SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Akamai

How Will the Internet Scale?

Executive Summary

Over the last several years, traffic levels on the Internet have grown 50% year to year. As the creation and consumption of media grows, traffic volumes will grow at an even faster pace. These growth levels have led some industry experts to predict that the Internet will eventually collapse under the load. As an example, an article in the June 7, 2007 issue of The Economist about new online video services notes that some of the companies that run the networks over which the videos travel are not confident that the Internet can handle the strain that will be placed on the infrastructure.[1]

In this paper, we examine the currently known approaches to delivering content over the Internet. We find that the centralized approach to content delivery practiced by many providers will not scale to the levels required for video, but that a distributed architectural model, such as that practiced by Akamai, can support hundreds of Tbps of traffic, which will be more than sufficient to accommodate future IP needs.

The paper is organized into the following sections:

- Scalability Requirements: How large will Internet traffic levels become?
- The Centralized Approach to Content Delivery: Why can't the traditional approach to content delivery scale to handle future traffic levels?
- The Distributed Approach to Content Delivery: What is it, and why is it better?
- Peer-to-Peer: Is P2P the answer?
- Distributed Delivery Architecture Requirements: What are the barriers to entry for developing an effective distributed delivery architecture?

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Table of Contents

Scalability Requirements 1

The Centralized Approach to Content Delivery 1

The Distributed Approach to Content Delivery 4

Peer-to-Peer (P2P) Systems 5

Distributed Delivery Architecture Requirements 6

Summary 8

About Akamai 10

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Scalability Requirements

Consider a viewing audience of 50 million simultaneous viewers around the world for an event such as a World Cup playoff game. An encoding rate of 2 Mbps is required to provide TV-like quality for the delivery of the game over IP. Thus, the bandwidth requirements for this single event are 100 Tbps. If there were more viewers or if DVD (at ~5 Mbps) or high definition (HD) (at ~10 Mbps) quality were required, then the bandwidth requirements would be even larger.

Is there any hope that such traffic levels be supported by the Internet? In this paper, we will show how the Internet can scale to handle such traffic levels, but only by using a highly distributed architecture. We will describe all the known approaches to content delivery, and we will explore the limitations, if any, of each. We will find that the centralized approach to content delivery is insufficiently scalable to deliver the traffic levels required for video, whereas the distributed approach can scale as required, while still delivering high-quality, on-demand content. Additionally, we find that "pure" P2P networks are insufficient to meet the coming traffic demands on their own, and that they must be integrated with a distributed content delivery network (CDN) in order to be effective.

The Centralized Approach to Content Delivery

What is the "Centralized Approach"?

The "centralized approach" to content delivery on the World Wide Web is defined by the following architecture components:

- A single to a few dozen data centers
- Dense deployment of Web and/or caching servers within these data centers
- Reliance on transit or peering relationships in each data center, and in some cases, a backbone network connecting the data centers.

The centralized approach is practiced in several ways. Some enterprises choose to host all of their enterprise Web sites in a single location, under centralized IT control. Commercial Web hosting service providers generally manage hundreds or thousands of servers located in one to a few data centers. Other CDNs have implemented architectures that rely on "dense" clustered deployments, with large numbers of servers in each data center, hosted in a limited number of geographic locations. These CDN providers include Limelight Networks, Mirror Image Internet, Level 3 (through the acquisition of the Savvis CDN assets), Navisite, AT&T, and Internap (through the acquisition of VitalStream).

Scalability Limitations of the Centralized Approach

Can the centralized approach handle 100 Tbps? The answer appears to be "no," as several facets of the architecture ultimately place an upper limit on the amount of traffic that can be delivered under this model. This upper limit ranges from tens of Gbps for enterprise or commercial hosting providers, to a few Tbps for the centralized CDN providers.

