baker, for participating. These are the kinds of discussions we have among ourselves as well.

And let me thank each of you for saying something about this process. What we're having here today is part of our ongoing experimentation and focus on improving the processes of the FCC, making them more open, having them be transparent, and facilitating a narrowing of the issues and a rolling up of our sleeves to figure out what the real questions are, what the facts and data are, how can we best solve them.

We're experimenting with many things here as we do this. This proceeding is open to the Internet and I think, as you'll hear from Joel and Julie, there will be opportunities for people who are participating remotely to ask questions. We encourage just vibrant, open dialogue, and these are staff-level working sessions at these workshops.

And so I thank you all for just saying what you think and helping Joel and Julie and the staff and all of us get our arms around these incredibly vital issues for the country.

And so, with that, thank you again to my colleagues on the Commission, the staff who put this together, each of the participants, those of you who are here and participating online, and, Joel, will you take it from here?

MR. GURIN: Yes, I will.

CHAIRMAN GENACHOWSKI: Terrific.

MR. GURIN: Okay.

CHAIRMAN GENACHOWSKI: Thank you all very much.

MR. GURIN: Thank you very much. Thank you, Chairman Genachowski. Thank you, Commissioners. And thank you for laying out so clearly and really so forcefully why we're here today.

We are looking at an extremely important issue that I think we all believe can affect the future of the Internet very significantly. And in looking at the Open Internet, what we are particularly focus on today is transparency and the Sixth Principle.

So I am going to speak just for a minute or two, and then turn it over to Julie, who will talk about our -- how we're going to go through our process today.

I think it's worth, since this is all about the Sixth Principle, I think it's worth actually reading it aloud for everybody. The -- what we are proposing as the Sixth Principle is that "subject to reasonable network management, a provider of broadband Internet access service must disclose such information concerning network management and other practices as is reasonably required for users and content application and service providers to enjoy the protections of these Open Internet principles."

So a few things about that: One is that the Sixth Principle includes both disclosure of information on network management and other practices. And I think it's almost impossible to talk about network management without talking about other issues like broadband speed, broadband performance, and issues that are related to the consumer experience.

But I do want to emphasize that the particular focus of this workshop today really is on network management principles and that whole array of issues that relate to how open the Internet is in terms of the consumer experience.

We are going forward and discussing the Sixth Principle with a lot of belief that transparency can be tremendously valuable, both for consumers and for ensuring a competitive marketplace. And we'll be asking a number of questions today.

Some of the kinds of questions that I think will get into include: What kind of information consumers really need about broadband network management practices and other practices. What kind of information do they have today? What tools do they have or could be used to get at information that is not currently disclosed? How should this information be presented? And how do these issues vary from wire line to wireless broadband service?

And finally, what are the appropriate roles of the public and private sectors in defining these disclosure policies?

So I do want to thank everybody on this panel. We are very fortunate to have such a remarkable group of speakers today, and look very -- very much looking forward to hearing what you have to say, to having a good a discussion.

I also want to thank, on my right, Julie Saulnier and Rebecca Hirsel, who put a lot of work into developing this workshop as well. And without further ado, Julie, please take it from here.

MR. KNAPP: Thank you, Joel. Thank you Mr. Chairman and Commissioners. We have an exciting lineup for this afternoon. It is -- as many of you all know, we also have a technical advisory process that has been working primarily focused on the Fifth Principle, with non-discrimination. But we've also integrated with all of the other areas of the rulemaking, including the Sixth Principle, and I've been working closely with Joel and his team. And sometimes the biggest challenge is taking very technical information and translating it in ways that are meaningful in this case not only for consumers and businesses, but for innovators and application developers, who really need to not just had the high-level information but to be able to drill down into the network management practices that the providers have put in place.

For this afternoon, we're pretty began with framing remarks from two senior officials from the FTC and CRTC, and we're very fortunate to have them with us this afternoon. And I'll introduce them momentarily.

Well, then have each of our panelists make a very brief presentation and to please try to keep their remarks to about five minutes. Then after a short break, we'll have about an hour for discussion, including questions and answers not only from Joel and I, but -- and also from the audience and those who are following online.

If you're here in the Commission Meeting

Room, write down the questions on note cards that are on a table near the entrance -- you'll be able to pick those up at the break -- and then give them to the person we'll have at the back of the room collecting them.

You can also e-mail questions to newmedia@fcc.gov or post comments and questions to Twittermarket with hashtag number oidiscuss.

On our website, OpenInternet.gov, you can find the agenda for today's workshop and bios for all of the participants, and you can watch live streaming of the proceedings.

With that, I'd like to introduce Chairman Jon Liebowitz. He has nearly two decades of a very distinguished career in public service. Of note, he served in the United States for 14 years in various positions, including 3 years as Democratic chief counsel and staff director for the Senate Antitrust Subcommittee; and 1 year as the chief counsel and staff director for the Subcommittee on Terrorism and Technology.

In September of 2004, he began serving

at the Federal Trade Commission as one of its Commissioners and he became Chairman in March of 2009.

It's my great honor to introduce Chairman Liebowitz. Thank you.

MR. LEIBOWITZ: Thank you, Julie. Let me just ask a quick question: Do you want me to sit at the table or should I go up? I'm neutral as to venue, you know, or platform, but as long as it's transparent to everyone.

MR. KNAPP: I think we've got you covered either way with the cameras, so just if you're comfortable there, that's fine.

MR. LEIBOWITZ: I'll stay here. Well, thank you so much for that kind and undeserved introduction, and I just want to also thank Chairman Genachowski and Commissioner Baker, Commissioner Copps, and Commissioner Clyburn, Commissioner McDowell, who I don't see here, but I know is peripatetically in and around for having this workshop and for inviting me to speak.

And, you know, I was thinking that --

and I don't consider myself to be a Washington institution despite what Julie has said about me. I will be remembering that when we play basketball together next Sunday.

But I will say that as I look around, this is really one of the strongest Commissions I have seen in my not-quite-lifetime in Washington, and it's also great to see the transparency here at the Commission, which had (inaudible) this I think we all know in the past. So we are following your work at the Commission with enormous interest, and very impressed with the energy and the vision that the Commission and that really you, the staff working to come up with ways to implement that vision are bringing to a notoriously, I would say, difficult and vexing issue.

It seems that you really have begun to cut through the sort of dystopian futures envisioned by each side of the Net neutrality debate if the other side's policy prescriptions are adopted, and there's no doubt in my mind that you're headed towards a reasonable, thoughtful, pro-consumer middle ground with your rule.

As everybody recognizes, transparency and an Open Internet are critical components of future broadband services to consumers and of further innovation in this incredibly dynamic and industry, and this is especially true, I think, on the content side.

But it's also clear to me that, absent some efforts by the FCC, those principles are not going to be certainties in the Internet of tomorrow.

Will carriers slow down or interfere with applications or services? Will some sorts of services be prioritized over others? If so, will consumers be told about any of this before they sign up, and, if they are told, will they be told in a way that they can understand it without having, say, a Ph.D. in electrical engineering? And at a time when they can make a choice whether or not to accept the terms. Will they be told if the terms of service change?

These are all important issues for consumers as they buy Internet access, and indeed they're sort of at the heart of what the FTC sees in a variety of other industries on an everyday basis. These questions implicate some of our core concerns: Whether consumers are told what they're paying for and whether -- so that they can make informed decisions, of course, and whether there's competition in the marketplace.

Let me start with the consumer protection component, which really boils down to what will people know and when will they know it. So, I suppose it always really boils down to what Howard Baker said in the Watergate Hearings. Okay. That was a really good joke, but it was an attempt at humor.

I am happy to see that the FCC has proposed to add a new principle of transparency to its four Internet Freedoms. With adequate transparency and disclosure, consumers are able to choose winners and losers in the marketplace. They can pick the technology, services, and companies that best fit their needs at the prices that they're willing to pay.

Providers that offer the best deals can grow and serve even more consumers. And this is true in any market, but I think it's especially true in a market like broadband, where consumers may not know what services they're buying when they buy it without adequate disclosures.

At the FTC, we've been working on issues that involve transparency and disclosure to consumers for many years. And these technologies themselves are very familiar to us. For example, in 2007, we investigated claims by Sprint Nextel that it was offering unlimited mobile web access to its subscribers, when, in fact, it was actually selling services that were limited by a monthly ceiling.

Now we ended up closing that investigation because the claims appeared not to be part of a broad marketing plan to mislead consumers, and Sprint acted quickly -- I think it was a Nextel plan, and they acquired this plan --

and then they acted quickly to remove the deceptive claims and to voluntarily make refunds to subscribers who may have been harmed by its failure to disclose the broadband access ceilings.

But, as staff noted in its letter to Sprint indicating that it was closing its investigation, we, and I quote, "take very seriously the sorts of claims investigated here and will continue to monitor such claims made by Sprint as well as those made by the rest of industry."

And transparency is important for another reason: Unless the FCC knows what ISPs are doing to manage their networks, how can you make the necessary distinctions between reasonable network management that allows ISPs to provide better services to their customers and the potentially abusive behavior that might harm consumers. I'm not saying anyone engages in potentially abusive behavior, but I'm saying to make that rulemaking as robust and as informed as possible, it's important that you understand what carriers are doing or what companies are doing.

Disclosure of network management techniques and plans by the providers is vital if we're to understand what industry best practices are. And it seems to me that ISPs should be talking about what they believe are reasonable network management practices and what they believe are not, if they want a rule that benefits them and the American consumer.

Let me be clear: I welcome the FCC's involvement in this area, and clearly, from the opening statements of all the Commissioners, you recognize that the principle of transparency and disclosure is enormously important. If you move forward with this rulemaking, and particularly with this provision, we're going to continue our close working relationship and leverage the expertise of both agencies.

The other principle that the FCC has proposed to add is non-discrimination. Broadband providers must, and I quote, "treat lawful content applications and services in a non-discriminatory manner, subject to reasonable network management."

From my perspective, some form of antidiscrimination language is critical to ensuring an open Internet, and reasonable network management is also critical to allow ISPs to serve their customers better by, for example, managing network capacity, stopping SPAM and spyware, something that our agency has an enormous amount of familiarity about.

Non-discrimination is a somewhat controversial proposal, and one that the FTC staff wrote extensively about in the wake of our Broadband Competition Workshop in 2007. In the report that we wrote on broadband competition, staff reviewed the record created by the two-day workshop and the evidence that existed at the time of the state of broadband competition in the provision of broadband services.

The FTC staff made what I believe in 2007 and still believe today to be a useful, a very useful contribution to the debate. Indeed, the Commission voted unanimously to issue the report. But let me just caution those who refer to the report in the current debate to do so carefully.

In particular, it should be clear from reading the report itself that while staff advised "proceeding with caution before enacting broad ex-ante restrictions in an unsettled, dynamic environment," they also recognized that a failure to act could be problematic.

Further, and I think all of us know this, and I -- when you're doing evidence-based public policy work, it comes with an expiration date. And it reflects the possibility that things can change after the work is finished. To some extent, that applies to the FTC staff's 2007 broadband report.

It has been nearly three years since that report was issued, which, as we all know, is an eternity in Internet time. Just as example, the report was completed before the Comcast Bittorent controversy that I suspect will be -- in fact, I was looking at the panel -- I know will be a topic of discussion today, and, as a result, it didn't really focus on the question of reasonable network management as a component to a balanced net neutrality regime.

Taken together, experience suggests a reason to question now, I think, what was a reasonable concern then about potentially premature and unnecessary regulation.

One of the critical predicates of staff's caution regarding pulls the road for the Internet in 2007 was the presence or the possible development of competition in broadband markets. As staff noted, competition provides the best results for consumers, providing the lowest prices, the highest quality products and services, and the most choices.

At that time, staff working on the report focused on disagreements over whether there was competition in the industry. It is a -- it's a stretch, but you could have read the FCC's statistics available at the time when the report was written to indicate that consumers in many ZIP codes had access to a very large number of broadband Internet access providers. But I would say it's seems clear now that a better reading of those statistics and an analysis -- a better analysis of today's marketplace is that broadband Internet access is, at best, essentially a duopoly in many markets.

Now that's not to say there isn't price competition, because I think there is a lot of price competition, at least right now while competitors are trying to sign up subscribers. But nearly all consumers buy their broadband Internet access from either their cable provider or their local phone company.

And more importantly, in the report FCC staff also foresaw -- and I have to say we were all hoping for -- a potential entry from a number of new actors who hadn't provided broadband access to consumers in any volume before. Among those were Wi-Fi, WiMAX, other broadband technologies. And what we've seen since then, however, is that the hope for competition has largely not materialized.

Most critically what seemed at that time to be a burgeoning movement of municipally sponsored broadband has never occurred, and, by the way, the first report we wrote, before we wrote our broadband competition report, was a report that raised questions about those companies that were going around from state to state trying to restrict -- pass state laws that would prohibit cities from offering broadband.

Thee fact that the broadband -- the fact that municipal broadband hasn't been terribly effective or hasn't been the competitor that people think is a good reminder that you really ought to let the marketplace work rather than trying to, even if you have the capacity to, sort of ratchet back anti-competitive or implement anti-competitive rules.

The WiMAX roll out also appears to have gone far more slowly than anticipated, and I don't think I need to say much about broadband over power lines. Probably the less said, the better.

On the other hand, mobile broadband, spearheaded by the iPhone and Google's Android, has grown faster, and it may become a game changer someday, but I think we all agree that day hasn't quite arrived yet.

So I would also say this: To the extent of the report was also skeptical about the need for rules, given the ability of anti-trust enforcement to ensure competition in these markets, it appeared to me, even at that time, that staff was overly confident about the ability of anti-trust law to deal with net neutrality-based concerns.

And since we released the report, my colleague, Commissioner Rush, has also looked at these issues, and has come to similar conclusions. To be sure, anti-trust is a useful vehicle for thinking about these issues, and the FTC has a unique and broad enforcement authority under Section 5 of the FTC Act to stop unfair methods of competition.

But it is reasonable for, and

appropriate for the FCC to consider the possibility that such enforcement may take too long and do -- not do much good in many cases.

And so, I guess I started to read, and I have not finished reading the Google-Verizon submission today, and to the extent that it suggests that the FCC shouldn't be doing a rulemaking which arguably it does and I think paragraph six I do take some issue with that.

I think this rulemaking is critical. I don't take any issue with the notion of the FTC staying involved in these areas, of course.

The FTC report recommended a cautious and thorough approach to any net neutrality regulations, and, of course, and that's the right approach to take. But to my mind, the FCC's initiatives, demonstrated by this workshop, your thoughtful NPR to begin the process last year does this job admirably. The rules that come out of this process have the potential really to benefit everyone, not just consumers, who need to understand what services they will receive, but also the businesses that provide these services.

And the need for predictability for these firms is critical, given the long-term investments made by many in the industry, including ISPs. It seems to me that the current proposal does a better job of addressing those concerns ex ante than an ad hoc approach that leaves me uncertain about what they can and what they can't do.

And clear rules of the road, by the way, are a much better defense against the public outcry over a questionable practice in an unregulated environment that will sometimes force companies, including companies in this space, to back track on important investment decisions long after those decisions have been made.

As David Cohen from Comcast recently noted, we are all obviously better off having clear rules. I will follow that admonition with another clear rule, which is that it's better to stop speaking before people are really tired of hearing about -- hearing you talk, and let me now introduce this terrific and interesting panel on Transparency and the Open Internet.

Julie, do you want to do that or should I do that? It's in my talking points, but I have a feeling that maybe you're prepared to do that, too.

MR. KNAPP: Yeah, I think --

MR. LEIBOWITZ: I think I'll stop right here, and you can turn it over to my colleague, from my our good neighbor to the North.

MR. KNAPP: Thank you very much. I look forward to continued working with the FTC very closely. The Honorable Conrad von Finckenstein has dedicated his 37-year career to public service, and I can only highlight a few of his notable achievements.

He served as chairman of competition and the head of the competition in Bureau of Canada for six years, from 1997 to 2003. He also led the drive to establish the International Competition Network and became its founding chairman until being appointed as a Justice of the Federal Court.

Mr. von Finckenstein served as a Justice until three years ago, when he was appointed to chair the CRTC. We are very pleased that he has joined us today, and look forward to his sharing his experiences on dealing with similar issues up in Canada. Thank you.

MR. von FINCKENSTEIN: Thank you very much. Thank you -- a very kind introduction.

I'm delighted to be here. I've never participated at an FCC Workshop, but I've ready many results of them, and I think it's a wonderful innovation, which we should probably copy in Canada.

I wanted to share with you today our views on network management. We've struggled with the same issues, and we've had recently made a major decision on it, which I wanted to explain to you.

The CRTC gets its mandate from the Canadian Broadcasting and Telecommunications Act. Internet Service Providers are subject to the Telecommunications Act.

And from -- aside from mandating access to wholesale customers, something which you don't do, we look at fair and reasonable terms of conditions. We monitor the activities, and we act on complaints about their business practices.

Obviously, with consumers and broadband connection becoming increasingly a part of our society and being essential for commerce, health, education, and the human government, these issues are coming very much to the front.

And the demand on the network capacity is rising, with many new applications, especially such things as real- time video streaming, gaming, HD video, and (inaudible) 3D.

We actually had an actual case, which pointed out to us -- out -- the need for some rules in this role, and so we dealt with that case on its terms, but then we held a major policy hearing on what you call network neutrality, which we call Internet Traffic Management, but it's all the same.

And basically, the question obviously

came up: To what extent should Internet providers, service providers, be committed to control or influence the users' (inaudible) of experience. To what extent can they protect the legitimate interests regarding the integrity of the networks without impinging on the use and the enjoyment of the Internet by its users.

We held a public hearing basically covering the six points that the FTC has announced, and we had numerous written submissions. We had online consultations. We had gotten more than 14,000 online comments, and we had a 6-day public hearing.

In October 2009, we then announced our decision, and that's what I want to share with you to now.

As background or short, when we started the proceedings and (inaudible) all Internet proposed service providers participated, it became clear that some are using ITMPs right now. Some are using only at specific times. And some are not using them at all.

And certainly the public had no idea who was doing what. For approach purposes, we found out the best way to do it is to, in effect, talk of two separate practices -- economic practices and technical practices.

Economic practices obviously as those (inaudible) charge consumers depending on their use and how much bandwidth they use.

Technical matters are such things where you use tools such as deep packet inspection to manage the traffic of networks by slowing down or prioritizing certain types of (inaudible) traffic or limiting heavy users. We have -- our approach applies both to retail and wholesale, but I'll only talk with the retail, because you don't have a mandated wholesale.

We -- as a result of the hearing, we firmly reaffirmed our often stated views that the first response of any ISPs to bandwidth demands should be make investment and increase (inaudible) whenever possible.

Secondly, we said that you may employ

Internet Traffic Management without our prior approval as long as they're consistent with the framework, which we established, which I will outline in a moment.

MR. von FINCKENSTEIN: And we said you should always prefer economic ITMPs over technical ones because they're obviously transparent. They're inherent in the effort (inaudible) they pay for. And certainly and then if you use economic ITMP, so you have to -- they have to be in a manner consistent with our Telcommunications Act, which basically has two principles: No unjust discrimination and no unreasonable preferences.

We will look at these matters ex post facto if there is a complaint and see whether they have done it properly, and the frameworks that we will use to value it are based on, as I say, these two concepts of unjust discrimination and undue preference.

Our Act states, "No Canadian carrier shall, in relation to the provision of telecommunications service or the charging of a rate for it, unjustly discriminate or give undue or unreasonable preference toward any person, including itself, or subject any person to an undue or unreasonable disadvantage."

Our framework works as follows: If we receive a complaint from a retail or wholesale customer about an Internet traffic management, we will use the following criteria. First of all, is it -- is the ITMP in question, is it designed to address a specific purpose and nothing else?

Secondly, is it as narrowly tailored as possible to achieve a desired result using the least restrictive means?

Thirdly, it causes as little harm as possible to the retail customer, application provider, or the ISP that's a wholesale customer.

Fourthly, network investments or economic approaches would not effectively deal with the same issue or achieve the same purpose.

And lastly, it is well advertised in advance. Now there's one exception to this rule, which is of ex post facto review, which is if you do anything to block the delivery of content or you slow it down to such a time that it becomes inherent -- incoherent and the content is degraded, then you have to come and get prior approval from us.

Those would be the most truly exceptional circumstances, and I frankly can't envision that who would ever give it. But if the ISP feels for traffic management purposes, they have to affect content or they have to degrade it so such in terms that it is no longer comprehensible, they have to come for prior approval.

Now we've (inaudible) learned about the issue of transparency that this hearing is about. Transparency is one of our four key principles. At the CRTC, we have four key principles. There's fairness, so a decision has to be fair. Diligence. They have to be as quickly as possible because time is money. They have to be transparent so people understand it. And they have to be predictable so everybody knows the rules of the road.

And since we apply it to ourselves, we obviously expect it from the people whom we regulate as well.

And as it was stunning that most of the consumer groups came before us were not even aware of the fact that their ITMPs and their being applied -- there was no information on anybody's website. It was not disclosed in the billing practices or in the contracts that ISPs signed with their customers, not even in the fine print.

We decided to establish the following rules: First of all, if you're going to use an ITMP, you must inform your users 30 days in advance before you can apply it. And this notification must be in plain language, displayed prominently on your website, and has to describe precisely in an intelligible form so that a consumer, an average consumer, could understand it. Number one, why is this ITMP being introduced. Two, who will be affected by it?

When will it occur, from what hours, et cetera? What type of traffic is subject to the Internet practice? Is it limited to a specific application or specific type or is (inaudible) to every user? And lastly, how will it affect the users' Internet experience, including the specific impact on speeds.

This is vital information so that the consumer can make an informed judgment and can decide whether he wants to vote with his feet and go to another provider or stay with this provider, because you're fully conscious that the people don't buy ISP services alone. They usually buy it in bundles, and that information may not be sufficient, but at least the consumer will have know and will have the choice whether he wants to break out of the bundle or go to somebody else and buy a new bundle or whatever.

We also dealt with the privacy because part of our mandate is privacy, and we found that any data that you collect or that you create for the purposes of administering an ITMP can be used for that purpose only and no other purpose.

It can also obviously not be resold or used for marketing or any (inaudible). ITMP information has to be as soon as you don't need it anymore, you destroy it, and you only use it for that purpose.

Now the mobile wireless, which as Chairman Leibowitz mentioned, is becoming more important. In Canada right now, until the data site is exempted from our regulation. We did that about 10 years ago. I mean it was far less important obviously than today.

And we firmly stated where we expect the mobile providers to act on the same principles as the ISPs. If not, we will revoke the exemption either partially or totally in order to deal -- give them exactly the same treatment as we have given the wire line ones.

So, in conclusion, our framework is intended to foster the environment for ISPs, application providers, and end users to have the utmost freedom to innovate. At the same time, we have to respect the legitimate interests of ISPs to manage their networks.

It was a balancing act, and first results would seem to indicate we have hit the right balance. This is the only decision that I've ever taken where everybody who participated has been laudatory at the end. Whether it was users, consumers, technology companies, ISPs, any other group, everybody said well, more or less, you got it right.

The decision came into effect 30 days after we made it, i.e., in November, and is now being -- we are carefully watching the websites of the ISP providers whether they have lived up to it. But it would seem they have got as always we basically are complaints driven. We wait for consumers to come forward and say, here's a problem, et cetera. And if there is, we will investigate it, and we will apply the framework that I outlined to you.

This is so much for the Canadian experience. Thank you for inviting me, and I look forward to your discussion.

MR. KNAPP: Thank you. I'm struck by the fact that we have stars and so little time. I hope you'll forgive me. We've got the full bios for each of our speakers up on the website. I can assure you that they are incredibly impressive, and are experts in all of the topics that they're going to talk about.

So without further ado, I'll just introduce the speaker, and we'll try to keep it to five minutes and move it through and so that we can get onto the second part of the session with the questions and answers.

Sascha Meinrath, you'll go first? Surprise. Keep them all on their toes. He's the director of New America Foundation's Open Technology Initiative. He has other titles as well.

Sascha, the floor is yours. Thank you.

MR. MEINRATH: Thank you, Julie. Well, I wanted to begin by thanking the Commission for granting me this opportunity to discuss transparency and data collection, and the efforts that I and many others have done over both the last few months and years, but also dating back many years and decades beyond that.

The Open Technology Initiative is a think tank within a think tank. We formulate policy and regulatory reforms that promote affordable, universal, ubiquitous communication systems. As you can probably gather by its name, you know my biases on this.

And we provide also in-depth, objective research, analysis, and findings for policy makers and the general public. And the nuts and bolts of what I'm going to talk about today are actually covered, in-depth, in our December 14, 2009, comments on Public Notice No. 24.

And in that 115-page filing, you'll find extensive technological documentation and both real-world options of our data collection efforts, as well as context from the research literature on what's been done previously and how problems that have been raised by some commentators and probably some of our panelists have already been addressed.

Now this second -- as we begin the second decade of the 21st century, I am utterly astounded by the level of ignorance presented to the FCC when it comes to the history of data reporting on the Internet.

And I say because first and foremost, the Internet was and continues to be a research endeavor. And as a research initiative, data collection and transparency have been at the heart of this endeavor since its very inception.

So many of the statements made about the issue being too complex or technologically infusible simply do not hold water, and many of the solutions have been around for years, if not decades.

So let me be clear: Almost all of the useful information that we would like to see made publicly available is already being collected by system administrators and ISPs.

And for the research and scientific community, the problem isn't how to collect information that would be useful to consumers and policymakers. The question is really how is it that the successful public data collection practices stopped in the first place?

And to drive this home, I found a memo from January 9, 1995, exactly 15 years ago this month, from Merritt, which ran NSF Net. And I'll quote very briefly from this memo. Now this was right at the time the NSF Net was privatized. And they said, "NSF Net performance statistics have been collected, processed, stored, and reported by the Merritt Network since 1988. During December of 1994, the numbers contained in Merritt's statistical reports began to decrease, as NSF Net traffic began to migrate to the new NSF network architecture, i.e., was privatized.

"Once a new architecture is in place, Merritt will be unable to collect the data needed to continue these traffic-based reports. The reports will be discontinued by spring of 1995."

Since then, as we have seen, year after year after year, there's been a steady removal of useful information from the public domain. The data collection requirements mandated under NSF Net gave way to voluntary adherence to these norms and the data collection and transparency that had happened was eliminated with disastrous consequences.

So when the story is written about how the United States went from number one in broadband service provision to our current appalling international standing, it's clear that were for ignorance has led to a series of unbelievably shortsighted policy actions and inactions.

And this loss of useful broadband information has systematically disempowered users of the Internet and allowed for the creation of increasingly dysfunctional markets.

However, all is not doom and gloom. The FCC has both an opportunity and the responsibility to rectify these problems. Now I'll focus the remainder of my time on two simple solutions to the problems that I've laid down.

The first is ATI's Truth in Broadband Advertising proposal, which has gained a wee bit of traction since we first proposed it, and the second is the Measurement Lab.net Initiative.

Now when we proposed a broadband nutrition label, as someone has called it, we were drawing both from the notion that consumers want meaningful information and that they have an increasingly diverse array of needs for their broadband connections.

Our Truth in Broadband Advertising proposal would ensure that specific information is made available to consumers, for example, upload/download speeds, uptime, latency and pricing, and that private industry had the opportunity to decide what levels of service they would guarantee to their customers.

In fact, this practice is already being done today on most business-class lines. They contain levels of service agreements. They can turn -- contain guarantees of available uptimes. And they enable businesses to make informed decisions about and, contrast at most, multiple service offerings.

The FCC has the ability to implement these disclosures uniquely, but is under substantial pressure to avoid meaningful information for consumers, and adopt a five- star rating that could utterly gut the intent of OTI's proposals.

Consumers want meaningful information, and the FCC has a responsibility to create clear disclosure rules that ensure customers have access to this fundamental information about their broadband service offerings in both the wire line and wireless world.

A five-star rating could be an information gateway, but only if the FCC also includes a clear mandate that the useful information underpinning this rating system is also disclosed and easily accessible to consumers.

Of course, in order for consumers to test their broadband connections and know about them, they first must have the tools to collect this information. The FCC should leave extensive effort to measure and collect fundamental data on their broadband service capabilities and Internet performances, traffic statistics, et cetera, across the country.

This effort should contain both nationally comparative information as well as a level of granularity that allows customers to compare offerings within their own neighborhood, amongst their friends, and across different Internet service providers.

Now within this context, the Open Technology Initiative has led a global coalition of researchers from Planet Lab and other institutes, along with a number of corporate partners, to pioneer a unique broadband measurement platform, the Measurement Lab.net Project.

And the Measurement Lab.net Initiative, or MLAB for short, is itself an open distributed server platform for researchers to deploy Internet measurement tools, and our goal is really to advance both network research, on the one hand, as well as to empower the public with useful information about their broadband connections as well as -- and to give policymakers the information they need to make informed decisions.

So every broadband measurement tool on MLAB is entirely open source, which means that if you have concerns about the objectivity of the tests, we'll give you the code, and please look through it and tell us where there's a problem.

MLAB is open to participation from all quarters, which means that if you want to be involved, if you have a measurement tool that you would like to put onto this platform, you can do so and you should contact us immediately. And all data collected by Measurement Lab is made publicly available under a creative commons zero license.

So anyone who wants to crunch these numbers or use them for any purpose whatsoever may do so. At its heart, MLAB is an open, independent, and transparent process for developing a suite of Internet measurement tools, for collecting data in an objective manner, and for ensuring that this useful information is made publicly available in a timely manner.

So the FCC has the opportunity to leverage these assets and create tools and information resources needed to empower consumers. So, here's my take-home messages.

First, systematic data collection efforts and the public release of these data have been a part of the Internet since its inception, and only ceased in the mid-1990s when NSF-Net was privatized.

Second, the best metrics for the data and the needs that we have for collecting this data and the processes for collecting this information have already been identified in a prototype system set up and fully operational.

And third, that the technological and scientific underpinnings have already been established, and, thus, this really boils down to an issue of the need for clear leadership from the FCC.

So, with forthright leadership, 2010 should be the year that consumers are finally empowered with meaningful broadband information, and the FCC started turning around broadband in the United States. Thank you.

MR. KNAPP: Thank you, Sascha. Our next speaker will be Jay Monahan. He is the vice president and general counsel of Vuze.

MR. MONAHAN: Thank you.

MR. KNAPP: Jay?

MR. MONAHAN: Thank you. I'm honored and pleased to be a part of this debate, to have been a part of this debate for a couple of years now. And originally became involved because we actually are a real company. We're that company I think they talk about when they talk about innovation and investment.

We are an online video platform that has raised three -- sorry -- $30 million in venture capital funding from Sand Hill Road in Silicon Valley to create and distribute a platform for delivering content; also some might consider competition for the established players.

About three years ago, we, like many of you, were reading probably online or your hard copy newspaper and saw an article about a guy named Rob Topolski that discovered what he thought was something strange going on with his Internet connection. As you all know at this point, he went to the AP Wire Services. They replicated some tests that he had done.

And that led to what we now know as the interference with the Bittorrent traffic by Comcast.

Now when that story broke, somebody in my office, which, as we pointed out at the Harvard hearing, our 20- person company above the Chinese restaurant in Palo Alto -- we're actually now above the bank, just for the record -- somebody in the office said well, hey, aren't they talking about us. And, in fact, that is exactly what was going on because our product, our content, which is licensed content at that time for many studios, was being delivered using the very technology that was addressed by Comcast's particular techniques at that time, which the Commission has since condemned.

Later in this debate, we heard that another carrier at least announced that it was beginning some tests to prioritize traffic, and, according to its public announcements and its website, it was going to classify certain traffic as time sensitive and certain traffic as non- time sensitive.

Well, the non-time sensitive traffic included all peer-to-peer traffic. The problem, of course, is that the technology that we used to deliver our licensed content is technically peer-to-peer traffic, and it was all classified by them as non-time sensitive.

Now the problem at the time, just to drill down a little bit so you get sort of the sense of how this -- how the rubber really hits the road, is that we had two features at the time.

One was what we called progressive downloads, which was a streaming-like experience where you literally hit a play button and then with a few seconds delay it would play, and, of course, the delay we were working on getting shorter and shorter and shorter; so it was essentially equivalent to what you think of streaming when you hit, you know, a video on YouTube.

We also were working on something called peer- assisted streaming, which also worked for HD content. Peer- assisted streaming works such that it -- you could -- or we could obtain the content from multiple sources and eventually some of the sources that would be used to deliver the content would come from peers who had already downloaded the -- whatever the video was.

So, in both cases, these were certainly products that we considered to be extremely time sensitive, and when we read that article, somebody in the office said, hey, isn't that us they're talking about.

And that is exactly what's going on here. The -- what I want to bring to this discussion -- obviously, the consumer disclosure is critically important, particularly for consumers who want to know whether a particular feature that we offer will actually work.

But when you are the competitor, when you're getting millions of other people's money to build a product, to run a product, to expand a product, we're not, from a business point of view, it's difficult to rely on having the next Rob Topolski out there, who's going to discover something that may be going on behind the scenes.

The particular technique, of course, that Comcast used was condemned by the Commission. There may be other techniques that may be perfectly permissible, but as a person helping to run a company and as investors, we still want to know exactly what those things are, because it may mean that a particular service is not feasible or it's not reliable.

Another reason for disclosure, which you appreciate when you are sort of hearing from the consumer directly, is that when there is an interference with the delivery of a product, like a streaming product, we know that consumers are very sensitive to the amount of time that it takes for pages to load, the amount of time it takes when you hit the play button.

And when there is something going on potentially behind the scenes that interferes with the delivery of the product, of course, the consumer says it must be us, meaning it must be the company's product, and it doesn't necessarily occur to them that it's actually somebody in a backroom, you know, forging packets on particular types of traffic.

So I think it certainly makes sense from our point of view in order to deal with our own consumers, we're certainly willing to take responsibility for the defects of our product or the infirmities of our own product. We should not have to take responsibility for interference that may be coming from other parties.

Were there adequate disclosure, we certainly would be able to note that certain performance issues have to do with your ISP and not with us. Obviously, those folks are going to be able to make their own decisions as to which ISP to use.

Now that said, as has been aptly pointed out, and I'm sure will come up further in this panel, we don't think that you can approach transparency issues without considering the fact that there is so little competition now in consumer choice. To that end, disclosure by itself is not sufficient unless you also have a network management tools that go with it to identify which practices are permissible and which are not permissible.

I wanted to also note that we appreciate that there are certain types of practices that go on that relate to security. There are certain types that go on that relate to emergency tactics that need to be taken by carriers. There probably will be less disclosure available in those cases. We understand that, but we believe that for the everyday practices about upload speeds and download speeds -- all right -- probably should say continuous upload speed and continuous download speeds, not just peak speeds, and any tactics or any techniques that are used that would slow down traffic or block traffic or delay traffic ought to be disclosed.

One another note that I think was -- a point that was alluded to earlier that this debate sort of-originally, there was a bit of a semantical debate about blocking versus slowing. And I think as a technical matter, what happened with the Torrent traffic with Comcast was, in fact, blocking.

But when you are offering a consumer device, it is equally important -- in some ways more important -- for us to know and for the consumer to know that traffic is being slowed as opposed to blocked because when traffic is slowed, even if it's still working, if your competitive alternative is working at full speed and you keep experiencing delay -- and just think about your own experience when you tried to load a page, and the page sits there and hangs, if it keeps doing that, you're going to say, well, there must be something wrong with the website.

MR. KNAPP: Jay, could you just --

MR. MONAHAN: You'd have the same problem --

MR. KNAPP: -- sum up?

MR. MONAHAN: -- yeah. I'm sorry. Okay. So, as we go forward in this proceeding, I'd urge the Commission to strongly factor in the point of view of the Sand Hill Roads and the other investment corridors of the world that we need reliability and we need disclosure at a pretty deep level in order to make investment decisions.

MR. KNAPP: Thanks, Jay.

MR. MONAHAN: Thanks. Sorry.

MR. KNAPP: It's always tough moving the clock along, because I know there's a lot of interesting information there. We probably can go on for quite a while.

I'd next like to introduce Parul Desai, who's vice president of the Media Access Project.

Thank you.

MS. DESAI: Thank you. And thank you -- is this on? Thanks for the opportunity to speak today regarding transparency and disclosure.

For over 35 years now, Media Access Project has promoted the values of the First Amendment by working to ensure that all people have access to an open and diverse media, which now includes the Internet, and that protects the free flow information and promotes universal and equitable access to media outlets and telecommunications services.

The Supreme Court in Reno v. ACLU has observed that the content of the Internet is as diverse as human thought. The Internet is an open and interactive medium, facilitating communication by anyone to and from everyone.

It is a medium that supports and enhances the free expression of citizens and serves as a vehicle for democratic governance and economic activities. One element that is necessary in ensuring that users are able to fully engage in these types of economic, social, and political activities on the Internet is making available relevant information regarding the Internet access services that consumers purchase.

Transparency involves more than just passive consumer interest. Transparency in providing interactive service affects the public's rights as participants in the marketplace of ideas, so the interests here are also citizens' rights and not just consumer interests.

Disclosure of management and practices that a provider utilizes is necessary since it allows all users to know what to expect from the Internet experience. Full and meaningful disclosure of network management practices and an explanation of how those management tools actually work would benefit both consumers and innovators and both passive and active users.

Disclosure would allow Internet users to know whether the services and applications they desire will work or be allowed on the network and would also allow innovators to determine whether their applications or services will work on the network.

Disclosure is also necessary to determine whether a particular management practice falls within any rule adopted by the Commission. The only way to determine whether a management practice goes beyond the goals adopted by the Commission is through active, clear, and conspicuous disclosure.

Finally, transparency rules are also critical because voluntary commitments are not as effective and sufficient as codified rules. Without actual disclosure requirements, service providers can choose whether or not to disclose certain information or hide the information they do disclose by making it hard to find or disclosing it in fine print.

Moreover, strong and effective disclosure requirements minimize the need for other, more burdensome regulations and makes it less likely that the Commission will have to deal with contentious complaint processes, as occurred in the undisclosed Comcast Bittorrent situation.

So, Internet access providers should disclose all practices that they engage in to manage congestions on their networks and practices they use to monitor, manage, or interfere with a user's Internet traffic.

Such disclosure must be clear and easily accessible. The information should be prominently, plainly, and clearly displayed on a provider's website. I would also ask the Commission to urge ISPs to provide disclosure in various different languages so that all Internet users have the ability to access and understand the information provided to them.

Prominent, plain, and clear language is especially necessary since many consumers may not read the disclosure terms if it contains legal jargon, small text size, or hidden or hard-to-find placement.

The Commission should not expect consumers to take the time to read language and terms that they cannot understand or cannot find.

On the other hand, clear and conspicuous disclosure can empower Internet users to make informed decisions regarding the purchase and use of a service.

The Commission has sought comment with respect to what disclosure requirements would be useful to Internet users and disclosure requirements that would be useful to content, application, and service providers.

It is important to note, however, that, in some cases, Internet users may also one day become innovators, so this distinction is not always clear-cut.

Thus, the Commission should establish disclosure requirements that empower all users, regardless of whether they are active or passive users, to make informed decisions about their service.

I think the model that the CRTC has is a good model for the Commission to look to in adopting actual requirements. I won't go into what has already been discussed.

But overall, I'd like to say that the current NPRM I think is a very good start to disclosure requirements, but it will be necessary to clarify and broaden the current language of the NPRM.

Thank you for the opportunity to testify today, and I look forward to working with the Commission on this issue.

MR. KNAPP: Okay. Thank you, Parul. Our next speaker will be Ron Dicklin. He's the co-founder and chief technology officer of Root Wireless, and I think we've got the presentation up on the screen, and we're set to go to, I guess.

MR. DICKLIN: Good. Okay. Good afternoon, and thank you.

Root Wireless is a venture-backed company located out of Seattle, Washington. We specialize in helping consumers make educated, well-informed decisions on their wireless purchases. Though we are a new company, we have over 80 years combined experience supporting emerging technologies, with a heavy emphasis on the wireless area.

