

1 MS. FARROBA: But is it inside the remote
2 terminal?

3 MR. LUBE: Again, it's going to be a site-specific
4 thing. It almost -- well, extremely seldomly would it ever
5 been in a cabinet, just because we all know how tight a lot
6 of cabinet space usually is. If you're talking about a CEV
7 or a hut it very possibly could be located right inside the
8 CEV or the hut.

9 The size of an ECS -- and let me make a very
10 specific point here -- again, like I said very early on, we
11 do not think it's prudent to pre-equip an entire RT line
12 capacity with a big, giant cross-connect field inside of an
13 RT whether it's a hut or a CV and certainly not a cabinet.

14 So the ECS -- when an ECS is requested by a CLEK
15 and we construct it on their behalf it is sized based on
16 what their need for pairs are, literally how many pairs to
17 each of the FDI's that are served out of that RT.

18 So I just want to make sure that you understood
19 that it wasn't, you know, an entire 100 percent capability
20 device that we're deploying with an ECS.

21 MR. ORREL: This is Barry Orrel with Quest. I'd
22 like to just add to on to what we were saying earlier. We
23 talked about the fact that it would be more costly to place
24 cross-connect fields in a controlled environment, if you
25 will, and I just want to expand on that a little bit.

1 If you -- the reason that the cross-connect fields
2 today are separate from the remote terminals is because
3 remote terminals contain electronics that require power,
4 HVAC, etcetera. Those kinds of things cost more than a non-
5 hardened box such as an FDI. Cross-connects don't require
6 power, they don't require HVAC. It's much cheaper to have
7 that box separate from the remote terminal, which is how our
8 networks got deployed to be the way they are.

9 To now require, say some time in the future, that
10 cross-connect fields be contained within the remote
11 terminal, the remote terminals are going to have to grow.
12 In addition, with the existing terminals that are there if
13 there is space to do -- to place those cross-connect fields
14 you're also going to have to account for the fact that the
15 investment that was made there is accommodating a certain
16 amount of heat dissipation and a certain amount of power
17 usage. Those kinds of investments will be lost.

18 MS. ROSEWORCEL: I think we're going to move on
19 now to quality of service issues. The second question we
20 have listed is what determines the feasibility of offering
21 various ATM classes of service such as constant bit rate,
22 variable bit rate or available bit rate over the shared
23 transmission facility?

24 Any volunteer?

25 MR. RANSOM: Well, I'll comment in the context of

1 light span. We provide the constant bit rate and the UBR
2 non-constant bit rate services right now at our site.

3 Typically, what controls the economics and so
4 forth is in general one tries to minimize the complexity of
5 equipment that's provided remotely. So rather than try to
6 make the remote terminal as though it is some sort of an ATM
7 switch with every variety of ATM service that you can
8 imagine, it's far more economic to try to do the minimal
9 functions required at a remote site, and then put other
10 equipment that's housed in the central office that can do
11 other kinds of services.

12 So in general we've found that the minimal will be
13 some sort of variable rate bursty kind of service and then
14 some kind of fixed rate services. We've looked at whether
15 it's necessary to expand that set to other kinds, but we've
16 found that those two pretty much handle the application and
17 for now that seemed to be the optimal choice in our
18 particular equipment.

19 MR. STANSHINE: Excuse me. CBR and UBR then are
20 your standard options?

21 MR. RANSOM: Yes. In light span those are the two
22 options available.

23 MR. STANSHINE: Does there -- is there an
24 additional charge if somebody wants CBR and UBR versus UBR
25 only?

1 MR. RANSOM: Well, now we're an equipment
2 manufacturer.

3 MR. STANSHINE: That's right.

4 MR. RANSOM: Oh, do you mean if you order a light
5 span with one capability or the other?

6 MR. STANSHINE: Yes.

7 MR. RANSOM: No. Those are not optional elements.
8 They come with the release of the ATM equipment.

9 MR. STANSHINE: It's a stripped down model,
10 basically.

11 MR. RANSOM: Yes.

12 MR. STANSHINE: Thank you.

13 MR. SACKMAN: The other thing I'd like to mention
14 is that at least CBR and UBR are models that have been
15 rolled out in very large-scale networks, It's not clear
16 that if you're going to deliver 220 million end points or
17 however many there are in the United States that you can
18 roll out any kind of intermediate style VBR service
19 effectively and PSTNs all CBR and the Internet is all UBR
20 and the rest of it's sort of in the middle.