The first limitation of a centralized architecture is in the number of data centers, coupled with the available exit/transit bandwidth from these data centers. The largest centralized CDN providers have at most 30 data centers in their network footprint. Using an aggressive forward projection of 30 outbound 10 Gbps connections from each data center, this results in a total capacity of only 9 Tbps. For this upper bound to be attainable in reality, the centralized CDN's network partners need to be able to support these traffic levels. This is not likely to be the case. Indeed, some carriers are already turning away paid high-speed transit traffic, citing the cost of the equipment needed to upgrade to 10 Gbps Ethernet connections[2].

The second architecture component that limits scalability is the reliance on peering relationships to deliver content into end-user networks. Given that the Internet is a network of networks, traffic will eventually reach a peering point, which becomes a bottleneck. Centralized CDNs have numerous peering relationships (often, hundreds) with other networks in order to exchange and deliver traffic to end users, and count on these relationships to help them deliver their traffic to networks where they are not deployed. While these relationships are valuable, peering with another network is a very poor substitute for deploying within that network.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Because of the high-cost, zero-revenue economics of peering, and the fact that peering requires the cooperation of directly competing entities, capacity at peering points can lag significantly behind demand, and this situation is unlikely to be satisfactorily resolved any time soon. Network economics dictate that money "enters" the network at the first mile, from connectivity and transit fees, as well as at the last mile, from end-user subscriber fees, so network providers have little-to-no financial incentive to invest any more money than absolutely necessary in improving "middle mile" peering capacity and related infrastructure. As such, it is unlikely that peering points will grow sufficiently to be able to handle anything close to 100 Tbps traffic loads.

Additionally, ISPs have a strong economic incentive to force traffic off their own networks onto peer networks as quickly as possible. This "hot potato" routing can further increase the traffic going through peering points, at a cost to overall Internet efficiency. Tier 1 ISPs, in particular, have very little incentive to peer well, as they forego transit revenue by peering, with little to gain in return[3]. This has been demonstrated in disputes between competitors that lead to sudden de-peering, with severe consequences for end users. This is evidenced, for example, in the well-publicized dispute between Level 3 and Cogent Communications in October 2005 that led to users from Cogent to be unable to reach Web sites on Level 3's network and vice versa[4], as well as Cogent's April 2007 de-peering of several smaller UK-based ISPs[5].

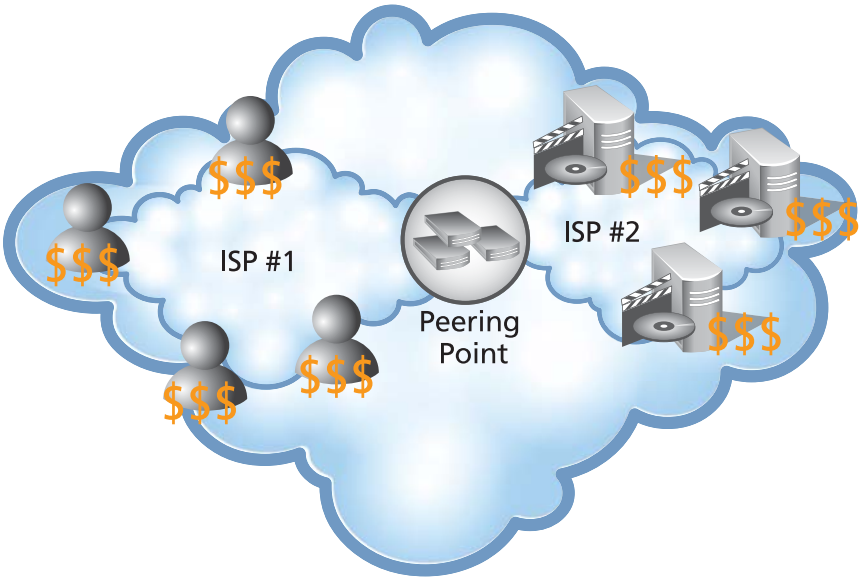

Finally, it is important to note that the large ISPs that supply connectivity to datacenters only deliver approximately 25% of Internet traffic to end users. The balance of the traffic is delivered by smaller ISPs that reach other networks by paying these larger ISPs for transit services. Because these smaller ISPs are not monetizing the content being delivered, and because they are paying upstream carriers for transit services, they usually cannot afford to have sufficient excess capacity, and therefore are unable to deliver growing amounts of traffic coming from outside their network reliably to their end users. Again, there is simply no substitute for deploying servers within a network—serving content locally would greatly reduce the upstream bandwidth requirements and related costs.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Quality Is a Problem