It's no secret that consumers can be overwhelmed when it comes to buying technology, like broadband Internet, for example. They need information to help support the buying decisions. You need to make sure that the product first fits their needs. They also to make certain that the product conforms as advertised in the marketing. They also need to make sure that their striving for the best value.

But buying technology can have some very inherent challenges for both consumers as well as the suppliers of the technology. First, it can be difficult in a fast and effective way to verify that the product meets a particular consumer's needs. There are some things that cannot be tested in a store or under normal purchasing environment.

Second, it's very common to have different expectations on how a product is supposed to perform. Usually, the more cutting edge or new the technology, the bigger gap in the expectations. Unfortunately, a lot of today's frustration over wireless broadband can be attributed to this type of mismatched expectations.

And third, the vast majority of consumers do not want to know the deep technical details to properly support their buying decisions. This requires an effort that the general consumer usually will not invest the time and understanding. They want unbiased information to tell them which product best supports their needs.

So with that said, anyone who spends time supporting broadband Internet will confirm that these exact same needs and challenges resonate loud and clear for the suppliers and consumers that make up this marketplace.

As it relates to wireless broadband specifically, Root Wireless is in the presence of building a business and helping consumers make informed, unbiased decision regarding their individual wireless data needs. Our backgrounds have enabled us to understand the challenges that face this ecosystem, and are investing our knowledge to develop tools that better align wireless data needs with consumers.

Our hope, if this is done correctly, it will create a win-win for everyone involved.

Given our experience over the last several years deeply and focused on this issue, we've learned a couple lessons while developing our tools.

First, the message needs to be simple. You need to understand the audience. There is no one consumer out there. There's some consumers that want -- that don't want to know the technical details, and there are some consumers that want to get very deeply involved in how this -- in how all this stuff goes together.

Second, you need to provide unbiased information. If the consumer thinks that there is somebody trying to get the upper hand, they're not going to pay any attention to the material you're posting.

And third, best used know rating methods. Consumers are used to understanding how certain products get rated. I think it was mentioned before here having some type of a star rating that consumers know, hey, four out of five stars is a pretty good service.

Secondly, and I think most importantly, is setting standards. This makes a scalable solution as wireless data will continue to be very dynamic over the next 5 to 10 years. And we need these standards.

Without proper definition on what is being offered, there will be a continuing mismatching of expectations. So, by setting standards, it helps better define the product. It helps -- enhances cross comparison of like technologies, and establishes a common and simplified message.

And lastly, moving forward, I think this all relates well to today's topic of discussions -- open dialogue around standards and better product definitions.

Other -- also being able to use ways of technology to help manage some of these expectations, and I think some of these other panelists will start talking about that.

And we truly believe that if this is done correctly, everybody will win. Consumers will have a better understanding of the products and services that they are buying as well as the suppliers of this technology will better align their customers with the products that they are delivering.

Once again, thank you for the opportunity to speak here today.

MR. KNAPP: Thank you, Ron. We'll next hear from Gerald Faulhaber, who's currently professor emeritus of business and public policy at the Wharton School and a professor at the Penn Law School.

MR. FAULHABER: I think currently professor emeritus means I'm still alive. Yes, that's pretty good.

Formerly, chief economist here, and I

will say this visit has been kind of exciting to see a newly energized FCC since the new administration came in. So.

Let me talk a little bit -- where am I here -- oh, okay. We got a little ahead of ourselves, didn't we? Okay.

This must be a Macintosh. Okay. What's going on here? There's some technical difficulties here. Ah, good. Okay. Excellent. Okay. Where does this fit in? As it was mentioned before, there's a number of proceedings in the National Broadband Plan -- Open Internet, and where does transparency fit?

Now I want to back up for a minute, I think I'm going to sort of reprise the Chairman's views about this. For me, the focus has to be -- and this is where transparency fits in -- has got to be on what I call consumer-centric markets.

And there are three features of consumer-centric markets, which are all critical and all important.

The first, of course, is competition.

Lots of guys slugging it out for the customers custom. The second is transparency, which is to say it's no good having a lot of players if you don't know what they're offering. So we have to have disclosure and transparency to make competition work.

And lastly -- too bad the FTC guy left -- which is what I call judicious anti-trust, anti-trust which enforces the competitiveness of the market, protects competition, but not competitors.

Now what's the role of regulation in all this? It's basically to enable all of the above. First of all, and this is a whole talk we could give, which is to say encouraging for competition, encouraging wireless broadband through the simple mechanism of getting lots more spectrum out there.

The technology is poised to take advantage of that, and I believe can bring competition to the broadband market in a finite time frame.

Last -- the second step to this is to

ensure that all Internet firms -- and I want to emphasize all Internet firms -- okay -- are transparent for customers.

Now the result is that once you do this, once you achieve this nirvana, okay, customers are in charge. They are now driving decisions. And the people who aren't driving decisions are firms. They have to respond to customers.

And it certainly shouldn't be regulators, and it shouldn't be pundits, by which I mean everybody at this table; okay?

This is what we mean by customer-centric; okay? Now successful competition requires transparency, and being an economist, of course, you know, competition drives a lot; okay? But you need transparency to do that, and I'm very strong on this; okay?

And it applies to all markets -- broadband ISPs -- in all features of their business, not just network management. Network management seems to be really important to the Commission these days, okay, which is a matter of history. But I think we have to look at all.

It also applies in my view or should apply to application and content providers. Okay. They shouldn't get a free ride here, okay. And backbone? We don't talk much about backbone networks, but they're in the game as well, and they influence what happens to customers. Okay?

Once you get this, put your faith in customers to make the right decisions. You don't have to tell them what to do. They're smart. They can figure it out. But you've got to give them the competition, and you've got to give them transparency.

Don't even listen to professors. Okay. We don't know what we're talking about.

And I am very much a hawk on this. I really think customers need to be in the driver's seat, and we shouldn't be telling them what to do.

There is, however, in transparency a market failure, and it's a well-known one. It's why we have the Federal Trade Commission. It's why we have the Consumer Group here, and that is information asymmetry. It's been a well-known thing in economics. This is a problem. You have to make sure that customers have good information.

And even in competitive markets, there may not be incentives for firms to reveal this. Okay? May or may not. I'm not saying -- but definitely this is an area where regulation should at least be involved; okay?

And I'm a hawk on this issue. I think it has to be done right, okay, but, you know, the words I'm hearing around here are good.

There's another implication, which is to say if there isn't a market failure, regulation is going to be worse than useless. And I'm a hawk on this as well.

So be careful. There's got to be a market failure you're looking at. If it's not a market failure, regulation is all downside and no upside.

Now what do we mean by transparency? We've said it a few times. This is a familiar example to American consumers. Credible, decision-relevant information has to be made available. It has to be immediately and easily available at the time of purchase; okay?

It has to be easy to understand. This has also been mentioned. Not ULAs, not buried three-deep on a website, does not require a J.D. to understand it.

Now I actually said this once to a J.D., who -- and she had a job for some period of time writing ULAs. And she said, "The job was to make sure even J.D.'s didn't understand them."

So, there you are. A little inside the sausage factory.

This is the example. We know what it is. We see it everyday in the supermarket -- well, every week -- in the supermarket, which is the nutrition data label. Okay? It's on every single piece of packaged goods sold in the United States. It's a standard format; has the basic facts. You have to work hard to not see it.

Now some of us don't. But you pick it up, and it's on the outside. Okay? Easy to see.

It doesn't have everything, but it's got most of the stuff you need to know. An interesting case is the FDA- FTC prescription drug information, which gets at a point that was raised earlier, and that is to say the important stuff is on the label.

Prescription drugs are at least as complicated as broadband, let me tell you. Okay. So you look at this thing, and it has do not operate heavy machinery. Okay. It's got the important stuff right on the label.

If you want to drill down, there's the package insert. Nice architecture for revealing information. I like it. Okay. We should be thinking about that.

I'm going to name some principles for assessing disclosure. Does it provide all the information customers need to make decisions? Does the customer have easy access? Is it clear and simple? And is it verifiable? Is it credible? Okay.

What info to disclose. Use a reasonable

customer standard. I mean we have this standard, for example, in consumer law, the reasonable consumer. We have an investor law, you know, reasonable investor for securities law. It's a fairly standard legal way to go about it. But here's the thing I read and really focus on it: Talk to consumers. Talk to customers. See what they want. See what's important to them.

We're all wonks up here. Don't talk to us. Yeah, we all have opinions, but, you know, what matters is what customers want. So you should be doing when corporations whose livelihood depends on knowing what customers want: Use focus groups. And talk to real people, people that file stuff with the FCC. That's, you know -- you don't want us.

MR. KNAPP: Gerry --

MR. FAULHABER: Okay?

MR. KNAPP: -- could you wrap up maybe in 30 seconds or a minute?

MR. FAULHABER: Ooh, I was hoping to talk less than Chairman Leibowitz. I'll do my best. Okay?

MR. KNAPP: Sorry. Thanks.

MR. FAULHABER: Incidentally, a hint on this: Network management may not be the most important thing on a customer's agenda. That's a working hypothesis. But I wouldn't do it.

Okay. Let me quickly go through this thing. I'm not going to do this one. This is the performance measure, which I think is more important than network management. You can look at it later.

Easy access. Application content providers. Adopt a standard icon. Require it be on everybody's website. You roll over. It's got the easy stuff; click it, it's got the more difficult stuff.

Do the same thing with broadband ISPs. I'd like to see this in the systray in Windows. Okay? You can opt out of it. You may not want it, but it will give you what you need to know; okay? Right there.

Okay. It can also be used to alert

customers to a botnet and attacks and stuff.

Here's the third principle: Clear and simple. This on the -- your left-hand side is credit card information label that's required. Can't read it? Most of us can't.

The right-hand side was an interesting thing they published in the New York Times. Some design people came up and said, "What would you really want to do?" Okay? That's it.

The other thing you should look at his financial reporting by publicly traded firms for the SEC, which should follow a reasonable investors' standard.

Fourth principle: Verifiability. A number of ways to get this. The FCC does not have to be the ones to do it. You could have an ISO 9000 standard. You could have outside auditors. There's a number of ways to do this.

Offcom in Britain has come up with some interesting ideas about how to do this, which -- in which they are very involved, incidentally. Okay?

Lastly -- this is the conclusion.

MR. KNAPP: You're doing --

MR. FAULHABER: -- necessary condition for competition to lead the customer-centric outcomes. With competition and transparency together, customers get to make the decisions. That's where you want to be. Transparency is an important part of that. Thank you very much.

MR. KNAPP: Thank you. Our next speaker is Nicholas Weaver. He's a researcher at the International Computer Science Institute in Berkeley at the University of California at Berkeley, and he's one of the leading developers of Netalyzr.

MR. WEAVER: Thank you very much. This is joint work with Vern Paxon and Christa (inaudible), also of ICSI.

ICSI is -- it is on. Is this better? The switch is on, and there is a red light. Should I turn up the audio levels?

SPEAKER: Yeah.

MR. WEAVER: Testing, testing, testing.

This better?

SPEAKER: Okay.

MR. WEAVER: Okay. Good. I'm actually doing something very gutsy. I'm a technologist who's doing a live demo.

To begin with, this is National Science Foundation- funded research, but all opinions are my own, not those of the NSF.

Our belief is we're a strong network transparency hawks, but we believe in doing one better. We actually want to discover what's up on the network. So, we've built a tool, which is designed twofold: To create a wide-scale survey so that we can actually build nutrition labels for ISPs without even having to rely on the ISPs doing disclosure; and also as a tool for individuals so that they can find out if there's anything interesting on their network. It could be misconfigurations. It could be deliberate mismanagement. It could be something local to their system.

And we try to be pretty much as

comprehensive as possible within the limits of making it trivial to use. Making it trivial to use means that it just runs in the web browser. So the user goes to our site, Netalyzr. We've had some 80,000 visits so far. Click the button. It goes to a back-end server, and they either allow or deny our signature.

And while that's busy running away, checking a whole bunch of different properties, I'll show you what the results are.

So this is a fairly typical network. What we do is we color code things and we minimize stuff when it's green, and we have convenient explanations.

But let's go through. This is from the Starbucks down the street. Starbucks -- well, you're behind a NAT. Very common for hotspots. Most individuals are behind NATs. But this NAT remember TCP ports, not listed on any bad blacklists.

What protocols are blocked in the network? Well, DNS is blocked, because they use

DNS to manage things so that you actually go through the hotspot login. This is a hotspot. You expect things like that. Likewise, on UDP, the DNS goes through a proxy, so we know you can't actually access anything (inaudible) DNS-wise, but at the same time when we do direct access, the results are sensible.

So it may proxy, but it's a correct proxy. There's no bugs on the MTU. There's subtle, little bugs that have driven us crazy with networks dropping packets. Latency is good. No outages. Bandwidth is good. This is an interesting one. Network buffering.

It turns out there's a design flaw in a lot of network hardware. It stores too much information, and, as a result, what happens if things don't work right under LUG. You get things feeling slower than they should. This checks for it, and we found out that this is quite an endemic condition.

As such, this is a case where you'd want really advanced traffic management techniques, because with advanced traffic management in the core of the network, you could fix this bug without having to fix all the buggy devices.

Is there an in-path HTTP proxy or cache? No, there is not on this network. We checked over a large number of ISPs. Caching is not common in the U.S. We did not see significant U.S. ISPs running in network caches.

We did see South African ISPs run in network caches. And this makes sense because if they work correctly, this improves latency, reduces bandwidth costs, especially when you're out in the middle of South Africa trying to get to U.S. websites.

A big set of tests is DNS. So does the DNS system work right? Will DNS-SEC work right? It works quite well.

We do a look up of a whole bunch of names on the client side, send them back, and do reverse on the server, because there's malcode that changes people's DNS settings.

More disturbing we saw a least one ISP

that wide open west, which does not return a valid value for www.google.com. Instead, it returns the address of their own proxy. Examining this proxy server shows that it will proxy search dot Yahoo, but it will not proxy mail.google or www.yahoo, and deliberate invalid requests give a reference to fishingwarningsite.com, which is a nonsense domain as far as we can tell.

So and ISP is deliberately manipulating DNS results, and we have no idea why. This is an example why transparency is important. This could be quite legitimate to protect users. It could be something completely illegitimate, and unless they disclose what they're doing, we have no way of telling.

And this is discussed in more detail in my written remarks, which should be on the website.

And finally, browser system clock, because those have issues, too. But to give you an idea of what more interesting networks are, this is what the hotspot looks like here. As you can see, there is a lot of red.

MR. KNAPP: We have no more time right now.

MR. WEAVER: The interesting thing is --

MR. KNAPP: How did we let that slide get through?

MR. WEAVER: -- (inaudible). This proxy is also -- has a security vulnerability in this, so we report that to the user. Oops. And it caches data incorrectly. Caches are all fine and good as long they work right.

And so the lesson of this -- that's just spurious Ameritrade -- it has bad reverse data a lot of times -- the lesson of this is quite clear: It is actually in the ISP's interest to do complete, clear, concise disclosure because it's going to be revealed by people like me anyway, but if you do disclosure in advance, it prevents these unfortunate surprises.

MR. KNAPP: Thanks, Nick. We'll now hear from David Young. He's Verizon's Vice President of Regulatory Affairs, and handles items before the FCC dealing with broadband and emerging issues. Thank you.

MR. YOUNG: Thanks, Julie.

MR. KNAPP: David.

MR. YOUNG: And it's a real pleasure to be here today. These are important topics, and they're also complex topics. And I think the previous presentation demonstrates just how complex these issues can be.

The Internet is a complex ecosystem, and user experience on the Internet is driven by the interactions of all of the elements of that ecosystem, whether it's the software on the device or the operating system; the hardware itself or consumer electronic devices that are integrated; things that are going on in the home network or the corporate network; things that go on in the access network; and then things that go on on the server side -- the applications and content and things that are running on the Internet.

And all of these things, together, contribute to the experience that the end-user has. And so robust disclosure of the things that affect that end-user experience are very important, and so all of the participants in that ecosystem have a duty to provide robust disclosure about the things that affect that experience.

Verizon does this for our own services, and we do it in a way that we think provides consumers with meaningful information that allows them to make informed choices when they're purchasing a product. At the point of purchase, we disclose all of the relevant information. And then after purchase, we provide a way for them to continue to get information that's relevant to their service and the things that affect their service.

I think, though, an important point that's come up a lot today is the fact that consumers should have and do have tools at their disposal to verify what they're getting, and things like Netalyzr are very detailed and provide a wide variety of relevant information.

But there are simpler tools out there as well that provide pretty robust information about upload and download speed and latency, for example. Those tools also frequently show a consumer not only what the result is for their test, but how does that compare with other people who live in the same area or who use the same ISP.

And so this is extremely valuable for them in figuring out, okay, how is my experience, but also in making market decisions as to what products and services they want to use.

There are security tests that look for vulnerabilities not only within the network and in the path, but also on the individual's own system themselves, where frequently those vulnerabilities exist.

And there's an ability to test for interference along the path.

I think that third party evaluation -- and I talked about it in these tools that are available -- but somebody who can look at and do side-by-side comparisons of these various things is useful to consumers.

I think Consumer Reports provides a good model for that in sort of the real world. They do compare Internet Service Providers, but, to my knowledge, they haven't done any real technical comparative testing of actual performance and what not. And that would be interesting to see or something like Root Wireless. That seems like an interesting concept as well for providing that sort of comparative information for consumers.

Because this stuff is so complicated and it's not easy to convey to average consumers, and it's not clear what the best way of doing that is, I think a process that allows all interested parties, whether it's service providers, content providers, consumer groups, academics, to come together and really look for what are the best practices. And this is the way techniques are developed in the IETF, for example, for the types of network management tools that are created for ISPs.

They're the types of things that are discussed in groups like, NANOG, the North

American Network Operators Group, that actually apply these tools to practices. And I think that a similar approach for looking at what are the types of things that are best disclosed to consumers. How best to disclose them, and develop a set of best practices in this area with respect not only to network management practices, but to things like speed and performance and those sorts of things is absolutely the right approach.

A group like that could also server as a sounding board for new ideas so that bad ideas don't get implemented in the first place, and then for identifying bad actors out there that are not sort of complying with what's seen as best practices, because it's in the interest of the good actors, the people who are trying to do the right thing, to weed out the problems, to fix them themselves, and then if the industry and the -- shining sunlight on it is not able to fix those problems, then there's a role for government to step in as a backstop and deal with those things in an anti-trust or FTC-type approach.

I'll let it go there, and look forward to questions. Thanks, Julie.

MR. KNAPP: Thank you, David. Next hear from Fernando LaGuarda. He helps Time Warner Cable develop and advance its policy positions that focus on consumer protection, competition issues, intellectual property, and telecommunications regulation. Fernando?

MR. LaGUARDA: Thanks. Thanks, Julie. Thanks, Joel. Thank you everybody for allowing me the opportunity to be here. This is an important topic, and thank you also to my fellow workshop members. I'm furiously taking notes here. I'm listening to your good ideas about this important topic.

Time Warner Cable's corporate mission is to connect people with information, entertainment, and each other, and to give our customers control in ways that are simple and easy. I think that lays out very clearly why we're interested in this topic.

Our business success is entirely

dependent on customer satisfaction with our services, and providing clear and comprehensive usable information to our customers is essential to our success.

That means giving customers the right information they need to control each their broadband Internet experience.

Powerful incentives already exists to ensure that we do that -- communicate clearly and effectively with our customers. We operate in a competitive marketplace. Our customers have ever-increasing expectations as to the quality, reliability of our services and their ability to control their experience.

In addition to the oversight of the Federal Trade Commission and various states, we also have content application providers, who also had the incentive and ability to monitor and ensure that what we are doing serves their customers, who are our customers as well.

As a result, we pay very close attention to the clarity and efficacy of our customer communication efforts. We give our customers clear information at every stage of their relationship with us, from the selection of plans to customer care and billing, to service termination. And we use a variety of channels to do that.

They include web content, advertisements, in-person assistance at local offices, toll-free telephone lines, online chat, and e-mail. And we continually refine our disclosure practices to make sure that they are simple and easy.

We're particularly attentive to the practices in connection with the network management issues at stake here. There is near universal agreement that reasonable network management is essential to protect the online experience, and we've always maintained a practice of transparency with respect to our management practices.

It's increasingly important in the competitive marketplace to do that. It's one of the many bases that customers have to make comparisons between their service providers in making decisions.

As a result, we don't believe there is a need right now to codify transparency principles to ensure customers can make informed purchasing and use it -- decisions. But if the Commission finds that some form of consistency in disclosure would help consumers, then there is an opportunity, as has been referred to here, to work collaboratively on best practices to make that happen.

The Commission has sought comment on the creation of a broadband clearinghouse to serve as a repository of broadband-related best practices, and recognize the value of that type of resource. It would be better serve consumers than the imposition of a mandatory, one-size-fits-all solution.

If the Commission determines to mandate disclosure, we think two important points should be made. First, disclosure rules should allow flexibility to adjust to marketplace and technological conditions. This is an evolving marketplace. Everyone refers to the importance of innovation.

Rigid disclosure mandates tend to produce boilerplate, as seen in the tariff model that prevailed for many years with regard to common carrier services.

Second, users should have as much information as possible regarding the conditions that apply to their use of content and application and online services as they do to their broadband service providers.

Search engines and others have the same obligation, and Professor Faulhaber made those points eloquently. If sunlight is the best disinfectant, as the Commission notes in the NPRM, then that sunlight should shine broadly on all parts of the Internet ecosystem and not be eclipsed in this important area.

Finally, while transparency is important, the Commission should focus on the provision of information to consumers. I think that there has certainly been widespread agreement on this panel so far that the purpose of transparency is to empower consumers, and that should remain the goal.

Imposing requirements upstream disclosure would be superfluous. I think the problems that have been referred to so far could certainly be discovered in disclosures to consumers as well as they could be by imposing upstream mandates.

And we believe that there are unfortunate consequences to security and the online experience if we make too much detail available to upstream providers beyond what consumers can understand.

I think that summarizes my written comments, and I appreciate the opportunity to make them here to you.

MR. KNAPP: Thank you, Fernando. Our final speaker is Joel Kelsey. There's Joel.

He's a federal and international affairs policy analyst for Consumers Union, which I'm sure most of you know is the nonprofit publisher of Consumer Reports. Joel?

MR. KELSEY: Hi. Thank you, Julie, and thank you to the Commission for the opportunity to testify today.

I'd also like to just quickly note that the workshop couldn't be more timely because if you're a subscriber or just a curious magazine rack browser, you'll notice that this month's issue of Consumer Reports is on bundled ISP providers, both cable, voice, and Internet service providers.

And there's also an editorial on network neutrality that you'll find on page four or five.

So, as a magazine that seeks to give consumers access to truthful information necessary to compare goods and services, we very much value and support strong and enforceable government disclosure standards.

And we believe that the Commission is on the right track. We were glad to see the

Commission take a comprehensive approach to transparency and the consumer information and disclosure NOI this past summer. And we're also very glad to see the addition of the Sixth Principle in the Open Internet NPRM.

We have long comments that we filed into the record on the Notice of Inquiry this summer that I would refer folks to for things like truth in broadband advertising, truth in billing, actual versus advertised speeds, surcharges, fees, et cetera.

For the next couple minutes, I'm going to focus just on network management practices, since that's what's mostly in the NPRM.

And I'd like to cover three areas. First, I'll give the consumer perspective on what the current landscape of transparency looks like and the need for disclosure. Second, I'll touch on the perceived tension between clarity of information and detailed disclosure. And then last, I'd like to give five examples of what we believe meaningful network management might look like, and offer two reasons why we believe the Commission should expand, clarify, and broaden the language that's in the NPRM right now.

So we very much agree with the Commission that transparency discourages harmful market behavior, and that network management practices should be disclosed to consumers. This is especially true when limited competition constrains consumer choice and dilutes the incentives among providers to provide an accurate representation of the services being offered.

Currently, broadband subscribers are confronted with confusing and often sometimes misleading information regarding service quality when making choices between providers. Internet access companies generally provide their subscribers with broad terms of service that give the flexibility to change their terms, their network management practices, and their level of service without any prior notice to their subscribers or to the public.

And again, I'll point to my written

testimony here and the many pages of appendices in the Notice of Inquiry this past summer that detail all of those different terms of service for many of the providers in the United States.

As a matter of just good consumer disclosure policy, we believe that the Commission should stop ISPs from describing binding terms and conditions within a multipage legal document in 8 point font.

In the context of an Open Internet, we believe that the disclosure of network management practices is particularly important for three reasons: First, it gives consumers -- and we've all talked about this on the panel today -- an accurate and ongoing representation of the Internet experience that they sign up for. Second, it ensures that network management practices are narrowly tailored to address a legitimate purpose and don't unduly interfere with the consumer's access to a best efforts network.

And lastly, it allows innovators and applications developers to make informed decisions about where to allocate their resources and how to design their applications. And I also will mention that we -- we're hawkish -- we're very hawk like on making sure that content providers disclose relevant privacy information to consumers as well. And there's a great debate over at the FTC about that.

Clarity versus detail. Being at Consumer Reports, we very much understand and value the need for concise, clearly formatted, and easily understood information design. In fact, we very much pride ourselves on writing clear and concise reports that help consumers make marketplace decisions.

And consumers must absolutely have access both at the point of sale and in an ongoing basis -- access to high- level information that plainly articulates factors like actual service performance rather than up-to speeds, and then also they must have general criteria regarding when a provider may block, prioritize, or monitor service use.

This information should be disclosed in a consistent manner across ISPs so that consumers can easily compare when they do, in fact, have a marketplace decision to make, and usually, as the Chairman of the FTC mentioned, that's only between two folks, mostly their cable provider or their telco provider.

However, the clarity of information provided is a question of information design, and it shouldn't be confused with the need for the Commission to issue requirements making thorough and detailed information about such practices available to the public.

Consumers are empowered when they have more, not less information on which to base decisions, and there's many organizations out there that will interpret that detailed information and try to give consumers access to unbiased opinions about what that information may mean for their end use -- or for their service and for the end use experience that they have.

And in practical terms, just because

consumers don't understand technical network terms today, doesn't mean they won't develop sophistication over time. It stands to reason that the vast majority of consumers weren't necessarily conversant with miles per gallon, caloric values, octane levels in gasoline, or APRs in credit scores before the federal government provided an industry-standard and mandated consistent public disclosure of such information.

So that leads us to what does meaningful disclosure actually look like. And we have a couple recommendations for the Commission.

We believe that the Commission should both require an easily accessed means for subscribers to learn about network management practices and how that might impact their end-user experience; and then also disclose a more detailed -- excuse me -- require more detailed analysis on the purpose, the message, and the congestion levels that are necessary to trigger a particular network management practice.

Here are five examples: First, any

limits imposed on a subscriber's upstream or downstream traffic -- this includes blocking, delaying, de-prioritizing, or prioritizing or intersecting traffic -- or excuse me -- inserting traffic into a stream.

Second, technical details of the method that was used to impose a limit on the subscriber's traffic.

Third, thresholds that trigger certain network management practices and an estimate of the percentage of users that were affected and the duration of the practice. If the practice is only implied -- used during congestion periods, let folks know what the trigger for that was and when the network management practice may stop or cease.

Fourth, any technology that inspects the content of the Internet traffic other than processing just basic addressing information -- this would be DPI or packet inspection.

Fifth, the differences in how the network is being allocated to different users -- to different uses, including "managed services."

This includes the amount of capacity dedicated to general Internet traffic and, if shared capacity, how that might be shared.

Lastly, I'd like to urge the Commission to clarify the language in the NPRM for two reasons: First, subjecting disclosure to reasonable network management we believe provides some ambiguity. For example, one interpretation of that exemption could mean that reasonable practices need not be disclosed and that ISPs would then be the judge of reasonableness. And they would only then have to be required to disclose information about network practices that they deem unreasonable or were, therefore, illegal.

That's clearly not the intent of the Commission, and we believe that -- we could go a long way in the final order to get rid of that ambiguity that might lead to that circular logic.

Second, disclosure obligation should not rest on the compliance of other consumer protections in the NPRM in particular. It's just not clear based on the language and the NPRM what falls under the umbrella of the Sixth Principle. That is, what actually requires disclosure.

We believe that the Commission shouldn't attach the requirements of disclosure to other consumer protections, but should craft disclosure principles as a standalone requirement and look at it in a different track related to all of the many marketplace reasons that you identified in the NOI this summer and make sure that it applies both to wireless and wireline providers.

With that, I'll wrap up and thank you.

MR. KNAPP: Thanks, Joel. Fantastic set of presentations. What we'll do we'll take about a 6-minute 30-second break, and we'll resume at 3:40 p.m. We should have the cards in the back. Have we got somebody designated to pick them up? Who's going to pick -- okay. Right over here. So make sure you get your questions in. E-mail questions to newmedia@fcc.gov, and I gave out the Twitter information and look forward to the closing hour. Thanks.

(Recess)

MR. GURIN: Okay. Okay. Well, thank you all for your presentations. This is really a very terrific and very thought-provoking session I think both for anyone who's new to the subject and for those who have been trying to puzzle through all of this for a while.

Julie and I were just talking at the break. I think we each have a few questions that we'd like to ask that I think are very general questions that would be very helpful to get some perspective from all of you, and then I know we have questions, both from the audience, from people following this online. And we will certainly leave plenty of time for anybody here to ask your questions as well.

But I wanted to just start with I think this whole question of reasonable and unreasonableness, and I think both to offer a little bit of a clarification, but then also to ask some questions that I'd love to hear your take on.

So just to clarify what the Sixth Principle actually says. What this principle says is that subject to reasonable network management, providers of broadband access have to disclose information concerning network management and other practices to enable users, content application and service providers to know what they need to know.

So this is I think potentially a little bit hard to follow, because we're using network management and a couple of different contexts here. And the intent here is that providers should disclose any information about the network management practices that would be of importance to consumers or other users, but that that disclosure in itself is subject to "reasonable network management practices," meaning, for example, that it would never be expected that a provider would disclose information that would somehow compromise the security of the network.

So, hopefully, that's a little bit clarifying. But the question that I think this raises -- because several of you talked about this kind of reasonable/unreasonable standard -- I would be very curious as to what your thoughts are on first, what are the basic kinds of things about network management that you think consumers and other users need to know to make intelligent decisions about how they use the network or what network they choose; and second, what would be, in your view, some important exceptions to that rule; that would be important to, you know, to not disclose for valid reasons.

And, Mr. Chairman, I thought actually, if you don't mind, it might be very interesting to -- for I -- it would be very interesting for us to hear if these are issues you addressed and how you address them in your process.

MR. von FINCKENSTEIN: We did. We, first of all, obligate ISPs to clearly and prominently do things on their website. And they have to be -- and they have to make references to disclosures and relevant marking in (inaudible) customer contracts or terms of service. So, if you as the customer look at it and you want -- there's a reference there, so you know where to go to them.

Secondly, the way we broke down -- I have here the actual decision. It says who's affected by the ITMP, when the management will occur -- so you know exactly the time period -- and what type of Internet traffic, i.e., application class or application protocol will be subject to it; how the ITMP will affect the user's Internet experience including the specific integral speeds. And that's really the (inaudible). What does it mean to a consumer?

You have to put it in normal English so the user can understands this means I will (inaudible) slowdown, interrupted, I can't view, et cetera.

And then secondly, we also made sure that in advance, 30 days, and not only new ones, but also amendments to existing ones. Say, are significant ones. And, of course, we take -- and we specifically exempt housekeeping ones or, as you suggest, what (inaudible) what dealing with (inaudible) reasonableness anything that would impair the security of the network -- would give information that could be used to harm the person.

MR. GURIN: Okay. And you -- and you're finding that that distinction is one that the networks are reasonably comfortable with, but that also is getting people the information they need in your judgment.

MR. von FINCKENSTEIN: It's the early days. I mean we have only enough --

MR. GURIN: Oh, sure. Sure.

MR. von FINCKENSTEIN: -- in place since November --

MR. GURIN: Yeah.

MR. von FINCKENSTEIN: -- but basically, everybody says, yeah, I think you've hit the main point. We will now see as we get complaints and get delve into them and through -- we'll obviously have to elaborate, put some flesh on these concepts, et cetera. But we -- by and large, people think we have covered the main areas.

MR. GURIN: Terrific. Thank you. Anybody else have a comment or a question on that point? Yes. Nicholas here.

MR. WEAVER: I think there's sort of also the need for two levels of disclosure. There's the high-level view for generic users and the more detailed view.

So, like, let's take an advantage of a policy actually I like as a customer is Comcast's fairness mechanism. A high-level description could be this enforces an approximate fairness between users on timescales of 15 minutes. Nice, sensible high-level description of what the goal is, and then the detailed is it operates by this thresholding of congestion and COS metrics and stuff like that.

So that I as a detailed expert can go and go, okay, the high-level description actually matches what's on the wire, and that this is how I test for it or this is why this is so subtle in practice, I'm not going to bother testing for it.

So there's I think two levels of

disclosure that's needed on a lot of things -- the normal disclosure and the geek disclosure.

MR. GURIN: Is that a term of art or? Great. Anybody -- thank you very helpful, yeah. Sascha?

MR. MEINRATH: I just wanted to add to this security through obscurity is a fascinating and easy way to undermine the basic fundamental security of your network.

And so I'd flip it on its head: I'd say, you know, there is almost nothing that shouldn't be disclosed because that's what maximizes the network security.

Now in terms of what should you exclude, I don' know, like, passwords to servers and things like that. But other than that, it's like if you're doing something, chances are people can identify it and have to tools to do that.

And if you're blocking ports or forwarding ports or doing, you know, DNS proxying -- all of that stuff should be disclosed, and there's almost zero times where it would be a good idea for security reasons not to disclose that.

MR. WEAVER: I strongly disagree when it comes to parameters rather than techniques. I believe that every ISP should do some outbound heuristics on Port 25 and dynamic blocking, and we see a lot of this. This is a good thing.

But I don't believe ISPs should disclose the thresholding involved and the detailed algorithms involved, because those can be used to game the system.

Likewise, the high-level description of the Comcast fairness is they're fuzzy on the parameters. The parameters are within this range, which means I as an expert can validate that this is what the general concept will do, but I as an evil person would have a real trouble gaming the system.

So there are -- I think there should be a distinction between techniques and detailed parameters in terms of disclosure.

MR. MEINRATH: Wouldn't you want to do that on -- I don't want to geek out too much in (inaudible) -- but you'd want to do that at the servers that are, like, the SMTP servers, not on the network itself?

MR. WEAVER: No. On the network itself, because you want to allow people to contact arbitrary SMTP servers so they don't have to use the ISP's SMTP server when their notebook moves from work to home, to this, but at the same time, you don't want them able to run SPAMbots.

So a lot of ISPs are doing heuristics that if you only contact one SMTP server and don't send too much, you allow it. But if you start contacting too much, it gets dynamically blocked.

MR. GURIN: Joel, you have a comment?

MR. KELSEY: Yeah. One just quick. Sorry. One just quick comment is that, you know, so a lot of this turns around what the Commission -- whatever the Commission decides will look like its reasonable network management test. And so, you know -- and then decide to kind of go forward on what's reasonable and what's unreasonable based on complaints that are brought before the

Commission.

So, from our perspective, we'd very much like to err on the side of disclosure, because then consumers will now whether or not they see a management practice that they'd like to bring a complaint before the Commissioner, not -- and then the Commission can employ its test.

And we've been huge fans of not listing black hat and white hat practices, but instead general criteria that kind of give consumers, Internet service providers, and innovators at the edge an idea of what would be deemed reasonable or not in the final order.

MR. GURIN: What would be reasonable for providers to do, do you mean, or not?

MR. KELSEY: Right. What would be a reasonable network management practice versus an unreasonable network management practice? And based on what that criteria looks like, one practice in a certain circumstance might be reasonable, but then in a different circumstance, it might be unreasonable, for example, on wireless or wired networks. In different -- you know, I'm sure you can come up with a million and one different examples of how a particular network management tool might in one case be okay, but in another case not.

So very much erring on the side of disclosure, I also very much like the idea of having a, you know, kind of simple format for most consumers to be able to look at and figure how it's going to impact their actual Internet experience. And then a much more detailed disclosure may be on a website or kept in a folder online on the FCC's website for consumers to access that folks like Consumer Reports, New America, and others can, you know, can go and look at.

MR. GURIN: Sure. Jay?

MR. MONAHAN: Just -- I think this is partly the same point, and that is that if the standard is narrowly tailored, which we believe is an appropriate -- striking the appropriate balance, there needs to be enough information to make for, if not a consumer, then one of the helpful organizations to make that evaluation.

By the way, for the record, we believe that some -- the term "strict scrutiny" has come up, and I believe the strict scrutiny standard is actually narrower, because it to me implies that there is one least restrictive means of doing it and anything else would be unreasonable.

We appreciate that there are different ways of doing network management, but given that the means chosen should be narrowly tailored to minimize the negative impact on traffic or content or whatever the case may be.

Also, I just wanted to note that while it's great that there are businesses, even startups, that are in the business of analyzing information and disclosing it for consumers, much of the information that we probably would want to see is not easily available. It may be available only anecdotally -- excuse me -- anecdotally initially.

We've tried some data collecting

practices ourselves, and found out it's quite difficult and it's hard to get very broad data.

So I think those outside groups are critically important, especially in the comparative analysis between competitors, but the starting point of the information has to be coming from the providers themselves, who have the best data and the broadest data.

MR. GURIN: Great. Thank you. I know we do want to get on to questions that we've been getting, but thank you. I think that was a good one to open this up, and let's continue with more discussion.

MR. KNAPP: Just --

MR. GURIN: Yeah.

MR. KNAPP: -- just to follow up on that same line. What are the areas with the development of the new measurement tools that you can't get at?

MR. WEAVER: Detailed fairness in conjunction with policies that only occur under the (inaudible). Those you can really measure with the long-term baseline monitoring. So, where you have measurement agents distributed on a lot of computers.

On the other hand, those policies, even when screwed up, are less disruptive for its users.

So things -- the more subtle the policy is, the harder it is to measure, but the less likely it's going to have a significant effect on a large number of users.

MR. KNAPP: Mm-hmm. Anybody else. Gerald?

MR. FAULHABER: Yeah. Let me -- let me raise a point, which we haven't talked about. I think David Young mentioned it, which is to say what are going to hold broadband ISPs responsible for, because certainly the latency, the issues that I run into depend as much on the fact that I'm still using an untweaked version of Windows XP, which has real limitations in its IP stack.

So I'm never going to buy a Comcast 50 megabits because I can't run at 50 megabits because of my system. I don't know that. Somebody had to tell me that.

So that means we -- the people we need to hold responsible -- we need to hold let's say Comcast responsible for stuff that goes on only in their network. But that is not the whole customer experience and somehow we got to close that gap.

MR. GURIN: Yeah. I think that's a great point and one we may want to talk about a bit more. We have some questions.

So the first question from someone here today -- I'll just read this verbatim. NSF-Net stats were cited. Those used application-specific measures, which involve essentially DPI. Do panelists believe, then, that DPI should be used in broadband measurement? And, Sascha, I think you were the one who talked about NSF-Net initially. So, start us on that.

MR. MEINRATH: Yeah. That's incorrect. I mean there was passive measurement. There's still passive measurement. A lot of this was actually used by -- there was time when ports and applications mapped much closer on to one another. And so that was used as heuristic for collecting these.

I'm not, you know, utterly familiar with the data collection methodologies from 15 years ago. I know that -- and this is the important part -- there were processes in place by which these methodologies were constantly evolving. And that certainly should be a part of what is implemented by the FCC.

You don't want to say, like, this is the end all and be all for all time. You want to have a process by which, you know, you can add new metrics, as, you know, octane becomes important; that that becomes something that is disclosed at the pump.

And we want that equivalent for what's happening, you know, on broadband connections as well.

MR. GURIN: Good. Okay. Yes.

MR. MEINRATH: Oh, wait. Also --

MR. GURIN: I'm sorry.

MR. MEINRATH: -- DPI is almost never a good idea, just because it creates data obfuscation on the behalf of the users.

So, no, I don't think that that would be a good solution to this problem.