21 We're talking about an ATM-type system that was
22 never organized to deliver \$29.95 a month or \$49.95 a month
23 service with variable bit rate across a large number of
24 interfaces.

25 So even though it's technically possible, whether

1 it's economical or logistically possible with actual human
2 beings that are available for the carriers to deploy today
3 is a completely separate question.

4 MR. STANSHINE: Okay. But AFC now is offering UBR
5 and CBR --

6 MR. SACKMAN: We also offer VBR.

7 MR. STANSHINE: Okay. The ability to offer that
8 again is as common -- does the stripped down model offer
9 both UBR and CBR? Does CBR cost more?

10 MR. SACKMAN: No. The product that has to do all
11 those features and offer them inherently and we don't charge
12 extra for that, the question is, is whether it's
13 operationally --

14 MR. STANSHINE: Okay. But it's all in the common
15 equipment, the ability to do a --

16 MR. SACKMAN: No. It's distributed throughout the
17 entire piece of equipment on every line card, on every piece
18 of equipment we make.

19 MR. STANSHINE: Okay. Thank you.

20 MR. LUBE: From the ILEC perspective, I agree with
21 the comments from the manufacturers about, you know, what
22 sorts of capabilities they build into their products and, of
23 course, we as the purchaser of those equipment products are
24 reliant upon what capabilities are in those products to
25 begin with. I mean that's just as a preliminary matter.

1 But there are some -- and there are costs
2 associated with that, as I think -- as I think my colleague
3 referred to, ultimately getting down to the most critical
4 issue of how many dollars per month you can sell this stuff
5 for?

6 Which, by the way, everything we deploy or
7 everything we consider for deployment, that's a very key
8 issue to us because the broad band market, as I'm sure you
9 all know, is a very price-competitive market.

10 So deployment decisions that we make in the
11 network organization are absolutely directly keyed to these
12 costs that we're talking about. So speaking of those costs
13 there's more to the picture than just the costs that are
14 built into the equipment that we buy from the DLC vendors.

15 Depending on how the network is actually
16 architected by the ILEC -- like, for instance, in the
17 project pronto network we have the NGDLC remote terminal in
18 the field but we have the scaled-down ATM switch that we
19 call the optical concentration device, the OCD, in the
20 central office. There's interconnections between those two.

21 When you start talking about other ATM quality of
22 service classes you immediately start talking about band
23 width. A network as it is deployed, physically deployed,
24 does not have infinite, unlimited band width. I mean we
25 know that technology does marvelous things with capacities

1 and stuff, but as a network is deployed it has -- especially
2 with respect to electronics in the network, it has limited
3 band width.

4 If you start talking about the provisioning to
5 thousands and thousands of subscribers of other quality of
6 service classes you have to look at the band width that
7 you've got available and what is the price tag to augment
8 that band width?

9 In our instance, if we're trying to modify our
10 platform, the project pronto platform, to accommodate more
11 band width, then that involves the addition of central
12 office OCD equipment, either the addition of other transport
13 capacity, whether it's additional fibers or electronics on
14 those fibers.

15 So there are network design issues that are above
16 and beyond the manufacturing design issues that are a key
17 element here.

18 MS. FARROBA: Well, as the network is currently
19 designed is it feasible to offer all of these different bit
20 rates?

21 MR. LUBE: Not in --

22 MS. FARROBA: Are you saying not right now?

23 MR. LUBE: Oh, I'm sorry.

24 MS. FARROBA: Go ahead.

25 MR. LUBE: Not in unlimited quantities. You all

1 may be aware of SBC -- the SBC ILECs just released a letter
2 to the carriers earlier this month announcing a CBR broad
3 band service over the project pronto platform.

4 The band width of the service we're offering and
5 the amount of that service that we can offer on an RT by RT
6 specific basis is limited, again because of these band width
7 considerations and cost considerations that I've been
8 talking about.

9 MR. FARROBA: Right. But I guess what we're
10 trying to get to is technically what are these limitations
11 and what do you have to do?