As broadband connectivity becomes more ubiquitous for end users, expectations of a high-quality, on-demand experience increase. This is true not only for multimedia content, such as a movie, but for large file downloads, such as game software, as well. Centralized CDNs face an increasing challenge in delivering such a high quality experience, due in part to the so-called "fat file paradox"—why does data traveling at the speed of light often take such a long time to make it across the country?

"Fat files" containing feature length movies in HD or DVD formats typically contain many Gigabytes of data. When considering whether latency matters for the delivery of such large files, many think that it does not. After all, who cares about 100ms of latency when downloading a file that will take minutes or hours to consume? The reality, however, is that latency is extremely important since it ultimately governs the rate of throughput and thus the total time needed to download the file. This is due to the underlying network protocols and the way they govern the transfer of data. In particular, a server can only send a certain number of packets before it must pause to wait for an acknowledgement from the client that the packets have all been received. Thus, one can only send a limited amount of data in a single round trip, meaning that the throughput is limited by the round trip time, or latency. Earlier this decade, the last mile constrained bandwidth speeds. However, that is no longer the general case—the distance between server and client is now the bottleneck.

Packet loss further reduces the amount of data that can be transferred per round trip since the underlying protocols back off and send less data when packet loss is detected. Since packet loss tends to increase when packets must traverse long distances, this means that high latency paths are even more likely to experience low throughput. The impact of this phenomenon is illustrated in the table below, which contrasts the theoretically fastest possible download times for a 4GB file (e.g., the size of a standard DVD) based on the speed of light[6] with download times given common network conditions.

Centralized CDNs delivering from far away face download times that can be significantly slower than content delivered locally.

Distance Between Server and End User	"Speed of Light" Network Latency (RTT)	Theoretically Fastest Possible Download Time Using TCP	Typical Network Latency (RTT)	Typical Packet Loss Rate[7]	Typical Download Time with Specified Packet Loss Rates	Typical Throughput	Quality
Local (100 miles)	1.6 ms	6.4 minutes	4 ms	0.6%	18 minutes	29.63 Mbps	High Definition
Regional (1000 miles)	16 ms	64 minutes	50 ms	0.7%	4 hours	2.22 Mbps	DVD
Cross-Continent (3000 miles)	48 ms	3.2 hours	90 ms	1.0%	9 hours	0.99 Mbps	TV
Different Continent (6000 miles)	96 ms	6.4 hours	150 ms	1.4%	18.3 hours	0.49 Mbps	Thumbnail
Emerging Markets (BRIC)	96 ms	6.4 hours	300 ms	2.0%	1 day, 1 hour	0.36 Mbps	Thumbnail

Because of the centralized CDN's limited deployment footprint, servers are often far from end users. As such, distance-induced latency will ultimately limit throughput, meaning that overall quality will suffer. In addition, network congestion and capacity problems further impact throughput, and these problems, coupled with the greater distance between server and end user, create additional opportunities for packet loss to occur, further reducing quality. For a live stream, this will result in a poor quality stream, and for on-demand content, such as a movie download, it essentially removes the on-demand nature of the content, as the download will take longer than the time required to view the content. Ultimately, "quality" will be defined by end users using two simple criteria—does it look good, and is it on-demand/immediate?

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

The Distributed Approach to Content Delivery

What is the "Distributed Approach"?

Akamai takes a distributed approach to content delivery and, over the last eight years, has deployed a globally distributed network of content delivery servers. This network currently (September 2007) includes over 27,000 servers, deployed in over 1600 locations[8] in over 1,000 networks, and over 750 cities in 70 countries around the world. These servers are located within those networks—there is no Akamai-owned network backbone that connects the servers, nor are the servers solely reliant on peering or transit relationships to reach end-users.