MR. GURIN: Fair enough. Thanks.

MR. WEAVER: I think, speaking as a network security person, I think DPI has been much and needlessly maligned. It is actually very useful for a lot of things, including network measurement. But it really, for network measurement and the like, it tends to be only appropriate on non-ISP networks.

But on the other hand, you can use measurements conducted at universities, companies, et cetera, to probe the other end. So, this is what we did with the reset injection stuff as we monitored our own traffic to probe the other side.

There are real cases where the ISPs should be using DPI, like, to be honest, would anybody object to the network recognizing this is known SPAM sent from known SPAMbot X. Cut off this person's network until they're fixed?

Notify them; clean them up. That's an example of DPI use that's very, very benign. On the other hand, you've got DPI use like some of the ad injection schemes, which I find totally nauseating.

So, it -- for network measurement, however, most of the stuff you want is actually lower level. The packet timing and passive analysis is often far, far more useful unless you suspect that a network is specifically discriminating against one class of traffic or another, and then you'd have to use DPI to validate that that's what they're doing.

But until then, mostly just headers are sufficient. Also, net flow alone is a huge source of information.

MR. GURIN: Okay. Great. Thank you.

MR. KNAPP: Again, just following on that same line, it -- what about the responsibilities between the application providers at the other end.

So, for example, I've got the transmission information and then I've got the application information. Should it be the responsibility of the application provider or if a filter you put, for example, on your PC to deal with the SPAM or viruses that come in on content or is that the responsibility of the ISP or should it be up to the consumer to opt in or not? Any thoughts on that?

MR. FAULHABER: Could you clarify that question?

MR. KNAPP: Yeah. In other words, I think the point that Nicholas was getting at is that if your ISP is getting into the packets and identifying the content and looking for SPAM or viruses, I mean I've heard some folks say, no, that shouldn't be within the realm of the ISP; that leave that to the filter you put on the PC.

So, are there -- I mean -- yeah. Okay.

MR. FAULHABER: I think the ISP could ask the customer, do you want us to do this, because there's a lot of things that are done better in the -- I mean I know this is end-to-end heresy, but there's a lot of stuff that you can do on the network which is going to be done better in the network.

And I think this might be -- at least some of this -- one of them. Just -- but transparency says ask the customer do you want us to do this.

MR. WEAVER: Also, there's a directionality involved. Directionality involved that stopping stuff coming into the customer's system is best done on the customer's system. Stopping bad stuff coming out of the customer's system has to be done in the network due to just simply how root gets work.

MR. DICKLIN: Well, I think from the wireless space, I mean the carriers control everything from certification of the radio modules on the device to certification of the hardware platform when it's done, even to certification of some of the applications that runs on there.

I mean you look at Apple and iPhone as

an example. I mean they have very strict regulations on what kind of applications can be provided on that device. Even if some of the applications aren't malicious, if you're not using certified APIs that they allow you to use, that application will not run on that platform.

MR. MEINRATH: Until somebody figures out a way to correctly identify what is SPAM in my in-box and what is not, I don't want anyone without my consent having the opportunity to drop e-mail that they may consider SPAM and I may consider vitally important.

And I think that's fundamental. This is why you have a network architecture is that you want to empower the edges of the network to make decisions over this.

Now in terms of the opt-in, I don't see anyone, anywhere ever saying that people shouldn't be able to provide services or applications to end users.

MR. GURIN: But -- if it's as an opt-in?

MR. MEINRATH: As an opt-in.

Absolutely.

MR. GURIN: As long as it's disclosed?

MR. KNAPP: If you'll forgive me -- for one more question. And this is for Chairman von Finckenstein.

So, the rules have been in place and parties have started to disclose their practices. Is this up online now that folks can look at, and I was just wondering if anybody has taken a look and part B of the question whether there's any disclosure provisions do you see anybody doing out there that you think would be a great model?

MR. von FINCKENSTEIN: No, I haven't see a great model. It is actually all up now, and we have a very live secondary market, because we mandate secondary ISPs to be able to purchase bandwidth at a (inaudible) price.

And so there is the internal industry monitoring, especially the secondary ISPs, very carefully monitor what the primary ISP puts on its website, because it affects them directly. And we want to make sure that's fully disclosed, et cetera, and in, et cetera.

Another thing: So far, we have never received a single complaint or suggestion that we should look into something that somebody is not following up.

But as I say, it's early days. But I count very much on those -- and on intermediaries, consumer groups, but also those secondary ISPs to make sure that the system is honest and people do what they're mandated to do.

MR. KNAPP: Okay. Gerry?

MR. FAULHABER: I think this is an appropriate time to make this remark having had that comment and heard from some of the ISPs here.

Transparency is in a good situation now, particularly with broadband ISPs, because I think many of the broadband ISPs, you know, have seen what happened to Comcast. They've got religion on this. Okay?

They know this is coming. They know they have to get with the program, and I think this is a good time for a cooperative solution to work with these guys. I'm not at all convinced that the application and content providers are anywhere near that degree of willingness to accept transparency.

But I think the broadband ISPs are, and I think that's a great opportunity for the Commission to look into something, which is a much more cooperative venture with their guidance, but, nevertheless, something where you can tap into the industry expertise on this, as well as from, pardon the word, pundits and academics.

MR. GURIN: Very good. Yes, and thank you. Very good point, also.

MR. KNAPP: Do you want to go into these?

MR. GURIN: Go ahead. Yes, so why don't you take that one?

MR. KNAPP: The Internet's end-to-end principles put the management of a connection into the end stations. TCP, HTTP, and other session protocols operate in the users PC and the content server.

Transparency must include these parts of the system. The question I'm going to take the liberty to modify, it says, will the FCC include end-to-end vendors -- I'll change it to should the FCC include end-station vendors in the process? Thoughts on that? Yeah. Go ahead, Nick.

MR. WEAVER: As someone who's looked at the Netalyzr stuff for the NATS, oh, Lord, I'd love a hammer to get the NAT vendors to cooperate, but I don't think a regulatory system would help given the number of vendors low cost, et cetera. I don't know what is the role of mandating the end systems actually apply the standards they're supposed to apply.

MR. KNAPP: Okay, David?

MR. YOUNG: I would take a different tack on that, and it's not so much that a regulatory approach is going to change that, but I think the end systems providers do need to be involved in the process of how you develop these practices, what the best practices are, what the effect of not complying with those practices are on the customers of these devices, and that, you know, that dialogue between all the different players in that ecosystem is beneficial and should prevent the types of problems that people are concerned about.

Now you can't bring everybody to the table, and there are going to be people who take shortcuts and hopefully the market will weed those people out. But I think that it is vitally important to have all of the various pieces of this value chain represented in the process.

MR. FAULHABER: I would agree. It was a little horrifying to hear someone suggest that here's yet another thing the FCC should regulate. I thought you proposed quite a lot of regulation in this NPRM.

But the notion of bringing these folks into this process from a transparency perspective I think David's right on. I think you absolutely have to do that.

MR. KNAPP: Yes, Sascha?

MR. MEINRATH: Let me add to this --

this sort of intersects with what Gerry was saying earlier -- that the best metrics and measurement tools are going to measure end to end, or at least end to the center, and, therefore, you will empower customers to identify when their operating system is actually the bottleneck as opposed to the ISP's network or when the router at the home is the bottleneck. Or you want to empower end users to identify whether it's application, their home network, their ISP, or something unknown.

And that kind of flips on its head the notion of you need to bring everyone into the room, but rather says, you should be in the room, and if you're not, just know that you will be held responsible if you guys are the ones that are messing around.

MR. KNAPP: Sascha, MLABS, as I understand it, does have at least one of the packages that tries to distinguish?

MR. MEINRATH: Yeah. The network diagnostic tool actually does exactly what I'm saying, and so, Gerry, would have gotten a response to his test that said your buffer on your operating system is the reason why you can only get up to 10 megabits per second, and until you change that, then it doesn't make sense to get, you know, Verizon's 50 megabits per second package.

MR. GURIN: Yeah. And that actually -- I'm sorry, Gerry, go ahead.

MR. FAULHABER: I just wanted to ask a question. I was on MLABS about six weeks ago, looking for the data. And I couldn't find it, which means maybe your website -- you didn't have it or your website wasn't transparent enough to tell me where to get it.

MR. MEINRATH: Right.

MR. FAULHABER: Is it there now?

MR. MEINRATH: Yes. It's actually stored on Amazon. There are links from the page. I think you have to go to the get involved --

MR. FAULHABER: Amazon?

MR. MEINRATH: -- section or something like that. Yeah, Amazon is providing us with free hosting of the 500 gig dataset.

MR. FAULHABER: Oh, so it's not on MLABS, then?

MR. MEINRATH: Well, we don't have 500 gigs to store the data.

MR. FAULHABER: But that's not where to look for it?

MR. MEINRATH: To the --

MR. FAULHABER: On your website?

MR. MEINRATH: Yeah. The link to the data --

MR. FAULHABER: Okay.

MR. MEINRATH: -- is on the website.

MR. FAULHABER: Okay. Great.

MR. GURIN: Actually, we have another specific MLABS question.

MR. MEINRATH: Okay.

MR. GURIN: Would MLABS consider opening up to transparently involve researchers from the Internet Research Task Force?

MR. MEINRATH: Absolutely. Absolutely. We have an open system whereby those that want to be involved just contact us. You can be involved. Those that have tools that want to host on MLAB can contact us and have our tools.

The source code is all open source. We're all about, you know, finding new people and partners that want to be involved in this.

MR. GURIN: Terrific. I have a question I'd like to ask on the format of disclosure and consumer information that I think actually relates to some of what we've just been talking about.

So we've seen a few different examples and models in this discussion about how information can be -- can be disclosed and shared with consumers, and they range from a kind of nutrition label for broadband to the kinds of things that Root Wireless is doing, where you really have, as I understand it, very localized maps of coverage in different areas.

And I think one of the interesting questions here, and this may relate more to things like broadband speed, but may also relate to network management practices, is that there are so many variables; that what you experience at home is determined not only by what the ISP does but also by the computer you have, by things that may happen more locally, and in other ways as well.

And I would love to get a little more discussion of how do you -- where do you think we are on the spectrum between nutrition label and local mapping or wireless coverage. Which of these -- or is this problem more similar to or is it somewhere in the middle or does it depend, Nicholas, as you said whether you're at the general level or at the geek level? Yes.

MR. FAULHABER: I've thought a little bit about this. In telecoms, we actually have a history here, okay? And one of the things we've learned is in things like networks, which are subject to lots of random variability, okay, depending on what other people are doing, and are time dependent. Are you getting on the network when everybody else is, which is when it really matters.

And we had in the old telephone world a

way to deal with this. We said the probability of blocking is 1 percent. Now it's a probabilistic -- oh, during the busy hour. It's probabilistic, time-dependent.

It's more complex with the Internet -- with broadband ISPs, but you can say things like in your neighborhood, or your city or something, you name the aggregation, at your service tier, okay, customers have received at least 10 megabits download speeds, 3 megabits upload speeds during the busiest hour of the week. Stop.

Notice it's a probabilistic statement. Anything you make -- any statement you make in this business without probabilities on it is wrong. I can guarantee you that. Okay.

It has to be place-specific. It has to be busy- hour specific, but you can do it. It's doable.

MR. GURIN: Great. Ron?

MR. DICKLIN: Yeah, I would definitely agree with what Gerry said there. In fact, if you look at what Root Wireless has really been spending a lot of their time over the last year has been just on this issue. You know, I come from an engineering background, and I continually fight some of our marketing people that the information we've been presenting is on the geeky side.

We've gone through many, many variations of how we present the information. In fact, one of our partners, C-Net, who's going to be posting the information for us, presently right now is -- has a pretty good investment of trying to make this information so it's very easy for the end-consumer.

On that notion, there's also people that want to dive deeper into these details, so having some kind of a progressive disclosure that upfront if somebody can get their information that they want very quickly, can get in, get out, make their decision. But if different people want to get to different layers, they have the capability of being able to dive down deeper.

MR. GURIN: Great. Thanks.

MR. KNAPP: Parul, maybe -- because you made some comments about this, too. Just what your thoughts would be on the kind of information that would be provided to consumers.

MS. DESAI: And are you talking about formatting or specific?

MR. KNAPP: The formatting and the.

MS. DESAI: Well, I agree with some of the comments that have been made here that there might be some thought put into something more high level for people who just want a generic idea of what's going on, and there's (inaudible) in something more detailed.

But I think it's important to have the detailed information available to people who want it.

But I do like the idea of something more high level and more generic for those who may not have -- or may not want the specific detailed information.

MR. GURIN: Good.

MR. KNAPP: Maybe we could also hear --

I'd like to hear from the service providers as well. Fernando and David, your thoughts on this? I mean --

MR. YOUNG: Yeah. I think it's an excellent topic because, you know, to Gerry's point, that in order to do this, it's not a one-time test that you would be using. You have to collect a statistically valid set of data, and that is going to provide a host of information about average, median, and various distributions of those tests.

And all of that tells you something about the user experience or the probable user experience on a particular service. And so, I think having a group of people who are interested in this topic -- the service providers, the consumer groups, the consumer people who understand how consumers process information and focus perhaps -- but bringing these people together to come up with what are the right ways of analyzing this data and communicating it to consumers is absolutely the right approach.

And trying to do it through sort of a notice and comment process is probably not going to get to where we want to be.

MR. LaGUARDA: I have to agree with that. I'll echo it, not be boring, but because I think it's right that this is something that in order to really ensure, to paraphrase Dr. Seuss, that someone is speaking for consumers, then we need to collaborate here and cooperatively work on best practices, because, otherwise, what's going to happen is we're going to end up with a process where the outcome either is static, shortsighted, or ignores the average consumer as opposed to the geek.

And I think it's -- that's important. And I also would say there's a lot of work that has gone into -- certainly, at my company, and I believe at others -- but I've experienced it by going to our site -- explaining to customers in an intelligible way the various aspects and attributes of the different service plans that we offer in ways that are simple and easy, as is our mission -- what types of applications can be used, what kind of demands you can put on your service.

I think those are ways that customers understand it that are friendly, and that -- we're best suited to intermediate, because that's a role that we play very positively, and I think our consumers or customers, subscribers, react positively to that, because they enjoy/use the service.

But I do think in the marketplace where there are different providers -- there are application vendors; there are others -- it's very appropriate, and, in fact, necessary to come together.

And in terms of where we are on this dynamic where this spectrum between, you know, a simple box all the way up to securities disclosures I think we're definitely not at either extreme. We want to get to a place where information is easier for customers to understand. That's certainly in the interest of providers.

But I don't think we're at a point right now where those definitions will have meaning to ordinary subscribers without some more work and cooperation.

MR. GURIN: Yeah. I think that's a great point, and I would say there is clearly going to be an educational component to this as well as the disclosure component. So. Joel? Yes?

MR. KELSEY: Just two quick points. One is that process is absolutely I think very important in bringing, you know, various different folks around the table to discuss, you know, what can and should the Commission do and where might it err on the side of caution is important.

But I think it shouldn't take away from the impetus or the need for the Commission to feel like it has to act here. When you asked where are we on the spectrum -- and I'll get away little bit from speeds and focus more on network management practices here -- and you can kind of go through our submissions in the docket and also in my testimony, but the examples in the terms of service are, you know, we may suspend your account, take any other action to prevent from -- I agree that you may suspend my account, take any other action to prevent me from utilizing certain account privileges, for example, home pages, or cancel my account without any prior notification.

I agree you may suspend or cancel my account from using any or all part of the service that posts content to the Internet or to engage in peer-to-peer file exchanges or other forms of file exchanges that violate the agreement of terms of use. Full stop.

So, you know, I don't want to take away from the need of networks or ISPs to manage their networks, but there needs to be some criteria there to let consumers know when they may get blocked or shut off or kicked off the network, when they may not be able to access a home page, when they may not be able to participate in peer-to-peer file sharing, et cetera, et cetera, et cetera.

MR. GURIN: And we're talking I think

about a different -- a couple of different kind of issues here --

MR. LaGUARDA: Right.

MR. GURIN: -- because we are inevitably going back and forth between performance issues and network management issues, because I think the consumer is both, in a kind of a meld.

But I would say that in addition to clarity, and, Fernando, this is what I was really kind of responding to on what you were saying, I think we should not underestimate whether it's the provider who does it or Consumer Reports who does it, or, you know, the FCC who does it, we should not underestimate the need to really educate people about what this all means as well as giving them data, because I think there probably is an understanding gap, because these can be fairly complex issues.

MR. KELSEY: Yeah.

MR. KNAPP: I had a question that came in over the Internet. What would happen if the Commission merely adopted two rules, one requiring transparency and another prohibiting anti-competitive conduct in enabling competition? Reactions? Would that be sufficient? Yeah. Joe?

MR. FAULHABER: No. Enabling competition -- I'm not sure what the questioner meant by that, but I think the Commission has to take -- I'm sympathetic with that point of view; okay -- the Commission has to be much more proactive in enabling competition. And my particular view here is wireless broadband and putting -- getting a lot more spectrum out. That's an absolute, fundamental -- a has to happen.

It's the only thing we're going to find, in my estimation, where we're going to see competition in the broadband market, and we can begin to let the market take care of itself. Transparency, absolutely essential.

Enabling competition, absolutely essential, but the Commission has to do something about that, and, frankly, you got to do it quickly.

MR. GURIN: Others? Yeah.

MR. KNAPP: Jay?

MR. MONAHAN: I'm going to chime in and say no, as well. When the Commission looked at the Comcast practices, some commentators said that this is really something that's more appropriately handled by the FTC, and my concern then and my concern now is that not every network management practice that hurts me as a competitor actually rises to the level of being a competition violation. Or it may be very difficult for me to make that case or extremely expensive for me to make that case here above the Chinese restaurant in Palo Alto.

So, I think having some rules about what the standard is for permissible practices as well as a way to understand what they are is critical. And I -- by the way, I think that the allusions to -- or the specific references to sort of a two-tiered view where there's sort of a simple consumer -- the thing that consumers mostly want that's more than nutrition label at one level of sort of disclosure and then the rest of it that's more like the prescription insert, because what's in the insert is what those of us are going to want to see who are going to evaluate is this a practice which is one, hurting me, and something that's permissible.

The consumer may just care about am I going to be shut down or is my speed going to be slowed at dinnertime.

So, I think approaching those as one pile of information, but two levels of disclosure in terms of the way it's displayed is probably a good way to think about it.

MR. GURIN: Great. We have more questions from the audience or from ourselves?

MR. KNAPP: Well, I have one.

MR. GURIN: Please, go ahead.

MR. KNAPP: Wireless, which many folks point to and say well, that's different. It -- how should we think about wireless in terms of disclosure? I know Canada treated wireless a little bit differently. Will the label look the same, but the information might have more variability to it? I mean thoughts on that? Yeah. Go ahead, Dave.

MR. YOUNG: Disclosure is just as important in the wireless arena, and Verizon Wireless makes an effort to disclose what a customer will reasonably expect to get when it uses our 3-G or 2-G data services.

And I would expect that we would have similar disclosure as we roll out our 4-G LTE. So, I think the principles are the same and that's that consumers have a need to know this information, and providers have a vested interest in giving them the information that they need.

MR. GURIN: Great. Yes, Ron?

MR. DICKLIN: I think on the topic of wireless, if we're talking either fixed wireless or mobile wireless, because one of the additional features when you start looking at mobile wireless, there's that where factor, you know? If my intention is to buy a device or that I'm going to use in multiple different locations, I need to have that additional piece of information to know that, hey, when I travel to these other areas, am I -- what am I going to expect to see in those areas.

I think you also see a little bit larger dynamic, and some of this can go to Gerry's point, where, you know, we've got some spectrum issues in here. You know, there's -- you go to downtown San Francisco at 12 o'clock noon as compared to downtown San Francisco at 12 o'clock midnight, you're going to see widely disparate types of service in there.

MR. GURIN: Yeah. Yes, Jerry?

MR. FAULHABER: That mobility issue I think is very, very important. There's certain dimensions that wireless stretches that, you know, wire line does not.

And I think the other thing we see because of mobility and other factors, wireless traffic is a lot more bursty, okay, and that means making probability statements is even more important. It really is a different space. My analogy and this is -- this is why it concerns me a little bit about regulation here -- what it looks like we're trying to do is regulate these same things the same, because we say, well, you know, there's wire line, wireless, what's the difference.

And I have the feeling this is a little bit like the 1930s when we decided, well, we're regulating railroads. Why not regulate trucks? Disastrous decision. They are very different, even though they're about carrying traffic.

I have the same concern about this NPRM. These are very different businesses, even though, from a customer's point of view, they're both carrying traffic.

But if we try to regulate them the same way, I think that's a bad decision.

MR. KELSEY: If I could just jump in real quick and say, you know, I think from a consumer point of view, it's a real head-scratcher that the same opportunity shouldn't be afforded to you on your PDA accessing the Internet as on your laptop plugged into a wall.

And so, you know, I absolutely think we need one national framework for both ways to access the Internet. I think that -- and I couldn't agree actually with David more, that the principles are the same, right? The principles here are the same and so, you know, looking at the entire NPRM, if you get the discrimination piece -- definition right; if you get the reasonable network criteria correct, then regardless of whether it's a wireless or wired network and you may have to take different factors into account based on how you're judging the network management principle, but as long as the test is correct and the discrimination definition is right, then I think the Commission has the ability to do that in a way that makes sense for consumers.

When it comes to the world that, you know, leaving the network management world and going into the consumer world of wireless, I think, you know, it's a little bit more nettlesome or a harder nut to crack, because you have things like switching costs or early termination fees, things like locked handsets that stop consumers from being able to take their -- the device that they've invested hundreds of dollars in from one network to another. You have things like price -- you know, pricing that's much different than it is in the wired world. You know, $0.20 for a text message after you go over your limit or however much per kilobyte after you go after your data limit.

And all of those things I think need to be disclosed to consumers in a way that they can actually understand so that they don't experience sticker shock once they get their, you know, $500 bill because their daughter went over the texting limit.

MR. GURIN: And that gets us I think a little bit more into the territory of the general Notice of Inquiry that we put out last --

MR. WEAVER: Right. That's correct.

MR. GURIN: -- August on disclosure. I see a couple of hands here. I think in order it was David, Nicholas, and Sascha, and then I think we're at about -- about at a wrap.

MR. YOUNG: Oh, yeah. I think I need to make it clear that when I talk about the principles being the same, my position is that this is not an appropriate area for regulation and that the appropriate way to deal with it is with an industry best practices. I think that wireless has actually set a good model out there with the CTIA's Consumer Code, and I think that that can be built upon. I think it's a -- provides a good model for the wire line broadband industry to look at as well.

And so, I just want to make that very clear that I think the principles are the same. I think that the value to consumers is important. I think that we all have an interest in providing that information and doing it in a way that's most meaningful to consumers, and I don't think that a regulatory process is the best way to get there.

MR. WEAVER: Just a couple thoughts as just a random consumer. I actually think there's a difference between regulation on disclosure versus regulation on behavior. I think conceptually they're very similar on disclosure requirements. We want the same thing, but behavior is a completely different bailiwick; like I would if I was a wireless provider, I would be so cracking down on peer-to-peer because of the differential costs.

The other thing that I think comes into play in terms of disclosure is there needs to be -- if there's ever a case of variable pricing -- so pricing per usage -- there needs to be a way for a consumer to impose a ceiling on total cost per month, and there should be a default ceiling of total cost per month that is some fraction, like 2X, 3X, 4X the base cost, and when you exceed that ceiling, you're cut off rather than run up the bill.

It's just -- why that isn't out there I don't know. That's the reason why I don't have a Droid phone. That's a reason why I don't have a mobile broadband card is they don't have that mechanism built into the system in a way that I know of.

MR. GURIN: Fair enough. And Sascha?

MR. MEINRATH: Sure. I will just kind of sum up, which is that, you know, consumer empowerment needs three things. It needs disclosure of service levels and guarantees that you as a consumer you need to know what it is that you're buying. It needs documentation of on-the-ground realities, i.e. what's actually happening in people's homes across the country. And it needs accountability for providers that promise one thing and deploy or give consumers something else.

And you only get there, you only get to that kind of a consumer empowered reality with government oversight.

MR. GURIN: Okay. Well, on that note, I want to thank all of the panelists. Thank the audience. Thank the people who've joined us online. I think this has been a very, very remarkable session really, and I know from our point of view I think that you've all brought perspectives and insights that I think will be very helpful to us as we continue on this path.

I do hear, you know, certainly a lot of issues that need thought and discussion and working through, but as you heard our Commissioners and our Chairman say at the very beginning, we do feel there's a lot of common ground beginning to emerge, and I think part of our role is to really foster that as well.

MR. KNAPP: Well, well said, and I just join Joel in thanking all of you. It was a great afternoon, and stay ready in case we get back to your for more questions.

(Applause)

MR. KNAPP: Thank you.

MR. GURIN: Thank you.

(Whereupon, at 4:33 p.m., the PROCEEDINGS were adjourned.)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

/s/Carleton J. Anderson, III

Notary Public # 351998

in and for the Commonwealth of Virginia

My Commission Expires: November 30, 2012

UNITED STATES OF AMERICA

FEDERAL COMMUNICATIONS COMMISSION

TECHNICAL ADVISORY WORKSHOP ON

BROADBAND NETWORK MANAGEMENT

Washington, D.C.

Tuesday, December 8, 2009

ANDERSON COURT REPORTING
706 Duke Street, Suite 100
Alexandria, VA 22314
Phone (703) 519-7180 Fax (703) 519-7190
www.andersonreporting.net

PARTICIPANTS:

Moderator:

JULIUS KNAPP
Chief, Office of Engineering and Technology
Federal Communications Commission

Panelists:

SCOTT JORDAN
University of California Irvine

K.C. CLAFFY
University of California San Diego

PAUL SACHIRICO
Cisco

PAUL LIAO
CableLabs

TOM SAWANOBORI
Verizon

JONATHAN ROSENBERG
Skype

Federal Communications Commission:

RUTH MILKMAN
Chief Wireless Telecommunications Bureau

SHARON GILLETT
Chief Wireline Competition Bureau

STAGG NEWMAN
Chief Technologist National Broadband Taskforce

JON PEHA
Chief Technologist

PARTICIPANTS (CONT'D):

WALTER JOHNSTON
Chief Electromagnetic Compatibility Division, QET

Other Participants:

ALISON NEPLOKH

MICHAEL GOLDSTEIN

RON REPASI

STEPHEN BUENZOW

SAURBH CHHABRA

RASHMI DOSHI

TOM PETERS

* * * * *

P R O C E E D I N G S

(10:03 a.m.)

MR. KNAPP: Good morning. I'm Julius Knapp, the chief of the FCC's Office of Engineering and Technology, and I'd like to welcome everyone to today's workshop, examining the technical aspects of broadband networks.

This workshop is being conducted as part of the technical advisory process for the Commission's open Internet rulemaking. The Technical Advisory Process, or TAP, was established to ensure that decisions in the Commission's proceeding on the open Internet reflect a thorough understanding of current technology and future technology trends. The process will provide an inclusive, open, and transparent form for obtaining the best technical data and insights from a broad range of stakeholders.

Today's workshop is an important step in this process. Our objective today is to lay a foundation for basic understanding for the way networks operate. The workshop will present tutorials by an outstanding group of technology experts on the relevant technical characteristics of various broadband delivery platforms and on current network management practices.

Each tutorial will be about 30 minutes long, followed by questions from the Commission's technical working group and members of the public.

While we have a great lineup of speakers, we recognize that there are a great many other stakeholders that have technical information to share. Today's workshop is just the start of the Technical Advisory Process. We will hold ex parte meetings with engineers and other interested parties to understand the range of views in the technical community on the issues that are presented in the open Internet rulemaking, to identify areas of common ground between the stakeholders, and clarify the scope of key differences.

In addition, SEC engineers from the working group will be integrated into other teams within the Commission considering the various issues that are raised in the open Internet proceeding.

Seated with me here at the table are senior level managers and engineers and technologists from across the Commission who are eager to listen, ask questions, and, most importantly, learn. Before introducing our first speaker, I'd like to take a moment to introduce them. There are some others who couldn't make it at the start but will be joining us as the day progresses.

If I can go around, and I'll start to my right and then we'll move the other way. Jon, just introduce yourself.

MR. PEHA: Jon Peha. I'm the chief technologist at the FCC.

MR. NEWMAN: Stagg Newman, on the National Broadband Taskforce to the technologist.

MS. MILKMAN: Ruth Milkman, chief of the Wireless Telecommunications Bureau.

MR. KIEFER: John Kiefer, Media Bureau

Engineering.

MS. NEPLOKH: Alison Neplokh, Media Bureau Engineering.

MR. JOHNSTON: Walter Johnston, chief of the Electromagnetic Compatibility Division.

MR. GOLDSTEIN: Mike Goldstein, Wireline Competition Bureau.

MR. CHHABRA: Saurbh Chhabra, engineer in Wireless Telecommunications Bureau.

MR. BUENZOW: Steve Buenzow, engineer in the Broadband Division of the Wireless Telecommunications Bureau.

MR. REPASI: I'm Ron Repasi. I'm the deputy chief of OET.

MR. KNAPP: There are a few housekeeping matters before we begin. I ask you that you place your cell phones into quiet mode and also please submit your questions either in writing to one of the attendants here or over the Internet. We can be reached both by Facebook and Twitter.

All right, with that I'd like to introduce our first speaker, Professor Scott

Jordan. Professor Jordan is concerned with the interplay between networking technology, communications policy, and economics. His research interests currently include communication policy, pricing, and differentiated services in the Internet and resource allocation in wireless multimedia networks.

He currently is professor in the Department of Computer Science at the University of California in Irvine, and during 2006, he served as an IEEE Congressional Fellow working in the United States Senate's Internet and Telecommunications Issue. I should mention, too, that the full bios are up on the website.

With the, Professor Jordan, I look forward to your presentation.

MR. JORDAN: Thanks, sir. Get an idea where I'm coming from, I've been a researcher in networks for about 20 years, focused particularly on quality of service issues both in wireline and in wireless. And then, as was mentioned, a couple of years ago, I spent a year on the Hill. And so that's kind of the intro into the policy side.

What I want to do today is I want to kind of lay the technical landscape that'll help give us an insight into the other topics that'll come up, so I'm going to start with a brief overview of Internet architecture. For those who have seen this before, you can tune out for a couple of minutes; for the rest of us, it will be a good refresher.

And then what I want to do is I want to talk about what's happened in the Internet with respect to congestion, and then I want to bring it around to this open Internet proceeding, and I want to get at how -- I want to focus on quality of service issues and how this may play out in different kinds of ways that quality of service could be offered in the Internet.

Okay, so with that, I want to start with one of the two ways that networking folks look at the networking world. So there's two pictures that we approach in everything, and you'll see these over and over again through the presentation. So the first one is this notion of layering. It's this abstract implementations of different functions. So if you view the user as being kind of off at the top of the diagram, the user's view of the Internet is through the content that they're reaching. And this community is fond of thinking of it as a content layer that contains all this wide variety of content out there in the Internet.

The Internet model actually combines this with what's called the application layer, and think of the application layer as being the protocols that are matched or kind of specialized to particular kinds of content and allowed those contents to be presented in different ways. So again up here there's a long list and should go on a much longer list of protocols that help make that content available.

The next layer down in the Internet model is where TCP sits. TCP is what is one of the mechanisms in charge of what e call congestion control. So think of this as one of the mechanisms that lets users know how fast they can transmit into the Internet.

Below that is the IP protocol, and that's the part that's in charge of determining, telling packets how to make it from source to destination, how to make it through the network.

And then finally, below that is what the Internet called the land link layer, and this part's in charge of telling users when they get to transmit, so who gets to transmit when. So the part that I want to stress here is that this protocol stack is fundamental to the way that networks are designed, but also we're fond to talk about this as an hour glass figure that it tends to be a lot of content, a lot of applications, and it narrows down to much fewer choices in the TCP and IP layers and then a fairly wide set of technology is down at the lower layers.

And what this means is it's as much easier to innovate up here and down here, both up and down. It's not that innovation doesn't take place in the intervening layers, but this is where it's a little more of a bottleneck. We will come back to this because this will be important as we start talking about quality of service mechanisms later on in the talk.

The other picture that we use all the time in networking is a geographical picture, a notion of what's where. So let me take the example last week of going onto the FCC web page trying to download some documents, so the FCC web server -- sorry about that -- the FCC web server's up here in this green blob. I'm sitting at my computer down at my university in pink blob. There are two things that happen in order to make that connection from the geographical point of view:

One is that we have to identify what sequence of networks or organizations I go through in order to get from the FCC to UCI. And so in this case it turns out it goes through Verizon, which serves as the FCC's ISP; go through Level 3, which provides backbone service probably to get across the country; and then through Cenic, which

connects all the California universities together.

And then the second step is then we have to figure out within each organization how you get a packet from the FCC web server to the gateway to pass over to Verizon's, how Verizon gets those packets across its network to Level 3, and then all the way through to the end. So this is the second fundamental way that we look at networks is in this geographical manner.

I want to zoom in on three of those providers: Verizon, Level 3, and Cenic. These three have decided to talk to each other in pairs. That peering that happens at this gateway involves service level agreements, and the service level agreements have as part of them descriptions of what traffic they'll transmit to each other. So this is going to be important as we get to talking about quality of service later because one of the primary ideas behind whether quality of service will be offered widely is will it make it into these service level agreements.

And, generally, there's two kinds of

agreements that are out there. One that is probably true between Verizon and Level 3 is that they agree to transport each other's traffic without payment. So it's an idea of I'll transfer your packets if you transmit my packets.

The other kind of agreement that's likely to be true between Level 3 and a smaller player like Cenic is that the smaller level player will generally agree to pay the higher level player a specified amount for transporting a specified amount of traffic. So this is where the traffic contract, or the traffic part of the contract, comes in. And again we'll come back to this idea in a little bit.

Okay, so how do these two things come together? If I take that layered picture and I take the geographical picture and I merge them together, then let's use again this example of I'm over here at the university and I'm retrieving a web page from the FCC. So what happens is that the FCC and my computer both have these particle stacks on from the application down to the lower layers. And the intervening routers typically don't have a full particle stack; they don't need to. Typically, the intervening layers only need these lower level layers.

So for the rest of the talk I'm not going to torment you with the differences between each layer, but there's a fundamental difference between the layers that I'm showing in pink and the layers that I'm showing in blue. The first distinction here is that the routers typically have just the pink layers in them.

Okay, so what happens is they go through this, the FCC web page gets separated off into packets, they get told how fast to go out into the network. IP decides how to packetize them and send them out to the next link. From there it determines what the next hop's going to be.

At that next hop it translates it back into a packet, it looks at the packet header, it figures out where that packet should go, and it sends it on its way, and it goes through router to router until finally it makes it to my machine and bound back up. So you get this complicated zigzag pattern that occurs that's, in essence, just representing the merging of this abstractness of layering combined with the geographical pictures.

Okay, so first thing I want to spend a little bit more time on here is the end-to-end principle comes up again and again in this debate, and there's all kinds of interpretations of it. So I want to be careful about what it means, and in the context of at least this talk, here's my picture of what the end-to-end principle means:

Same picture I just had, source on the left, destination on the right, multiple routers in the middle. And what I said was the typical thing that happens at the router is, comes up to the IP layer, that's good enough to know where the packet's trying to go, and this layer, this box, decides how to route it.

There is a question of, well, what if the router does more? What if it goes all the way up to maybe application layer and it makes some decision there and sends it back down? And what the end-to-end principle says is that this shouldn't happen. So the end-to-end principle says if you can implement some functionality in the blue layers, you implement it in the blue layers and you do it just at the two end points: The source and destination.

If you can't, if there's no way to do it that way, you implement it in the pink layers, and you implement it t throughout all the routers, throughout every hop. So that said, it's just a principle, doesn't mean it always followed, and also this was created at a time when we just had the technical aspects in mind. Didn't have the economic aspects in mind, didn't have the legal aspects in mind. So in some way you can view this open Internet proceeding as a really nice way to revisit this concept and to try to bring in the economic and legal pieces.

The reason that I want to also focus on this distinction between these lower layers in pink and these upper layers in blue and this interface between the two is it's useful in its economic context. The lower layers are really expensive to put out. I'm sure the carriers will comment more on this as the day goes on. So we see few providers' high barrier of entry, like this is the typical economic discussion. The upper layers at application of content, many providers' low barrier of entry, and so it's really this distinction between these two that is the pertinent distinction.

It's also pertinent in terms of thinking about competition. A lot of the focus of the proceeding is on trying to encourage competition. So I'm going to use this example a few times as we go. You have cases where there's an ISP that's offering the service, so he'll pick up classes in example of offering their voice service. The voice goes over IP. In order to make it good quality, they use prioritization mechanisms in their network for their void traffic. And so, in essence, the service that represents both applications voices the application and represents a prioritization that's happening down at these lower two pink layers.

Another example -- we have a speaker from Skype later today -- so Skype offers a voiceover IP product, also. They were just offering it above the green line, but they're not a facility SPS provider. The time may come -- it hasn't yet, the time might come -- where Skype may approach Cox and say that we would like to get quality of service mechanisms from you, all right, and we'll talk about availability and how that happens later on. But the main thing to note here is that competition is going to be reliance on whether these mechanisms are available in order to have a equal kind of offering between these multiple providers.

Okay, with that I want to shift a little bit to talking about congestion. So I want to start with a different picture, and here I want to think about what the demand is for end-user versus time. So if I was going to plot how much traffic is going out from my laptop connection during the day, so what would it look like? So what would the picture look like?

So I just made this one up, but this gives you a rough idea: It's going to be very bursty. It's going to be times when my demand is very high; there's going to be times when the demand is very low or zero. And because of this there is worth noting the difference between the highest level at which I need traffic demand and what my average case is, averaged over time.

And the reason that's going to matter is in terms of thinking about what the subscriber pays for, so what they're not paying for, they're not paying, generally, for this key capacity 24/7; they're paying for something that is closer to the average with some understanding that there's going to be burstiness.

Okay, so how does this show up in terms of networking technology? So let me again start with this one user in pink. Let me add a second user in yellow, and what the network has to do is it has to transmit these combined. So if I just add these two together, I get this green curve, and the main thing to note -- two things to note -- so one thing to note is, again, it's extremely bursty. There's big variations. The next thing to note is there are times when one user wants a lot, the yellow user; the one user doesn't want much, the pink user. And so there's an advantage to sharing this natural advantage that it's much more efficient to multiplex or to merge these users together because there are times when one user's demand offsets another user's lack of demand.

And so that shows up in terms of again thinking about what the average demand and what the peak demand is on the network, and asking the question, if you're an ISP, how much capacity are you going to put into this network? But you have to make a financial decision, and that financial decision's going to be based on the network engineering. And so the thing is are you going to put enough capacity to satisfy the peak demand, and the answer is no. It's going to be very expensive, and it's not going to be used very often, so you're going to put in a capacity that's going to be somewhere in between the two. So maybe it's something like this level.

The next thing to note is then it means that whenever there's demand that's greater than the capacity, then that's what we refer to as congestion. So this is where the congestion comes from. It's this notion that there is, by design, periods of time when the network's congested because the demand is greater than the supply. The thing that's important here is by design and this period of time is measured, typically, in seconds. We know there's rush hours during the day, but these times where it's congested don't last for hours on end; because of the burstiness it lasts a very short amount of time.

Okay, so that's the end of the background. Now let me jump into techniques. And I'm going to cover this at a very broad level because I expect that a lot of the speakers following today will get into more details on this. Okay, so how do you deal with this? The

simplest way I can deal with it is whenever there's congestion I just take those packets that are coming into the router, we chew them up. Congestion means that at some point the queue overflows. What do I do? I just throw the packets away.

So it would be nice to think that the networks did something a little more intelligent, but at this level the simplest thing is just throw the packets away, let the application deal with it later.

Okay, the second technique I'm going to go back to this geographical picture, and let's suppose that that link's the one that's congested right now. So a second technique is I can say, well, I'm going to try to find another route. I'm going to try to route around it.