12 You talked about having to add additional capacity
13 such as with the OCD, etcetera. What is it that's
14 determining the feasibility right now of adding that new CBR
15 offering? I mean what is it that you have in place that
16 allows you to make that feasible? Did you have to upgrade
17 something? Is there enough capacity now? What are the
18 different components?

19 MR. LUBE: Okay. Well, the basic components of
20 the project pronto network that we're talking about right
21 now are the NGDLC remote terminal, since it has an Alcatel
22 light span 2000, just as a for instance. Well, that
23 actually is the predominant vehicle we're using right now.
24 The OCD in the central office and the fiber connections that
25 are between that remote terminal and the OCD.

1 The capacity that is installed in our network
2 right now -- of course, it was originally put out there on a
3 UBR basis because the primary objective of pronto was to
4 reach the mass market with Internet access. But capacity-
5 wise with what's deployed in the network today, we are able
6 to accommodate some degree of CBR service. But like I said,
7 it's a 96 kilobit service which allows voice to be carried
8 with DSL by the CLEC if they want to do that.

9 It is not unlimited in terms of the number of end-
10 users in a particular RT. In other words, in a typical RT
11 we have 672 end-users that we intend to serve DSL. I'm
12 talking about one of the light span type RTs. All of them
13 cannot have CBR and all of them cannot have unlimited
14 amounts of band width on CBR because the pipes or the fiber
15 paths back to the CO will become exhausted.

16 If you do some things to the network to increase
17 the amount of that band width you can take back from the RT
18 to the CO then you're adding to the OCD costs. You have to
19 add more OCD ports, which may trigger you to a second or
20 perhaps third OCD box or switch in the central office.
21 Those are the kinds of impacts I'm talking about.

22 MR. STANSHINE: If I understood you correctly
23 then, it's a matter -- you don't need new features. It's a
24 matter of bonding either additional interface cards or
25 higher speed interface cards on the OCD, the same thing at

1 the RT or in the loop electronics. Possibly more -- but
2 it's more stuff rather than new kinds of stuff?

3 MR. LUBE: Well, let me clarify one aspect of your
4 statement.

5 MR. STANSHINE: Yes.

6 MR. LUBE: In our architecture we are using the --
7 the optical card we're using at the OCD and the optical card
8 we're using at the NGDLC remote terminal is the highest
9 that's available. It's OC3C or OC3. It's the highest band
10 width available to us today.

11 MR. STANSHINE: I thought there were OC12 cards
12 available on the light span.

13 MR. LUBE: There is a light span 2012 that I'm
14 sure Dr. Ransom can describe more eloquently than I can, but
15 basically the DSL capable channel banks that are within the
16 light span 2012 are still fed at the maximum OC3 rate.

17 MR. RANSOM: That's with any detail that you want.

18 MR. STANSHINE: But that in a nutshell is correct?

19 MR. RANSOM: That is correct.

20 MR. STANSHINE: Okay.

21 A PARTICIPANT: -- per channel?

22 MR. RANSOM: Per channel bank? Because there's
23 channel banks in --

24 MS. FARROBA: I'm sorry. Just a second. This is
25 Kathy Farroba, for the record. I think we need to just make

1 sure we identify ourselves and then also would you speak up
2 a little bit.

3 MR. REILLY: Dave Reilly with Rhythms. Just to
4 clarify on that. The OC3 per channel bank, is that on a per
5 channel bank basis, meaning there's multiple channel banks
6 in the light span or is that per light span?

7 MR. RANSOM: The OC3 -- and this is Neil Ransom
8 from Alcatel -- in the light span system and whether it's
9 light span 2000 or 2012, there is an OC3 which is shared
10 across multiple channel banks and, in fact, can be shared
11 across multiple light span systems.

12 So the OC3 capacity is shared across the entire
13 light span, be it multiple channel banks, and can be shared
14 across multiple light span systems.

15 MR. McNAMARA: Let me add one thing to that.
16 Actually, there is a way to handle a little bit of
17 additional capacity. They can, indeed, be shared across
18 multiple channel banks but it actually is available on one
19 single channel bank assembly. You can actually isolate an
20 OC3 down to as few as about 200 lines, but that requires
21 additional transport capacity, a lot of additional money to
22 beef up this capability.