Akamai's deployments vary in size and scope, from just a few servers at small local ISPs, to hundreds of servers in high-traffic networks. For the most part, Akamai serves content from within (i.e., "on-net") these thousand networks, directly to end users, or to end users of 2nd or 3rd tier ISPs through downstream transit connections when necessary. Akamai's distributed network deployment places Akamai servers within one network hop of 90% of Internet users.

The Distributed Approach Removes Scalability Limitations

To service growing end user bandwidth demands, broadband providers are spending heavily to build out their high-speed last-mile networks. As a data point, Verizon expects to spend $18 billion to reach 18 million homes with its FiOS consumer broadband service, and AT&T plans to spend $5.1 billion to reach 19 million homes with its U-Verse consumer broadband service[9].

Unfortunately, the aggregate bandwidth demanded by end users within ISPs will quickly overwhelm existing uplink capacity and upstream peering points, which makes the ability to deliver content from within edge networks critical. Akamai's goal is to be able to support 100 Tbps of traffic. This could be done by deploying 20 servers (each capable of delivering 1 Gbps) in each of 5,000 locations within edge networks. Additional capacity can be added by deploying into PCs and set-top boxes. Ultimately, a distributed server deployment into thousands of locations means that Akamai can achieve the 100 Tbps goal, whereas the centralized model, with dozens of locations, cannot.

The Distributed Approach Solves the Quality Problem

Akamai's distributed approach to content delivery largely avoids the fat file paradox, as content no longer has to travel cross-country, or overseas, to reach an end user. By delivering content to end users from nearby servers, round trip times are reduced to just a few milliseconds, enabling higher throughput and reducing the opportunity for packet loss.

As a practical example, consider the delivery of a console game file to a consumer in Korea, which has the highest level of broadband penetration in Asia[10]. By serving content in-country, Akamai can deliver a 100 MB file in approximately 27 seconds. In contrast, a centralized CDN relying on servers located on the West Coast of the United States would take upwards of 27 minutes to deliver that same file. Assuming that this game file was actually 1 GB in size, the difference grows significantly—the distributed approach can deliver the file in 4.5 minutes, while the centralized approach forces the user to wait just over 4.5 hours. Even if the centralized CDN were delivering the file to Korea from servers in Singapore, those download times would be cut in half but are still significantly higher than those seen with the distributed model. Although we are talking about a game file download here, if this were a video download scenario, the distributed model would be able to support an HD-quality experience, while the centralized approach would result in a TV-quality experience at best.

Reducing round trip times from tens or hundreds of milliseconds to just a few milliseconds means that higher bitrates can also be used for the encoding of media, such as movies. As such, the higher throughput rates enabled by the distributed approach enable content providers to offer higher quality products (such as DVD or HD videos) to their consumers.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

As noted earlier, "quality" will be defined by end users using two simple criteria—does it look good, and is it on-demand/immediate? Akamai's distributed approach to content delivery is designed to ensure that video content will look good now and in the future, as encoding rates increase, and to ensure the immediacy of on-demand content, removing the multi-hour waits for a video or file to be completely downloaded.

Peer-to-Peer (P2P) Systems

What Is P2P?

A peer-to-peer (or P2P) computer network relies primarily on the computing power and bandwidth of the participants in the network. Peer-to-peer networks generally do not have the notion of clients or servers, but only peer nodes that simultaneously function as both "clients" and "servers" to the other nodes on the network. This model can be thought of as the logical extreme of the distributed approach to content delivery. However, the P2P model also has a number of shortcomings that limit its utility for real-time content delivery, as we will describe in the sections that follow.

Can P2P Scale?

Taken to the logical theoretical extreme, P2P networks could be thought to be infinitely scalable, since each node that joins the network increases the capacity of the network. However, this theoretical extreme ignores the real-life constraint of available bandwidth at each peer (i.e. client) and the significant imbalance between uplink and downlink speeds on consumer connections. As an example, Comcast offer Boston-area consumers two packages: 6 Mbps download/384 Kbps upload[11], or 8 Mbps download/768 Kbps upload. Similarly, Verizon FiOS packages range from 5 Mbps download/2 Mbps upload to 30 Mbps download/5 Mbps upload[12].