So maybe Verizon says, ah-ha, we can send it through this router instead. We'll eliminate that congested link. This is very similar to if you're driving on the freeway, you hit some congested segment, you say is there a better route? Is there some way I can drive around that traffic? So that's a very common mechanism. It has a limited ability to deal with congestion because it doesn't get rid of all of this congestion. It depends on how many possible routes. And it could be that all routes get congested, which you're all familiar with during rush-hour driving.

Another way that you can do it is you can say, well, when a user tries to form a connection, we notice that there's no good way to get to the destination and so it might say the network's too busy, come back later. So you don't see us too much in the Internet today, but we're all familiar with this from telephone networks where occasionally, by design, hopefully rarely, there's not enough capacity and you get a message that says, "Come back later." This really more appropriate for real-time application, not so appropriate for data applications, but expect to see it for real-time applications because of that.

Another way to do it is we can say that

maybe there's content that I'm trying to get that I can get from somewhere else. So this is the idea of content distribution networks, or caching, and so in this example we may say, okay, rather than getting that web page from the FCC, maybe Akamai, as an example of a content distribution provider, has cached that content, and I'm going to get it from Akamai instead and so avoid having to go through all of this part of the network. And that's another way I think of is in essence avoiding congestion.

A third kind of technique at a general level is instead of trying to avoid it I'm going to delay traffic. So this is more appropriate for applications that can tolerate just a few seconds of delay, and you can think of is that tiering does this, right? When you subscribe, you subscribe at a certain maximum band width. If you go over that band width, then what happens is that content that would normally go out then if you had a higher limit, will get delayed a little bit and would be transmitted just a little bit later. The next time it's congested again it'll be a delayed, and that delay is probably a few seconds. So if applications can tolerate that, that's a reasonable approach.

TCP, I won't go into the details, but TCP is doing something similar. It's like a freeway on ramp signal, and what it's doing is it's saying: We're going to hold you up until the network's a little bit less congested. A few seconds later and then we'll let you get on. And it tires to pace traffic in order to do that.

The fourth technique which is particularly important for this discussion is differentiation, and I'll call it differentiation at this level rather than discrimination, although both terms get used. The idea of differentiation is during these times that are congested, I'm going to look at all the packets that are trying to be transmitted during that time. And what I'm going to do is I'm going to identify some of those packets as high priority, and those packets are going to get through right now, transmitted

without delay. And from all of the rest of the packets that are low priority, what I'm going to do is I'm going to pick some and then I'm going to use those previous techniques. And then I'm going to throw them out, or I'm going to delay them, one of the best previous techniques that I had.

So the idea here is I want to intelligently choose which packets are high priority and which packets are low priority.

Okay, important here is how is that done? So how is quality of service done? There's in essence two ways to do this: One's a reservation paradigm. The reservation paradigm is I'm going to take the band width that's available, I'm going to segue then off part of it, I'm going to let high priority traffic into that part, and I'm going to try to ensure that the high priority traffic doesn't hit congestion.

How am I going to do that? I'm going to somehow limit the amount of traffic going in. So the common example that gets thrown around in the terminology is "toll lane," more accurately so we

have a couple of toll lanes out in California where I come from where it's not just a toll lane but there's a price that's posted at the entrance to the toll lane, and it changes according to the time of day. And as you get further into rush hour it gets more expensive, and the reason is it's trying to limit the traffic going in there so the congestion doesn't happen on that site. So that's one of the paradigms.

The other paradigm is a priority paradigm. I don't reserve bandwidth but I stamp packets at the priority level, and those packets get some kind of preferential treatment when they hit the routers. So think of this as sort of like priority mail. It goes into the same system, but somehow behind the scenes the network makes sure it gets there quickly than regular mail does.

Okay, so far those techniques are not in and of themselves either reasonable techniques or unreasonable techniques. It depends on how it's used, and this is where the fun comes in. So this is where I'm going to slow down and pay a little bit more attention to this.

When these techniques were envisioned by techies designing Internet protocols, the view that we had was we can see a day when the Internet's going to try to do everything from very uninteractive applications like e-mail. Internet was designed for e-mail and file transfer. Non-interactive, I'm happy if when that e-mail ascends it gets there within a few minutes. All the way through to then in the '90s web browsing became more popular. Web browsing: I type in the URL, I hit Enter, I'm happy if the web page shows up in a few seconds, to streaming of radio and TV, to making phone calls or doing video conferencing. Over here in order to have it be what I consider decent quality it has to get through within a few tenths of a second; otherwise, performance breaks down.

So the techie folks said, okay, this represents a range of applications. We then envision the Internet will eventually try to do all of this. And so what we're going to do is try

to create some notion of worst performance to better performance, and we're going to try to match them to the application. So e-mail doesn't need to get there very quickly, and so we're going to design a mechanism for that, all the way through to voiceover IP and video conferencing these good performance, we're going to clear the mechanism for that.

That's again just purely just a technical approach. The question is, how is this going to work in reality? Well, they created some standards, but the standards just say: Here's how ISPs can talk to each other and make these decisions. They don't tell the ISPs what decisions to make, that's up to the ISPs to decide.

Okay, so what does this mean? It means that now this is how we get to this debate. It's a question of how is quality service going to be used? So examples -- and these are all being in current discussion -- could be as in the previous slide I do it based on the application.

Applications that are toward the left side that are more interactive need better quality of service. And so you get examples like Cox's current trial where they're saying certain applications are low priority, certain applications are high priority.

We get examples where it can be this in the service provider. So again, picking on Cox today, their voiceover IP put their subscribers gets priority over voiceover IP of other subscribers.

The third way as it could be done based on source or destination, it could be the traffic coming from certain places and going to certain places get higher priority. This is what got Comcast in trouble. And I said I don't think we need to pay as much attention to that now than we would have might have had to previously. And the other could be payment that an ISP gives preference based on either the subscriber or the application provider providing payment.

Okay, so here's the way that I want to

tie it together. I want to tie it together by focusing on three ways the QoS can happen, so Example No. 1 we do it based on the application. ISPs look at the application and say more interactive applications get higher quality of service, and let me in this example propose that they do it for free. And if this is one of the examples that's given in the NPRM in a form that's been discussed, so what does this mean?

So tracing it through means I transmit some packets that it hits my ISP, that ISP looks and determines which application I'm using, so that's de-packet inspection. So let's suppose that I'm doing a voiceover IP call, they notice the voiceover IP call, they've made the decision that that's high priority, and they assign it high priority. And it'll be high priority all the way through their network in green. Then it's going to hit the next provider, so in this example it's going to hit Verizon, and Verizon's going to have to make a decision: Do they honor that priority or not?

So if I want to continue it in this way, I could say, well, maybe Verizon does the same kind of thing. They look at the packet, they see it's a voice packet. Perhaps they've also decided that voice packets get priority, and so they give it priority and it gets priority through their network, and on through the rest of the providers.

Okay, so what's the good news and bad news here? The good news is it's fair in the sense that all voice applications are getting equal treatments. The bad news is because I've said it's free, you can ask whether there's incentive here for the ISPs to do this, or what's the incentive? All right, and that's an economic argument that bears more consideration.

A second problem is because it's the ISPs looking at the packets and determining what the application is, we hit a problem that we don't yet have any standards for how to do this. This isn't something that the Internet engineering taskforce or other bodies have done.

And the third question is we don't know

when it hits that gateway between, in this example AT&T and Verizon, if Verizon's going to do the same thing that AT&T does. So we don't know whether that quality of service will be peered or not. Okay, so it's a way to do it, it may work. There's a few problems, and these problems are things that should come up in this proceeding.

The second example, proceeding brings up the notion of managed services so, doing a little bit of guesswork, isn't yet pinned down what this would mean. Let me focus on, in this example, let's suppose Verizon offers managed service. So what does that mean with respect to this picture? So what it might mean is it might mean that both the source and the destination have to be in Verizon's network. They're both Verizon subscribers or Verizon's providing some of that content.

In that case, what may happen is Verizon may make a decision that they're going to provide quality of service to this managed service, and they might make a decision that they're not offering it to the rest of the world, all right? Unknown. It's not addressed. So it's a second way that quality of service can work. It's going to solve the problem within Verizon's network, so it's useful, right. It'll solve the problem with respect to real-time applications like voice or video within one ICP's network. But if it's not peered, right, if it's just within one network, then it means deployment's limited. We have to have both the subscriber and the source in the network, and it's limiting competition in the sense that it's not widely available, it's not widely deployed.

So third example, a lot of times you hear discussions of user choice. I want to give a user choice example, so user choice example might be that I'm back over here, it's my computer. I decide what I want to get high priority, so maybe I make the decision similarly that I want my voiceover IP to get high priority. It goes into my ISP's network. The ISP then the question is when it hits their router, are they going to obey the priority, the wish that I expressed? And my guess is, well, it depends on what my contract with them is. It might depend on whether I'm paying them for that.

So let's suppose that it's in my contract, and they honor that priority, so it's going to get that kind of priority service until it hits this gateway. Now again we have that question: What's going to happen when it goes over into Verizon's network? Because I'm trying to talk to somebody at the other end, and so we have a similar question: Is Verizon going to give priority to that? And this is where that notion of service level agreement comes back in that I mentioned back at the beginning.

In this example, this is AT&T and Verizon. They already have a contract that they signed with each other that dictates the terms where which they exchange traffic. The question is, will that service level agreement eventually have written into it something about quality of service?

For instance, if they're equal peers and AT&T and Verizon say we'll transmit your packet if you'll transmit mine, will they also then say I'll honor the quality of service marking that you put on your packets if you honor mine, which is what in this example you would hope to see. And then we'll go through the rest of the networks similarly.

So items to note about this, if it's done this way, hopefully it's done standards-based, because there are standards for this unlike standards for DPI. We still don't know how widely it's going to be available, and the main reason for that is we don't know if it's going to be done in a way that peers into their service level agreements.

Okay, I'm going to leave you with a couple of challenges, so I wouldn't be a university instructor if I didn't leave you with some questions to think about. So here's my list of questions to think about:

So will quality of service be available

not just within a single network band intent? So I want that picture, not this picture, all right? I want it to be available widely deployed. So what does that mean?

I think what it means is as you go forward and you're trying to pin down what managed services means, however it's defined -- I'm not suggesting a definition -- make sure it doesn't discourage end-to-end QoS, right? It should encourage this.

The same challenge different picture, back to this picture where I said, okay, here's an example of Cox offering vertically integrated service; here's an example of Skype maybe eventually wanting quality of service from an ISP. So same kind of idea is that in order to have it available end to end, that's what's going to give it the competition, So this idea is we want to encourage competition. We want to encourage competition also in this voice and video services. So that's why you want it available end to end.

The second challenge. You will note

that until now I haven't said anything about wired versus wireless, and yet there's a big question: Is wireless different? So here's my one-minute answer to this: Wireless networks are very different. They're very different down here. All the challenges they face are considerable challenges, they're special challenges, they're different down here at the lower layers. They're not different up here. They operate the same as wired networks up here. So whether they're different depends in large part about whether the regulation of these, if there is to be regulation, is done in a way that recognizes this boundary or doesn't recognize this boundary.

If it recognizes this boundary, and it says: What we're concerned about is quality of service mechanisms that are implemented down here should be available to the rest of the world up here, then wireless is the same. My contention. If it doesn't recognize this boundary and it says; Okay, there's a class of services that meet quality of service and the application, and it's all rolled into one, then it may be the wireless is different. And that bears further investigation. That has to be thought about more.

Okay, last challenge. I've had a version of this slide forever. I would be remiss if I didn't put it up, so you all know this story: Telephone network is designed for voice service; cable TV network is designed for video; cell phone network is designed for mobile voice; Internet originally designed for e-mail and file transfer, right, varied data or (inaudible) applications.

We all know where this is going. At some point everything becomes over IP. So why is this a challenge here? Well, we know that we may have to wait for Congress to some day rewrite the whole Communications Act and merge all of these titles. It's not something you can do here. But what it does mean is when you come back to managed services, make sure that however it's defined at least doesn't hinder or run into problems as everything becomes over IT. And I think that's going to be a tough thing to do.

Okay, a quick summary, layering matters. I think this boundary between infrastructure and applications matters. I think it needs to have more attention. Congestion happens, and so the question is subscriber buys what? I think that needs some more attention in here.

Techniques aren't by themselves reasonable or unreasonable. It's a matter of how they're used, and so the focus should be on the use.

Quality of service, I think there's now wider agreement some applications need it. Let's make sure we don't discourage it. When looking at reasonable or unreasonable, it's a matter of what the practice is. I think the big one is do what you can to make quality of service widely available end to end. I think that's really the big one out of this.

And that last light, of course, convergence always the ultimate challenge, try to do the best you can to make sure it works with that.

That's the end of the formal comments. I would be remiss if I didn't say that, of course, I do have some policy opinions on the issues to how to make this happen, but I will refrain from giving those today.

MR. KNAPP: Thank you, Professor Jordan. That was outstanding. We've had a number of panels where Walter, John, Stagg, and others here have co-moderated it, and we got into the habit of giving homework assignments, but when we have professors, I see it still works the other way around, that we get the homework assignments.

Just a quick question: How do you think about the interplay between congestion and management techniques and prioritization versus motivation to improve bandwidth on the networks so you need to do less of that?

MR. JORDAN: Yeah, I view these as in constant tension. There's been a debate in the 20 years I've been in the field about whether capacity would eventually be plentiful enough that, in essence, bandwidth was free. It hasn't happened in 20 years. My personal prediction it won't happen in the next 20 years, continues to scale up, but every time we scale it up people come up with new applications to use it more -- but which is a great thing, which this is what we want.

So I don't ever see that it will completely go away. What that means with respect to congestion management techniques, they need to be there, they need to be done in a smart way. They need to be done in a way that makes the applications work the right way, depending on how they are.

I think of this proceeding brings in two ideas: The positive side usually approaches quality of service, and the negative side is usually referred to as "throttling" or "degrading." I view them as two sides of the same coin and that they should be wisely used so that it represents the best overall utility for the user base which you can. No way to get around that.

MR. KNAPP: Questions from the colleagues? Stagg?

MR. NEWMAN: Could you put up one of your five bubble charts?

MR. JORDAN: Geographical ones?

MR. NEWMAN: Yeah. Yeah. My understanding now is more and more of the traffic of the Internet is actually bypassing several of your bubbles there. So, for example, Google will now do direct connections to the large ISPs, or Akamai will actually there is providing the pipe, et cetera, and therefore, you know, if I want -- if I'm a Google or a content provider and I want really good service, I am then getting my traffic directly to the local ISP. How do we think about that part of the problem?

MR. JORDAN: There's two aspects. One of that is caching; the other part of it is what you briefly refer to part of your question is, if I want a high quality voice service that parallels the same kind of calling that we have in the phone network, then one option, which is common for cable providers, is I run it through, an IP network within here, and then I don't go the rest of the way. I go off into the phone network. And the reason is, is I don't have end- to-end QoS.

So going forward, what do we expect? At some point in the future we expect ever to be over IP; we expect a phone network as a circuit switch network to disappear. The reason is, eventually if you can get QS end to end, it's more efficient to just do all over IP. And so we want to get rid of those kinds of detours through other networks to the extent that they're not economically efficient.

Caching, I think, is a little bit different because it's not really aimed so much at the what I did at the very interactive left-hand side of the spectrum. It seems that the intermediate applications where you're not trying to get the delayed under a few tenths of a second, you're trying to get it down to something a little bit higher than that in interest to avoid really long delays.

So there the thing is just do one other picture, the layering picture matters here. Caching's all up here, so personal belief, it doesn't represent the same competitive concerns as quality of service does, which is down here. I am much less concerned about caching and the anti-competitive problems that may corrupt (inaudible). I'm much more concerned about mechanisms that are put down here.

MR. NEWMAN: Well, going back to your bubble diagram, though, suppose I'm -- the bottom one says Comcast, okay, and they're in the business of delivering video, obviously. They may want to do caching in their own network to delivery their video well and say Disney, you know, doesn't caching get into the heart of that when more and more video becomes IP-based services?

MR. JORDAN: Yes. It depends on whether the video is going to be provided just within one provider or it tiers off into the Internet. So I'm going to cross over that boundary from technical to policy for a moment -- forgive me -- so one way to handle this, which I don't think is in the NPRM right now, is you say that if the source and destination are both within one ISP's network and it's video, given them the option of calling it Title 6. If you call it Title 6, then I don't think you necessarily have to be subject to the rules that should apply if it's peered.

If the video is coming from somewhere across the Internet, then I think these issues that are concerned about an open Internet proceeding matter; if it falls within one provider, I think it can be handled in Title 6. It's a personal policy rec.

MR. NEWMAN: I've got a whole series of questions.

MR. KNAPP: Any other questions? Sure. Go ahead, Ruth.

MS. MILKMAN: A factual question, I think. To what extent is quality of service available today across different networks?

MR. JORDAN: It is not.

MS. MILKMAN: Not at all.

MR. JORDAN: No. Well, I'll let the carriers feel free to disagree with me if I'm wrong. What I get from them and what you can ask them later is it seems to be the shorter term strategy as use it within their own networks to provide their own services makes sense, like there's a good business case for that.

I think there should eventually be a business case for doing it end to end; I don't think it's on the horizon yet, but I think you got to face it now because it won't be that long, and the time to address it's now.

MR. PEHA: I guess my question is now related to both of the last two. Certainly, if you go through five networks and you want a quality of service guarantee, then agreements across networks matter. Given that not both everything goes through five networks per stage comments, and I don't know, maybe you think a guarantee is important, maybe you don't. How important is it to have quality of service just within one network even if the networks don't get these agreements across networks?

MR. JORDAN: I'm going to see putting it in one network is a starting point, but I really want to see it incorporated into the service level agreements. And I don't think it's unreasonable because they're already talking about traffic specifications in those, and since the ultimate objective is widely deployed, enable all kinds of applications and enable competition, I think that's the long- term way to do it. Otherwise you're stuck with having a much more limited number of application providers that own the infrastructure and are able to make the combined solution. You want, I think, a more -- you want both: You want the ISPs to be able to do it; you want the non-facilities-based folks to also be able to do it.

MR. GOLDSTEIN: In your diagram, Cenic isn't really a tier of Level 3, they're really buying transit. Do you see those service level agreements including a requirement on this smaller

ISP to be providing a piece of the end-to-end?

MR. JORDAN: I'm going to duck that question. Service level agreements terms are kept very closely held, and so we know not much about them, and I suspect your next speaker will make a point of saying that we need more information on that in order to be able to answer your question.

MR. JOHNSTON: In the business environment service level agreements on IP private traffic have been worked out for a decade now. Do you think they serve as a good model for what might be a generalized service level agreement? I mean in the competitive environment in the -- if I buy an IP business service, they have specified the parameters that I can measure what I should expect, bandwidth requirements, peak, average, are typically specified in the contract. And it's been worked out in an environment in which companies compete for business.

Is that a good foundation for looking at what we need to do in the future?

MR. JORDAN: My instinct is that the

focus for good reason is on the access network. It's on the access network because that's where the competition is limited. As soon as you pass over that point, then my instinct says that you don't need to worry about the same competitive concerns because here's where competition kicks in, I hope. And therefore my instinct is you don't need to regulate peering agreements because it doesn't fall on the access side. There's a big asterisk.

MR. JOHNSTON: Where these are -- what I'm referring to were the business side IP contracts, I'm a small business or large business, I can go to AT&T and it will give me an IP service with guaranteed quality service only across their network. Or I can go to Verizon and get a similar service level agreement, contractual. If they don't meet the service level, I actually get money back. And how that's specified, it's been worked out over the last 10 years.

MR. JORDAN: So here's my temptation:

My temptation is to say that they are a subscriber just like a consumer in house is a subscriber. They have a different contract and different package. They're buying a much bigger package, they're paying a lot more money for it. But I'm tempted to say you're a subscriber, and as long as you're a subscriber and you write the rules as pertaining to subscribers, it, I hope, should work out similarly.

So just like you're not going to dictate detail for subscriber agreements for consumers, you're not going to dictate it for businesses either. The part where it gets a little more complicated is if they're on the other end of the network. I won't dive into that at the moment.

MR. KNAPP: Thank -- go ahead, Stagg, one last question.

MR. NEWMAN: Oh, somebody else then. I've had my share.

MR. KNAPP: Go ahead, ask.

MR. NEWMAN: Okay. Go back to your -- now, your layered diagram with your nice bright line. Yeah. Okay, so what percent of the traffic today is TCP versus UDP versus, say, Akamai doing proprietary solutions at the third layer there?

MR. JORDAN: Wonderful question to ask your next speaker.

MR. KNAPP: Our next speaker is Professor KC Claffy. Because she has a very tight schedule she wasn't able to make it here, but she -- we're very grateful -- was willing to present from the West Coast. And KC, are you there?

MS. CLAFFY: Yeah. Can you hear me okay?

MR. KNAPP: Yeah, terrific. KC Claffy is founder and director of the Cooperative Association for Internet Data Analysis, or CAIDA. And am I pronouncing that correct, KC?

MS. CLAFFY: I tend to use CAIDA.

MR. KNAPP: CAIDA.

MS. CLAFFY: Yeah.

MR. KNAPP: Okay, based at the University of California's San Diego Super Computer Center, an adjunct professor in the

Computer Science and Engineering Department at UCSD. KC's research interests include measurement analysis and visualization of Internet workload, routing, topology, performance, and economic data.

CAIDA seeks through the collection and curation of strategic Internet datasets and freely available tools and analysis methodologies to improve the scientific integrity of network research and to promote more informed engineering, business, and policy decisions regarding the Internet infrastructure. KC, I think we're ready to start and, hopefully, we'll do a good job of keeping in sync.

MS. CLAFFY: Yeah, okay. The webcast is really great. This is the first time I've given a talk and then seen myself come back to the Internet on the webcast, so we'll -- great Internet metric -- Internet performance metric. And I really appreciate you guys letting me do this remotely.

I'm actually busy writing an NSF proposal this week on Internet economics, which is the biggest challenge I've ever done on NSF proposal writing. And then I just got slammed with a bunch of deadlines, so I'm going to try to do this remotely as best I can. Fortunately, as Scott set this up perfectly, he gave a fantastic talk on technical background.

I'm going to focus on the historical context, but still coming from a technical architectural perspective. So hopefully, it'll fill in some gaps for those who may not have been there through some of this, like I've had a front row seat, at least from the academic perspective.

This graphic I have in the front is our latest attempt to try to do some visualization of the inter-AS topology, including some semantics associated with the topology, like how many routes are exchanged between AS's and how many tiering sessions are happening. But to be honest, I'm not going to talk about this slide and this graphic anymore. It's basically eye candy on the front; it's not even ready for public release yet, the tool that's doing this. But it does give an

illustration of the sort of things that we are trying to think about at CAIDA and it gives a relevance to this -- conversations happening here. So next slide.

And Walter is a (inaudible) over here on the side got me on Skype as a back channel so he can -- I may give him so instructions or if you have questions in the -- that you want to interrupt, please -- and let him know or you can just speak up and I'll slow down. That'll also slow me down in case I start talking to fast to be understood.

So just as a brief summary, I have quite a few slides with even more words on them and some of those slides are really more background information. I won't go over ever bullet on every slide, but leave them there for you to consult later or ask me questions about later. So what I really want to make sure I leave you guys with today is what do we know. And, of course, from an academic perspective that is going to be fairly limited. What do you know about traffic management? But I am going to be able to provide you some historical and architectural context that I don't think you've gotten, at least from reading the documents that have been available to me. So hopefully, that'll help.

Now the punch line: I'll just let you know now and it's not going to be a big surprise to you and, in fact, I've seen it mirrored in other countries making decisions on this exact topic lately, is that we do have a conflict we're trying to resolve in the future. In the short term, the best approach, I think, is going to be the thing that the FCC has already been trying to push for in the last couple of years: Reactive basis based on the circumstances that have arisen.

But the key is going to be transparency and providing objective data to justify any sort of traffic management needs, as well as transparency in the process of executing the traffic management. I'll talk about that more later, but that's a key point.

In the longer term, I think it's great.

Again, the FCC has already started to do this to build a -- not only a technical advisory function, but really an interdisciplinary advisory function on how to manage the traffic management regulatory function, whatever that turns out to be, if there is one. Because I think that part of the problem that gets us into trouble, and I see this in a lot of other areas of -- policy advice that are happening right now, including an ICANN -- is that we try to segment the technology, the policy, and the economic advise because they're all very complicated and we think it'll be simpler if we go at them from one angle at a time. But it turns out to be kind of a recipe for failure because the one thing we've learned about the Internet is that it connects everything, including these different disciplines in trying to think about our problem. Next slide.

So let me say a bit about my background. I haven't been studying as long as -- God, I'd say about 15 years. No, I guess it is almost 20. But I started in grad school studying this. In fact, while I was a graduate student I did the very cheeky thing of trying to publish an interdisciplinary paper before I even knew what the discipline was I was studying.

So I wrote a paper that is quite in the eye of this tornado called "Mitigating the Coming Internet Crunch." This was before the Internet even privatized. So there was really a single organization that operated and it was a government organization -- indeed, the National Science Foundation -- operated a general purpose national backbone for the country. A general purpose in terms of research and education support because really it wasn't supposed to be the case that Apple was supposed to be -- was to use NSF's net to communicate with IBM, although that is what ended up happening because of the limitations of the technical architecture in preventing certain traffics from getting to certain places and the technical capabilities that were implementing the architecture at the time.

So we had a single, you could argue, a

single operator, administrator, of the -- itself. But even back then we had very many different independent administrative operators of regions of the networks. So in the NSF's case it was an explicitly designed hierarchical architecture where the NSF operated or heavily contracted out the operation, contracted out to a partnership of academics and private sector to operate the backbone.

But then they had regional networks explicitly designed to attach sort of up the chain, upstream we call it today because it happens in a commercial world too to some extent, upstream to the backbone. And then we campus networks, which we would consider like end sites today, would attach to these regionals. So I was at UCSD and we had a regional network in California. It was a decade later to become AT&T. And then there was the (inaudible) backbone at the "top."

Today we have several backbones you would consider at the "top." Today we call them

Tier 1. And, of course, today they are competitively provisioned. But at the time, they were more or less cooperatively provisioned, although obviously everybody still stressed about resources within their own domain, so it wouldn't be fair to say that everybody treated each other's resources as their own. It wasn't that cooperative. There were budgets to be worried about, for sure, just like academics always are worried about.

But it wasn't just -- it wasn't the sort of architecture -- economic architecture in terms of competitive provisioning that we have today, and yet we still had the exact issues. We still saw video coming. Cornell had applications some of the oldest folks in the room may remember called "See You, See Me," which basically allowed (inaudible) video, allowed what my mother uses Skype for today and it worked in 1992.

And so several of us, including the PI, the director of the backbone architecture himself, realized that this wasn't going to get better;

this was going to get worse. We were going to have congestion problems on the Internet. And I'll talk a little bit more about what we recommended in that paper later.

But just to make the point that these are not new questions. We're asking the same questions that we were asking 15 years ago. The context is different; I'll talk a little bit about why. And I can't say that we've solved the problems yet, although certain domains we have. And I'll talk about that later too.

And then my real dissertation work, which I graduated the year after I wrote that paper, co-authored that paper with several other folks, was on technical Internet traffic: Using public traffic data, whose collection was mandated by the U.S. Government because it was traffic on the NSF net backbone and that was a U.S. operated and regulated -- sort of regulated -- network.

It was a thesis that I did in 1994, right before, literally months before, the backbone retired and privatized, essentially. The government decided to privatize the network. And you couldn't reproduce that thesis today unless you're in Japan, where they have quite a bit of cooperative data sharing going on and, in fact, somebody reproduces a similar type of traffic characterization, a little piece of it, in Japan a couple of years ago.

Okay, next slide. I've given this slide to some FCC folks before. It's very busy. It's a historical perspective. It's just a timeline of events that have happened in the last four years. And I'm not going to go over all of these bullets in detail, but the big point -- okay, I see the slide isn't up yet on the webcast, but -- I'm on slide 4.

The big point -- the big bullet point is this architecture, the Internet architecture that we're using today was introduced first in 1966. I mean the TCP/IP architecture. Technically, I mean the IP architecture or really the underlying packet switching architecture that lead to the IP architecture. And the important thing to know

about this architecture is that the proposal that Larry Roberts wrote when he decided to submit this to ARPA at the time, was titled "Toward a Cooperative Network of Timeshared Computers." And that's essentially the same architecture that we're using today.

It's an architecture that was designed for a cooperative network. That is a network of maybe different administrative entities, but operating under the same overarching administration. And 10 years later, we came up with a routing architecture to match, which took advantage of this explicit structure in the network that was built, the structure that I mentioned earlier.

That structure no longer exists today in any real sense, and I'll talk about that a little bit later. The big problem, because we are still using the same technical routing architecture to support a network that is fundamentally differently shaped and I mean shaped in a way that the routing architecture is no longer a good match

for. So that's a problem that most people studying inter-domain routing in the Internet understand today and yet none of us have a path forward to solve the problem because that's a pretty big problem to upgrade the routing architecture of the Internet. That's not something that any given ISP, or indeed, any given country, could pursue on their own. So really, that's another talk.

But just to kind of drive on the point that we are using really 30-year-old, 40-year-old architectural technology that was a design for a completely different, not only economic and political context, but really quite different technical usage at the time. And we're suffering. And in fact, you could argue, all of our problems are rooted in this incongruity between the architecture we're using and the demands that we have put on it.

So I'm going to skip the rest of this slide and go to the next slide where I focus on -- I have two slides here: What has not changed since let's say the mid-'90s when the government decided to really privatize the network and let the private sector take it as far as it could go and what has changed. So what hasn't changed?

I mentioned in the last slide that the network architecture has not changed; still using an architecture built for a cooperative environment and relatively low bandwidth stochastic applications. That is not video; very important to keep in mind. The architecture is basically a misfit for video and to the extent that we're using it for video now, we're putting sort of layers of tax on top.

That's the academics. We call it tax, but I think Verizon would call it innovation and I would probably call it innovation, too, if I were Verizon. But if you look at it from a purely architectural perspective, it's kind of ugly what we have to do to get the IP architecture to do video.

The routing architecture also hasn't changed. That's a bigger problem as far as I can tell, although, obviously, I think the network architecture is a problem, too, because the routing architecture -- we are headed for a sky is falling kind of scenario eventually. We don't know when.

We can't really model how the -- what the failure mode will be, but we do know that we're going in the wrong direction because the shape of the network, again, is evolving naturally sort of by natural market forces or economic forces of how can -- connect together in a way that is not good for the -- is not allow the routing architecture to be efficient.

What else hasn't changed? Well, the addressing architecture, unfortunately. We're still using IP, addressing architecture designed for the military in the '60s, and we are almost out of these addresses. Like literally in about two years is the recent apocalypse time stamp, we're not going to have any more addresses to hand out in the current political regime of the way that we hand out addresses.

And we're going to -- most likely what the address registries have voted for in the last couple of years, although by quite of a jagged consensus, I've heard -- privatize the -- how people just buy and sell addresses with unclear regulatory framework for that, and again, hard to regulate that on a per country basis.

So that hasn't changed and yet we're going to sort of change the political architecture around it because the technical architecture is too brittle to change. We have failed to execute the IPv6 transition thus far. I can do a whole other talk on that, but I'm not going to now.

The transport architecture hasn't fundamentally changed, although you see more experimentation now. BitTorrent is coming out with a new type of PCP to take better advantage of the network as they would say. The naming architecture hasn't really changed either, the DNS architecture, although this is a pretty complicated topic. I won't go into too much detail here. Depending on how you measure it and what aspect of the architecture you're talking about, you can make arguments that the naming architecture has changed.

But, in fact, the biggest change is to the DNS architecture of the Internet are literally about to happen in the next 12 months. The ICANN, and in collaboration with a lot of people that have helped, are trying to deploy security DNS (inaudible) for the first time as a route of the DNS, as well as, expand the DNS to include many more top level domains.

So instead of.com and.net and.org -- not instead of, in addition to those, you'll be able to have.verizon and.ibm. And initially the most important party will be internationalized domain names, like having.china in Chinese characters. So we're trying to support a fundamentally different alphabet, different -- many different alphabets that, again, was not considered in the original architecture and yet the same architecture, fundamentally, will be used to try to support those (inaudible). And then, of course, we still find that we have the same sort of economic aspects to bit transport. That is it's still -- basically impossible to move packets around over long distances to a wide variety of places. And I'll show another data slide in a -- I'll show a data slide in a second.

We still have a problem that end points receive and send, not to their own knowledge, much unwanted traffic which makes a lot of pricing models, including metering, tricky. I'll talk more about that in the next slide. But the bottom line is a lot of things in terms of China -- Internet as a critical infrastructure, I have a list of these things I call the fours S's, are all still in progress; quite a bit of progress.

Okay, next slide. This next slide is just a data point of how much unwanted traffic we do see. So here's a data point of a network that we measure that is easy to deal with on the privacy side because there's actually no host on this network and yet this network gets 20 meg -- up to 20 meg, and you see it's growing over the last 2 years. And part of this is Confikr, and part of this is just allowing more traffic in so we could see how bad it was because we were filtering in to (inaudible).

But you can see here -- and this is a large chunk of address space. I don't want to imply that every -- every IP address that doesn't have a machine on it on the Internet is getting this much space. This is a big chunk of address space and yet it's several megaseconds -- 20 megaseconds -- that turns into about several gigabytes per hour, well over a Comcast 250 gigabyte limit.

Now, Comcast is obviously dealing with these -- sort of the pollution out there, some people call it background radiation of Internet traffic, by modulating their metering techniques to sort of give a bit of margin of -- to deal with this. But this is something that we have keep in mind.

There is a lot of crust out there on the Internet that we haven't really even fully characterized yet and it's always there and it's growing. I can't argue that it's growing faster than the legitimate traffic, for some -- legitimate, but it's certainly growing and it's something to be looked at further.

Okay, next slide. And then, of course, just an obligatory quote from the (inaudible). What also hasn't changed is that there's always going to be areas of the network -- Scott drove this point home, previous speakers have driven this point home -- there's always going to be areas of the network where there are resources to manage. Of course, yes, if we get a one- shot capital infusion to get everybody fiber to the home, we're going to be in a much different world than we're in today. And I would love for that to happen, and I don't see it happening in the current economic climate.

But even if that happened, and, of course, we always have the wireless environment to compare it to, wireless for the foreseeable future, due to physical limitations, is going to have, at some points in the network, resource contention. So an obligatory here -- this problem isn't going away, although I do want to -- I do want to clarify something that Scott alluded to, which is it's a really different problem in the wireless versus wireline/fiber world.

Okay, what did change? I've got to pick up the pace here because I'm a third of the way through and I'm halfway through my time. Again, a lot of words on this slide number 8 and I want -- I don't need to go through them all individually, but just to give you confidence, it really -- it is a different world than it was in 1993. It's a different world in a lot of important ways, despite the fact that we are still using, to a certain level, fundamentally the same technical architecture to support this quite different world.

So okay, the industry structure has changed. Obviously it's inverted from the telephone whereas the telephony world, you know, it started out in the private sector and then eventually was turned into a regulated monopoly.

The Internet is interestingly quite the opposite of that. The Internet started out as a public sector activity. In fact, for the first 30 years of its life it was essentially publicly U.S. Government-funded, maintained, operated, managed, code regulated, although it's obviously a different kind of regulation that we think of today.

OPERATOR: Please pardon the interruption. You're conference contains less than three participants at this time. If you would like to continue, press star 1 now or the conference will be terminated.

MR. KNAPP: It's more of that background noise creeping in. You still with us, KC? She --

MR. JOHNSTON: The one thing that hasn't changed is the complexities of (inaudible).

SPEAKER: Oh, here it is. You've got to turn it (inaudible).

MR. KNAPP: We'll pause for just a moment to reconnect.

MS. CLAFFY: Can you hear me now? (inaudible) folding from the telephone system there. Are you --

MR. KNAPP: KC, are you there again?

MS. CLAFFY: Yeah, I'm here.

MR. KNAPP: Okay.

MS. CLAFFY: Are you guys there? I feel regulated. Okay, now I've lost a couple more minutes, so back to slide 8. All right. So what did change? Almost everything but the network architecture and the routing architecture changed. Right. The industry structure changed, the efficiency of bandwidth provisioning -- just miracles have happened that none of us would have imagined even in 1990.

Data processing and storage has obviously followed Moore's Law, so that's also exponentially changed. The way that we do access provisioning, the way that we do tiering, the way that we handle -- provisioning, the way that we do address provisioning is about to change. Pricing models are changing under our noses. Data access, you could argue, changed dramatically right after 1994; very little data is leases from providers for legitimate reasons -- with privacy, as low as competitive concerns.

And if data is released, it's to a select few researchers under strict NDA. Researchers are not allowed to publish anything that is not approved by the provider, for again, a lot of historical important reasons. But it's a huge problem both for network science, as well as network policy. And of course, the uses of the network have changed over the years. Although, we saw hints of it in early 1990's, we realize even now, it's only just begun what people are going to want to use the Internet for.

Okay, what did we recommend -- next slide, slide 9. What did we recommend in 1993? So I offered this hint at the beginning that I studied this 15 years ago, right, but I studied it in a world where I could pretend that some sort of cooperative solution wasn't completely fantasy, which, if you read the paper today, it would look

like complete -- but I suggested that you use existing fields of the IP packet.

Now, something that many folks might not know is that there are actually bits in the original IP packet designed by the Department of Defense funded researchers that specify quality of service. There are three bits, which means you get eight -- two to the three -- you get eight levels of service. And so I suggested again, with some other folks, a management professor and the NSF, well, why don't we just get people to volunteer to use these bits and we get ISP to volunteer to respect these bits?

Now, the only sensible thing I said in the paper was we did warn folks that even that was a hack. Even using the existing technology built into the architecture was a hack because the current architecture was living on borrowed time. If we were going to try to do things, like real-time, non- stochastic, heavily streamed video, or not even streamed, heavy bandwidth applications like video are living on borrowed

time.

Now, of course, the solution is that we proposed it never got traction nor did any other solution that was presented in the next decade. And in fact, the IETF worked on this for-- you know, a bunch of engineers, really smart engineers, worked on this for a decade. And the thing that we managed to do most effectively, I think, was convince the operators that we were just completely out of touch with reality; not just the academics, but even the engineers that tried to work on this.

Fortunately, we're starting to get it now. It's -- may be a little slow on the uptake, but the academics are starting to understand now. You heard at the research panel last week, Dave Clark and NSF talked about the future architecture efforts; obviously there's a very long-term thinking now. We don't have a deployment plan, but we are starting to recognize the importance of the interdisciplinary aspect of this.

Now, I also need to point out, although

Cisco, I'm sure, is going to get up and point it out because every time Cisco is in a room with me and I talk about the failure of QoS, Cisco jumps up and says no, no, no, no, QoS works. It works inside a given ISP; intra-domain we call that, within one administrative domain.

In fact, I totally agree; the quote "reasonable traffic management problem within a single domain with solved and deployed black decade." And Cisco is selling products to many different enterprises that manage their traffic within their networks perfectly well with current technology.

It's not a problem, intra-domain. Across ISP's, that is with -- when you have multiple AS's with different administrative operators and in particular, if they're competing with each other, it's not solved at all. There is nobody doing this. Maybe Cisco will get up and say the world is different than I understood it to be and I'm just an academic. But the last time I checked, it was not solved and it's not solved for completely non-technical reasons.

Next slide. So this is -- I'm going to skip this -- basically skip this slide, but it's the slide that I gave in February to another FCC panel on what should be the broadband conditions on the -- on BTOP money. It's -- BTOP money. So I, again, tried to drive the point home that Scott promised you I would drive home which is we know so little about critical infrastructure conditions. But I find it rather frightening that we're trying to regulate it at the moment, frankly. And I'm sure you all find it even more frightening.

And so I was just making a list up here of the things that we don't know. And again, the reason I can tell you that -- what a -- is, is because as a network scientist, as a researcher for the last 20 years, it's a huge problem for science. I mean, probably the Climate-Gate scandal has really hit close to home for me because I'm afraid that something similar is going to happen, although I don't have to worry so much about not providing data to people because we don't have too much data to worry about FOIA for.