23 The key thing to remember is that the services
24 that the ILECs are offering today is basically UBR. We've
25 found we can very effectively support at least 2000

1 subscribers on a single DS3, which is the core of our
2 network. We rarely exceed more than about 50 percent
3 capacity on a single DS3.

4 If you compare that to what you can do with CBR --
5 let's say we're trying to give someone 500 kilobit CBR, in a
6 single DS3 you could put 90 customers in there and it's
7 completely filled as opposed to 2000.

8 MR. KIEDERER: Charlie Kiederer with Verizon. One
9 point I want to make and I think it brings together maybe a
10 lot of what we've heard today.

11 You know, we as incumbent local exchange carriers,
12 have probably built the most efficient voice network in the
13 world. Now what we're trying to do is to force-fit data
14 onto it.

15 Whereas, if you had your druthers or had the
16 opportunity for a desert start or had the proper economic
17 incentives you might do it a different way. We would not
18 necessarily use NGDLC and we wouldn't be talking about
19 whether you use a light span 2000 or a light span 2012
20 because there would be other ways of doing this.

21 So I think we need to keep that fact in mind of
22 where we've been, where we're going and what we're trying to
23 do with this. You know, we're trying to force-feed advance
24 services onto what started out to be a platform that was
25 meant to efficiently serve voice traffic. So we need to

1 keep that in mind.

2 With regard to the CBR, the issue did come up that
3 to the degree you have many, many, many CBR customers served
4 off of that RT, depending on the pipe that you have going
5 back from the RT to the CO, there is a possibility that you
6 certainly can exhaust that pipe. Then you're faced with
7 putting in more broad band transport capability from the RT
8 back to the CO.

9 I also don't want to minimize the fact that
10 whatever you do at the RT and whatever services you decide
11 to offer, whether it's one, three or six, all of them have
12 implications in the EMS' and in the back room operation
13 support systems that have to drive all of those service
14 provision inventories.

15 MS. ROSENWORCEL: You know, building off of that
16 question. If we could just look towards number seven, which
17 is, what options are available to the ILECs if the capacity
18 of the shared transmission facility nears exhaustion? What
19 are the options available if interoffice transport, just for
20 comparison, nears exhaustion?

21 MR. LUBE: I think it's very perceptive to jump to
22 number seven because it fits in with the description of the
23 changes that I was describing that we would need to
24 accommodate large amounts or larger amounts of CBR and that
25 type of thing.

1 Let me preface really quickly though, I mentioned
2 a minute ago during this same discussion the price-
3 competitive nature of the overall broad band market even
4 with respect to the other technologies like cable modem out
5 there.

6 The type of thing that we're talking about in ILEC
7 maybe having to do to upgrade its network with more capacity
8 is the very time of thing that has to be looked at very
9 seriously because of the economic impact of that additional
10 cost. Is it going to price DSL out of the broad band
11 market? That's a very critical concern to all of us that
12 are trying to sell it.

13 But, anyway, to number seven specifically, there
14 were basically three different scenarios and there may be
15 variations of these three, but I think there's probably
16 three major in my mind scenarios that we would use in the
17 project pronto architecture.

18 Before diving into those, just to add something to
19 what Charlie said, he mentioned or made the comment about
20 when trying to use this network we may have to add more
21 transport capacity back to the CO.

22 I cannot overemphasize the fact that it goes
23 beyond just the transport back to the CO. What's happening
24 inside of the CO is also a very big cost factor here. As
25 I've said several times already, we use this OCD in the

1 central office and the more OC3 connections or paths that we
2 bring in from the RTs and have to terminate on OCDs we need
3 more OCD port capacity.

4 OCDs are port capacity limited. So we may have to
5 buy a second OCD full of cards with more port capacity. So
6 these are major, big-ticket cost items.

7 Okay. Now to get into the three -- and, by the
8 way, I think this will answer the question that was asked by
9 the Rhythms engineer just a little bit ago as to whether the
10 OC3 is per RT or per channel bank?