Assuming that average end-users in a P2P network have broadband service similar to Comcast's basic service level, uplinks will get saturated 15x earlier than the downlinks, meaning that the number of nodes sharing content needs to be at least 15x greater than the number of nodes downloading it. If the number of downloaders is the same as the number of uploaders in the P2P system (as it often is by design), then this means that the average download capacity in a P2P system is limited by the average uplink capacity, which is typically too small to support even TV-quality traffic.

For live P2P streaming of high quality streams (encoded at high bit rates), the imbalance and saturation problem creates a significant challenge in effectively delivering a full high quality stream to all participants. The instantaneous available upload capacity ultimately becomes the constraint. While it is theoretically possible to serve 100% of the streams from the peers, one will sacrifice performance, which results in lower stream quality. To achieve better performance, one needs to sacrifice efficiencies—the amount of content served from peers. Ultimately, it means that the P2P model requires a hybrid approach, where a distributed CDN is used to "seed" files into the P2P network and act as a backup for the delivery of content when it is not available from peer nodes, or when peer nodes have insufficient capacity to distribute it across the network.

This is especially true in the "flash crowd" scenario for on-demand content, which occurs when there is a sharp increase in the number of instantaneous downloads that outstrips the supply of peers in the network—there are lots of downloaders, but not enough uploaders and not enough upload capacity. In this case, some of the content can come from peers, but a distributed CDN would be required to service the large majority of the requests.

P2P's Role in the Distributed Approach to Content Delivery

Over the last ten years, various protocols, networks, client programs and related companies have appeared on the peer-to-peer landscape. Some have gained an "underground" following, mostly trading on the illicit exchange of commercially available multimedia and software. None have demonstrated that a truly successful long-term business model can be built on P2P. As an example, in March 2006, VeriSign spent $62 Million to acquire Kontiki, a provider of peer-to-peer technology, in an effort to enter the broadband content delivery market.[13] However, in December 2006, they announced that they would be building out CDN infrastructure to support their Kontiki P2P technology.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Akamai believes that there is a role for P2P as part of a distributed approach to content delivery. For certain types of content, such as background downloads of large files, content providers may feel that quality and performance are not as critical as reducing costs for that download. Examples include the delivery of software programs, where a low-cost completed download is more important than the speed of the download. In this scenario, P2P-augmented content delivery would fulfill the content provider's requirements, with Akamai's distributed CDN serving a seed function and providing the required control mechanisms.

P2P Needs the Control Mechanisms of a Distributed CDN

Akamai's customers demand the quality of service and control that they receive with its massively scalable backend control system and global network of edge servers. For P2P to realize its potential for enterprises, Akamai believes that the only effective approach is for client-side software to leverage an existing massively scalable backend control system. By applying the control mechanisms of its distributed CDN to the P2P technology acquired from RedSwoosh, Akamai is able to apply unique mapping capabilities, enabling delivery of content from the optimal location.

In addition, incorporation of granular authentication capabilities and digital rights management are necessary to ensure that business rules specified by content owners can be enforced. This is especially important for distribution of media and software, where P2P systems have typically been used to gain illicit access to content. Furthermore, advanced features including media management and content syndication give content owners control over the delivery of their content by partners and affiliates.

Finally, reporting capabilities ensure that content providers continue to have access to the metrics that are important to their business, such as download completion and origin offload rates. Content providers can use these metrics to easily calculate ROIs against alternative solutions, such as distributing physical media or building out additional server infrastructure. Additionally, reports on such metrics as content popularity, user locations, and traffic over time enable content providers to track the effectiveness of marketing campaigns, plan future activities, and otherwise gain visibility into their extended e-business infrastructure.

Distributed Delivery Architecture Requirements

The obvious question from the examination of these architectural requirements is why aren't there more in-market examples of the distributed delivery architecture? The fundamental answer is that it would be extremely difficult to replicate what Akamai has achieved across these three areas. Thus, it will be extremely hard for other companies to create effective distributed delivery architectures.