But it is this very similar problem about openness in science and how do you justify results as legitimate if you can't show the data and you can't reproduce the work. Okay, so I'll push on that point later and I pushed on it at the beginning, the sort of need for openness in transparency, and it just sort of permeates every aspect of the industry. But I've said it a million times, so next slide.

Surprisingly, I -- slide 11 now. Surprisingly, NTIA actually put in some language -- some of you guys are acutely aware of this I think -- put in some language that tried to address some of this transparency issue. So here's the list of data that (inaudible) BTOP money are suppose to provide. I don't think any of this is is implemented. I don't -- I kind of have my doubts that any of that will be implemented; that is that we'll ever see this type of data the next few years. But it's certainly a stake in the ground and I think it's worth looking at. And they did ask for my input, although this isn't exactly what I sent them. It's not too far off, actually and I'm impressed that it got in there, whether or not it gets executed.

Next slide; okay. So what is our big problem that we need to worry about right now? I would say given the historical context of sort of the architecture for what we're trying to use it for now, including the fact that, again, we're using the cooperative architecture in a competitively provisioned environment. The biggest risk that we have is that, again, we sort of don't want a heavy handed regulation, but we also don't want to incentivise strategies that induce artificial scarcity or that promote artificial scarcity. So we've got these sort of elephants in the room problem.

We've got, again, a platform that still exhibits natural monopoly of economics, or duopoly in this case. It's sort of from an economic perspective they're called monopoly (inaudible).

So 13 years after, we swore that it wasn't going to be true in the 1990's (inaudible); it just still seems to be true. We still haven't been able to get out of the -- to have facilities-based competition be really a success story. And we've got not only the facilities' owners, but the spectrum allocations that they control turned by some decreasing number of facilities owners.

And so the third problem is that, unfortunately, despite the things -- they seem to be moving in the right direction and even more so in other countries. We -- lack of transparency prevails in the industry. Again, good historical reasons, but poses a big problem for China to regulate.

Next slide. And then, of course, from a legal perspective, this is my only legal slide because I'm not a lawyer, although I've talked to a bunch of them lately, but when we people try to talk about what legal status Internet traffic has, the bottom line is it doesn't have any legal status right now. Of course, there's ECPA which talks about privacy, but not about the equality of traffic or the rights of traffic to not be discriminated against. And anyway, ECPA has exceptions that, you know, you could argue allow the Comcast bitTorrent stuff and people have argued it. I don't like the argument, but just not for this talk. And in any event, the ECPA restrictions are not preempted by the Telecom Act.

So the significance of the political regime in which you treat traffic matters--and, of course, none of this has been resolved in courts yet, and, of course, even resolving it in court is not plan A. So that's a -- yet another sort of elephant in the room we're trying to deal with. But that's to the technical side.

What do we know from a -- slide 14 now, Walter. What do we know about QoS -- or is it 14? Yeah. What do we know about QoS technology, economics, and transparency? Well, unfortunately, again, it's going to remind you of Climate-Gate a little bit, the inability to study real networks has lead to unresolveable scientific contradictions, in particular related to this issue of QoS. What are the cost and benefits of using QoS to support tiered service?

So you've got two papers in the field: One written by some folks at Internet 2, the U.S. R&E document to support the academic community, apparently based on their own economics, although they don't provide data about their network I should point out, and there's the citation for it. And they claim that QoS is just a waste of money and they'll never use it. It's not the way they should {inaudible) their network.

So you've got another paper by, surprise, AT&T, insisting that QoS is critical. Now, that paper doesn't -- not only has -- does it have no data, but it actually doesn't even talk about AT&T's network, specifically, per se. It said some simulations of something and then it says that QoS is critical. AT&T is, surprisingly, given its level of research investment, is not actually publishing any papers on the economics of its network. Although I've seen papers lately about it, behavioral advertising of people using AT&T's network, which is kind of interesting. AT&T is fairly concerned, I would infer, about publishing information about its network.

Okay, so the bottom line, slide 15, is that scientific researchers have not solved this problem. Well, one could argue it's not a scientific problem if you get empirical grounding to your work when, in fact, the data dearth, or the data sharing dearth, is a policy problem.

Now, I'll make a side note because it does sort of--as I was thinking about this talk, it did sort of remind me of a project that DHS is funding right now to develop a framework for ethical consideration; how to determine what kind of network research, which includes a lot of traffic collections, so it really comes down to the same common denominator as traffic management.

What kind of ethical concerns should go into determining whether a network -- given a network research project is okay to do? Now for example, UCSD, when we do traffic research, we go

through our Institutional Research Board because it's considered human subjects research, although it did take us a bit of time to convince them it was human subjects research. They get it now.

The IRBs are coming along. And very similar concerns I think that when we finish with this report that we're writing, I think it will of interest to a lot of folks in the room. But, unfortunately, it's not going to happen probably until after the February deadline.

Next slide. There's a great over-honest quarterly report -- quarterly earnings report by Cogent to its investors where they basically admit that they're not just competing with other ISPs, but they're competing with the Postal Service in terms of movement of data around. And so they need their transit pricings, their bit transport prices, to fall in order to be as competitive as the Postal Service, which I thought was a brilliant illustration of our problem because the Postal Service is down by a charter that dictates profit minimization while the carriers are bound

to a -- completely the opposite charter. Another problem we haven't really dealt with head on.

Next slide; okay. So the rest of my talk -- it's just five minutes left; I'll go really quickly through these -- is case studies internationally. So I happened to work with some Japanese researchers, like I said, because they really put the U.S. to shame in terms of what they've been able to publish in the research literature on traffic characterization because they've got industry self-regulation models that we just don't have in this country, I think partially motivated by the massive ownership -- government ownership of some of their telcos.

So last year they published -- and again, a few companies got together. I don't think it's every company, every ISP in the Japanese market, but the big ones got together and essentially established guidelines for packet shaping, so to figure out what is the right balance between privacy and needs for doing network management. They specify in this document, which you can go find here, that it's not legally binding, it's only based on the industry "consensus," and again, it's not every single provider.

Next slide. Now, their big point is that -- well, they have several big points and the main ones are summarized on this slide. So this is the most important slide about this study, although I have a few more in here. Packet shaping should only be implemented in exceptional circumstances.

In response to some important network operational problem like suggestion, must be substantiated by objective data. And that -- they really put a stake in the ground there. They basically said that before you can even do traffic management, you must publish data that justifies the need for it.

Now, they admit that the terms of work for the traffic management, just like the FCC has admitted, just like other countries have admitted, you can't really do a rigorous clean algorithm for what traffic management is. Okay. So they provide case studies to help evaluate it. And one of their case studies includes content examination, looking for -- for example, looking for (inaudible) payload they determined is not being reasonable for a couple of reasons. You can't really do it accurately for a single user and you certainly cannot do it at scale for all users, given current technology.

It's also not the business that ISP's are in and I'll get to that sort of criteria that they use in the next slide or slide after next. And then they talk about there's other reasons to support traffic management, like helping users out themselves with security. For example, there are security issues with several P2P applications. But they specify that that kind of traffic management should be explicitly requiring informed consent of the user. And that's the stake in the ground that they put.

Now, next slide. The Japanese have a fairly strong notion of secrecy. They call it

secrecy of communications. I think we call it privacy here, but it includes a lot of things in the packet, including the headers and even the traffic volumes. And they consider infringement to be anything that falls under these -- this information that is used to advance one's own or affect another's interest against the parties of original communication. So that's the stake in the ground; they also include in the guidelines.

Next slide; okay. So what are the three criteria they use? And the interesting thing about these three criteria -- I really like these criteria even though that you can tell that it's translated from Japanese and it's a little clunky in terms of the words. It really matches kind of nicely the one that Canada just released last month. Anyway, there are three criteria.

Legitimacy of purpose. It must be in the nature of the business of the ISP's, so it -- for example, managing congestion.

Number two they call necessity of action, which boils down to it must be supported

by objective data.

And number three is the validity of means, which in the Canadian model they treated -- or they -- something like minimum amount of impact necessary to accomplish the goal. And here they call it method aims only at objectively demonstrating necessary.

And then they say, you know, these things only come into play if there's not user consent. If there's user consent, okay, well, then you open up a whole new set of things that you can do. But again, in Japan, that -- a competitive access environment, which we don't have here.

Okay, next slide. Okay, and then they also point out, as does the Canadian study I read, or ruling I read, that there's been quite a bit of need for further study. Slide 21, the increase in video content, everybody recognizes, is going to require more need.

I want to point out again this architecture wasn't built for that, although we're doing, again, what the academics would call hacks to the architecture. We really need to think about a fundamentally new Internet architecture. I know that's a line you've heard from NSF in the past, but there really is something behind it.

And then what are the impacts of the packet shaping on the access networks? This needs to be studied. What sort of applications specific packet shaping should be okay? Because, again, they rule it out as not okay to do applications, specific packet shaping.

And then paid-for content, which the Japanese guidelines identify as problematic given that ISP's themselves are expanding into content and necessarily so, so they sort of red flagged that one. They also recognize that even though Japan's made more progress than most countries on this, information sharing among protocol players -- among players regarding the packet shipping implementation needs to be improved.

And then they will point out that many P2P protocols are actually improving their efficiency and sort of reducing the impact on the networks themselves. But that all needs to be studied; not all of them have done that.

Okay, and then the last case study, just a second one and there's only one slide on this one, is Canada, which Harold, fortunately, let me know about a couple of days ago. And he wrote this nice blog entry and if you haven't read it, it's maybe the nicest blog entry Harold's ever written about the FCC. It basically says that Canada copied what the FCC did because it made so much sense. No, that was the FCC response to the bitTorrent Comcast thing.

But another interesting point about the Canadian decision is that not only is the -- similar to the FCC's reaction to this problem last year, but it's quite congruent with Japan's articulated guidelines from last -- from May of last year. Now they call it slightly different so their criteria are minimized harm, again, transparency of need, make sure that it's clear, and -- I didn't write it down here, but

transparency of execution.

So you also have to make it clear to the user what sort of traffic management techniques are being done and, just as the Japanese specified, make sure it's narrowly tailored. From a technical perspective, you would probably see efficient, meaning it doesn't do more than it needs to do to the traffic in order to accomplish the congestion management goals.

Now, of course, the big challenge, Scott alluded to this, too, is ultimately about defining what is reasonable. That's not a technical problem. But some of the stakes in the ground that you can see have been put, including by Canada and Japan, targeting content, not reasonable; targeting specific applications, not reasonable; not basing techniques on quantifiable data, not reasonable. So I think those are pretty good stakes in the ground. And frankly, as near as I can tell, FCC is already following them.

All right, next slide I'm going to skip. It's just a bit of terminology, taxonomization there; you can use it if you want. And then let's go to the last slide, which is, again, an executive summary on what you've said. We all get it. The reality, someone needs to pay for the infrastructure. The current economic models are not going to get us all fiber to the home in the next 10 years. Something needs to change. But in the short term, we do have a -- sort of a navigation of conflicts to manage and I think the best way to do it is continue on the same path that we've pushed, which is transparency and objective data obligations.

As Andrew is fond of saying, this quote below is from Andrew, it's a repeat, to evaluate claims about the need for additional revenues, we need solid cost data, which we don't have, and a dynamic model of the industry, which we don't have.

Academics is working on it, but, you know, we also need some additional work on the data sharing models. Some -- the nearby case studies I suggested, Japanese and Canadian, suggested the FCC is asking the right questions and is still waiting from the facts base. So when I reread the NPRN this weekend, the facts base is all over that thing, so you guys are really interested in facts. Suggestions and explanations for pricing changes.

And I should argue I really -- this really hits close to home because my own discipline, Internet research, is in the same situation as QoS is really in and that Climate Science is really in, in that it could meet its demise, if marketed in the absence of legitimating independent objectives for review. So the last slide is just a set of references that I used in making the talk and, hopefully, I can hear some questions now.

MR. KNAPP: Thank you Professor Claffy. Questions. Go ahead, Walter.

MR. JOHNSTON: KC, you mentioned Internet 2 had not supported quality of service in the Internet. About three years ago, there wasn't a summit between the ITU and the IETF over the convergence, I'll say; it covered many issues. But one of the biggest points of departure between their two perspectives was on a need for quality of service with the IETF arguing that the -- it was all about best effort and the ITU arguing that inter- domain quality of service was essential. How do we go forward with these two important stakeholders disagreeing?

MS. CLAFFY: Yeah, that's a great question. By the way, all of the questions -- best questions I've ever heard from the FCC. You guys are on top of this stuff. Unfortunately, the ITU versus IETF is, again, in technical terms, really a political economic argument. Of course, ITU is going to say that inter- domain QoS is essential, just like when I presented in 1992 to the -- to AT&T research.

Everything I knew about the Internet, which didn't take long, I was a graduate student, it took about an hour, they looked back at me and said, but how do we manage traffic on this? And I said, well, that's really not part of what the

Internet does, the Internet architecture itself.

So when the IETF says, well, the Internet's best effort, you don't get inter-domain QoS. They mean that from an architectural perspective. That's what the network architecture can support inherently. When the ITU says inter-domain QoS is essential, they mean that from a business perspective or from a user needs perspective.

They don't mean that from an architectural perspective and, in fact, the ITU had no solution then and has no solution now for having implement QoS inter-domain across many, you know, as Scott was trying to say -- as Scott was saying, end-to-end QoS. Neither the ITU nor the IETF has a technical solution to support inter-domain QoS that has survived in the economic environment.

Now, the IETF did try for 10 years. The IETF hasn't ignored this problem. They worked on (inaudible), they worked on (inaudible), they spent the last decade on it and they provided technology, just like the original 3-bit and the IP header has technology to support inter -- even inter-domain QoS if you want.

That's -- it's not a technical problem; it's a -- it's an economic problem, it's a market problem. It's a -- it's also an enforcement problem, that is, how does MCI -- or how does Verizon know that Sprint, in fact, respected the QoS that was specified in the packet when it was sent? So these guys need to enforce it on each other, never mind what the FCC would need to do, and they don't have a mechanism for doing that right now.

So yes, that biggest departure, it's not just between the ITU and the IETF, it's also between the IETF and itself, frankly, because we all recognize inter-domain QoS is "essential" and yet not supported by the current architecture. And by that, I mean not that it's a technical architecture, but the politically economic architecture in which it is placed. Did that help?

MR. JOHNSTON: Yes, thank you very much.

MR. KNAPP: Well, I -- yeah. Now I've got some hands going up. Go ahead; Jon go first and then Stagg.

MR. PEHA: Hi, KC, it's Jon Peha. I want to -- you've called for more information and I'd like to drill down on that a little bit and see if you can help us prioritize. Imagine that I have a specific goal, which is not to learn everything about the Internet, but to help make an intelligent decision on the open Internet proceeding. And let's say that the network service providers wake up tomorrow morning and decide that they're anxious to share some kinds of data with us to help us make that decision. What do you think are the, you know, one or two most important kinds of data that we should be hoping for to help us make this decision?

MS. CLAFFY: Well, again, I think that the -- we should let the burden be on the ISP to provide data that will justify the need for whatever pricing decision they think they need, right. So this classic example is when Comcast backed off of the traffic management, the sort of opaque traffic management they were doing with bitTorrent last year, and reverted to doing metering. And then they promised that they would provide the end user with a tool so that they would know when their own personal traffic loads -- levels were getting near the quota that they were set, 250 (inaudible), I believe. And then they never made the tool available.

So they implemented this pricing mode that really required transparency into the network on both the provider side and the user side in order to rationally respond to the pricing model. And then they never made the tool available.

Now they have released -- I'm notified they have released this tool in -- last week, I think in Portland for testing, although their webpage said there's not a -- contradicting information on their webpage. First they said there's not a clear determination of what they're going to do; it'll depend on the trials in

Portland. And I had not -- I mean, I have not seen anybody using this tool. I don't have any connections in Portland, if anybody has been able to use this.

But they -- I've also seen they say they're going to release it the first quarter of next year to other folks. So certainly you could argue that traffic data and cost data are the main types of data that you're going to need to determine reasonable traffic management approaches. But I wouldn't argue that a blanket request for data on service providers is going to help as much as putting the burden on the service provider to demonstrate objectively whatever the traffic management proposal that they have in mind and why it's needed. Because the FCC is going to fight a losing battle if you're going to try to say, okay, every provider needs to provide just some kind of data so that we come up with some aggregate sort of recipe for what sort of traffic management is okay.

Keep the burden on the people that run

the networks. They understand the networks; they understand the provisioning problems. They just have strong counter incentives to making any of that information public. We need to help them with that.

We need to give them ways of making that information -- of sharing that information that protects the privacy of individuals that reasonably protects their commercial interests and yet that enables the customer, as well as the regulator, as well as everybody else, potential consumers of that service, to understand what they're doing and why. Does that help?

MR. KNAPP: That was helpful. Stagg has -- this is our last question I think.

MR. NEWMAN: Okay, a question and a homework to go with it then. KC, you argued that -- you had some great slides about network -- or about architecture, not just network architecture.

So the -- I'm told now that if you look at the average user accessing the Internet from their home, much of that content is not going over the Internet, as sort of Scott depicted it, the layer of hierarchical Internet, but much of it is bypassing that and going directly from the content provider to the Internet access provider by Akamai or Google Connections or what have you.

Would you say that traffic is still using the conventional network architecture or because the middle of that is no longer the Internet, it's actually going over a different architecture?

MS. CLAFFY: That's also a great question. I would argue -- and now I'm not aware of Akamai building out like fiber infrastructure. I know Google is doing that, but I thought Akamai was still mostly about putting sort of caching infrastructure, as Scott was saying, that it's PCs that can hold files, that can hold data, so that it's closer to you and you don't need to go across the Internet to get it.

But it's perfectly reasonable for them to sort of expand in the direction of building out their own backbone infrastructure because everybody who can afford it has done it, right. Microsoft has done it; Yahoo has done it. Anybody who sells content and has to deal with the peering pricing -- opaque peering pricing policies of providers is going to build up their own network and just write it off as the cost of doing business.

I think we do need to consider that still the Internet. If it's using TCP/IP and, you know, I think even Google has talked about a more efficient version of TCP/IP within Google's own AS, but I think it's just talking about peering that happens over TCP/IP, you still ought to consider that the public Internet. I wouldn't consider that bypassing.

It's certainly bypassing, say, the providers where there seems to be -- or is more likely to be congestion, but that's not just an issue with Google and Akamai. Anybody is going to try to bypass their peering connections or modify their peering relationship such that they improve their performance. I mean, that's the name of the game, really, is manipulating your peering relationships to optimize your performance and your cost. Does that help?

MR. NEWMAN: Yeah, it leaves about three more questions, but let me stop and just ask if you can help us. Your slide, I think it was slide 5, you listed about six different architectures: The naming architecture, the network architecture, transport architecture, routing architecture, et cetera. Could you either send us or point us to expand the definitions of each of those architectures and good illustrations? Because I think we really need that deeper level of understanding when we have to go from principles to, you know, embodying that in real policy.

MS. CLAFFY: Sure, absolutely. And I can try to write this up as more of an essay, these slides, and I'll definitely put more details into that. Thank you very much for having me.

MR. KNAPP: I should add, too, for both you and all of our speakers, Stagg said three more questions, but I'm sure there will be a lot more.

So we'll be in touch.

MS. CLAFFY: Yes, please send e-mail. I'm happy to answer questions. I think Stagg had a question that Scott promised I would answer and I didn't catch the question, so send that one in an e-mail, too, and I'll answer it.

MR. KNAPP: Yeah, I think it was more, and I had the same one, too, similar to what Jon asked.

MS. CLAFFY: Oh, okay.

MR. KNAPP: Yeah.

MS. CLAFFY: What data would you want? What's your --

MR. KNAPP: Yeah, what data would you -- what would you be looking for? What would you do with it, without getting down into all of the nitty-gritty details of how the networks are -- every packet going through --

MS. CLAFFY: Yeah, I would encourage you to look again at the NTIA list. There's a couple on the NTIA list which drives this point home in terms of -- each -- peak and average utilization of access links, which I think is where the claims are that the congestion is happening. But, you know, I better answer this question after I listen to the four industry speakers this afternoon. Are we going to have a panel this afternoon?

MR. KNAPP: More presentations, but today is only the start. There will be much more dialogue as we go ahead.

MS. CLAFFY: Okay.

MR. KNAPP: Terrific, thank you. Our next speaker will be Paul Sanchirico, who is vice president of Engineering at Cisco. He's vice president of the Service Provider Systems Unit for Cisco, and in this position he leads the organization responsible for developing Cisco's service provider's solutions and architectures.

The organization's work spans, IP NGN systems, manage services systems, mobility, WiMAX systems, cable systems, and video IP TV systems. I'm not sure there's anything left out.

Before Cisco, he spent many years with Bell Corp where he was the position -- managing director of Bell Corp's Business Broadband Consulting and Engineering practice. And as soon as you are ready to roll, we're all set.

(Pause)

MR. SANCHIRICO: Good morning. It is a privilege to have the opportunity to talk with you today. Thanks very much for that chance. I'm here representing Cisco Systems. We're the worldwide leader in networking and we just recently celebrated 25 years in existence and in many ways we have grown up as the Internet and IP networking has grown up.

Today what I would like to share with you is to talk a bit about innovation, talk about congestion, where it happens in the network, what we see as the customer and service provider expectations that are placed on us as an infrastructure supplier, and talk about the basics of network management and then summarize.

So the Internet and the IP network has been a tremendous platform for innovation since its inception. It started with e-mail and bulletin boards and since then has expanded to web browsing, the carriage of voice, the carriage of music, gaming, file sharing, video, IP TV, and it's been able to deliver these services and applications over an increasingly diverse set of end user devices; whether or not it's a smart phone, it's a PC, it's a television, right.

And increasingly, people are getting a variety of applications on that same device, right. So no longer is there one application, one device; there's a tremendous amount of innovation that has happened in terms of the services and applications that are supported on this IP network and in terms of the devices that people receive those services and applications.

Those new services and applications, though, have placed different and unique demands on the networking infrastructure, right. And so it's important to characterize how those -- what those new demands are. So for example, when you put voice on this network, one of the key requirements that voice places is that the one way delay should be less than 150 milliseconds, right, and that number is arrived at through mean opinion scores on what's the satisfactory level of quality, right.

That's a very different kind of delay objective then what we started out with on the Internet in terms of the transmission of files or the transmission of e-mail. We're also seeing the transmission of broadcast video on top of IP networks. And so this stresses the network in unique ways as well.

An obvious one is the amount of bandwidth that's required in order to deliver the video signal. And it can vary from 2 megabits to 15 megabits, depending on whether or not it's a standard definition signal or it's a high definition signal; whether or not it's encoded in MPEG 2 or it's encoded in MPEG 4.

But equally important is the fact that the packet loss requirements are far more stringent than you see in other services. Right, the rule of thumb in the industry is that you should have no more than one dropped packet for a two-hour movie. Why? Because a dropped packet translates to macroblocking on your television set.

Let me calibrate this for you. The DSL forum previously specified that you were allowed to have one dropped packet every two minutes before that DSL line retrained to a lower speed. That's satisfactory for web surfing, but clearly not satisfactory for the delivery of video. None of us would tolerate a video service where there was macroblocking on our television set every two minutes.

All right. And then what we're seeing is telepresence, right, a real revolution in how people are communicating, eliminating distance. Right, and this stresses the network in a variety of ways as well. Not only does it deliver -- require more bandwidth, but it requires very stringent delay requirements and it's symmetric, right.

So as we see an explosion in innovation

and the services and applications that are delivered on this IP network, we're seeing a tremendous growth in the demands that are being placed on this IP network. All right.

So one of the topics that has come up this morning and has come up in previous discussions is this notion of managed services. And so I thought I would start with a definition of what a managed service is, right. And so I offer two definitions of a managed service. One is kind of a classic business service definition where the service provider manages, that is provisions, monitors, and fixes business communication services, including devices on the customer premises.

Another definition is where the customer pays for a particular level of service, sometimes implied by the application or service purchased. Best effort service is an unmanaged service. All right. Application services like voice or video generally have an implied set of transport service requirements. For example, as we just went through, voice services have a latency requirement that has to be met or the service quality is unsatisfactory; right. So to offer a voice service, there's an implied level of transport service that's delivered.

Now the environment that we live in is an incredibly dynamic environment, right. What we do at Cisco, we produce something called the Cisco Visual Networking Index, and this data and how we collect it is available to you all and the public on our website. But basically, what our studies are showing is that there's going to be a cumulative annual growth rate of traffic on the IP network of 46 percent between 2007 to 2012, right.

So you're going to get to the point where in 2012, about half a zettabyte, or 1 trillion gigabytes, of traffic will cross the global network. Now, that's not to say that there's a small amount crossing today because there's not. There's about 15 exabytes, right, which is about a billion gigabytes, right. So there's a lot of traffic that's crossing the network.

And so there's a -- if you look at this, this says, well, that means that on average I'm doubling the capacity of the network every two years, right, to give you a sense of it. Right, so this is a very dynamic environment. Now this is -- the other thing that you have to look at is kind of the volatility of that traffic.

Right. One of the things we're also witnessing in the studies that we're doing is there's an increasing amount of volatility. The peak-to-average ratios are growing, right, on the order of, you know, 5 to 7 percent a year. So this is a very dynamic environment. And at one level, right, this tremendous growth that we're seeing is a proof point positive that there's a tremendous amount of innovation happening on top of the Internet today.

Now let me -- the position that with some data from our development organizations, right. Advances is ASIC technologies have really enabled significant increases in density of the chips and, therefore, have lead to improvements in costs, right. So for the routing engines that we produce, right, what we're seeing is a 23 percent cumulative average price decline per year to carry a gigabit of traffic. Right.

Now it turns out that that's not the entire cost that a service provider incurs in carrying traffic. A significant part of it has to do with the optics and we're not seeing the same drop in price in the optics world. The complexity of long haul, DWDM optics at 40 gig and 100 gig is really seeing that those prices are staying pretty flat as far as we can tell. Right.

So juxtaposition this 23 percent cumulative annual price drop with the 46 percent cumulative annual growth rate that we're seeing in traffic and you see that -- our customers, the service providers, are running into a challenge. Now, the other point about the traffic characterization to note is that about 10 percent of the Internet users make up more than 60 percent of the traffic, right. So not everybody is using the same amounts of traffic all the time, right. And so that becomes important as we talk later on about traffic network management.

So congestion happens. Where does congestion happen in the network? Here's a diagram of the network. There are customer premise on the left-hand side, there's an access network, and aggregation and -- network. There's a service delivery center, a data center out of which people might access web pages or get access to services, maybe that's a video head end, and then there are peering service provider networks.

Congestion happens in a number of different places in this network. One of the places it happens is because of this device source sync mismatch, right. So I've got a smart phone that's downloading a web page from a very fast server, right, there's a device mismatch there and there's going to be congestion caused by that device mismatch. Another place in which congestion occurs is that between the premises network, where you might be running a wireless network at 54 megabits per second and it connects to a DSL line, right, or a HFC plant, right, at far slower speeds. Right. So that's another point in which congestion might occur.

Another place is that the aggregation network is typically over subscribed, right. And by that I mean that the aggregate bandwidth of all of the pipes going to the premise environment is greater than that going into that network, right. That's because that access link is dedicated to specific subscribers, right, it becomes shared once you get into the network. And so there's an oversubscription that happens there.

You get focused overloads, right, those of us who have spent time in New Jersey know that if Bruce Springsteen concerts go on sale at noon, right, there's going to be a focused overload to the ticket site. Right. So there are focused overloads, right.

And then there are inter-carrier linked bandwidth differences, right. So all of these cause congestion in the network. The other aspect

of this is that remember the traffic is very bursty, right, and so it makes a great deal of complexity, right, in how it is that you operate these networks in order to meet all of those services and applications in the unique and new demands that they're placing on the infrastructure.

So customer expectations as we understand them. This is taken from a filing that we signed onto along with a number of other companies to the European Regulatory Commission. Basically, customer expectations are that there's continuous innovation; that their Internet mailboxes aren't going to be overstuffed with spam; that a user's online experiences aren't going to be unfairly degraded or disrupted due to congestion from the online activities of their neighbor; that the services they pay for are not going to be susceptible to a variety of online attacks or other threats from malware of various kinds; and that the market will continue to offer new and differentiated services and content.

Customers also expect choice. They want meaningful information about the nature of the broadband services and the practices of their providers. All right. So with that as backdrop, as an infrastructure supplier, right, what demands are placed on us is that we are asked to enable scale with a lower total cost of ownership; we're asked to enable the delivery of a fundamentally better user experience. In our mind, that establishes the need for network management because the alternative is impaired access to content services or a rise in consumer costs.

So now we're at the stage where I'll talk about the basics of network management. So network management is used for a variety of purposes. It's used to ensure the quality of service is maintained as demands skyrocket, right, the current environment that we exist in; it's used to protect against bad actors or bandwidth hogs; it's used to prioritize network management traffic such as routing tables, right, if there's a problem in the network that network management traffic has to get through, such that the right corrective action can be taken.

You need to maintain a network security to contain the proliferation of spam, spyware, worms, or other malware. You want to make sure that there are parental controls on content, you want to hamper unlawful dissemination of intellectual property, and you want to be able to enable new services. So those are the purposes of network management.

There are a variety of tools that are -- that we have enabled our customers with. I've listed some of them here. I'll not have time to go through all of them, but quality of service treatment, traffic shaping. Typically traffic shaping is applied for business customers rather than residential customers. Virtual private networks, access control -- deep packet inspection, content delivery networks, and visual quality of experience technologies.

I'm going to focus on the quality of service treatment and the deep packet inspection today. At another time, happy to go through more of the tools here in more depth for you if you so desire.

So enabling quality of service; so step one is to really provision the network to accurately calculate the required bandwidth for all applications plus their overhead, all right. Step two is to classify or mark the packets with a specific priority denoting the requirement for the class of service from the network. Is this voice traffic? Is it video traffic? Does it require low latency or not?

And then also to define and enforce a trust boundary, right. So typically that's the edge of the -- part of the network that the provider controls, right, where they can trust that the markings are true, right, that, yes, this is a voice packet and deserves a higher quality of service because of it's application requirements.

All right. And also to then schedule, assign the packets to one of multiple queues based on classification for expedited treatment

throughout the network, right, and then use congestion avoidance for data. So those are the steps that are typically taken with quality of service.

I also want to point out that the tools that we -- that in the network the congestion doesn't happen on fiber links. Congestion happens in the routers, right, and the queues where the congestion is managed, right. And so we provide a variety of tools to our customers just so they can manage that congestion, right.

So if we have a particular link capacity, we might -- and we also give -- the service provider uses those tools as they see fit, given the unique circumstances that they face in their network, the type of traffic they're carrying, the volume of customers that they have. But we also give it some design guidance and so I've just picked a generalized kind of bandwidth scheduling design principle here to illustrate some points.

Individual carriers for individual

circumstances might tune this a bit differently. But taking a look at the voice traffic or video traffic, what you might do is decide, well, those require low latency queuing, right, because they have specific delay requirements.

And so you might decide that I'm going to allocate a third of the link capacity to those low latency queues, right, to those applications. Then what I might decide is for the rest of the traffic, what I will do is I will allocate say maybe 90 percent and I'll -- for the rest of the data for that traffic and then the remaining 10 percent I'll reserve for routing and another network management overhead.

Now, in the traffic, what we used to do 10 years ago or so, is that we applied priority queuing, right, which basically said I'm going to service all of the packets that are marked with the top priority first. And when I'm done, then I'll get to the second priority and when I'm done with that, I'll get to the third priority.

What we found was that we were starving

out, right, entire classes of applications. So we haven't used that for 10 years, right. Instead, what we use is something called weighted fair queuing or class based weighted fair queuing, which basically says I'm going to allocate a specific portion of the bandwidth for particular classes of traffic, right, and then I will -- then that context, fairly allocate the bandwidth among the traffic that's coming in, that's marked in that way.

All right. So in this example I've taken my voice and video traffic, I've put them in the low latency queues because those applications require that, for the rest of the traffic, I put it in a class based weighted fair queuing algorithm and I managed the scarce capacity on that bandwidth. And I reserve a particular amount of the bandwidth for network management traffic.

Now, the question has been raised, well, how much money does quality of service really save you? Well, if you don't have quality of service, what you have to do is fundamentally design the

network to the level of service required by the most stringent application that you want to support.

So -- but oftentimes, that is a small percentage of the traffic, right, voice or high priority data, for example. And so what we -- what this chart shows you is that some of the studies that we've done and have seen demonstrate that I can get at least a two times bandwidth gain by using quality of service over networks without quality of service, right. So it's a real economic benefit for using these network management tools and meeting the quality of service requirements demanded by the innovative set of applications that people are putting on these networks today.

Now -- oops. The queuing algorithms that we've talked about so far manage the front of the queue; that is which packets get transmitted first. We also have an issue of congestion avoidance that we have to deal with, right, and that is really about how you manage the tail of the queue and that is which packets get dropped first when the queue buffer overflows.

Right. And so the tail drop method says the queue fills up, I drop the next packet that comes through. But what we found was that that didn't work so well because what happened was we'd get the network into an unstable state. Because what would happen is you'd get a bunch of TCP traffic flows starting at different times, but then the network capacity would be reached and they would all be dropped and they would start again.

So you end up with this synchronization thing that happens and you end up with waves of traffic coming through the network and destabilizing the network. All right. And so instead what we implemented was something called weighted fair queuing -- I mean, weighted random early detection, which basically says I'm going to drop packets before my queue is completely full, right. When it's approaching getting full, I start to drop packets and I drop the lowest -- I can do this in a class-based way or not, right, I might drop the lowest priority packets first or I might just start dropping them randomly.

And in that way, I avoid this TCP windowing synchronization problem and avoid the instability in the networks; all right. And so these are some of the things that we have developed over our 25-year history, spending billions of dollars in research and development in order to deal with the tremendous growth that we've seen in the Internet and the tremendous diversity and innovation that we've seen in the services and applications that are supporting it. And we've developed them because we have to because if we don't the competition will, right.

Now, another tool that we make available is something called a deep packet inspection technology. This is really nothing more than providing a more granular view of traffic data than I can get currently in the routers, right, and this line between what's in deep packet inspection and what's in the routers blurs, right, but -- over time.

So an example is I might use it to get visibility about which subscribers are infected with a virus and which ones aren't. I might use it to get visibility about the average bandwidth per subscriber and look at what are the services that they're running over that link. I might use it to get average subscriber bandwidth for one service or look at the total number of subscribers.

This technology is really about understanding the type of traffic, not really understanding the content of the traffic, right. And so it's a tool that's made available to users.

Another example of how this tool might be used is maybe I want to take a look at file sharing, right, and so I see that, okay, this is the amount of file sharing traffic on my network, maybe I want to break it down; maybe I want to take a look at what kinds of file sharing traffic I have available.

Right. So it gives the service provider

more visibility into what's actually transmitting over their network. And so the -- and inspection engine really enables you to take a look at specific flows, specific service, specific subscribers, take a look at some global controls, policies, and really it's about how do I ensure that bandwidth is distributed fairly between network flows according to their assigned subscriber service and global controllers policy. You know, you could do this at a TCP level.

The problem is that sometimes you end up with bad actors which might open multiple TCP sessions, right, and there -- and that way consume more of the scarce network resource than they are allocated, right. And which is not a problem if there's not congestion on the network, right. None of these tools are applied if the network is lightly loaded. These tools are applied only when there are congestion situations, whether or not it's weighted fair queuing or whether or not it's weighted. Right.

And so this is a tool to ensure that

user's online experiences aren't being unfairly degraded due -- or disrupted due to congestion from the online activities of others. So that's a bit about the DPI technology.

So just to summarize, all right. So the Internet and the IP network has been a tremendous platform for innovation, right. I think that we've seen that in the tremendous growth in the services and applications that are running on it and the proof positive that that's being successful as the tremendous growth in the traffic that we're seeing.

All right, congestion happens. Congestion happens in the network because of speed mismatches between the home network and the access network, between inter-carrier links because of focused overloads, because of over subscription, right, device source synch, mismatches, variety of reasons that congestion happens.

All right. The customers' expectations are really driving the need for network management, right. They want a high quality of service and they want it cost effectively. And there are many network -- intelligent network management tools that the market place has driven us to develop and make available to our customers. So that summarizes the presentation that I had prepared.

Thank you very much for your attention. What questions do you have?

MR. KNAPP: Thank you. I suspect we're going to have more questions than we can deal with in a handful of minutes. Have you got any sense, and I know this is going to vary from one provider to the next, of how often -- there's a thirst for data relative to congestion and how often it happens, and how often tools like this might get applied. Do you have any sense, as a provider of the equipment, what are we talking about? Ten percent of the time, 5 percent of the time? Any way to characterize this?

MR. SANCHIRICO: I can't sell a router without quality of service capabilities, right. I think that, you know, the ability to provide quality of service is an essential element and I -- it's -- so I don't know, right. All I know is that for me and my products, I can't sell it without this capability. Right. It's that important.

MR. KNAPP: I assume, though, you get feedback from your customers about their requirement. I mean, why --

MR. SANCHIRICO: You know, my impression is that because of the bursting nature of this traffic, right, it happens all of the time, right. So for example, when I log in to work from my home, I have a wireless network at home; it runs at 54 megabits per second, right. I have a DSL line, right, and that DSL line is slower than 54 megabits per second. There's congestion that's happening right there. Right. Every single time I access my website, right, it has to wait. So I think it happens on a regular basis. I don't think it's an exception condition.

MR. KNAPP: Okay, go ahead.

MR. NEWMAN: Paul, if you look at the

combination of things that you're selling to the service providers, caching, depacking, queue management, et cetera, et cetera, if none of those were available, how much more capacity would they need in the network or, you know, how much are you saving of the total amount of capacity? Realizing that that's a little bit of a flaky way of thinking about it.

MR. SANCHIRICO: Well, you know, as the chart showed in the presentation, at least two times more capacity is what I would need, right. But that doesn't account for --

MR. NEWMAN: That was queuing, right, and then you add up caching and DPI and all of those?

MR. SANCHIRICO: I don't know. So that -- I think that would establish the floor, right.

MR. JOHNSTON: You know, you said you can't sell a router without quality of service. We've heard early discussion about the need for, or the growing need for inter-domain quality of service. One of the things the Internet's never been great on is getting capabilities implemented across domains. How do we approach the problem of a multiserver environment where quality of service is desired and how do we achieve some rationalization of what that means?

MR. SANCHIRICO: I think that's a policy and business question. Cisco's providing the tools that allow that to happen.

MR. JOHNSTON: We're open to suggestions.

MR. KNAPP: I have a question from the audience. I'm going to read it exactly the way it was presented. It says, "You list one of the goals of network management is to control bandwidth hogs. Who decided that? Is that inherent in the design or creation of network management tools? And should other mechanisms for addressing bandwidth hogs be preferred to use of network management?"

MR. SANCHIRICO: So the fundamental issue is, is my behavior affecting your experience? Right. And you know, I purchase a particular amount of service from you and during times that are lightly loaded I might -- you might let me use that. Right, but when it's congested, I want to make sure that I'm preserving your experience and Walter's experience, and Stagg's experience, and Bill's experience, and Paul's experience. Right.

So there's a greater good, right, that's at play here, right, is to make sure that one person that's using more than their fair share of the available resources, doesn't adversely affect what the experiences of the broader set of people.

MR. KNAPP: Let Jon have a shot. Stagg, you'll be on deck.

MR. PEHA: Weighted fair queuing assumes the existence of multiple classes of traffic. Can you say more about how many classes there are in a typical system? What -- what are -- what differentiates the different classes that we would see in a Cisco switch?

MR. SANCHIRICO: Yeah. So there's a balance here, right, between -- the technology can

enable a large number of classes, but you run into operational complexity in -- from a service provider's point of view and how many different classes you want to manage. Right. So, you know, typically we'll see four, right.