11 To add to what Dr. Ransom replied on that, the way
12 that we're deploying the 2000 -- the light span 2000 in our
13 network -- there's typically going to be three DSL channel
14 banks in an RT. Within that RT those three DSL channel
15 banks are chained together such that we use a single OC3
16 coming back into the CO for that entire RT.

17 Again, that RT we're looking at in that
18 configuration is probably 672 end-user -- not probably --
19 we're looking at 672 end-user capacity.

20 MR. STANSHINE: Per three -- is that 224 per
21 channel bank or is it --

22 MR. LUBE: It's 224 per channel bank, yes, sir.

23 A PARTICIPANT: -- or looking forward?

24 MR. LUBE: I can't speak for the vendors on that.
25 I'm sorry. But, anyway, to continue on, what we would have

1 to do as one of the options of dealing with the question
2 asked in number seven is to unchain the three channel -- the
3 three DSL-capable channel banks within the RT. What that
4 would do, we would physically be bringing three separate
5 OC3Cs back to the central office.

6 Now just to give you some metrics on my OCD
7 capacity comments, consider the fact that we may have 20 or
8 let's say up to 20 -- not in every wire center, but 20 RTs
9 within a central office, a wire center area. If I triple
10 the amount of OC3s that I bring back into the central
11 office, I triple the amount of ports that I need on the OCD
12 which specifically is what triggers me into a second OCD
13 switch or a third OCD switch in some instances.

14 So, anyway, that would be my first option, to
15 unchain the channel banks. That would require more fiber,
16 more OCD.

17 In addition -- and I may be incorrect about this,
18 but I believe that would even involve taking some of the
19 channel banks out of service, but I would have to defer to
20 the folks that understand how -- we've not done that so I
21 don't know what exactly that would cause.

22 The second option is really a variation of the
23 first. It would be to unchain the three DSL channel banks
24 and instead of using three times the amount of fibers to
25 come back to the CO, a technological possibility -- I don't

1 consider it economically feasible, but a technological
2 possibility is to put something like DWDM on the fiber so
3 that I can carry all three of those OC3s or OC3Cs back on
4 maybe the same fibers.

5 There are problems inherent with that because DWDM
6 is generally fairly large and uses a lot of power and
7 there's heat dissipation issues. So we regard that as not
8 being a viable alternative in the RT environment. So I'd
9 have to say my second option is really probably no option at
10 all economically to us.

11 The third option, which is really the worst, would
12 be to add another NGDLC, a whole new RT structure.
13 Obviously, that requires all of the costs of the additional
14 RT, the additional fibers to feed that, and, of course, the
15 common problem of the additional OCD capacity.

16 MS. FARROBA: Thanks.

17 Actually, we'd like to hear about some of the
18 other network architectures that are out there, but -- well,
19 I had a question, as well. I'll just throw it out there and
20 then, Jerry, go ahead with yours.

21 But on the upgrading of the capacity, you said it
22 would require additional OCD ports. I wanted to know if
23 you're upgrading capacity because it's due to converting the
24 constant bit rate instead of the UBR that's out there, why
25 that would take additional OCD ports?

1 MR. LUBE: Well, I think for the numerical example
2 that my colleague from Bell South, I believe it was, gave a
3 minute ago, you know, if you're looking at CBR at -- I'm
4 just trying to think of an example. Well, I think you said
5 like 500 kilobit CBR and I think you said 90 customers. I'm
6 not sure what --

7 I think we figured that if you're looking at --
8 oh, I forget what our calculations were -- I could run
9 through those in a minute or figure those out here in a
10 minute, but what happens is if you have individual users
11 with CBR with band widths in the hundreds of kilobits per
12 second, that literally consumes the available band width on
13 the facility going back to the central office, you know, on
14 the OC3C.

15 In an OC3C, I think our estimates are that you get
16 about 135 megabits per second of useable DSL band width
17 coming back. I suppose if you had one megabit CBR you could
18 put 135 customers on there and then it's gone.

19 If you say, well, you know, I could go ahead and
20 unchain my three channel banks, I still have a port capacity
21 for customers of 224 per channel bank. If I only get 139 --
22 excuse me -- 135 CBR customers on 224 port capacity hardware
23 then I've stranded some of our hardware out there. That's a
24 cost, another cost to me, another cost that affects the
25 price competitiveness of DSL and the bigger broad band

1 market. There's just a lot of issues like that.