Technology

As we have discussed earlier, it is critical to be at the edge of the network to deliver the best performance. To get there, one needs to have highly distributed, highly fault tolerant software. While the development and implementation of such software is difficult, the resulting distributed system is dramatically more fault tolerant, cost-effective and powerful, than centralized or standalone systems. In addition, efficiently organizing the interaction between multiple systems and subsystems is not easy, but once in place, it allows for near limitless scalability.[14]

One of the defining design principles of Akamai's network is the assumption that a significant and constantly changing number of component or subsystem failures are occurring at all times within the network. This means that, at all levels, the network is designed to be highly redundant, self-monitoring, and self-healing, incorporating many of the ideas used in recovery-oriented computing. To build such a system successfully, one must adhere to design principles around redundancy, reliable messaging and decentralized control. Having invested over two thousand man-years of research and development into this very task, Akamai has developed, tested and proven the technologies that make its vast network operate with extraordinary intelligence, efficiency and performance. Akamai has further demonstrated its leadership in this area with over 50 issued patents and approximately 60 pending patent applications, many focused on the concepts underlying the operation of a distributed computing system on the Internet for content delivery.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

In addition to content delivery, a significant investment has been made to monitor, report, and log customer usage of such a distributed platform. On any given day, Akamai is logging and storing 22 terabytes of data, generated by 200 billion end-user content requests. Real-time access to this data is provided to thousands of customers through a user-friendly portal.

While it would not be impossible for a competitor to replicate Akamai's technology, it would require a massive undertaking.

Customers

Content providers use Akamai's services because it guarantees improved performance and unmatched scalability for the delivery of their content. In order to achieve this improved performance and scalability, one needs to have a distributed content delivery network, with deployment within edge networks. As noted earlier, to get to the edge, one needs to have the technology that supports the distributed computing requirements. Additionally, to get the deployments within edge networks, one needs to have the customers (content providers) that have the content that end users on these networks want to consume. This "chicken/egg" issue is an important consideration—most centralized CDNs simply do not have the customer base to facilitate server deployments into edge networks and without these deployments, customers tend to choose Akamai's distributed CDN, rather than the centralized CDN of a competitor.

Networks

With massive amounts of bandwidth available at the last mile, Akamai's distributed approach is the only practical way to scale to support the demand generated by broadband end users. Because the Internet is comprised of over 13,000 individual Internet Service Providers, peering plays a critical role in moving traffic between them, as discussed earlier in this paper. By delivering content from servers deployed on-net, Akamai avoids the congestion and resultant limits to scalability introduced by upstream peering points. Contrary to popular belief, the distribution of Internet content delivered to end users does not reside with a few large players. In fact, there is a very long tail of providers to ensure that over 1 billion people in 75+ countries have access to the Internet. According to data collected by Akamai, the top 10 end-user access networks account for only 30% of Akamai's traffic, while extending that figure out to the top 1000 networks accounts for 90% of Akamai's traffic.

Because most of the leading content providers are Akamai customers, this has led to a symbiotic relationship with edge ISPs, as they invite Akamai server deployments into their networks to help them scale their upstream connections. It is important to note that, because many of these edge ISPs do not sell collocation services, other CDN providers are not able to replicate Akamai's distributed deployments. Said another way, it's hard to be in all of those networks because you can't buy your way in—those providers only want your servers in their networks if you have the content their subscribers want.

The distributed approach mandates building business relationships with as many of the thousands of edge ISPs as possible. This takes time and skill, as creating thousands of network relationships and then physically deploying and maintaining servers in these regional ISPs is a considerable challenge for anyone looking to solve Internet scaling issues. Akamai has nearly a decade of experience in the operational demands inherent in operating in the thousands of locations that comprise our globally distributed computing platform.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

Summary

The increasing amount of video content coming online has raised concerns about the Internet's ability to scale to support this burgeoning demand for rich content. In particular, the centralized model of delivery, offered by CDN providers such as Limelight Networks, Level 3, Internap, Panther Express, CDNetworks, and others, will face several key challenges outlined in this paper, including limited outbound datacenter capacity, peering constraints, and the throughput degradations caused by round trip latency (a.k.a. "the fat file paradox").