And so there might be, you know, low latency queues for voice and for video and maybe kind of a best effort and a higher tier service, right, maybe for business customers, right, for which there might be an SLA contract. All right. The service providers can give you a -- probably a more detailed answer on that. Right. The technology is pretty sophisticated, but you quickly run into operational issues.

MR. NEWMAN: Paul, you made the point that voice over IP video services of a certain nature streaming video, telepresence, need to be offered over -- as a managed service or over managed service. So let me be more specific.

Would you say that many, if not most, of the applications that I as a user are going to be accessing over my Internet access pipe, from my

Internet access service provider, that first provider, will in the future require that provider to be providing managed services to those applications?

MR. SANCHIRICO: I don't -- so clearly there's an important set of services and applications that will -- that imply a higher level of transport service than best effort. I don't know if it's going to be most or more, right. You know, it depends on how you define that, right. I mean, if I --

MR. NEWMAN: I guess most or more is not the important point. The point is there are a reasonable number of applications that will require that first access provider to be doing those over managed services.

MR. SANCHIRICO: Yes, absolutely. They're a reasonable number and they're an important set.

MR. KNAPP: I've got another question that came in over the net. "Nice to hear technology presentation. Existing inter-domain quality of service signaling exists, but networks do not use them. Who's asking for them?"

And I'll add to the first part. I mean, we had a statement earlier that it's not being used. Do you have any sense for the extent that QoS is being applied today? And then the second part of this, who's asking?

MR. SANCHIRICO: So you're asking do I have any sense as to what extent the quality of service is applied between networks?

MR. KNAPP: I think that's -- yeah.

MR. SANCHIRICO: Yeah. I don't know. Right. I know the tools are there. (inaudible) can give you a better answer than that.

MR. KNAPP: Okay. Any other questions? Well, we're pretty close to being on schedule, which is very important just before lunch.

Paul, thank you very much. I really appreciate your presentation and thanks for dealing with some tough questions.

MR. SANCHIRICO: Sure.

MR. KNAPP: And I'm sure we'll be

following up. Thank you, Paul.

(Whereupon, at 12:20 p.m., a luncheon recess was taken.)

A F T E R N O O N S E S S I O N

(1:32 p.m.)

MR. KNAPP: Everybody take your seats, please. I hope everybody had a good and relaxing lunch.

We've got a lot ahead of us this afternoon, four more presentations. The first will be by Dr. Paul Liao. He is the president and CEO of Cable Television Labs. In this capacity he heads the research and development organization that's responsible for chartering the cable television industry's course in technology.

Prior to joining CableLabs, he held the positions of vice president and chief technology officer for Panasonic and was the president of Panasonic's R&D Company of America. In these positions he was responsible for technology direction and R&D for Panasonic in North America.

And I'll just mention a couple of other things here. It's a very distinguished career, noted with many honors. He is the recipient of the IEEE Millennium Medal and a past president of the IEEE Lasers and Electro Optic Society. He's a fellow of the IEEE, the American Physical Society, and the Optical Society of America, and was part of the Leadership New Jersey Class of 1989. And above all that, he is a full member of our Technological Advisory Council.

So, Paul, welcome, good to see you again. Thank you.

DR. LIAO: Thank you, Juli. And thank you for inviting me to come here and talk about this. So when I talked to Walter, he said that I should talk a little about how cable networks operate and plan their networks. And so I'm going to try and do that.

These are the three concepts that I'd like you to remember as a result of this discussion. The first one, and really the bottom line, starting at the bottom there, is really the cable industry was really blessed with a network that it can easily evolve to revolutionize the customer experience. That is, they've been lucky enough to have developed a network over time that

is easily upgradeable for them to really enable them to take advantage of this new broadband era. Together with suppliers like Cisco and others, this has been made possible.

Now the three concepts that I'm going to emphasize here are first, that the planning and operation of their network is really in some sense all about spectrum management, which is a concept that the FCC is very, very familiar with. Fundamentally, going back to the very beginning, the idea was that they would take the spectrum of free broadcast and capture it and put it through a wire, a protected wire. But as a result they've had to manage that spectrum within that wire. And things like spectrum reuse is a very important part of their planning.

Digital compression and transmission, that's what it's all been about, everything from digital TV on up. And that means multiple services operate on the network that they built over these years. So things like multiplexing, statistical multiplexing, those all go into that.

And because the market is really exploding and the services, the number of services and the kinds of services are exploding, constant investment is required because they have not only this tremendous market demand, but also competition. And this competition requires them to have evermore innovative solutions to figure out how to use that spectrum evermore effectively.

And one of the advantages they have in terms of managing spectrum compared with, say, the FCC, is that they can do that because it's their plant. And for the FCC, you have to go through all these hearings and all the rest and have these workshops. So they don't necessarily have to do that.

So first a little bit of history. Cable TV was launched simply simultaneously in Arkansas, Oregon, and Pennsylvania. And probably the guy that first came up with cable in cable TV -- because in some of these initial launches they used the twin lead antenna wire. It was CATV, community antenna television. And it was basically launched because somebody wanted to watch television and the nearest television station was hundreds of miles away, or, in fact, in the case of Oregon, where a guy named Leroy Parsons established the first system, his wife wanted to watch television. So he put up an antenna because he lived outside of -- 125 miles away from Seattle. Seattle had a TV station, so he put up an antenna and ran coaxial cable and amplifiers. So he was the first to use actually coaxial cable and amplifiers, the two fundamental concepts of that time, to bring a signal 125 miles to his home. And then, of course, he let his neighbors do it, and then it became a business, right. 1948, similar things happened in Arkansas and Pennsylvania.

And this develops over time, but the real big next big event was distribution by satellite to the individual cable systems. And this happened around 1972, and when HBO was launched. So this meant it was no longer just getting the over-the-air signals, but also what we now call cable programming. And because it was satellite they could distribute it across the whole country, essentially, and bring it to all the cable systems.

Around the end of the '70s, early parts of the '90sXXXSICXXX, addressable set-top boxes were launched. This was very, very important. This was a very important innovation because that allowed cable customers to have a wide, wide, wide range of commercial premium cable services.

Then in 1991, the first HFC, hybrid fiber coax network, was launched. Before that there were the efforts -- full service network and so on sort of trials that went on. Digital set-top boxes were finally launched in about 1995 or so, and then HDTV. Toward the end of the '90s, we saw DOCSIS. And you can see that things are accelerating there. DOCSIS, the first digital transmission of data over coax. DOCSIS 1.0. By using DOCSIS, because you could now transmit data over coax, you had packet cable, which is digital voice. DOCSIS 2.0, and now DOCSIS 3.0, which can

allow hundreds of megabits per second communications to homes.

So we had a period that was from in the '90s, where there was a tremendous explosion of, as you can see, of the '80s and '90s, a tremendous explosion, which was really the wiring of America. And that was a wiring that was based on coaxial cable and amplifiers and analog network. And then when you got into the '90s and up until now, the cable companies invested something like $150 billion since around 1996 to transform that into a hybrid fiber network.

The important thing here is that this sort of structure was created. You have a fiber ring connecting a master headend, which might have, you know, 50- or 100,000 subscribers. But then out the hubs. From the hub, you have multiple fibers going to nodes. Okay? And you can see how this thing gets built out. You might have a single laser here, initially, feeding these multiple nodes. But then you might split it and get multiple lasers to fit multiple nodes.

Because there are multiple fibers in here you can split it again.

And I'll talk in a little bit more detail about that. But fundamentally, going from the analog to this world by going from purely coax to a hybrid fiber coax you could begin to have an explosion of services. So over this single cable you had analog TV, digital TV, switched digital video, video on-demand, circuit switched telephony -- which would be later replaced by packet switched VoIP, though some systems are still running circuit switched telephony -- broadband Internet access, and as well as various business, voice, and data services.

Now, the other thing I'd like to really emphasize is this system, these systems, are not -- there's no sort of typical system. There are now over 1,000 cable companies, and, like, tens of thousands of cable systems. So even Comcast has many, many, many, many systems, and a lot of them are different. So the architecture itself, for example, going from a hub, this might serve 2,000.

And those might be 50 to 100 homes, to 200 homes. There's no simple way to describe this. And so you need to understand how that thing is working.

So here's the way the spectrum is managed. All right. A typical cable system will go out to 750 megahertz. There are some cable systems that go out to a gigahertz, depending upon their local situation. And they will divide it up. Because television was originally 6 megahertz channels, this is typically divided up into these 6 megahertz channel slots. So from 50 to 750 megahertz, you have these 6 megahertz channel slots.

Above this point, it's digital. Below this point, it's analog. Channel 3 is here, all right. So there are lots of restrictions on which channels they can turn to digital and which channels will have to remain analog because of these are the analog channels and in some communities they're required to carry those channels as analog.

Below 50 is where the return path is.

And the return path is channelized in a different way, but up to 6.4 megahertz, as an example.

Now, this portion of the spectrum is all broadcast. So the same signal goes to every home on the network. All right. It's all broadcast.

But these channels are digital, sort of, interactive services that -- where we really need to reuse the spectrum. So if you look at the spectral reuse, you have the programming which will be received either by a backbone network through satellites, or off-air if it's being retransmitted local programming, and then distributed through a fiber ring to the hubs through the nodes, and then you could add one fiber to each of the nodes, each one having a different set of the digital services. And that way you can reuse that digital spectrum.

So, again, a portion of the cable spectrum is broadcast through all the service groups. These are typically the broadcast TV services, analog or digital, and a portion of the spectrum is reused for the narrow cast or the interactive services.

So if you look at the total use of the spectrum, and this is just a typical example, if you go down to the household level it would seem that in the broadband case for DOCSIS 3 we're allocating four channels downstream for DOCSIS. All right. They might be getting, say, 100 channels of linear content, 16 channels of switched digital video, 4 channels have been allocated to video on-demand, 4 channels have been allocated to broadband. But if you look at the actual delivered bandwidth from the cable system itself, considering that these are the service group size, how we're reusing the spectrum and then how those slots are delivered from each of the hubs, then in terms of the total delivered bandwidth, because of, again, of spectrum reuse, the percentage of the delivered spectrum that's delivered is split up this way. So really, in some sense, in a typical cable system, in this example cable system, 14 percent of the total delivered bandwidth is allocated to broadband, the rest up here. So you can see the great benefit of doing this reuse of the spectrum.

So this is the overall planning process, then. So for each of the services, whether it's analog TV, digital TV, whether it's standard definition or high definition, video on-demand, switched digital video, high speed data, or telephony, or the business commercial services, they've got to consider how to use this single pipe, and divide up and use that spectrum. So what they look at are trends in the operational performance metrics that Paul Sanchirico talked about in some detail. They look at demand trends and forecasts, including what they're expecting for new services.

So right now a big concern is 3D-TV. What's 3D-TV going to do in terms of the requirement to allocate some of this spectrum to 3D television? HDTV has been a major effort over the last few years because HDTV uses so much more spectrum. They look at the competitive trends. What are their competitors, whether it's satellite or telephone, are doing, and then forecasting those trends. And of course, as you all know, they need to make a return on their investment for their shareholders.

So there are lots of factors that will influence whether a particular plan is feasible or not. First of all, whether -- and this will vary greatly by region to region, system to system, the penetration of the services. In some areas of the country, some services may be more popular than in other. The bandwidth that's required for each of these services. The opportunity cost associated with this versus the cost to re-split the nodes and reduce the service size. For example, if they were going to remove an analog channel and increase their bandwidth for broadband, there's an opportunity cost associated with that. There's always a negative service cost associated with that because -- so we'll get customer calls when an analog channel disappears. Of course, the capital expenses for both the network as well as the client devices, that means the set-tops that

are deployed in homes. A huge effort, a huge expense goes into the changes that must be made to the business and operations support systems. And, of course, there will be service disruptions whenever you make these changes in the spectrum. And finally, there will be regulatory restrictions. So the combination of all these factors influences exactly how they choose to plan and then operate their network.

So just some examples of the performance metrics. For analog broadcast, they need to look at the signal quality, including the picture quality. All right. So with digital broadcast, it's signal quality and picture quality. For example, how much compression do they do on the standard definition channels? Switched digital video, blocking and timeouts; video on-demand, dropped sessions; broadband, the most typical thing that's monitored essentially continuously is CMTS port utilization. When they reach a certain point in which the total bandwidth capacity of the port is reached, they will then consider splitting a node. Okay. And they're able to split that node because going from the hub to the node there are multiple fibers, dark fibers, and they can split those up, and then split the node just by adding a new (inaudible). Switched digital telephony, these are the sort of traditional telephony metrics; likewise for the packet switched voice. And for business voice and data services, the service level agreements that they have with the -- whether it's for a metro ethernet over DOCSIS or a primary rate replacement, they'll have these service level agreements. And based on those, the cable companies will modify their networks.

So, for example -- this is just a completely hypothetical example -- but they may have 85 channels and they would like to go to 90 in the analog. They would like to have digital TV go from 30 to 250 channels, especially maybe in HD, and 100 to 250 in standard definition, right. So they will look at what bandwidth is required for -- if they want to increase their video on-demand, whether it's standard definition or high definition, they would look at penetration and what the peak usage rate would be, allocate some bandwidth according to the statistics of that usage. For the broadband Internet they might allocate more channels to go for DOCSIS 3. In the case of DOCSIS 3 you're using bonding between the multiple channels, and so you get some allocation of that. There's the voiceover IP, which is using that DOCSIS bandwidth, right. So this DOCSIS bandwidth is being used for both the broadband Internet access, as well as the telephony, as well as the business services.

And so they'll allocate the total bandwidth. In this case it comes to 1.3 gigahertz, which is more than they have, about twice. So they'll have to go back and rejigger this. So they will do this and finally come to some best understanding of how they can do this.

The point of this chart is really to say this is not a trivial exercise. There really is a spectrum scarcity that exists in the cable system.

The beauty of the cable system is that they fundamentally send that RF through the optical cable, they've modulated the optical light with the RF carrier frequencies, and then they can just simply detect it with a diode and amplify it and send it out of the coax. That's the beauty of it. The downside, of course, is that they have a limited spectrum.

So it's all about creating bandwidth. And they create bandwidth by analog reclamation. So you have some companies like Comcast, which are being very aggressive at going from analog to all digital. Now, in some sense, it's not completely all digital because their basic tier is still analog because they don't expect their customers to completely give up those at that basic tier. Switched digital video, it's a way to make use of the fact that not everybody is watching every channel. And so for those channels that aren't watched frequently, you might as well just switch them on when they're being watched.

Advanced techniques like advanced

compression, the cost associated with this is replacing the set-tops, right. We've already talked about reducing the service group size through node splits. And if you look a little further down the line, there will be advanced modulation and protocols, DOCSIS 3 being one of those. Okay. And, of course, you could expand the spectrum simply by expanding the spectrum, go out to a full gigahertz.

On the upstream side, the problem with the upstream is that the area of the spectrum between 0 and 50 megahertz is quite noisy. So it's really important to maintain that extremely well. Possible equipment replacements may be necessary. There's advances in the modulation; DOCSIS 3 did that. And eventually you might have to replace the drops if the noise is substantial. So all this is to manage that spectrum.

Now, if they're going to provide a competitive broadcast service, it's more than just Internet access. And this was already mentioned to some degree. They've got to provide the full set of services: DHCP, TFTP, you know, time of day, these things, domain name service; all these things need to be done. They manage their network right out to the set-tops. So a digital set-top, for example, a 22A set-top has actually an IP address associated with that for the management. Every voice EMTA, the DOCSIS ver for the packet voice service has two IP addresses, right. Actually, every modem has two IP addresses, one of them for the public Internet access and one of them for managing that modem. There's an IPDR usage reporting that can be tracked by the CMTS from the cable modem. And, of course, there's a whole billing and customer support system that's required.

In addition to that there's the whole list of necessary services that customers are expecting: E-mail, web service, you know, some people will require virtual private networking, SIP trunking, firewalls, the whole thing, CALEA obligations, right, anti-spam, anti-virus. So when you talk about operating the cable system,

every operator needs to provide this whole list of services.

And they need to do this considering the fact that the usage is growing. And to do that they've been increasing the residential speeds, in part to meet the competition, in large part to meet the competition. And it's quite interesting: The data rates are doubling every 21 months, faster than once every 2 years. Okay. An increase of a factor of 10 every 6 years. Right.

This is the dial-up, downstream and upstream. First the usage here, it's growing at about a 40 percent rate. This is actually accelerating somewhat. Paul showed a somewhat higher number, but these are the measurements that we have. During the last period from 2007 to September 2009, it's growing at a compound annual rate of about 40 percent. Prior to that it was a little slower, maybe around 30 percent.

The interesting thing here is that the top quartile of the subscribers used 114 times more data than the lowest quartile, okay, and 180 times more data in DOCSIS 3 systems. So these are the systems with the highest bandwidth that use even more. The heavy users really dominate the consumption:

Percent of the users use 20 percent of the total consumption, 5 percent use 50 percent of the total consumption, and 20 percent use 80 percent. So I guess this is the canonical 80/20 Rule.

So they need to manage all these users and things, and then look toward the future, where they're going, and DOCSIS 3 is playing the big role here. We're going from the age where it was simply high-speed Internet access and circuit switched voice to once we got to DOCSIS 1.2 and 2.0 we began to add voiceover IP, right. Here we are today, where video is now a primary -- video over the Internet, online video is a major contributor to the bandwidth. Business services are expanding very rapidly, using that same network that they invested $150 billion in, using that same network to provide these business services. And within the next three to five years you'll see a transition to IP video. So whereas today the video is provided by QAM modulation, that we expect that that will change to IP transport through their network because that can really allow a full convergence of applications.

So I'm going to stop there and see if you have questions. I wanted to just quickly go through that and see what kind of questions you might have.

But here's the three points that I wanted to emphasize. That for the cable companies, operating, planning, and managing network, it's really all about that spectrum management. Their services, they take advantage of the fact that it's digital, that allows them to multiplex their services together, all on that single wire. And that they need to continually invest in new technologies to improve the performance, and especially to increase the spectrum efficiency.

MR. KNAPP: Great. Paul, that was

terrific. Aside from dealing with the spectrum issue, is there anything unique relative to the capabilities or implementation for traffic management on cable systems versus other delivery platforms?

DR. LIAO: I don't think so. I think that the main restriction that they have is they have this limited spectrum, okay. And so they will optimize the use of that spectrum just like any other provider. Beyond that access network, the rest of the, you know, beyond that node back, it looks just like any other provider.

MR. KNAPP: If the analog portion is going to be with us for a while yet, how does the location of that help or hinder the availability, as some of that starts to wind down, the availability of more spectrum that can be used for broadband? And how do you -- is it a matter of the individual carrier making the choices of how they split that up?

DR. LIAO: So, you're talking about whether you changed the nodes, the spectrum split.

MR. KNAPP: Yeah. I mean, how does -- so the spectrum coming back effectively.

DR. LIAO: So right now there are multiple reasons for why that split is set at 50 megahertz. And one of the primary ones is what you're alluding to, I think, is the fact that channel 3 is at, what, 52 megahertz. And that produces a lot of interference, obviously. So it's really not possible to move that split from that plane as long as channel 3 is a broadcast channel creating tremendous amounts of interference. If that were to go away, that would help.

But then we need to go in and change the amplifiers. So there would be a major cost associated with that. Companies are investing all the time, so it's possible to do that. And if it's necessary, you can do that. They would have to play that off against other possibilities. So, for example, you can get more effective and efficient modulation yet still. So one of the things that would happen with DOCSIS 3 is we went from, fundamentally, from DOCSIS to DOCSIS 2, is we went from 16 QAM to 256. And there are still more efficient ways that are still in the lab, to see whether you can get, you know, 1028. There's things happening in Europe that look like -- that are quite attractive. So there's a bunch of new technologies that can do that. At the moment, they're taking channels and binding them together to make DOCSIS 3. And there's a lot of overhead associated with that binding. Maybe you could just do a, you know, a wider spot of spectrum and modulate it in a different way.

So these are all technologies that are in the lab and working. So they will be, whether you move the split, whether you use more advanced modulation technologies, they will need to be done on a case-by-case basis.

MR. KNAPP: Other questions?

MR. GOLDSTEIN: Along those same lines, how much of an effort is it to increase that 750 up to a gig or higher?

DR. LIAO: What's involved in that? Is

that what your question is?

MR. GOLDSTEIN: Yeah, right.

DR. LIAO: Well, again, it's just a capital investment thing. The technology is there. Cox has got a substantial part of their network at a gigahertz today. You need to change the amplifiers out. In some cases the coaxial cable may itself need to be placed out. So it'll be, again, a question of choice, a careful analysis of whether that's the more effective thing to do, or do a node split, or to even drive fiber deeper into the neighborhood.

And for each -- I can't overemphasize the fact that for each of these 1,000 companies, they make the decision in a different way because they optimize it according to their particular geography, their particular topology, their particular customer base. So customers in a rural area may have a different kind of requirement in terms of what services that they wish to have versus those in a metropolitan area. So it's not a single size fits all.

MR. JOHNSTON: Paul, you heard Scott Jordan talk about quality of service brought out to the end user. Now, that's not included in DOCSIS 3, obviously, but is that an issue -- cable way of just looking at the cable industry in terms of how quality of service interfaces might be extended and into the home environment?

DR. LIAO: Quality of service is part of DOCSIS. In fact, that's why they were able to deploy packet cable voice.

MR. JOHNSTON: But it's not brought out to the end user. It's internal in terms of your link itself. So Scott talked about quality of service perhaps being made visible to end user applications and services. Right now that's not done, to the best of my knowledge. It's not visible to the end user. The effect is, but they can't control it directly themselves.

DR. LIAO: That's right. They can't change the parameters in the routers. But as Paul Sanchirico said, the tools are there because the service provider, in this case the cable provider, can go in and adjust the quality of service. And they do that for voice. So they give voice traffic priority. In the case of Comcast they also do a thing where they manage, you know, the priorities of the heaviest users so that they don't destroy the user experience for those who are normal users.

But these types of -- the tools are existing in the DOCSIS technology for measuring, let's say, latency or speeds, you know, and all of that. And the companies measure those things, not -- the thing that they really use to determine how they do the planning is the port utilization.

MR. PEHA: So, as you said, to do capacity planning in spectrum management you routinely collect port utilization and per user traffic volume, I think you said. Just wondering, is that enough? If you took those measurements and broke them up by hour of the day, say, so I can tell the busy hour, is that enough information that you might be able to predict and inform a subscriber as to what their performance experience might actually be? Or is there other data that you would have already collected and available that might help with that kind of prediction?

DR. LIAO: I'm not quite sure how to answer the question. What I will say is that Ofcom did this study in the past year where they looked at various user experiences in Europe. And they compared various technologies. And in the cable systems, which are very similar to the cable systems here, the user experience is very, very close to the so-called advertised rate. And so that -- the times -- of course there will be times when the user experience will really suffer the, you know, sort of the peak times when there's a special event and everybody's on. That will happen.

But by continually upgrading the networks, typically by splitting the nodes, by watching the port utilization, when the port utilization reaches a certain level, they will split the node. By doing that they've been able to stay ahead of the degrading the user experience. And I think all of the, sort of, the metrics in terms of customer surveys and so on have continuously shown the cable services to be very -- to be evaluated quite highly in terms of the broadband service users' experience.

MR. NEWMAN: Paul, what percent of -- in the cable networks is it common for the broadband Internet service, not for the video services, for a lot of the content to be stored locally? For example, a YouTube video gets hot or something, do most of your operators employ local caching?

DR. LIAO: That's a question -- I don't know the answer to that, Stagg. I believe that's the case. But a lot of the content providers will, you know, make arrangements with Akamai or something to assure the user experience. But I don't know the answer to that, Stagg, sorry.

MR. DOSHI: Paul, maybe take you to a, sort of, three to five years' time window, where you talked about IP- TV. Can you talk a little bit about what the decision criteria are going to be in terms of allocating bandwidth, or spectrum as you call it, for IP-TV versus what people have called over-the-top IP services that perhaps consumers may want? And how would your clients make a decision between the two?

DR. LIAO: So that's not a technical issue. That issue is strictly one of making sure that their customers are happy. Because their customers are getting a spectrum of services. They're getting voice, they're getting video, they're getting data. And the transition to IP video is being driven more by the fact that they want to keep their customers pleased and, you know, enthused about their services. So by going to IP they can begin to get the convergence of traditional IP services and bring them to the consumers at lower cost than if they were to stick with the normal cable telephony technologies.

The way they will allocate this split between how much of that single pipe they will allocate for, let's say, IP video as opposed to some other service will depend upon what the competition is doing, what their consumers are demanding. And I think that's the thing that they need to keep track of, is what consumers want.

Does that make sense?

MR. DOSHI: Well, is it purely --

SPEAKER: (inaudible)

MR. DOSHI: Right. I'm assuming it's a business issue, basically you're saying. But in terms of this bandwidth or the spectrum are there any rules? Like currently, your spectrum is divided between analog, digital, STV, and broadband. All of the digital and STV probably go to IP-TV or would the --

DR. LIAO: You're asking will it all be IP-TV at some point. I think it's safe to assume that, as somebody who has watched IP over the years, that there will be nothing in the world that won't be IP, you know. It's just -- that's just -- IP is gobbling up everything. So I can't imagine that it won't be. But it will, again, it will happen differently on different systems. And it will happen at a timescale that's really dictated by what consumers are requiring.

MR. KNAPP: Paul, you answered a while ago about looking at different options for improving efficiency and so forth. Is there any look at dynamically managing the spectrums so that -- because you have times where, for example, the NBA games are on or, you know, or any other sport. And they tend to be particular hours of time or blocks of time. And then times where those channels, at least from the viewer's end, seem to be unoccupied. Is anybody looking at a way to maybe dynamically manage the spectrum so that you can take advantage of capacity that's not being used for video at the time?

DR. LIAO: So, today there's sort of two course -- there's a (inaudible) course management, which is one that takes time, to change the allocation of the different 6 megahertz channels. Of course within that 6 megahertz channels even today, or within, let's say, four channels that are bonded for DOCSIS, and some manufacturers even bond 8 channels together for DOCSIS, those things are all, in a sense, those services are all statistically multiplexed against each other. And if you use QoS priority you can manage the services within that bandwidth. As I said, they're working in the Lab now to work on more advanced things which take even wider swaths of the spectrum. You could take the whole, you know, 750 megahertz and get around 5 gigabits per second, and manage it all as one swath, in some sense.

And so all that is technically possible. And today what I can say is that today they're already doing that for voice and Internet access. Those two are already dynamically, in some sense, dynamically allocated one to the other because they're using the same DOCSIS transport mechanism.

MR. KNAPP: Paul, you've been very good. We've had you up here with all of these questions, and I just got three more that came in. One we've asked. Let me just see.

Earlier we heard a presentation from Cisco that explained that broadcast quality video requires no packet dropped in over two hours plus (inaudible) less than a 30- millisecond latency. Is this accurate for a cable's pay-per- view on broadcast channels? Is that similar metric?

DR. LIAO: I assume it is, if Paul said it was. Paul, was this your question?

MR. KNAPP: I have one other here. And I promise this will be the last one. Advocates of Internet regulation have claimed that cable companies refuse to allocate sufficient spectrum on their cable plans to Internet, instead preferring to use more of it for TV channels and non-Internet VoIP, which are more profitable. These groups claim that cable companies should be forced to allocate more of it to the Internet. What is the cable industry's perspective on this issue and on the FCC's power to regulate such plans?

And I understand you can't speak on behalf of the entire industry.

DR. LIAO: Well, what I can tell you is this. Is that, let's say, a DOCSIS 3 system, all right, cable companies will -- they may offer three, four, or five tiers of services, okay. Typically probably more like five. So what that means is that if you're on a DOCSIS 3 system you can elect to subscribe to a service that is 10 megabits per second guaranteed service or, you know, that the advertised rate would be 10 megabits per second, 50 megabits per second, or 30 megabits per second. Now, DOCSIS 3, if it has two bonded channels, each channel being about 40, that would be an 80-megabit thing, right. So if everybody chose 50, that would fill up that DOCSIS 3 path pretty fast. The reality is people don't do that. And even if they do choose 50, they don't use it 100 percent of the time. So you get this typical multiplexing of statistical usage.

So the important thing -- what I'm trying to say here, however -- so the important thing is what they do. What they do do is they watch very carefully, you know, on basically a daily, hourly, almost by minute basis what the utilization is. And if the utilization starts to come up, you know, north of 50 percent or so, they immediately start planning to split the nodes. All right. At some point they've ordered the equipment, and well before it gets to 100 percent utilization, they have already split the nodes or added spectrum or whatever is necessary, whether -- so there's all these different ways they can do it, right.

They can add more channels, they could split the nodes, they could put more bandwidth, they could go to higher modulations, you know, which means buying more equipment from our good suppliers. All right. So there's all these choices for how to increase the efficiency of the bandwidth and how to increase the efficiency of that spectrum. Each one of those things will be made on a case-by-case basis, but they will be made well in advance of when the customers require it. The point is that the cable systems have to keep up and meet the demands of their customers. They have to keep up because if they don't, their competition is going to take their business.

And what I've tried to emphasize today,

it's all about managing the spectrum, and there's multiple ways that they can do that and they increase the spectrum that they have. And at the moment there doesn't seem to be any problem with the, at least out to the foreseeable future, for what they have. And, you know, they have even the option of pulling fiber deeper into the network. So there's just multiple ways to do this.

MR. KNAPP: Great. Thank you, Paul. The presentation was fabulous.

And please give him a round of applause.

(Applause)

MR. KNAPP: Next up we have Bill Smith, who is the president of Local Area Networks at AT&T. He's responsible for all network-related operations, including over 75,000 employees across the company's domestic footprint. This includes construction and engineering, installation and maintenance, U-Verse field operations, and operations planning for wireline and wireless facilities.

Prior to his current position, he was

executive vice president, shared services in charge of mass market and enterprise operations, corporate real estate, procurement, regional wireline planning, and business planning and integration.

And I'll stop there. The rest of the bio is online.

Bill, as soon as you're ready. You're in good hands with Rashmi.

MR. SMITH: (inaudible)

MR. KNAPP: Yeah, it is.

SPEAKER: Bill's we're going to talk about Verizon.

MR. SMITH: Good afternoon. We'll just switch the presentations, just for the fun of it.

SPEAKER: Maybe announce the next merger. (inaudible) system.

SPEAKER: Not only will you not be doing the right presentation, but Paul will speak for you.

MR. SMITH: Okay, we should be -- I'm going to try and walk through a lot of the same principles, I think, that we've talked about already today and touched on today. In fact, it's kind of amazing how close a lot of our statistics are.

But this is just kind of a highlight of the topics that we'll touch on. Let me say that we'll be delighted to come back and, you know, it was -- one of the things that was a little difficult to understand exactly what level of depth. So, we stand ready to come back and talk in any more detail that you'd like on, you know, specific techniques or things that we've implemented in our network.

But this is kind of highlight. Let me start with a little bit of an overview of the network. And as you can see, it's a global network. This is really just the layer 3 MPLS network that we run most of the IP traffic on for both consumer and business traffic. And as you can see even if you look at the U.S. map up there, this is just a major backbone route. These are the 40 gig OC768 routes, so obviously we've got other routes into some of the areas that aren't covered on this map.

We're looking in some of these routes to go to 100 gig. And as soon as we've got the clear path to get the technology there, then some of these routes will probably be upgraded to 100 gig.

We have 25 major backbone nodes around this. We carry, you know, a few facts: 18.7 petabytes of data on every business day. We have almost 1 million route miles of fiber, about 932,000 route miles of fiber. We have 3,900 MPLS nodes around the world in 163 countries, and we have 38 data centers around the world.

So, I think you can kind of see the scope and scale of what we're dealing with here. In trying to deal with some of the parts of the world where, you know, maybe the rules are a little different or the way people play by the rules are a little bit different. But I think the important thing here is, a lot of this is the scale and trying to get the economies of scale.

And particularly one other key element,

I think, that we try and do is kind of take advantage of inter-segment leverage between business and consumer networks. You know, some of the things that we'll talk about today in terms of implementing QoS, some of our business customers have extremely tight QoS requirements that were held to. You know, we could do that two ways. We could build a separate network for business and consumer kind of best effort. That wouldn't be very economically efficient, in my point of view, when we've got the tools and capability to manage those within the network itself.

This chart is amazingly similar to others that you've seen today. I think the numbers may vary just slightly, but, you know, it kind of depends on the timeframe. The bottom line is we have a huge amount of growth. I think Paul said it very well, you know, the part that gives me pause in managing my capital budget is, unfortunately, traffic growth is going a little bit faster than the price performance on the equipment we're buying. Or, you know, equally you have to remember the embedded price performance of the equipment we have in the network. So that creates a challenge.

You know, our data would suggest that 1 percent of our customers drive about 17 percent of the usage, 5 percent drive 41, 10 percent drive 57. And our numbers are pretty close. The last set of data I saw, Paul was pretty close. I think we had about 14 percent of our customers that drive 80 percent of our traffic.

So, as you can see, that's a similar challenge to what others -- I think one of the things that I've been pleased to hear in the discussion today is, there is a combination of technical and commercial solutions here. And I think ultimately that's probably what we're going to have to deal with. Because throwing capital at this problem is frankly not one that is attractive for most businesses like mine. You know, it's funny, I was in a discussion just a couple of days ago about a particular route that we had turned up a new OC 48. And immediately as you opened it up, you know, it immediately fills up because it, you know, becomes one of the better peering points. So, you know, the challenge that I have in dealing with my boss and our shareholders is, okay. You spent, you know, X-millions of dollars on this. Where's the added revenue that comes from it? And so that's a challenge that the network operators are providing today.

This is just kind of an illustrative chart showing how we do manage all the different types of traffic on the network. You know, so underlying we've got this integrated, shared network. But on top of that, we're implementing a variety of service from, you know, Fortune 500 companies, government agencies, you name it across the board as well as interfacing into, you know, legacy circuit switch networks.

So, our challenge here is really to take this network and boil it down in such a way that each one of the individual applications is getting the kind of capabilities from this network that they're looking for and to do it, frankly, in the

most efficient way we can.

I think what we've learned over the years is, that you know when we started, IP networks were kind of an overlay network. Quite often they ran on top of ATM networks or frame networks. You know, more and more they've become the core backbone. And so instead of having, you know, a circuit based network that has IP running on top of it, we're quickly becoming an IP based network that, you know, has a little bit of a circuit element still embedded in there. But largely the shift to core IP is pretty dramatic here.

You know, Scott did mention one of the things. I really loved the chart he had this morning showing the different networks and how we connect across the networks. I think that's a great illustration of how, you know, the Internet really is a connection -- an interconnection of private networks, largely done through private agreements. The other part that is just an added facet of complexity to that, it's in the traffic balance as well. So it's not only are you agreeing to interconnect with another network, but it's also pretty important what the traffic balance and issues are there. Because the way Internet routing goes, with kind of a hot potato routing scheme, you know, whoever you are you try to hand that traffic off as quickly as possible.

So, if you've got a significant imbalance of your traffic -- and, you know, frankly some of our peers, you know, may be we may get 10 to 15 times as much traffic from them as we send, we pay to deliver that traffic back, really, into our network. So, it's another nuance of the way network peering works.

And you say, well, gee, what's -- how does that matter? Well, what matters about it is the way it drives cost into the network, and how network operators share cost as we implement those capabilities.

So, this is kind of the next phase. And this is a fairly elementary chart, so I don't want to, you know, offend anybody with -- but the key point here I think is, if you look at all of these traditional technologies here on the left, virtually every one of those had some individual QoS mechanism within them.

In fact, I was thinking about it on the trip over here last night. And of all the data networking protocols that I've been involved with through the years. From frame and ATM and Ethernet and you name it, I can't think of a single one that doesn't have a mechanism for priority management within it.

So, if you look at all those legacy protocols on the left hand side, they really inherently had some ability to manage prioritization. When you put all of those onto a single backbone, this one being 40 gig, you know, with our plan to go to 100 gig, it becomes even more important that you have an ability to implement some sort of mechanism for prioritization across those.

Now, I would say, you know, we've had a lot of discussion today. I think one of the first speakers talked about, you know, the ability to do inter network prioritization. And that's something, frankly, that we're very open to working with other providers on. That I think the hard part is not, can you do that, can you pass the information between networks. I think, you know, our engineers know how to do that pretty effectively.

The difficult part becomes kind of enforcement and policing. Because, you're counting on everybody to "play by the rules." And so, then what happens if they don't? And, you know, I remember this question coming up 20 years ago in some ATM prioritization discussions. When you have carrier- to-carrier interfaces, you know, you have to have the ability to make sure that when people mark traffic as high priority, there's some enforcement mechanisms to make sure that it's treated that way. And it is, in fact, legitimate high priority traffic. I think that may be one of the things, ultimately we have to work on.

You know, I think Walter, you question

around exposing some of those things to the consumer. You know, I think part of the issue is how do we then have the consumer in a position to make sure that they're using high priority for traffic that absolutely requires it. But it's not just indiscriminately being used for everything that comes in the network. Because if that happens, then obviously the whole kind of concept breaks down pretty quickly.

But if you look at our network, you know, this is kind of how we do it. And again, we've got a global network that has a lot of information in it. And these are just some examples. But it's a shared IP MPLS backbone that has, you know, everything on it from, you know, basic Internet access to things like GETS. And so what we do is, we basically have a logical segmentation of that network so that we can meet ultra high priority traffic such as GETS, as an example. We can meet real-time needs for VoIP and other real-time applications. We can serve needs for enterprise customers. You know, I'm frankly

-- I'm kind of amazed at the stringent requirements that some of the customers are putting on us in the enterprise domain today in terms of things like latency when you used to think, you know, 100 or 200 milliseconds used to be pretty good. Now, we see customers that are demanding, for example, 20 milliseconds of latency on, you know, large spans of the network.

So, for us to be able to meet that requirement, we've got logical partitioning. And then obviously, we've got a significant amount of our network capacity allocated to what we'll call, you know, public data. And I think Paul Sanchirico did a great job of kind of describing how that works today. But, again, each one of these really designed to meet some specific requirements in the network.

And this chart -- and this is one that I think is important, but even this chart doesn't cover all of the facets. This talks about kind of the different needs of the network, and -- or the different needs of applications and what they look for from our network.

So, you can see on things like bursty -- and this is also intended to kind of show current applications, things that are kind of developing, and then long term or managed network applications. So, you can see things like today, e-mail which may be fairly high on the bursty side, may be fairly high on the ability to tolerate delays in the network. You know, if I send a big PowerPoint file or if I receive one, you know, I don't really care if it's done real-time or if I get it a few seconds later. No big deal.

But you look all the way through to high quality audio, and then ultimately to real-time high definition video services, you know, all of the applications running on this shared infrastructure, many of them are looking for different characteristics from the underlying network. And the way we are able to implement that is by using some form of QoS capability.

In fact, it was kind of interesting, I

was in this room yesterday, and a CYSRIC communications services interoperability and reliability council meeting, and one of the topics of discussion there is that we've traditionally looked at e911 services from landlines and from, you know, cell phones. But now the question becomes, well, how do we move that into domains like text messaging or e-mail from a handheld smart phone? How do we make sure that an emergency message gets a priority treatment in the network? And so we're, I think, seeing a kind of a new area of opportunity and or needs from applications of the network. And I think the underlying capabilities exist, we've just got to figure out how best we implement them.

And you know, Stagg you asked a couple of times about things like caching and how does that change the requirement. And it -- I think depending on the application, it can have a significant impact. I mean, obviously the example you used, the hottest YouTube video that, you know, suddenly the -- I think the one this week was the kitten or whatever that I saw about. But, you know, that can help a little bit. We don't decide to do that. It's kind of the content provider that decides whether to invest in the caching infrastructure for that. Because it's not really up to us to decide whose content to keep, and are we keeping the freshest version of the content and so forth.

It can help by moving that closer to the end user, but as all -- we'll look in a minute, there are some -- still some kind of choke points in the network, if you will. However, you know, some aspects of the content here, like maybe some of the, you know, streaming high-quality video, if it's real time -- you know, Paul's example, best example is, you know, telepresence is an example. Obviously that's not an application for caching. So, I think it can help the network, but it -- I think it requires more study to see how much of it is discretely concentrated.