2 MR. ROSENWORCEL: Copper Mountain.

3 MR. REISTER: Yes. I feel like -- you asked what
4 determines the feasibility, it's the chips, of course, that
5 are in the equipment and it's the vendor of the carrier's
6 equipment choice.

7 A couple of points. One is that you can provide a
8 wider array of classes of service today. I mean there are -
9 - boy, there are OC3 SARs that can do quality of service for
10 -- on a variety of classes with queuing for, you know,
11 10,000 individual ques for \$80 or \$90 dollars for the chip.
12 Obviously, by the time the vendor goes and integrates it
13 into their equipment and has their profits they charge the
14 customer. But it's certainly feasible today and has been
15 feasible for quite some time.

16 I think that the issue here though is we keep
17 talking about this quality of service and CBR per
18 subscriber, which is a really crappy -- in my opinion --
19 crappy way to do it because, obviously constant bit rate is
20 what it sounds like, which is, you are permanently
21 allocating that band width on a per subscriber basis.

22 So, you know, it's commonly called like a circuit
23 emulation service because you can use it to create the
24 equivalent of a T-1 from Point A to Point Z.

25 But what would make a lot more sense, is, if you

1 did the constant bit rate in the ATM Lexicon you called it a
2 virtual path, which in English is like a tunnel that you can
3 put many lanes into. So now either the incumbent carrier or
4 a competitive carrier could say, "I would like 10 megabits
5 out to this particular location" and the only thing that the
6 incumbent carrier would have to architect for is that 10
7 megabits.

8 At that point now, how much traffic each
9 individual subscriber going into that tunnel takes up, does
10 not matter a wit to the incumbent carrier architecting their
11 fiber, it's just that 10 megabits and there's a variety of
12 things you could do for the individual subscribers feeding
13 into that.

14 You could certainly provide real-time services,
15 burstable services. You could just divide the 10 megabits
16 using weighted fare queuing. There's a lot of schemes you
17 could do. But you've set that tunnel so that the outer
18 bound is that 10 megabits and now the CLEC knows exactly
19 what it's paying the ILEC for.

20 It's the same concept when a carrier is putting
21 equipment in a collocation space and they've got a DS3.
22 There they've got 45. The nice thing is if one carrier was
23 more successful than another in a particular market area
24 they can say, "Hey, I would like to go from 10 megabits to
25 20 or whatever." Obviously, the carrier that was managing

1 the facilities would then charge correspondingly based on
2 the 155 megabits that was available.

3 In terms of question seven, to upgrade it, you
4 could -- I mean the fiber itself, assuming it's single-mode
5 fiber, doesn't really care whether it's OC3 or OC12. What
6 does cost is the interface card that would be on the RT
7 device and the interface card on the OCD. OC12 certainly
8 does cost substantially more than OC3.

9 Presumably though, if you were running out of
10 capacity on it, it would be because you had so much service
11 and if you had that much service hopefully you would be
12 getting revenues that would correspond to that type of a
13 band width upgrade.

14 Then, finally, I point out that ATM is not the
15 only way to carry over fiber or maybe I should say SONET.
16 There was the mention of DWDM, which is currently I agree
17 too expensive in this environment. But there's gigabit
18 ethernet and other fiber-based technologies that over the
19 coming years certainly hold the promise of substantially
20 increasing the band width available.

21 MS. FARROBA: Jerry?

22 MR. STANSHINE: Just to get a couple of things on
23 the record with Alcatel and AFC. Your CBR is virtual
24 circuit and not -- or virtual VC rather than VP, is that
25 correct?

1 MR. RANSOM: Well, the current offering we have is
2 a virtual circuit where we're adding also the virtual path,
3 but that's a virtual path going to a single customer.

4 I think what Copper Mountain was asking for was
5 whether there was a way to share a single CBR virtual path
6 across multiple customers and that's not currently scheduled
7 for this product.