Akamai's approach to content delivery was implemented in anticipation of these scaling issues and relies on a distributed architecture to circumvent these limitations. Distribution through more than 1,000 service provider networks in thousands of locations worldwide offers delivery of digital assets at superior levels of quality and performance, with unmatched scalability. While TV/DVD quality video is acceptable to consumers today, going forward, the "does it look good?" metric will be defined by HD quality video, due to growing availability of broadband and the wider adoption of Internet-connected home media systems.

Akamai's distributed delivery approach bypasses peering constraints and reduces latency so that the true potential of broadband is released to consumers. While peer-to-peer holds promise as an ingredient of a distributed delivery system, it requires a CDN to function as a control mechanism to ensure correct mapping of end users to files, authentication for digital rights management, and enterprise-level reporting systems.

For nearly 10 years, Akamai has focused on the development of the world's leading distributed content delivery system—purpose-built to meet the needs of the industry's most discerning customers. This investment of time, money and effort is designed to scale to meet the demands of Internet delivery today and in the future.

SUBJECT TO REQUEST FOR CONFIDENTIAL TREATMENT PURSUANT TO 47 C.F.R. §§ 0.457 AND 0.459

About Akamai

Akamai® is the leading global service provider for accelerating content and business processes online. Thousands of organizations have formed trusted relationships with Akamai, improving their revenue and reducing costs by maximizing the performance of their online businesses. Leveraging the Akamai Edge Network, these organizations gain business advantage today, and have the foundation for the emerging Internet solutions of tomorrow. Akamai is "The Trusted Choice for Online Business."
For more information, visit **www.akamai.com**.

1. http://www.economist.com/printedition/displaystory.cfm?story_id=9249193
2. http://www.datacenterknowledge.com/archives/2007/Jan/23/some_isps_spurn_10_gige_video_peering.html
3. Norton, William. "The Evolution of the U.S. Internet Peering Ecosystem." 19 Nov 2003
4. Borland, John. "Network Feud Leads to Net Blackout." C-Net News. 5 Oct. 2005
   http://news.com.com/Network+feud+leads+to+Net+blackout/2100-1038_3-5889592.html?tag=nefd.top
5. http://www.renesys.com/blog/2007/04/rough_day_at_cogent_part_1_cog.shtml
6. http://nanog.cluepon.net/index.php/Traceroute_latency#Propagation_Delay
7. As derived by empirical data from Akamai's network monitoring systems
8. A location is a (datacenter, network provider) pair
9. http://arstechnica.com/news.ars/post/20061205-8356.html
10. http://www.oecd.org/document/7/0,2340,en_2649_34223_38446855_1_1_1_1,00.html
11. http://www.comcast.com/shop/buyflow/default.ashx?FormName=ProductDetails&ProductID=20185
12. http://www22.verizon.com/content/consumerfios/packages+and+prices/packages+and+prices.htm
13. http://news.com.com/2100-1030_3-6048990.html
14. Experience with some Principles for Building an Internet-Scale Reliable System
    Mike Afergan, Joel Wein, and Amy LaMeyer (WORLDS '05, San Francisco, CA, December 2005.)

©2008 Akamai Technologies, Inc. All Rights Reserved. Reproduction in whole or in part in any form or medium without express written permission is prohibited. Akamai and the Akamai wave logo are registered trademarks. Other trademarks contained herein are the property of their respective owners. Akamai believes that the information in this publication is accurate as of its publication date; such information is subject to change without notice.

AKAMWP-0108

Akamai Technologies, Inc.
U.S. Headquarters
8 Cambridge Center, Cambridge, MA 02142
Tel 617.444.3000
Fax 617.444.3001
U.S. toll-free 877.4AKAMAI
(877.425.2624)

Akamai Technologies GmbH
Park Village, Betastrasse 10 b
D-85774 Unterföhring, Germany
Tel +49 89 94006.0

www.akamai.com