You know, when I looked at the last traffic stats that I saw, and I -- this is not something that I look at every day in my job at AT&T, but I do look at occasionally. You know, there are some entities or some network locations that have got large amounts of volume and we go -- you know, there's a lot of our traffic is coming from those sites. But, you know, then you have to look underneath.

For example, let's just take YouTube. You know, we may see a lot of our traffic coming from YouTube, but it doesn't mean it's exactly the same YouTube videos all the time. It could be, you know, a dispersed amount of content coming from that location. And that may or may not determine whether it's financially viable for YouTube to disperse that caching or not.

So, any event, that's something that we're -- you certainly looking at as a way we manage our network. I think this chart is really just kind of designed to show a fairly basic configuration of how we manage an IP-TV or U-Verse location, for example.

So, we have our IP backbone. But we

basically come in from our head ins. And if you look at most of the traffic here from voiceover IP, the network management, and the IP-TV stream is really a logical network that is built and managed, if you will, very much like what we would do for a commercial enterprise customer.

So, it's segmented, managed in terms of the characteristics all the way down to the VRAD or the special purpose video DSLAM, if you will, that sits in the neighborhood. At that point in time, we're trying to get all this information that the consumer, our customer, has asked for. We're trying to bring it in and converge it into that pipe to get down to the customer. So that's kind of where, you know, we'd implement these capabilities.

And our job is to try and make sure that we do this in such a way that the customer has a great experience. If they're running their cell phone over a femtocell from their home or their, you know, watching TV or having a voice call or surfing the Internet, you know, our objective is

to make sure that they're getting the user experience that they expect from us.

Because believe it or not, most consumers have no idea what we're talking about here today. They just want the doggone thing to work, right? And so I look at it as, you know, personally -- you know, if I'm sitting there watching the Super Bowl over IP-TV and my large screen TV in the den, I don't want one my kids to start a music download or something of that nature and cause a disruption in my Super Bowl. And frankly, it doesn't need to. Because we've implemented the kind of capabilities -- implementing QoS and managing the network in such a way that we're able to deliver both the person downloading the music video, and the person watching the Super Bowl. Both get the experience they're expecting, without a problem.

And frankly, I don't think it's an overstatement to say that, you know, we couldn't be in the video business. You know, we couldn't be competing with Paul's companies if we didn't have the ability to do this. Because we're taking, you know, a VDSL connection to the home that is, you know, in the order of 25 megabits, give or take. And we're able to deliver multiple high definition TV channels, we're able to deliver voiceover IP, and we're able to deliver Internet access. And the way we do that is fundamentally using network management techniques that allow us to manage the applications in the way they need.

So, if you look at kind of going deeper into the network, there are basic network principles here that kind of have been around for the 30+ years that I've been here. You know, we have the traditional threats, cable cuts, hardware failures, software defects, plan maintenance configuration changes, traffic surges, hacker attacks. And we respond to those through a series of principles, and you can see here. You know, we first of all try to utilize all available resources. So, if there's a suddenly a flood of traffic on the network, or if we loose a cable in a particular area, obviously we want to reroute

traffic to other facilities as quickly as we can.

But we also want to make sure that we have our traffic filled with calls that have a probability of getting to their intended destination. So, you know, probably -- and these principles work in the legacy public switch -- telephone network, they work in the IP network, largely. So, if you look at, you know, one of the examples that some of my colleagues like to use, is during -- after the 9-11 attack when some of Verizon's offices in lower Manhattan were destroyed or severely hampered. We didn't deliver calls into those offices because that tied up a resource across the country that was never going to be put in service. So, the basic principle of network management is that you try and understand whether there's a probability of being able to deliver that call. If the answer is yes, you attempt to deliver it. If the answer is no, you don't tie up the asset.

If you look at -- you know, even the fundamentals of signal System 7 were built to do that.

So we try and look at some similar capabilities as we manage our IP network. And so if we do get in overload, we try and give priority to connections that have the best probability of being successful.

So, we really then try and manage traffic congestion and maybe, more importantly, look at inhibiting it's spread. So those principles basically come into two strategies. One is, expansive. So if we do have additional resources that we can bring in to solve a problem, we do that. If on the other hand, we don't have those additional resources, then we try and stop the problem as close to the source as we can.

You know, one of the things KC mentioned this morning, the growing amount of -- I forget the exact term she used. I think it may have been "malicious content," but don't hold me to that. But I think it's important, these are the kind of threats that we see in our network on an average month. And I think it's -- you know, it's pretty telling. You know, we have about 10,000 planned maintenance events. About 300 significant cable cuts, you know, we may have 345 mass traffic events of one type or another; 49,000 hardware failures, on an average basis. Obviously this varies a lot by the time of year, but four natural disaster events of some type.

Go to the top, 205 software bugs, 1.3 million network changes, but we have almost 40 million malicious attacks on a monthly basis. So, I think again, if we weren't able to try and respond aggressively to that, I can't imagine how our network would function. Because it's amazing how many people around the world, either because they have nothing better to do or because they have an absolutely malicious intent, are launching attacks on our network.

So, you know, if you look at the way we used to do it and the way we currently do it, it's not dramatically different. The left-hand side of the chart kind of shows the way the PSTN worked. And we've talked about expansive controls of

rerouting when a failure occurs. We've talked about protective controls when -- you know, call gapping, for example, where the classic example is the radio station contest that gives, you know, $10,000 to caller number 20. You know, that creates -- it always has created havoc in our network. So we had to learn a long time ago how to deal with that.

We throttle calls to an end office. The example I just used with the southern Manhattan after 9-11 is a good example there. And then last but not least, we selectively prioritize for things like GETS and wireless priority service.

If you look at the IP domain, the details of the tools and techniques may have varied slightly, but the fundamental principles are the same. So, in the expansive side, we will change the weightings on a route so that we can start moving traffic around. And this is something we do on a maintenance basis as well. If we have planned maintenance occurring on a particular route, we can change the weightings and

gradually bleed traffic off of that route.

On the protective controls, we look at access control lists. And I think, you know, Paul had -- I think both Pauls, frankly, had some elements of this in the presentations: Access control list management, distributed denial of service protection, and SMTP spam blocking.

So, those are the kind of approaches that we use to try and manage this. And so last but not least, you know, convergence. And I think the discussion came up earlier today as well that convergence exacerbates this need. Because we are truly, you know, trying to bring our network so that it behaves as one. You know, our objective is to make it -- make the network smart enough to recognize me, whether I'm interacting with it on my iPhone, whether I'm sitting in front of my TV in my den, on U-Verse, or Internet connection, et cetera. So we're converging all of those then, and it means that we've got to be able to manage the network to offer those different capabilities that each application is looking for.

And that's done basically on kind of a dynamic basis. So, some of these things are principles that we implement on a longstanding basis, some are done more on a dynamic basis.

But with that, let me wrap and see what questions you have.

MR. KNAPP: Bill, that was great. QoS, is -- does it kick in only when we're approaching some congestion level? Or is it just applied to all of the traffic?

MR. SMITH: It's not -- we don't set the network to implement QoS in non-congested times. So it only basically -- it may be there, you know. The traffic is marked appropriately. But under times of non-congestion, you know, nothing every happens.

MR. KNAPP: And what would be a ballpark? A metric for congestion? I mean --

MR. SMITH: You know, I would rather get a good answer for you than to throw something out. And the reason I ask that -- or say it that way, is it -- you know, I could envision providing a stat from a lot of different angles. I mean, how many events per day, how many seconds per day, or of the overall, you know, 18 petabytes of data on a daily basis, how many bytes get affected by congestion?

I think a good answer to your question would require some characterization of all of the above. So rather than giving you --

MR. KNAPP: Fair enough.

MR. SMITH: -- something out of my -- out of the air, I'd like to get back with you better --

MR. KNAPP: You have the answer, Jon? What's that?

MR. PEHA: I was going to ask a different question, but. I was a little surprised, I want to follow on Julius to make sure I understand.

Are you saying that when there's no congestion, it doesn't matter that you do these things? Or are you saying that at somewhere -- wherever the line is drawn between no congestion and congestion, something kicks in that requires differential handle?

MR. SMITH: Well, I mean, if you look at some of the techniques, for example, Paul talked about with managing queues, if the queues never fill, it doesn't matter. Right?

MR. PEHA: So, it doesn't matter. But it's not like the mechanisms are different. It's just that they have no impact.

MR. SMITH: No, they're not --

MR. PEHA: Okay.

MR. SMITH: -- exactly. They don't kick in until you experience congestion, but they're there if you need them. Right, exactly.

MR. PEHA: Right. So --

MR. SMITH: And to the other point, you know. The alternative would be just to build so much capacity that you never got there. But we've talked about the financial impracticality of that solution.

MR. PEHA: Right. So I guess related -- so you have multiple queues with multiple classes and you allocate bandwidth, particularly you said you -- accordingly. And you favor -- I don't know if it's connections or classes that you said that make more effective use of resources. Can you say more what you mean by "effective uses of resources" and how you determine that?

MR. SMITH: That was -- that answer was favoring routes that more effectively use -- when we talked about --

MR. PEHA: Okay.

MR. SMITH: -- the way we change route, you know, if -- so for example, if I've got the ability to route traffic directly from, you know, Washington to New York, I would favor that route rather than Washington to Atlanta to Chicago to New York. So, it's a more effective use of resource.

MR. PEHA: I understand.

MR. SMITH: Make sense?

MR. PEHA: Yeah.

MR. KNAPP: Any others? Walter.

MR. JOHNSTON: Yeah. I was just

interested, you threw up on your mosaic there of network technologies and domains, you had Ethernet up there. Could you talk a little bit about how Ethernet is important both as a technology and possibly as a service?

MR. SMITH: You know, I think Ethernet is to the future of our network what T1s were to the past. I think Ethernet is becoming a tremendous, you know, access transport. You know, it -- people debate, you know -- OSI protocol purists may want to debate about what you call it, but, you know, base-based services are becoming very prevalent in our business.

And I -- you know, when I talk to my employee groups, I say, look, you know, you better get up to speed, you better get yourself trained, we'll help you, but, you know, you need to understand that Ethernet technologies are the future of our business. Like I say, as T1s were to the past. So, whether it's, you know, business customers looking for Ethernet access, which is very prevalent today, you know, both in enterprise and even in school systems. If you look at a lot of the school systems that put networks in using e-rate funding, you know, many of those were based-based networks.

But now, even if you move into the wireless domain, most of the, you know, 3G kind of equipment for the wireless world is based on Ethernet transport from the cell site. We've fortunately gotten to the point where instead of going out and, you know, augmenting T1s as hard as we can go, we're trying to put, you know, Ethernet connectivity into those locations so that we've got, you know, the sufficient capacity to handle 3G and then ultimately 4G wireless.

So, it's a huge and growing part of our network.

MR. JOHNSTON: And what does that mean in terms of quality of service metrics or techniques you might use?

MR. SMITH: Well, that's been one of the big challenges, I think, Walter. And I'm not the expert in that, so I'll say that right up front.

But, you know, not only in quality of service, but even in some of the operational stuff that I am closer to is, you know, most Ethernet protocols were not originally designed for long-range end-to-end network. In fact, you know, the protocol is designed to run from the corner of the room there to the corner of the room there. And so that has been one of the -- frankly, one of the challenges that we've had and probably has kept Ethernet technologies from rolling out even more quickly, is building the hooks into the network to manage it effectively.

But, you know, to get into more details on that, I'd have to get somebody that's a little more deeper in that than myself.

Unless Paul wants to handle that question.

(Laughter)

MR. KNAPP: Steve, go ahead.

MR. BUENZOW: Yeah. I had a question about the statistic on the users and the amount of traffic that they use sort of universally; the universal thread through a lot of the presentations here that it's a small percentage of the users that create a large amount of the traffic.

And I was wondering if that -- if you have data behind that to tell you whether that's a function of the user's patterns or the applications that they're using. In other words, is it a few users that are adopting, you know, things like video on demand that are using that versus the large number of users that are using less bandwidth intense of applications or is it, you know -- is it just that certain users are on more often using, you know -- watching online videos or whatever constantly?

MR. SMITH: I think it's a combination of things. I would say, when we first started looking at this or when I first started looking at it in my -- in a previous assignment, it was probably four or five years ago, and at that time it was almost exclusively, and I guess I can say probably illegal, peer-to-peer traffic.

You know, so we did some deep packet studies just to characterize it, and that was kind of the finding at the time. That's no longer the case. I think there are some of those things still there. But I think what we find is, you know, more and more it's people who just choose to do more and more, you know, high bandwidth, high rich content applications. And, you know, a lot of people I talked to say that, you know, they virtually never watch TV live. They, you know, stream their favorite shows over the top and, you know, that that's basically what they do to get their content. So, I think that characteristic has changed pretty significantly over the last few years.

Now, I will say that there's still occasionally the situation where you find that, you know, the parent who is the primary holder of the account has no idea that, you know, their 14-year-old son is, you know, running some peering point out of the bedroom. That still happens, quite honestly, quite a bit. But I think there's been a pretty dramatic shift of the nature of the usage.

MR. KNAPP: Stagg?

MR. NEWMAN: Paul, in effect, made the point that, you know, if you look at coming into the home, 4 percent of the bandwidth on that pipe is devoted -- or 3 percent to Internet access, which is really what the focus of the NPRM is all about. And 14 percent of the bandwidth across this network, arguably.

Your diagram showed over that pipe, you know, voiceover IP as a managed service. Your -- video IP-TV is a managed service, and then the Internet access is another service on the Triple Play. And I think there was something else on it.

What is your comparable sort of allocation, if you think about it? And more importantly, how do you manage the question of congestion, you know, given you're dealing with a 30 or 40 megabit per second pipe? Paul said it's a digital pipe; his is a several gigabit per second.

How do you manage the problem of the IP-TV congestion versus the voice congestion versus the Internet access over that pipe?

MR. SMITH: Well, our numbers, given that it's a -- you know, kind of a fundamentally different architecture for us. It depends a little bit on what service our customer signs up for. So, we have or will soon have even a higher speed data connection. I don't know if we've announced it yet, so I probably shouldn't go a whole lot more in detail.

But we offer multiple tiers of Internet access. And let's say, for example, that we're going to have one that's hypothetically well north of 10 megabits per second. And let's say that that customer is getting a U-Verse offer from us that they have a 32 megabit per connection, as an example. In that case, we may have as much as 50 percent or close to 50 percent of the bandwidth to the home dedicated for just pure Internet access.

So, it's kind of up to the customer to choose, you know, what package or what feature set

that they want. Now, the way we manage that is basically, we implement prioritization at the VRAD because we know that, you know, the customer for voice and HDTV or video, they're looking for real-time, you know, capability from us whereas the Internet access, if there is a congestion point and something needs to be buffered and something needs to go through, we'll buffer the Internet access momentarily and get the video through.

Now, you know, it's --

MR. NEWMAN: So, what does that do to the -- let's -- so I'm trying to understand that. So the customers subscribe to IP-TV and then they have two HD sets. So that's, you know, 16 or 3 --

MR. SMITH: Right.

MR. NEWMAN: And they've subscribed to Internet access. So, the kid in the room is watching the video over the top from our next speaker. You know, how does that work?

MR. SMITH: Well, I have never -- you know, I was just thinking. I have never seen a kid situation, gotten a service complaint, I've never heard of a situation where there's been a problem. So, I guess the issue, Stagg, would be the assumption is that I'm -- if I'm a U-Verse customer and I've got three HD streams all running on different programming, so I'm not getting any sharing there, and then suddenly we start trying to blast down a multi-megabit per second file, then conceivably there could be congestion there.

But, frankly, it would take a very bandwidth heavy application in conjunction with multiple HDs operational to cause that kind of an application disruption.

And I'd have to get -- I'd have to see if we can collect stats on that. There may be a way that we can query or study and see, you know, how often do we see that kind of a situation occur. But frankly, I'm not aware of it being an issue.

MR. NEWMAN: Okay. So, maybe -- here -- tell me if I heard correctly. I mean, one, you're saying, in general, you see that as more a hypothetical than a real problem. Then, in effect, you're offering two types of service: The video service, which is comparable to what I would get over cable TV, and then, in effect, a best effort -- you're advertising a best effort, which can peak at maybe 32. But it is a best effort service. So -- and you're -- that's what you're based --

MR. SMITH: Right --

MR. NEWMAN: -- (inaudible) the customer. Okay.

MR. SMITH: And that's kind of the way we try and -- our peaks, now, we don't have -- for the raw Internet service, we don't go quite as high as you're talking about, but conceptually, yeah. You're -- the terms of service are that, you know, you're buying from us a video service, and a voiceover IP service, and an Internet access service, and, you know, the Internet access piece is best effort, which we think is going to be pretty darn good, you know, virtually all the time.

But if in the event of a congestion event, just like my Super Bowl example, you know.

(Applause)

MR. KNAPP: We'll take just a five-minute break, so that we're not going three hours non-stop. But that also was -- since I'm being benevolent in giving the five-minute break, we want everyone to be disciplined and be back here at -- we'll make it 10 after.

Thank you.

(Recess)

MR. KNAPP: Can I have everybody please take their seats again?

I'd like to introduce Tom -- and I know I'm going to mess up the last name, Tom -- Sawanaboy?

MR. SAWANOBORI: Sawanobori.

MR. KNAPP: I got close. In parens it just says "Tom."

So it says vice president of Network Technology Strategy for Verizon. He's responsible for technology strategy for all of Verizon's business units in his group; leads technology strategy, wireless technology strategy in planning, standards, and regulatory support; provides wireless strategy and planning direction for network headquarters and field support. And his organization also leads capital budget allocation for the corporation and manages the intellectual property portfolio. In his previous role as vice president for network planning for Verizon wireless, he led technology direction, planning, and evolution of the radio and core network. I'll stop there. Very impressive bio.

And, Tom, look forward to your presentation.

MR. SAWANOBORI: Okay, thank you very much and good afternoon.

So, I'll spend a little more time talking about the wireless network, but initially I will talk from Verizon's perspective on both wireline and wireless networks.

So, pleased to be here this afternoon. And just a little bit of an introduction. Some of the similar comments I'll make about our networks, so obviously we didn't collaborate before here, but Verizon's networks overall for both landline and for wireless networks communicate with more than a 100 million people each day. So, that's a tremendous amount of volume of traffic: Over 1.7 billion text messages exchanged; 400 million e-mails; many, many petabytes of video- streamed. I think I want to emphasize also a lot of potential threats, over 5 billion potential threats monitored. And that's upon every day and obviously delivering a high volume of phone calls as well today. So, that's on both the wireline and wireless networks.

And, you know, I think I think our premier advanced networks here most people are familiar with: Our FiOS network, which is really our fiber-to-the-home solution, and then our wireless network.

So, our FiOS network today is deployed in over states where Verizon operates or currently past just over 14-

1/2 million homes today with plans to extend that beyond in 2010. We've invested a significant amount of capital in the ground over the last several years in order to deploy all that infrastructure and to provide both broadband and broadcast television to our customers, and TV is available today in about 14 states.

On the wireless side, we operate both 2G and 3G networks. Our 3G networks today cover over 280 million people. We've, again, invested a significant amount of capital over the last few years to build out that network and deploy both the coverage and capacity needed.

And we're currently trialing 4G technology using LTE technology -- long-term evolution -- in both Seattle and Boston. So, we announced plans to get that up and running in 25 to 30 markets next year -- by the end of next year.

Okay, so some of the other presenters talked about bandwidth demand, and this is a chart from Infonetics and Alcatel-Lucent from several years ago, but I think it points to some of the same characteristics that people have been looking at.

So, if you look at peak data demand, there's been a significant increase over the last several years, and, you know, wireless -- when we talk about some of these other technologies, they've been around for many, many years. Wireless just celebrated its 25th anniversary last year, so you think about that, and in a very short -- relatively short period of time we moved from 14.4 services, which seemed pretty adequate when we were just getting going on wireless and shortly thereafter moved up from -- I think people thought it was fairly impressive to have 56 kilobit -- and then today consumer demand, traffic, those kinds of applications are driving the desire for higher and higher peak traffic.

So, today our FiOS, for example, we have tiered offerings, as some of our other competitors do, on the order of 15, 25, and 50 megabits per second downstream and similar -- and then we have

-- we're working on advanced plans to advanced that even further, which are obviously competitive.

So, if you think about bandwidth demand, and other people put that in -- I think Bill put it nicely in terms of multiple dimensions, but those -- the different traffic and applications that people want to run have different requirements in terms of bandwidth, but also latency. So, obviously e-mail and text have lower requirements for that. Gaming, depending on how interactive that is, has a higher requirement for bandwidth than latency. Photo sharing is getting more and more popular, especially with some of the social media and social networking capabilities, and uploading of photos is becoming more and more intensive, video and music streaming obviously more intensive.

Bless you. So, shifting gears a little bit more into the wireless focus, if you look at actual traffic on our Verizon wireless networks, on a monthly basis over the last several years our

2G traffic -- and what I mean by 2G is our 1XRTT traffic, which is based on CMA technology -- was capable of about 144 kilobit peak service and utilized by 60 to 80 kilobits per second average. So, it doesn't sound exciting by today's standards, but when we launched that in 2002, that was pretty good for the day and people were glad to be able to get mobile Internet service and be able to take their laptop and move around and go to airports and be productive at hotel rooms and so forth. So, we have seen that -- you know, that grew fairly slowly on this relative scale, but it was fairly impressive for the time.

Now, over time obviously technology advances get additional spectrum, and we actually saw 3G traffic obviously growing at a much higher rate, and that's a combination of things, right, both the capital investments that we made, technology becoming available here.

We utilized EV-DO -- which is evolution data optimize -- which is our 3G technology that we chose to deploy, and deploying that broadly we

started the rollout in 2004 and completed that over about primarily a three- to four-year period. So, we still have a few nooks and crannies that we're continuing to equip throughout the network. But for the most part, it's nationwide. It's got over 280 million people covered. But when you think about the growth of data services, particularly in wireless, they're dependent on a lot of different things, both the footprint, availability to the customers, the right devices being implemented, and so forth, and so the data traffic is really starting to grow quite significantly -- over 500 percent recently. And when you think about it, I didn't even put 2009 data up here, but with the proliferation of new devices like today's devices -- Smartphones, BlackBerrys, and so forth -- we're seeing a continuing exponential growth in data traffic. So, those things in combination with some of the other applications that I mentioned earlier are driving our -- along with competitive pressures -- are driving our plans to go to fourth-generation technology.

And I think one of the things that is similar about some of the other presenters, but different in many cases with wireless -- we have a lot of competition, and I think you've heard it, we've said it many times, but, you know, we have four national competitors -- Verizon, AT&T, Sprint, T-Mobile -- and then in many cases you have regional operators and then you have some others like Leap and Metro PCS as well, not to mention some of the resellers. So, there's quite a bit of competition here. We're all competing for those customers and their usage and selling them services. So, part of this is, you know, driven by capacity and needs to serve it more efficiently; part of it's driven by the competitive nature; and part of it's by our need to serve the consumers' applications.

So, that kind of brings me to technology, and when I refer to "today" here, it's referring -- these now I'm referring in this context to average speeds in megabits per second on the upper charts. So, today using our 3G technology, our EV-DO traffic or our EV-DO technology -- excuse me -- we're able to deliver roughly 1 megabit per second, and I'll talk more about, you know, what the range of expected speeds are.

When we go to LTE, we're expecting to get 5 to megabits per second on the downlink and 2 to 5 megabits per second on the uplink. So, it will enable more and more applications. By the way, those expectations are based on some of the trials of testing, simulations that we've been doing, and so forth; and they vary based on -- they're going to vary based on where you are in the cell, how much we're building, how much traffic there is, and so forth, and other characteristics I'll go into in more detail. So, your performance is obviously going to vary, but we do think that that's a reasonable expectation to be setting with customers and to be communicating at this point.

Latency is going to be about a quarter

of what it was on the -- or what it is today on the 3G network, so that's important for certain applications, particularly for certain business applications and things that require low latency, which I'll talk about in more detail.

So, generically, network management, some of the things have already been mentioned. Internet is a network of networks. I think Scott Jordan -- Professor Jordan talked about a number of different networks linked together, but I think they're -- when you look them in aggregate, there's many, many networks -- tens of thousands of them -- and the speed and quality and latency are all defined by how many hops, how far the packets have to be transmitted across some of the network capacity and the variation and traffic obviously, and then obviously the complex nature of the applications.

I think it's already been mentioned before, even with the very high-capacity networks and technology that we have, it's not possible, it's not economically feasible to build capacity to serve all network peak traffic, which is beyond what we would expect on an average day or month.

One of the things that makes it more challenging is this whole issue of cyber tax and malware, because that creates unwanted traffic that we still have to deal with. So, we have to still figure our how to manage that and get rid of it, so it doesn't impact our customers, to the maximum extent that we can. And, you know, we -- as engineers, our folks like to work with engineering forms like the IETF, in order to figure out how to manage the write technical solutions for these. So, we're working across industries, across standards groups, and so forth, in order to make sure we're using common approaches wherever possible.

Okay, talking a little bit more about latency, I think it's already been defined by others here, but obviously the time it takes for a packet to travel from a certain source to a certain destination -- that applies for both wireline and wireless networks, but I'll talk a

little bit more about how it's different for wireless.

So, you know, there's a certain amount of buffering and scheduling, and I think some of us have seen that, trying to utilize video, for example, over an Internet connection. But, obviously, it varies by application. So, a sensitivity to latency for VoIP -- voiceover IP -- and gaming is going to be a lot tighter than something like web browsing or e-mail reading or something like meter reading, which maybe could be in several minutes or even hours. So, the latency requirements are going to vary quite a bit, and that's why it's important for us to be able to have the capability to offer managed services.

Okay, so I promised you I'd talk a little bit more about what's different about wireless, so wireless is definitely more complex an engineer to manage, and, you know, why do I day that? Okay, first of all, we have a very dynamic environment. If you think about it, the traffic -- first of all, we're engineering, you know, putting in cell sites to locations where the traffic is going to be variable. It's moving around. Some of it's stationary; some of it's pedestrian; some of it's going to be moving at high speeds, and that requires, you know, radio propagation, radio resource management, and so forth, handover and managing traffic engineering across constantly moving and changing traffic patterns and while it's growing at the same time. You know, that's something that engineers have perfected over the last 25 years, but it's still -- you know, it's still -- there's always optimization and tuning required, and particularly because the traffic's always moving around, engineers are having to look at, plan the network, provide capacity into it, and provide initial coverage into it, then tune the parameters, and so forth.

The other thing is that the radio signal environment is changing. So, when I go outside and step out to the street side, that radio signal environment is constantly changing and fluctuating, and as soon as I turn the corner, and even if I'm walking or whether I'm in a taxicab or I'm on the train, it's constantly changing. So, that's makes it even more difficult, especially when you're trying to handle some of these advance services. Most of our voice today is handled by circuit-switch connections. When you think about trying to handle voice or other real-time data services over an environment that's constantly and dynamically changing, it makes it more difficult. Then you've got things like interference. Signals in wireless are important in terms of the signal-to-noise ratio or signal-to- interference ratios. So, it's not just how much signal strength. Typically, people want to look at their phone or their device and see how many bars do I have; well, if I have four bars, then I'm in good video conditions. Well, yes, possibly, but probably not necessarily, because you also have to manage interference and have relatively good signal-to-noise ratios. So, I know that many people understand this. But a lot of the consumers that we are having to deal with don't really understand it. It's managing the appropriate signal-to-noise ratios that's important, and that can be dynamically changing as well.

And then you've got, like I said, mobility management, which is our customers do expect to be able to place their phone calls, make their data sessions, and so forth, and really desire to be able maintain their connections. They want to have their applications run whether they're playing in a taxi from here to the train or getting on the train or whether they're sitting in the hotel or moving about in a vehicle. And more and more, the applications, I think, are going to be riding over mobile systems as we go forward. It just makes sense.

One of the other big challenges here is the interface band with this constraint. There was a little bit of discussion earlier with Paul about limited spectrum availability on cable. Well, he was talking about hundreds of megabits --

or megahertz available. We're talking about tens of megahertz, depending on your particular operator, how much spectrum they may have, and how it's being utilized.

So, if you think about what we're utilizing in wireless today with cellular PCS systems, you're talking about a couple of tens of megahertz available. So, your bandwidth is definitely constrained, and we're sharing that pipe, if you will, or that spectrum resource, across all the users, where both users are using voice or they're pulling down data, whether they're doing something in between, text messaging, all of these things, picture messaging. You know, people are using much different devices. All you got to do is turn and look around at what's happening in the room, on the subway systems, and what you're doing in your workspace. And more and more, the workspace is being time-and-place shifted as well. So, spectrum is a limited resource, much more limited for wireless and a fraction of the bandwidth compared to wireline, in particular fiber systems.

So, we operate fiber systems as well. And you think about the amount of spectrum you have available and a piece of fiber, and, relatively speaking, it's very small. So, it's another constraint that we have on wireless.

And, of course, the shared-user access, it's more of an issue in wireless because as you're moving about and accessing the wireless, the cell site, the more capacity that's being used by one user, the less that's available to others. Now, if that sector of cells is not congested, that's not so much of an issue, but obviously it's difficult for us to engineer and manage for that because we don't know what the traffic's going to be like on any particular cell or sector at any particular time of the day. We have to spend a lot of time and effort looking at the trends in our traffic data, how it's growing, and then when to add capacity.

Sometimes capacity can be provisioned -- additional capacity can be provisioned in a short

period of time, meaning days or weeks, depending on quickly we can identify the capacity need and how quickly we can order equipment and install it. Sometimes if you need additional cell sites in there, you acquire a cell split or additional cells that can take months or years, particularly given some of the owners' zoning processes and so forth.

So, you know, it's a difficult situation to manage, but we do work very hard at it, and I think ourselves and our wireless competitors, you know, pride ourselves in doing so. We do think that that's -- at Verizon, well, we believe that that's a key differentiator as a network quality, so I'll talk more about that in a minute.

So, all these factors impact performance and latency, and they obviously impact customer satisfaction, because if a customer can't access the network, make calls, and do what they want to do, then they're not going to stay with their service provider.

So, briefly, I want to talk a little bit about the RF situation, and I think it's important that people understand that there's a lot of complexity in the algorithms that are being implemented such that -- this isn't something that engineers decide to do, but more of what has been perfected over a number of years and is probably similar for both -- more advanced with as you move from, like, 2G technologies to 3G and so on.

But what we have here is a plot of the relative signal strength and time. So, what we're selling is over time and in a matter of a very short period of time. The signal strength fluctuates quite a bit, so it's moving up and down; there's peaks and valleys all the time. And during this fast -- this relatively short period of time, we're presuming that the customer wants to have -- is making a data request and asking for data to be sent to them. So, in the upper left you can see that the best time to transmit that is when the signal strength is good, so there's little kind of white square boxes up there which show that that's when the SNR is really good.

That's when the best time would be.

And on the order of milliseconds, the RF scheduler in the cell site has to figure out how to -- when's the right time to send data to that person as they're moving around -- moving through the city or driving down the freeway and sending data to that user dynamically as decided by the network infrastructure in conjunction with the mobile device. So, in order to maximize a performance, that's what some of the algorithms in the infrastructure do along with the device interaction.

The other thing is if you look at the bottom right, the worst time to transmit is when you're in a deep failing condition, okay? So, if you send that data in a deep failing condition, probably the likelihood is that the customer won't get the device -- won't get -- or receive that properly or there could be errors in it or could be delays. And I think all of us have experienced some of that, and we try to minimize that as much as possible, but sometimes that happens.

Now, if you look at the top right, many times these RF schedulers will actually adjust what they need to do, so if you look at that light blue box, that's indicating -- in that time sequence the data cue is actually ready here, but if you look straight below that, we're in a failing condition. So, what the algorithms try to do is delay the packets when they can be sent at a higher modulation index and a better modulation type, resulting in better throughput for the customer.

And the other thing I'll mention about RF systems is they utilize different modulation techniques to different times and different circumstances, so that's the other complexity is that -- and that's not something that we decide a priori; that's something that's decided dynamically by the systems depending on the sophistication of the system and the technology. So, it might use more sophisticated and better algorithms or modulation techniques. When you're closer to the cell site you have good, strong signal strength, but when you're at cell edge condition or you're in a higher interference environment, it may not be able to take advantage of that.

Okay, so what are some of the other factors that affect wireless? Well, there is a strong interdependence of the devices in the network and also the applications. It seems pretty apparent that that's the case. But if you think about it, most of us just think about the phone and we dial when we want to dial and we press the SEND key and we expect it to work, and we don't really think about what's really inside that device, right? It's a complex set of radios in there, along with -- depending on the sophistication of the device, it's basically a minicomputer in the device, because it's also asked to do web browsing. It's also asked to do e-mail. It's also asked to do some advanced services as well. Some of the new Smartphone devices are doing -- are accessing the web and doing things like Facebook and other capabilities.

So, those devices have to be fairly sophisticated, and while we support openness, we do have a device certification process to ensure that when people dream device, it's not just a dumb device. The device has to have a fair amount of intelligence, sophistication, and meet the radio -- the type radio parameters that are specified by the technology and also by our technical requirements in order to do that. So, we have both our own certification process, and since we announced our Open Development Initiative a few years, we've established a third- party certification process as well.

I think I already alluded to supporting a bunch of different traffic and services simultaneously, so we do do this.

Now, today in wireless, this is all done over the same share of radio spectrum and access, so voice, data, messaging. Now, some of them might be done a little bit differently on different technologies. Today, most of our voice is carried on our circuit-switch network, as well as our text messaging. Broadband data is handled primarily on our 3G data network, but also can be utilized on our 2G networks. And in the future, we'll be able to offer another technology as well with 4G.

So, we manage all that. We're trying to make sure that customers have the best experience with that. And also things are going to be -- special techniques are required to implement things like wireless party service, which Bill mentioned earlier. So, that requires, you know, reserving or prioritizing access to the technology. So, that's some unlimited circumstances, but we do have -- those capabilities to manage the network are required if we're going to provide those kinds of capabilities.

We need to be able to prevent harm to the network. That's already been mentioned, but one data point I want to talk about briefly is we have to put security gateway spam filters in. A lot of people on the outside that are kind of the bad actors that someone referred to want to send spam to mobile devices, and, in fact, on a typical month we get 30 to 40 percent of our incoming text messages are actually spam, and I think our customers -- occasionally that gets to a customer and they don't usually care to get it. So, we try to filter that out to the maximum extent possible, but that's not always easy to do. You have to have technology that tries to figure out, you know, to identify that there's a lot of -- there's a broad range of numbers or addresses or so forth that somebody's trying to send information to, so that kind of indicates that it might be spam, and you don't want to be filtering out traffic that's legitimate traffic. So, it can be a very difficult thing. And sometimes we see storms or attacks of text messages, and sometimes that can be 80 or 90 percent of that particular day's traffic. So, it's a challenging thing to manage, but we do spend some time and effort and money doing that in order to protect our network.

I briefly mentioned our Open Development

Initiative. I think we announced that about 2 years ago, and we do have over 30-some devices available today, so those range from things like metering, telematics, health care. So, if somebody wants to develop their own device using or running their specific application, they can come to us and work through our ODI process and get that device certified and be able to run their applications on our network.

Okay, I was asked to talk briefly about some of the key factors that -- or metrics that we utilize in wireless networks. Some of the things that are utilized for voice in particular are ineffective attempts. Some people would call that blocking, but not necessarily means blocking; it just -- in our context it means any failed attempt, so it could be for radio channel blocking or it could be because the signal strength is inadequate or for any other reason. So, ineffective attempts and lost calls -- we measure those on a daily basis or an hourly basis, but the data is always changing, because the situation is very dynamic. So, you know, we try to engineer -- we use that information to engineer the networks and to make sure they're the highest quality possible, and if we see an area that has a high lost call rate, we'll try to investigate and find out if there's an equipment problem; is there a problem with some of our transport; is there a problem with capacity; and so forth. So, it takes some constant engineering, monitoring, and so forth to ensure that quality is maintained as high as possible. And data that even gets more complex, because we look at other parameters that just fail connection attempts which are analogous to ineffective attempts in voice and drop connections. You also have to look at things like task success, so once you've established a data connection, did the session stay up and was the person able to conduct the task they were trying to do?

And then we've had things like throughput as well, but throughput measures for actual customer data are difficult to gauge, because some people want to do e-mail and web browsing, and so they may require only a few kilobits per second, where somebody else may be trying to stream a video and that may require megabits per second. So, measuring throughput by itself doesn't mean that just because there's low throughput being generated on a particular site or sector doesn't mean that that's necessarily a bad thing. It depends on what the customer's requesting, how many customers there are, and so forth.

And then in terms of statistical performance, we look at statistical performance in our network core platforms. Are they serving the capacity that we expect now? We look at them to make sure their utilization isn't too high, and we look at the performance overall on a cell site basis as well. We use both statistical data coming out of the network and then we also utilize drive testing, so I think we're one of the few carriers who does their own internal drive testing. We do that on a regular monthly --

quarterly basis over the same route. So, we actually drive our own systems to make sure they are performing as we expect them and to focus in on the problems that we have and to continue to refine and optimize those systems, and we also compare that to our competitor systems, so other people I think commonly subscribe to external third-party drive-testing data as well to help optimize their systems and to know how they're doing against the competition.

Some of the data or the information that's available to the customer when they're choosing a wireless broadband provider and plan and device are -- I listed here on the chart. So, if they wanted to get a data service, obviously they're going to looking a certain kind of device? Do I want to get a Smartphone? Am I interested in a laptop service? Am I interested in something more of a BlackBerry device, or what have you, or do I just want some phone and occasional web browsing.

They're going to look at our price plans in terms of the cost for that, how many megabits they expect to use on a monthly basis, and so forth, and then they'll choose the appropriate price plan. In the future we may have tiered price plans as well.

Devices that I've mentioned -- there will be a variety of different devices, and depending on what kinds of products and service they want to utilize on it will help guide their choice.

Coverage -- obviously, they're going to want to know what technology and device they could use to carry out the kinds of services they want, so they can look at what our voice and messaging coverage is on our 2G network; enhanced data, which is our 1X or 2G data network; broadband, which is our 3G data network -- and they can look at it both on a national basis and they can drill down into their local region by ZIP code and see what the geography is that we cover. That includes roaming services, so we typically -- we do as well as other carriers -- have roaming agreements so people can look at when they're going to be utilizing our network and then when they're going to be roaming.

So, all these choices or else information is available on our website for folks to look at, and then we'll -- in wireless we typically give information that represents where we expect the average or typical user performance on a broadband system to be. So, for example, for our 3G network, which is EV- DO Rev A, which is what we have today, 600 kilobits per second to 1.4 megabits per second. Five to 800 on the uplink is what a typical user should expect to see. And we kind of validate that based on actually doing file transfer testing and so forth under both mobile and fixed conditions.

Okay, I think a lot of the other presenters have already talked a lot about these tools, but I'll just briefly mention potential tools that we're looking at in order to manage service quality.

Quality of service, which has already

been defined and described by others, so there's -- you know, as we get more and more applications, we're probably going to have the need to utilize some of these service capabilities. We utilize quality of service in a very limited way today on the wireless network, primarily on our push-to-talk service, which rides on our 3G data network, so it's really sort of voiceover IP over the 3G data network. But that's important to do in order to make sure that the voice packets get transmitted as efficiently and as quickly as possible, because during all those different conditions, all those mobility varying signal conditions and so forth, it's obviously important to a congestion, but it can also be important under other conditions when the radio channel is varying.

And then you've got things like traffic shaping, which are used for optimizing bandwidth and so forth to make sure that we're allocating bandwidth appropriately and fairly for our customers.