8 MR. STANSHINE: Okay. NAFC?

9 MR. SACKMAN: We only do virtual circuits today
10 and we haven't been asked by anybody to do virtual paths.

11 MR. STANSHINE: Got ya.

12 And for SBC, if the network has to basically
13 absorb substantial extra cost for CBR service is there
14 anything about the technology that would force you to burden
15 the UBR customers with this or could you really concentrate
16 the cost recovery in your tariffs on CBR customers?

17 MR. LUBE: Well, first of all, I have to put the
18 disclaimer out that I'm not a pricing spokesperson. I
19 really don't know how we would differentiate the pricing
20 between the UBR customer and the CBR customer.

21 I do also want to clarify that as deployed, our
22 network does have a limited amount of CBR capacity. I do
23 want to restate that because I don't want there to be a
24 perception that it's not possible to put any CBR over the
25 deployed network today.

1 If I might also just branch out for just a second
2 into the permanent virtual path, the PVP. I don't recall
3 that you had a question specifically about PVP, but that's
4 another capability that a lot of the carriers have desired
5 in the proceedings, in the state proceedings, that I have
6 been in around the country.

7 PVP is a particularly onerous issue to us right
8 now for at least two distinct reasons. I think you need to
9 understand this part of the technology to see where we're
10 coming from.

11 First of all, if you had multiple -- even multiple
12 carriers out there saying, "Gee, I think I'd like a 20
13 megabit PDP" and another carrier comes along and says, "Gee,
14 I think I'm going to do really well in this service area, so
15 I'd like a 30 megabit."

16 I mean you get to a capacity management issue of
17 which CLECs get have what sizes and what do you do when the
18 next CLEC comes along and you don't have any more size left
19 or any more capacity left? That's an issue that we see in
20 terms of the band width available in the OC3C that we have
21 today.

22 But also I mentioned, and I think Dr. Ransom
23 referred to this on behalf of the light span, the current
24 product that we have now -- and please correct me if I
25 misstate this -- but I believe we have a capability of one

1 PVP per channel bank.

2 So if there is CLEC No. 1 that would like a PVP
3 and one of my RTs, it literally instantaneously wipes out
4 one third of my customer or end-user capacity for DSL in
5 that RT. I think there may be some other developments in
6 future product releases, but there's limitations as we
7 understand even in those further developments.

8 MS. FARROBA: Okay.

9 I think the gentleman from Rhythms had something
10 he wanted to say and then we might want to move on.

11 MR. REILLY: Yeah. This is just a question
12 concerning the capacity issue. I think it was Bell South
13 that had mentioned they could get 2,000 users on an OC3, but
14 yet there's only 672 customers that could be served out of
15 an RT, a next-generation RT. So the numbers don't add up
16 from a band width capacity issue.

17 MR. McNAMARA: Well, that's actually for one light
18 span system. You can actually serve 2,000 customers from a
19 light span. But I wasn't specifically even talking about
20 RTs. I was talking about ATM in general.

21 MR. REILLY: Right, right.

22 MR. McNAMARA: And, actually, 2,000 customers per
23 DS3 and not OC3.

24 MR. REILLY: Right. So there's an OC3 and 672
25 customers that can be served out of one next-generation DLC.

1 There should be a lot of capacity left to do other types of
2 QOS services, UBR, ABR, VBR --

3 MR. McNAMARA: I wasn't really referring to light
4 span or any channel bank in general.

5 MR. REILLY: Right.

6 MR. McNAMARA: I was just talking about the
7 capacity of UBR traffic on a given circuit type.

8 MR. REILLY: Right. And that's what I'm saying,
9 if you have an OC3 and 672 customers that are doing UBR
10 there's a lot of capacity left.

11 MR. McNAMARA: If I could --

12 MS. FARROBA: Just a second, please. I think the
13 gentleman from Nortel --

14 MR. EDHOLM: I think -- Phil Edholm from Nortel.

15 I think it's really important to step back and
16 understand the architecture of the networks in place today
17 are based on -- it's the good statistical nature of data.
18 So data innodes are very bursty, whether they're web
19 browsing or doing e-mail.

20 As you move up in the network, the natural
21 smoothing effect of gathering up multiple users tends to
22 reduce that burstiness. So at the edge where you have
23 dedicated band width, the burstiness is handled by dedicated
24 band width. As you move farther up in the network, the
25 burstiness is handled by spreading it out across lots of