I know I'm over time, so I'm going to try to wrap here, but -- so, as far as Verizon goes, we are looking for common ground here as you've looked towards this proceeding. We believe in an open Internet, and the customer should be able to use the devices and applications they want. We think that flexibility in the policy is key. We don't really believe that rules are needed for this, particularly on the wireless side. But if there are some rules, obviously, developed then we think they should be very prescriptive, because they're going to need to be flexible, and in terms of flexibility we think the providers need to be able to flexibly manage their networks, do it reasonably consistent with what the customers are asking for and not unreasonably discriminate in ways that harm other users, that industry collaboration and consensus should dictate what some of these best practices are. Competition's always going to drive us to do the right thing, we believe, and providers need the freedom to offer a high-quality managed service in addition to kind of the best-effort services that are available.

So, with that, I'll end my prepared remarks and take your questions.

MR. KNAPP: Thanks, Tom. Fades and interference are dynamic, so -- and you had the slide up there about the packet essentially waiting for what I'll call a clearer channel, for making it simple. And maybe this isn't the -- yeah, that's the one -- the time to get into depth on this, but I think this is really important to understand as we go forward. Whether you continue to transmit and then you get a response back from the handset that says here, I got 80 percent of it, but there's 20 percent you've got to -- I know I'm over-simplifying.

MR. SAWANOBORI: Right.

MR. KNAPP: Can you comment a little bit more on that process and -- because this obviously is something that's a bit different for wireless networks as opposed to -- and understanding how -- the cue forms and how the prioritization is done is I think something we'd really like to understand.

MR. SAWANOBORI: Okay, so in this environment, and it depends -- the level of sophistication of depends on how -- the particular technology that's being implemented, but in this case what the infrastructure providers and the technology are trying to is to try to optimize when the data packets are sent, because the radio conditions are changing very dynamically as you mentioned. So, that's not as -- just to be clear again, that's not something that we decide ahead of time; that's something where the algorithms are dynamically measuring how quickly that's happening, and if for some reason the signal strength from the mobile comes back and it's in deep fade, then -- because there are signal measurements being made back and forth between the mobile and the base station, then the base station might decide a few milliseconds or microseconds to delay and then send the packets. So, it's a little bit different whereas I think typically an

IT network -- and I'm not necessarily the expert on all this -- but there might be a lot of retransmissions. So, you'd rather -- it's probably more efficient to wait and delay for a few milliseconds and to then transmit that data when the signal strength is good than to keep sending retransmissions which are ineffective and effectively create overhead.

MR. KNAPP: Mm-hmm.

MR. SAWANOBORI: So, I don't know if that answers your question, but that's the kind of thing that's being done, and I think that's one thing that makes it kind of unique for wireless.

MR. KNAPP: Yeah, I think we probably will need to follow up and take down a bit more on this more than we can today, but thank you, that's helpful.

Yes, Stagg.

MR. NEWMAN: Tom, I want to ask you a similar question to what I asked Bill Smith. That is, you currently, today, offer video services that are not part of your Internet service, right, your V CAST, and you're offering, obviously, voice and push-to- talk, and then you offer Internet access over my -- whichever one of these is a Verizon device. How do you handle congestion among those different services both today in your 3G network, and how do you envision that when you move to an LTE network with a I assume probably a core IP and IMS? How do you handle the congestion in those -- among the different services you're offering the same RF interface?

MR. SAWANOBORI: Okay, so your question is today -- first part of it was today how do we manage congestion amongst the different applications.

MR. NEWMAN: Right.

MR. SAWANOBORI: You know, I would say at this point, it's -- today it's primarily best effort. So, if there's a number of different applications we run simultaneously, whether it's V CAST video or web browsing or what have you, basically it is a best-effort type, you know, environment. However, what I'll say is it's a little more sophisticated than that, because the radio scheduler is constantly looking at it primarily from a radio perspective what's the most efficient way to manage it, so if you've got multiple devices and you're doing -- let's say you want -- you're browsing the web because you're checking some e- mail back at the office or -- and you happen to be looking at a video at the same time that somebody else is doing something down the block, both those things are going to be interacting, but the network doesn't look at -- we're not discriminating based on applications; we're just -- it's just looking at what's the most efficient way to deliver your traffic the other person's trying to get and so forth from a radio management and efficiency standpoint.

So, I'm not sure if that answers your question, but --

MR. NEWMAN: Do you -- how do you -- when -- okay, today your voice is actually not offered sort of in the data mode, right? It's not voiceover IP.

MR. SAWANOBORI: Our voice is primarily on circuit switch. It could be voiceover IP going over our networks, right, but that's a small percentage of the traffic. The vast majority of it is circuit switch today.

MR. NEWMAN: When you go to your LTE and voice, say, in 2013 goes to voiceover IP, how do you -- do you envision having to provide a lot more qualitative service to the applications on how to envision handling that?

MR. SAWANOBORI: Yes. When we get the 4G services, you know, I think that's still under study, that's still under evaluation, because, frankly, we're still planning how we're going to implement voiceover IP, and we're working with industry standards, (inaudible) one of them is -- we also have an industry initiative called One Voice where we're specifying how it's going to work. So we'll be working over the next several months and, you know, next year or two to define how we're actually going to manage these different services, because essentially they're going to all

be IP services, and they're all going to run over the same IP network, albeit wireless. So, we really haven't decided yet. We'll still be looking at how we're going to manage that.

You know, initially, I suppose the networks won't have a lot of congestion. We'll plan to implement this on our 700 megahertz spectrum. We're going to implement 4G there, so -- but we still have to look at how we're going to -- what our policies are going to be and how we're going to manage multiple applications, so it's something we're still planning on studying.

MR. KNAPP: Other questions? I've got one here. Well, let me pass this, actually this came in over the net. It's from a wireless Internet service provider. There's actually six questions here. And I think I'll just -- I'll take the liberty of boiling it down to two and reiterate that --

MR. SAWANOBORI: Thank you.

MR. KNAPP: -- we can't -- we won't -- we don't aspire to answer all the questions today.

Can -- you had on one of the charts a -- typical speeds for EV-DO Rev A. It was 13 I think it is, maybe two more ahead.

The question was actually posed in the context of capacity for a cell site, but that gets even more complicated, so I won't pose it quite that way, but can you compare the data speeds for LTE and what's expected typically?

MR. SAWANOBORI: Right, so, this is what we expect on a 3G network today. As I mentioned, and we're expecting, based on, you know, trial testing, simulations, and so forth. So, we don't have a lot of live traffic, because we haven't rode out these networks and don't have real commercial use on them, obviously, yet. But we're expecting to have downlink speeds on average in LTE, the way we're implementing it, so it depends on how much spectrum; it depends on how you implement advance antenna techniques; it depends if you provide all the necessary backhaul.

If you're in -- you know, obviously some radio conditions are going to be better than

others, but on average, we expect to have 5 to 12 megabits per second on the downlink and 2 to 5 megabits per second on the uplink when we go to 4G. So, it'll be analogous to these sorts of numbers. Obviously, performance is going to vary, depending on the application you're running and what the device is and the radio conditions and how much variability, but on average that's what we expect.

MR. LIAO: -- hand-held mobile type device as opposed to, say, a high-gain fixed directional device.

MR. SAWANOBORI: Right. The question was does that apply to a mobile device as opposed to a high-gain -- right, if you're in a fixed wireless condition where you have, say, a better antenna configuration or higher gain, that would lend itself to a better performance, correct.

MR. KNAPP: And the second part of the question goes to what happens to ensure that one or a couple of consumers don't consume all of the available bandwidth in a particular cell site? In other words, how is it parceled out?

MR. SAWANOBORI: Well, right now, today, as I mentioned, the algorithms are going to use -- the algorithms that we use in the network are going to decide how to handle that traffic most efficiently, but if one particular user is using a lot more bandwidth, then they potentially could be consuming a lot of the available bandwidth and others could suffer. So, we currently don't do anything proactively today to deal with that situation. But over time, that's something where we could be implementing some tools like some of the traffic optimization in order to, say, to adjust all potential users in that sector such that they could -- you know, there's more a fair allocation to others. So, a way of adjusting or something like that could be utilized in the future to address situations like that.

MR. KNAPP: I think I saw a question. Tom did you still have --

MR. SAWANOBORI: That was my question.

MR. KNAPP: That was your question.

Walter?

MR. JOHNSTON: I have a question on your Open Development Initiative. If I remember, it was about a year ago, you announced your first -- somewhere around a year ago, you announced your first certified device, and I think it was a meter reader if I remember. How successful has that program been? How many devices have been qualified since the program's been underway?

MR. SAWANOBORI: Okay, so I don't know the exact number, but I know that at least -- you know, within the last few months over 30 different devices have been certified, so those range from things like meter-reading devices; telematics devices, which are in vehicles, to -- recently we had a health care tablet that was certified. So, there were a bunch of applications and people -- not necessarily with huge volumes, but a special application. We wanted to be able to facilitate that sort of thing.

So, it's going very well. We're very pleased with it, and I think some of our customers are very happy that

we have that of initiative in place for them to be able to bring, you know, be able to utilize mobile -- you know, use their applications in a mobility environment.

MR. KNAPP: I'm going to end the questions here. I'm sure we could go with more.

Tom, thank you very much. Really appreciate your presentation. It was a great job. And we're close to on time. Our final presentation for today is Jonathan Rosenberg, who is the Chief Technology Strategist for Skype. As the Chief Technology Strategist, he is responsible for Skype's overall architecture and technology strategy. He's most well-known for co-author of the Session Initiation Protocol, or SIP. And has authored many of the standards that are in wide use today within the telecommunications industry.

And, in the interests of time -- while he gets wired up or wirelessly connected -- I'll turn it over to Jonathan just as soon as he's ready.

MR. ROSENBERG: Can everyone hear me?

Hello? Okay, all right, we're good.

Thank you all. Good afternoon. I thank you very much for giving me the opportunity to talk to you today about Skype and how we use the network. And, obviously, our perspective here is decidedly different. Unlike most of the other presentations we've heard today, we're not a network. We don't have facilities, we don't have MPLS networks, we don't have backbones and so, really, what I'd like to talk to you about today is how we use that best effort pike that is given to internet applications like us.

And the kind of things we do, to deal with the fact that sometimes there's congestion, sometimes there's problems with the network, and how do we deal with that fact?

So, a little bit of background about us and what we do. I think, hopefully, everyone is aware of the core services our software provides. We're in the business of enabling the world's conversations. We provide communications software that runs on PCs that allows people to talk, chat, have presence updates, and see video with each other. And we've been quite successful at that. We have a pretty decent number of registered users these days. This is a somewhat old statistic -- I think about a quarter old -- but as of the end of Q3-09, we had 521 million registered users, which is a pretty decent number.

And they're all over the world. As an internet service we run everywhere the internet runs, and the internet runs everywhere. So this covers the United States, to Europe, to Asia, and places I've probably never even heard of.

Our core functionality is to provide high quality voice and video calling to use across the internet. And in support of that, in Q3 of '09 we did 27.7 billion minutes of voice communications for Skype to Skype. Between users on PCs that are running out software application. And if you look at the total count of the minutes that we've provided over the lifetime of the software, it's in excess of 200 billion minutes, so far. So, a fairly good volume of communications happening.

And if you put it all together, actually, that represents almost 8 percent to the international calling minutes worldwide, across all providers that enable communication services. So a lot of people using this -- connecting across international boundaries is one of the key applications Skype is about. Keeping in touch with friends and family that are far away. And so, a lot of times, that overseas family members that this provides a great way to closer point of contact.

Perhaps most interesting is actually this statistic: That if you look at all the Skype to Skype calls -- these 27.7 billion minutes I talked about in Q3-09 -- a full 34 percent of them were actually using video. And this is a number that has been rising, frankly, meteorically. And I don't think we're alone in this -- everyone is seeing the rise of video -- largely, frankly, for streaming video. For watching YouTube, Hulu, this kind of thing.

Here we're talking about person-to-person real time communications. These numbers are growing tremendously and we see that trend only continuing. And that, obviously, has an impact on the consumption of bandwidth on the network, as video is an entirely different story than audio. And with increasing quality web cameras and computers, more and more bandwidth will be potentially utilized to enable these kinds of high quality conversations.

So, even on the mobile space -- so, we have an iphone application that people may be familiar with -- I think this, again, this is slightly old data -- that application has been downloaded 6 million times. It enables voice communications presence and Instant Message and chat on WiFi networks only. Today -- and one of the most popular iphone applications of all time.

And I think the lesson to be taken away with here is that there is extreme consumer demand and interest in these types of communication services that just run anywhere, that work on the internet -- wherever the internet is. That provide new types of communication capabilities that uses are greatly desired.

So how do we do that? And I'm going to spend most of my time on the first two of these things. The core of our business is really voice -- is enabling high quality voice communications between users. And there's really only two disparate paths that communications can take today, Skype to Skype calls -- where I'm going between users on two different computers.

We use a voice system called SILK. This is a voice codec. If you're not familiar, this codec technologies are ways that we take digital voice that comes in via the microphone on your computer and we represent it in a format which is suitable for transition over an IP network. SILK is a Skype codec that was developed internally and it's one that we've actually opened up to the community, allowing everyone open access to the technology. And it's actually -- if anyone's not done it, by the way, please go home, download the application, make a call -- it's free -- it sounds great. And one of the reasons is SILK uses what we call wide ban and super wide ban modes, which use a higher fidelity of voice encoding than what goes on the public switch telephone network -- everywhere in the public switch telephone network limits the amount of bandwidth that gets -- of the frequency that your voice uses, based on studies done a long time ago.

But if you sort of broaden the spectrum, just like high def TV brings that extra level of crispness, this voice technology bring an extra level of crispness for voice communications. And that's what happens when we go for Skype to Skype calls. When you're calling out through the public switch telephone network -- we have services Skype Out and Skype In. Those use these -- an industry standard codec, G. 729, that runs about 8 kilobytes per second.

The reason I'm spending time on this and I'm going to talk more about it is that these are the types of things that consume the researchers

on the network. The techniques we use for transmitting voice and video largely define how we use the network. And the fact that we're using these codecs -- and SILK, in particular, I'll talk about more shortly -- are highly relevant.

On the video side, we use a codec from a company called OnTube that's called VP7, and it's similar in its overall structure to many other industry codecs. And we use that for both screen sharing -- so there's a feature where you can select a region of your screen and you can send it to other participants to see -- an addition, of course, to the webcam -- are both using that codec.

And then there's IM, and Presence, and File Transfer which, relatively speaking, actually, don't even count. When I spoke to our software engineers who build the rate adaptation tools that I'm going to be talking about in a moment -- they said it doesn't even show up on the radar of how we allocate bandwidth across these different applications.

So it's really about voice and video and how those things consume the services on the network. So we have a bunch of techniques that we've applied over time that are very unique to an application provider, on how we deal with the network. And there's two, in particular, that I want to spend a bunch of time on to help you understand what we do.

Media Elasticity and Media Relay. So, you have to keep in mind that as an internet application provider, we have no control over the underlying network. We send a packet in, and it comes out the other side and sometimes in makes it and sometimes it doesn't. And sometimes there's more bandwidth, sometimes there's less bandwidth. Sometimes packets get lost, they dropped, they get reordered, they get delayed, they get jittered. All kinds of problems happen. And since we don't have access to any of these quality service mechanisms that could be used for the managed services that we've been talking about today, we have to deal with what we get. And the way we do that is, we're elastic.

We're adaptive. We look at what the internet is giving us and we adapt what we send, in order to be responsive to it. And that adaptation actually takes into account a whole bunch of different constraints on the system. It's not actually even just the network. It actually includes things like the computer. Some computers aren't very powerful. If you've got an older machine and you plug in a high quality video web cam to it, it just doesn't have the horsepower to do high quality video and codec. So we take these kinds of factors into consideration, as well. Again, we want to give customers a choice to use whatever machines they have at their disposal, whatever cameras they plug in, and try and make the best of it.

All right, so we're adaptive to the constraints of the network, bandwidth, packet loss, delay, jitter -- in addition to the computer -- and try to work within those constraints to optimize the experience so that users walk away

with a great communications and conversation.

So, in that regard, this is where we've done a large amount of innovation. If you look at -- just focusing for the moment in our voice codec. So I'm showing here a table of how other industry standards for voice communications deal with varying network constraints. So, G. 711 -- if you're not familiar -- this is the lingua franca of digital voice transmission. The entire public switch telephone network is built off of this. This is 64 kilobytes per second, that's why everything in the telephone network is multiples of 64 kilobyte per second -- it's all, ultimately, because of this. It's widely used in voice over IP systems. It's used by us, it's used by application providers, network providers, every -- this is the common codec. But it's pretty beefy, it's 64 kilobyte per second. G. 729 is another industry standard codec that runs much lower, at 8 kilobytes per second.

The SILK codec that we use is actually adaptive. If you look at these two numbers over here, they don't change. The minimum and maximum bandwidth used by these systems doesn't vary with the network. The assumption is that, well, we're going to engineer the networks so that there's enough for this thing to go through, but that's not how we operate. We're adaptive to what the network provides us. And we go monitor and probe the network -- and I'll talk a little bit more about what we do there -- and then actually adjust the rate of our voice compression over a pretty broad range, actually. Going as low as 6 kilobytes per second -- which doesn't sound great. At that point we're narrow ban, we're back down to the same kind of low fidelity voice you get on the telephone network, but you can see that that rate is actually less than even what the low rate codecs are elsewhere.

Up to 40. At 40 we're what we call Super Wide Ban. Where now we have the crispness and fidelity of voice that allows you to really hear nuances in speech. Then if you -- if anyone has kids that's tried to communicate with kids

over a phone call, it's sometimes harder. It's sort of easier in person. One of the reasons for that, again, is that loss of fidelity in the higher spectral range of the voice, so we go all the way up to try and capture that stuff. And we adapt. As we measure what rate is available on the network, we'll move up and down to deal with it.

Another common technique that's been used in the voice over IP industry for a long while is sound suppression, which is that -- in most cases -- in a conversation, both parties aren't actively talking all the time. It depends on the conversation. Some heated arguments, it may be both of the parties speak a lot, but typically only one person goes at a time. And so, why bother to transmit information when there's nothing to say? So that kind of technique -- sound suppression -- is actually built right into the systems that we use, rather than something that's an optional add-on, as it often is in systems that use G. 711 and G.729. So when there's nothing to send, we just go right down and we just send less, and it's built in a way that makes it work really well and, frankly, hard to notice. So we're trying to be a really good citizen and only send when we have useful data to send.

And I have already mentioned this one, too, but we're adaptive to the capabilities of the computers themselves. So when there's a CPU on a machine that can't handle the higher quality voice or video, we tone down the rate of compression in order to deal with the limited constraints of the computer.

So, adaptation, elasticity, meeting whatever the network throws at it, that's our over-riding philosophy in this. And the way we do it is based on feedback. And, again, this is different than the kind of techniques we've been hearing about today. From our perspective, the internet sort of looks like this cloud we've brought. Like I said, we put packets in and something comes out on the other end. And our ability to control is limited to the observations that we can make as applications that ride on top of this IP network. There's no API or interface that tells us a priori that the cable modem is only able to support a particular value of tiered service. There's nothing that we can talk to the DSL modem that will tell us, here's the amount of QUOS or packet loss I can request for this call. All we can do is put stuff in and stuff comes out the other end, and we can measure it. And we do.

So, in a typical voice conversation here -- actually showing a complicated scenario where we have one computer that's sending and they're actually in calls, perhaps a multi-party voice call, or perhaps just several calls simultaneously, with a number of other computers with different media types. With audio -- perhaps this one has audio and video, perhaps this one is just Instant Messaging? And each of these computer cooperatively sends feedback back to the sender of the stream to help them figure out how the network is treating their packets.

Measurements of things like delay, which we can now measure a one way delay between here and here. Or packet loss, detect packet loss and feed that information back.

And this type of feedback system isn't actually entirely unique to Skype. The core protocols that have been in use in many voice over IP systems for a long time incorporate feedback techniques, so this idea of utilizing the N to N measurements to adjust the rate of transmission is something that's actually been in the core voice standards as technology called RTP -- the Real time Transfer Protocol -- for a very long time. And we're using very similar techniques here.

So what do we do with this feedback and what are some of the properties of our rate adaptation? I'm not going to go into details of exactly how it works, and what the mathematical properties are, and all that kind of stuff. But let me help you understand some of the key facts of the properties of the rate adaptation systems that we have. And this applies to all media, by

the way. So, voice, video, everything I'm talking about covers the whole range.

First of all is, we throttled down our rate of transmission based on available network resources. We actively measure what there is and when there's less, we reduce our transmission rate. So it's not this greedy thing that tries to consume as much as possible and boot other people out. We actually just try and figure out what there is and fit it within the constraints of the network that is available.

And perhaps most importantly, the approach that our software takes is actually fairly conservative. One of the metrics that used in the industry for this is, fairness with regards to TCP. As we heard this morning from Scott, TCP is one of the underlying foundational protocols of the internet. When you go to a website, when you send an e-mail, all those things are being transported by TCP. And it provides several service, one of which is reliable transport, but more importantly it provides congestion control.

And a fairly cautious congestion control, at that.

We saw those curves where things jump up and down -- it has to do with TCP's sort of gentle probing of the network where it looks for additional bandwidth and it moves gradually up, but when it sees a packet loss it jumps back really fast. Additive increase, multiplitive decrease. That's the core technique that's actually allowed the internet to survive by being very careful with the way it applies congestion control.

TCP also has some nice properties around Fair Share Bandwidth Allocation. Interestingly, without explicit Q management controls and the types of things we were talking about this morning -- just through end-to-end probing at packet loss -- studies have shown that when you have multiple people who are sending TCP traffic into a queue, in a router the way the protocol itself works tends to sort of roughly even out the amount of bandwidth that each of them are allocated. I say roughly because it is quite rough. It actually depends on a lot of factors, including the number of streams, their relative round-trip times -- between one and the other you do not get perfect allocation. But given uniform RGT, it actually -- round-trip time -- it actually does a pretty good job of evening that stuff out.

What's interesting is when you look at the type of -- and TCP rate adaptation is actually quite poor for voice because it has these dramatic swings that you saw previously and voice doesn't like that. Voice likes to have smooth -- in fact, if you can get it, constant rate is perfect. But you don't want to vary it too wildly. So we have a rate allocation technique that is appropriate for real time communications but, actually, it's not as aggressive as even TCP. It'll back off a little and let the TCP traffic have more than its fair share, as if we had been a TCP stream in the same situation.

So we're very polite to the network with the type of consumption of resources for real time communications traffic that we take.

We do this jointly across all media streams. It's not like we do this just for voice, or just for video, we look at the whole picture and try and optimize the whole thing so that across all of the streams that we're sending, across all the participants we're communicating with, we're not using more bandwidth than is available in the network. So, another important property of the rate adaptation techniques that we apply.

We also try and adapt quickly. Things happen, you know, and all of a sudden as we've talked about, someone goes -- some kid goes in the bedroom and starts downloading a huge video or movie, or something like that. And all of sudden the network gets congested, so we -- like TCP and other rate adaptation techniques -- look at these things, we detect them, and we back off in response to those congestion events, to make sure that we're not contributing to the overall congestion in the network.

And we go pretty low, okay? We're sort

of a scavenger some times. We'll take what little tidbits the network will throw at us -- please give me your kilobytes per second. And, including video alone, as low as 4 kilobytes per second. And I had to triple check that number. When the engineer told me this number I said, you must be kidding me. We take video as low as 4 kilobytes per second -- you mean 40? There's a zero missing there somewhere. But, no. Now you're not getting good video, it's more like a series of stills, like a minute apart or something like that. It's not really video any more.

But our system's were designed to be adaptive. To take, again, whatever bandwidth is available from the network and make the best use of what we can get and sometimes it's pretty darn low. In fact, relative to some of the other folks who work in this space, we think we're actually a pretty good -- you know, I'm familiar with lots of voice over IP systems, I've been working in this industry for a very long time. I'm familiar with many of the standards and techniques that are used

by application providers and core network infrastructure providers, and very few of them do this level of rate adaptation for their real time multimedia traffic.

And so I think I would argue that we're amongst the most network friendly users of real time communications traffic, due in large part to these network friendly rate adaptation techniques. So that's sort of our rate adaptation story. And that story talks about what we do when we have communications between two users. And in our ideal world, communications streams -- audio, video, and Instant Messaging -- do go directly between users, right? We're not unique in this, by the way. Voice over IP systems were always designed to try and keep the real time traffic -- the voice and the video, in particular -- following what we call a direct path between the caller and the called party, in any call. This means that as we look inside this internet we'll see routers, and switches, and peering points, whatever the IP routing protocols say is the

shortest path between this IP address and this IP address -- that's where we'd like to have the media traffic go. Any detours it takes for additional processing that would not be on the direct path mean extra delay.

And delay is the enemy of real time communication systems. There's this number -- I forget who showed it -- something like 150 milliseconds -- this is a human property that people have gone and studied that say, you need that one way delay to be under 150 milliseconds, in order to have an interactive conversation. You start to go over, like, 300 milliseconds, 400 milliseconds, you're not having a real time conversation on the phone. You're leaving like voicemails for each other.

It's really hard to have a rapid fire conversation. That low latency need -- a human need for communications -- is why people get on airplanes to still have face to face conversations. Because when you're face to face you can have an even lower latency conversation that you can't get on a phone call. And there are just fundamental limitations there having all to do with the speed of light and things that you just can't fix.

So we want to -- people who build VoIP systems -- us and everyone else, like this direct pure to pure path for media traffic but, unfortunately, it doesn't always work that way. The real internet is not as simple as a bubble with connectivity between any pair of IP addresses, the reality is that there's lots of things that get in the way. And one of the biggest enemies for real time communications is the NAT -- Network Address Translators -- that hopefully you guys are familiar with what this is. They are IP address translation systems built into lots and lots of products. Residential NAT functionality, in particular, resides in home routers. It's not an advertised feature, but when you go to Best Buy or Fries and you come home with a home router from Linxus, Nechier, Belkin, anyone, they have this NAT functionality into

them. It's what allows you to have multiple computers hiding behind a single IP address from your service provider.

And those NATs have the property of sort of forming a one way gate that allows traffic out, but not traffic back in. And that's good, actually. It's fine for web browsing, for e-mail, for video streaming, anything where you go to a server and you ask it and it sends stuff back to you. But when you want two end points to directly send each other stuff, like we want to do for reasons of lowly instant media, it doesn't work so good.

So, unfortunately, these things aren't just in these boxes people going to their home to buy. They're starting to show up in the core of IP networks, too. And this is due in no small part to the near exhaustion of IP addresses that we're approaching. There's lots of different camps on what the future will hold as we deal with this problem. I think it's fair to say that either way you go, we're going to see more and more of these NAT things to deal with, either translation to IPv6 or carry a grade IPv4 net to deal with IP address exhaustion. Either way, we've got more nets coming our way and this is only going to make the problem worse.

And as a consequence of that, you can't always get direct media paths. Sometimes it's blocked. You can't work your way through it. This solution that everyone in the industry uses -- and Skype is not alone in this -- is to use relays. Each client that wants to communicate basically contacts a central point that has direct internet connectivity and that central point acts as a relay for media traffic, for voice, for video, for whatever the traffic is, from one place to another.

And relays are bad because they increase the latency of the solution. So why am I talking about relays? The reason I'm talking about relays is because, in the case of Skype in particular, our relays are part of the software that's in the PC client self and as a result certain computers

-- not every computer -- but certain computers that have really good open internet access will be used as relays for voice conversations between other pairs of users. And that mean that more resources are being consumed. So that person who's acting as a relay -- media will come in and it will turn around and it will go back out utilizing the resources of the access links.

We don't like that. I wish we didn't have to do that. It's bad for us. It's bad for access providers -- it's mostly bad for us because it heavily ruins the experience that our users have when using real time communications. So we work like crazy to not have these things in the loop. And all kinds of really fancy detection techniques and logarithms have been implemented in our network so that we only every use these things as a absolute last resort. But some times we have no choice.

So that's it. Our relay service are provided by other computer that run Skype. We try and minimize it to always come back to this direct connection. Sometimes some router, for example, will permit protocols like UP&P that might allow us to get the right connection. Certain NATs are more friendly and can be set up to allow direct connectivity -- we do that whenever it's possible, but sometimes we cannot.

I already talked about that -- that it worsens the overall user experience.

When we do have to do it, we try not to send the traffic very far. It's not good for our users, it's not good for network providers, either. You know, if I'm in a conversation with someone down the street and our media has to go through a relay in China, it's not a very good experience, in addition to consuming lots of network resources. So we deploy logarithms that actually try and select nearby relays to the greatest degree possible, including information on what countries and IP address blocks these things are in. So we try and minimize the consumption of resources for this traffic.

We also put hard limits on the usage of

our services when these relays are in use. Video is just hard capped. You will not get a very high quality video experience when we have to relay a call because we're afraid to consume too much resources from these relay computers.

So, in summary, you know, we are -- I think as folks understand -- one of the leading providers of voice and video communications on the internet today. We're an application provider. We're not a network provider. We don't have MPLS networks, or routers, or things of this sort. We're not a network. But we also have our own forms of congestion control and bandwidth management. That's a key message I want to get across today.

We carefully control the consumption of network resources for the benefit of our own users, as well as to make sure that we're a good citizens on the internet, as an internet application provider. As a consequence, we're highly adaptive. We go and actively measure and deploy the best in (inaudible), latest logarithms for detecting available resources, in terms of bandwidth latency, and adjusting the operation of our system as a consequence. And we do so conservatively.

We try and -- we believe we're less aggressive even than TCP, in similar condition, so we don't consume more than our fair share relative to TCP at bandwidth. We do use other computers connected to the internet as media relays that will cause consumption of traffic in and out of those computers over the access lines, but we only do that when absolutely necessary. We try and minimize the amount of bandwidth and we put hard caps on the amount of bandwidth those things consume.

So, I'm sure there's questions, so I've left a little bit of extra time for Q and A. Fire away.

MR. KNAPP: Thank you, Jonathan. Since I've mostly gone first, I'm going to defer to anybody else who wants to go first this time, but I do have a question -- just in case nobody does.

Go ahead, Rashmi.

MR. DOSHI: I guess I'm probably similar to you, being the nice guy using the network and how would you use quality of service capability that Professor Jordan discussed, for example? In fact, they used Skype as a -- do you think that would be part of the useful thing, or is it that you solved your problems and you really don't need anything more from the network.

MR. ROSENBERG: Absolutely we could use it. Nothing is in absolutes. All the work we do still results in cases where users have a poor experience. There is congestion on the network and congestion happens in different places. And if we were able to have access to quality of service somewhere in the network -- it doesn't have to be everywhere. A lot of time congestion tends to be focused on the access links, so being able to have an impact on requesting particular slots in the uplink or downlink DOCSIS schedulers. These things would help solve some of our problems on one part of the network.

And then our rate adaptation will deal with some of the problems elsewhere in the network. I and we don't believe that, you know, you need to have a complete end to end QUAS solution in order to improve the overall experience that the user has. All right, just a little bit at a time. Every little bit more helps, especially when you combine it with the kinds of techniques I'm talking about today.

So, yes, if we were able to have access to APIs that would allow a computer to inform key congestion points -- that this is voice, we would like to request that there's levels of QUAS streaming, we'd absolutely love to have that. Today we do not have that anywhere. It does not exist.

MR. NEWMAN: I'm almost afraid to ask this question because I'll show my age. But almost 30 years ago, when I was about 12 years old at Bell Labs -- we evaluated a technology that used voice detection -- silence detection -- to save bandwidth at the network. The business case was easy because there was one Bell system and so we could do the total evaluation. We killed the project because it didn't save enough Cap X to warrant the additional expense. And a few years later we looked at something that would improve voice quality by using more bandwidth, but the revenue potential wasn't there.

Now we're faced with this dilemma that, for example, SILK uses as much bandwidth -- up to 40 kilobytes -- as it can grab. But if you're not grabbing your own bandwidth, so to speak -- and Paul said video is coming which goes up to 5 kilobytes per second for adoptive rate video?

MR. LIAO: About that.

MR. NEWMAN: 5 -- the video -- yeah, whatever it can take, the new video. How do we deal with this in the sense that these are great services, you're trying to be good citizens, but ultimately the Cap X burden is born by the operator, is it through QOS and paying for the QOS, or what?

MR. ROSENBERG: Pay -- our users pay for

that internet access and that internet access allows them to send data at the rates that have been promised on the service contracts that are associated with those services and so, as far as we can tell, they're already paying for those things, so I guess I don't see the contradiction here.

MR. NEWMAN: Can you verify that most users are paying for variable byte rate access, so are you saying that we need to go to a pricing scheme where the pricing is much more based on consumption of bandwidth?

MR. ROSENBERG: I think of you as -- you know, carrier pricing is carriers pricing. And that's their discretion to deliver different tiers of service that meet their customer needs. I can only describe the kind of things that, if they were offered, would be good for Skype to take use of. Again, the more bandwidth, the better. Having higher bandwidth means that we can make better use of best effort for better quality video, which is something our customers are

definitely demanding. So, sure, if higher tiers of bandwidth are available and our customers can purchase that from their service providers, that definitely enables us.

If there was some kind of services that made available quality of service treatment in a variety of different ways. If users could pick applications or, somehow -- and, again, I don't want to propose specific pricing techniques, that's not our discretion to do that. But services that become available, that applications like Skype can use to request quality of service from the network would be valuable to applications like us. However the carriers choose to price them.

MR. JOHNSTON: We seem to have this metaphor before us of good actors and bad actors, although anybody who's seen a recent film from Hollywood knows that the world is full of bad actors. But --

MR. ROSENBERG: Just a recent one.

(Laughter)

MR. JOHNSTON: And you've spoken in the context of a good actor on the network.

MR. ROSENBERG: Yes.

MR. JOHNSTON: But your software actually goes on somebody's device, whether it's a PC or a wireless, handheld device. Is there a good active metaphor in the context of what conduct should be for applications -- and on the computer it's almost -- I'd say the analogy in my mind is the computer is the FiOS of the world and the handheld device is actually the wireless condition. Is there -- are there good conduct that an application that seeks to utilize resources should undertake in terms of handheld devices?

MR. ROSENBERG: So you're asking specifically about mobile devices?

MR. JOHNSTON: I'm going on an iphone, going on a Verizon handheld phone. You're sharing space with other applications, some of them from the vendor, perhaps some of them from other suppliers. What would a good actor do?

MR. ROSENBERG: So, I think it's the similar situation, but much more extreme, right? As we heard, you know, the wireless are fundamentally -- by, like, laws of physics, more constrained than wired networks. And that reduces even further the amount of available bandwidth makes even worse the kind of latencies we see. And so it requires even more aggressive techniques around rate control. And, again, the great thing about many of these techniques is they have a pretty broad span of controls. TCP itself, for example, has proven to be fairly resilient in its ability to shuttle down its rate -- going pretty low, and reasonably high. It does suffer in the sort of foreign extremes of these things, but these types of broad rated adaptation techniques are powerful.

You have to be careful, too, because in the olden days it was sort of a one to one match. It had a handset and it had a particular network access and it was either good or bad. But we see things like this lovely little MyFi access point that our friends at Verizon make. It's great. You put this thing down on a table and it puts WiFi on one side and then a, you know, (inaudible) or a 1XRTT uplink on the other side. So here, this computer thinks it's on a WiFi network, all right. But, in fact, there's a constrained bottleneck downstream.

So then this argues again for -- and I think this kind of thing is going be more common as we get, you know, devices in the home that wander around the WiFi or PCs with (inaudible) cards plugged in. These sort of change the nature of what's a handset and what's not a handset. In terms of usage of network resources. So, the thing that ultimately works across all these different things -- you know, these types (inaudible) rate control and congestion controls, in concert with the ability to ask the network for quality of service treatment. Or inform the network of the type of resources it requires to consume, so that that can interact properly with the type of schedulers and quality of service treatment mechanisms that exist downstream from the point of origin.

MR. JOHNSTON: Thank you.

MR. KNAPP: Go ahead, Saurbh.

MR. CHHABRA: Yeah, I was thinking about the quality of service that you might manage on your own network for your own service. You mentioned about the video feed if it goes down to 4 kilobytes per second -- it's just, you know, one frame or one picture every so often -- I would imagine the audio drops as well at the same time. Would you rather not drop the video feed and allow the audio to be at a higher quality and then, basically, effectively manage --

MR. ROSENBERG: So this is -- I don't want to go into detail, but this is exactly what our logarithms do when I talk about how they jointly manage re-adaptation. We take into account the overall available bandwidth in the network, the type of services that the user is seeking to consume -- audio, video, file transfer, lots of peers, one peer -- and we sort of do a global optimization to pick relative prioritization of that traffic.

And this is something, of course, we're constantly evolving as we get feedback from our users about what constitutes a quality call and what doesn't constitute a quality call. Okay?

MR. KNAPP: Go ahead, Tom.

MR. PETERS: I would like you to elaborate a little bit about your plans in the mobile space. Right now you're supporting the iphone with your application. Do you have plans to support other phones and the limitation to WiFi, is that impeding your progress with global development?

MR. ROSENBERG: Okay, so obviously I'm not at liberty to give any specific product plans, so I won't. It's fair to say that mobile's important to us because mobile is important to our consumers. More and more people are taking their communications on the go, they're utilizing a variety of devices for those communications and wherever our users are, we want to be there, too.

As you mention, a lot of times that is limited to WiFi and is that a barrier? I mean, I have to admit it is. I mean, so I think everyone is quite well aware of the iphone situation with the Skype application today -- that we were initially limited to WiFi. AT&T very graciously agreed to lift that restriction so we could start to work over 3G and see how that goes.

We have yet to be able to deliver that capability, as we're still not able to get access to those capabilities in the underlying platform SDK provided by Apple. Rest assured that we would be quite eager to make that service available once the capabilities are exposed to us.

MR. JOHNSTON: You don't have to answer this question, but I have to ask it.

MR. ROSENBERG: Okay.

MR. JOHNSTON: I'm trying to be fair to you. To the extent that you have a voice application on a voice phone and there are public safety implications about what the user expects or believes, do you start entering the domain

where -- you know, for the rights that you advocate -- you also take on some responsibilities?

MR. ROSENBERG: There comes tough questions, like you said. So, you know --

MR. JOHNSTON: (Inaudible)

MR. ROSENBERG: I mean, I can sort of repeat our position on this. I mean, we're not a telecom service provider, we're a communication application. To your point exactly, we're an application that sits inside a device that already provides such services and we believe that consumers should utilize the features of their service provider to access emergency services and things like that. Ours is about enhancing communications, it's not a replacement for traditional telephony. Okay?

MR. KNAPP: Great. Jonathan thank you very much, it was a wonderful presentation. Let's just give a round of applause.

(Applause)

MR. KNAPP: And we don't have all of our

speakers still here, but we have a few that stayed and I just wanted one more round of applause for everybody, they did a fabulous job. Just a couple of points before we close. I mentioned this when we started that this is just the start of the process. There was no way we could pack in everybody into one day. There are many other people that we want to talk to, and hear from, and have dialog with. And we'll be doing that through the ex parte process with the team here.

In our public notice we had listed a contact for setting up meetings and our e-mail address just -- if you're interested in meeting, reach out to us. And, you know, if you don't reach out to us, we'll be reaching out to you.

And I do want to reiterate a point or two that came up during the day. We understand, this is not just purely a technical issue, but we really do need to understand the technical piece of this. We have set up a process here at the agency to integrate with the examination of the economic, and the policy, and the legal issues. We know it's a multi-dimensional issue and we're trying to be responsive to that.

And, lastly, I want to thank Stagg, and Jon, and Walter, who did the lion's share of the work in putting this together. Also, Zenia for helping out with the --

(Whereupon, at 4:38 p.m. the PROCEEDINGS were adjourned.)

* * * * *

CERTIFICATE OF NOTARY PUBLIC

I, Carleton J. Anderson, III do hereby certify that the forgoing electronic file when originally transmitted was reduced to text at my direction; that said transcript is a true record of the proceedings therein referenced; that I am neither counsel for, related to, nor employed by any of the parties to the action in which these proceedings were taken; and, furthermore, that I am neither a relative or employee of any attorney or counsel employed by the parties hereto, nor financially or otherwise interested in the outcome of this action.

/s/Carleton J. Anderson, III

Notary Public in and for the
Commonwealth of Virginia
Commission No. 351998
Expires: November 30, 2012