



## The Never-Ending Rush Hour

Internet Traffic Growth Requires  
Continual Investment in Capacity  
and Innovation in Network  
Management

Prepared by Jason Kowal  
August 9, 2007

Analysys Research Ltd.  
919 18<sup>th</sup> Street NW, Suite 220  
Washington, DC 20006  
Tel: +1 202 331 3082  
Fax: +1 202 331 3083  
[research@analysys.com](mailto:research@analysys.com)  
[www.analysys.com](http://www.analysys.com)

# The Never-Ending Rush Hour

Internet Traffic Growth Requires Continual Investment in Capacity and Innovation in  
Network Management

## Contents

|     |  |    |
|-----|--|----|
| 0   | Executive summary                            | 1  |
| 1   | A Primer on Internet Traffic                 | 4  |
| 1.1 | Origins of the Internet Traffic Explosion    | 4  |
| 1.2 | Founding Principles of Internet Architecture | 6  |
| 1.3 | How Networks Handle Traffic Today            | 7  |
| 2   | The State of Internet Traffic Today          | 9  |
| 2.1 | Internet Traffic Drivers                     | 9  |
| 2.2 | New Applications Taxing the Internet         | 10 |
| 3   | Keeping the Traffic Flowing                  | 14 |
| 3.1 | Investment in Capacity is the First Step     | 14 |
| 3.2 | Smarter Networks                             | 15 |
| 4   | Policy Considerations                        | 16 |

The author wishes to acknowledge and thank AT&T for their support for this project. The views expressed herein are the author's alone.

## 0 Executive summary

---

*“The Internet represents one of the most successful examples of the benefits of sustained investment and commitment to research and development of information infrastructure... [I]t started as the creation of a small band of dedicated researchers, and has grown to be a commercial success with billions of dollars of annual investment... At the same time, the industry struggles to find the economic rationale for the large investment needed for the future growth, for example to upgrade residential access to a more suitable technology. If the Internet stumbles, it will not be because we lack for technology, vision, or motivation. It will be because we cannot set a direction and march collectively into the future.”*

From “A Brief History of the Internet,” by Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff, Internet Society, last updated April 5, 2007 <http://www.isoc.org/internet/history/brief.shtml>

---

### **Rapid growth is standard operating procedure on the Internet**

Throughout the history of the Internet, dealing with very rapid growth has become standard operating procedure for network managers, engineers and planners around the world. The companies which operate the Internet’s core networks, or backbones, have managed consistently robust and sometimes incredible traffic growth by carefully managing available resources, augmenting network capacity, directing traffic flow, and taking advantage of technological advances to expand effective capacity. Although this process is now accepted as the norm, it should not be taken for granted.

In order to cope with traffic growth, networks must be continuously monitored and prepared for the next upgrade. Service providers, software companies and network vendors must constantly invent new ways to meet enhanced expectations for speed, security, and accountability.

Network operators apply three primary methods to cope with the perpetual growth and upgrade cycle:

- Continuous monitoring of traffic flows to track network performance and prepare for the next stage of network growth
- Planning for and investing in additional capacity long before existing networks are over-utilized
- Applying intelligent traffic management systems and practices to make the most of existing resources (for example, many network managers have suggested that providing differentiated classes of services for delay-sensitive traffic will be essential for meeting new capacity demands)

### **Online video will require a new wave of network upgrades**

Today the Internet arrives at a new threshold, driven primarily by online video. Little over a year ago, large scale online video distribution was still largely an experiment with mixed results. Now web sites like YouTube – as well as dozens of imitators – are flooding the Internet with over 100 million user-generated videos a day.<sup>1</sup> Combine this with the video offerings of traditional media outlets like Disney as well as myriad other new video services and the traffic is piling up rapidly.

Despite the current advancements in online video distribution, the real challenge – delivering real-time TV-quality video over the Internet to a large number of simultaneous viewers – has yet to arrive. To handle this level of traffic increase, backbone operators will need to reach a new plateau in network scale and traffic management capability. Even if they do, there is still the potential for a traffic jam in the “last mile” of the Internet, which provides access from subscriber homes and buildings. Most existing broadband access technologies today can not reliably deliver these new bandwidth-intensive applications at the quality expected by consumers. Additional capacity and routing intelligence closer to or even onto the end user premises will be a prerequisite. While major operators in the US and abroad are actively rolling out the next wave of ultra-broadband (as well as TV

---

<sup>1</sup> See “YouTube hits 100m videos per day,” BBC, July 17, 2006

services of their own), the process requires significant and sustained levels of investment associated with high levels of risk.

### **Policy decisions need to encourage broadband investment**

Given the continuous need for expanding network capacity and intelligence as well as the recent explosion in online video, there is a role for policymakers in stimulating the next wave of Internet evolution:

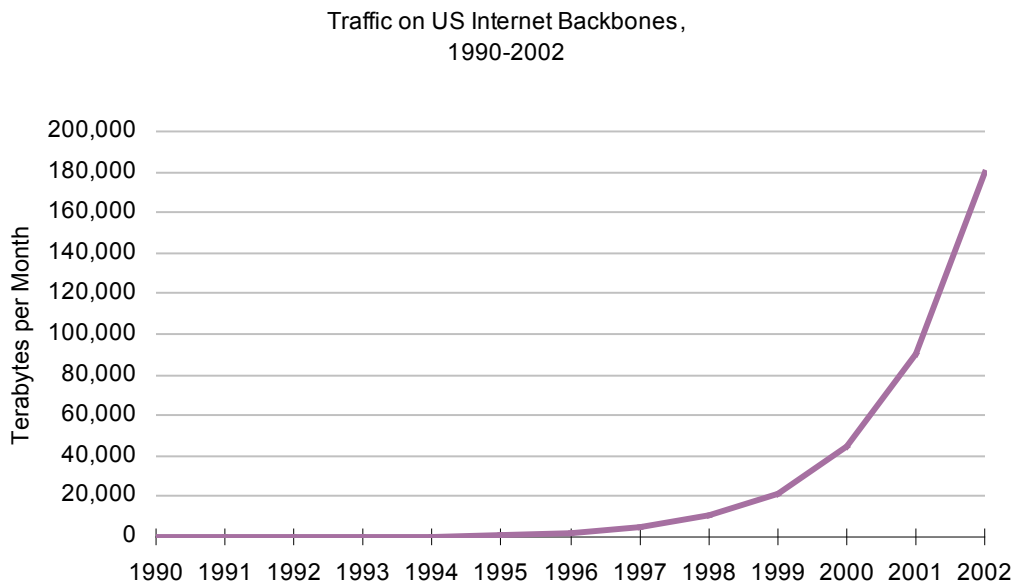
- First, careful attention should be paid to any new regulation which might adversely impact the business case for internet investment or set preference for one business model over another.
- Second, policymakers should consider methods to encourage network investment, particularly for the last mile of underserved locations.
- Third, regulators should not inhibit Internet service providers' flexibility to experiment with new traffic management technologies and strategies in order to deliver new services with the quality and security options that Internet users desire.

By following these simple principles, policymakers can help provide a framework in which network engineers can keep the Internet running reliably and create the additional network capabilities required to meet the needs of future applications.

# 1 A Primer on Internet Traffic

## 1.1 Origins of the Internet Traffic Explosion

Since the world wide web was first popularized in the early 1990s, Internet traffic has grown rapidly and without pause. During one period, between 1995 and 1996, many researchers believe traffic grew 1000% in a twelve month span (see Exhibit 1 below).<sup>2</sup> Several factors, including the already deployed but unused capacity, rapid investment in new infrastructure and the heroic efforts of network engineers (as well as the architecture of the Internet), combined to keep Internet traffic flowing.




---

**Exhibit 1:** *Traffic on US Internet Backbones, 1990-2002 [Source: Andrew Odlyzko, University of Minnesota Digital Technology Center]*

---

<sup>2</sup> See "Internet traffic growth: Sources and implications," Andrew M. Odlyzko, University of Minnesota Digital Technology Center. Note the figure above displays the average of high/low estimates developed by Dr. Odlyzko. Also, although no reliable data exists for the period 1995-1996, the figure assumes a steady growth curve for consistency.

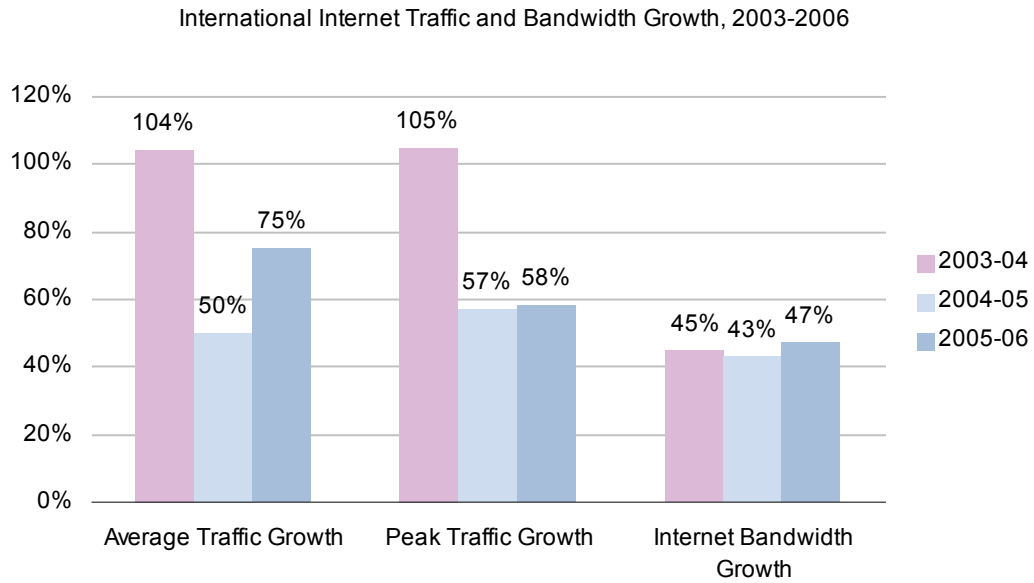
The incredible growth rates of the mid-1990s were exceptional because the Internet was experiencing its first wave of mass market popularity. Since then, however, global Internet traffic has continued to grow at an astounding pace, typically increasing 75% to 100% or more each year (see Exhibit 2 below).<sup>3</sup> This has put the Internet's architects into perpetual crisis mode: in order to cope with traffic growth, networks must be continuously monitored and prepared for the next upgrade. Furthermore, despite the fact that the Internet is still relying on protocols developed in the 1980s, service providers, software companies and network vendors must constantly invent new ways to meet enhanced expectations for speed, security, and accountability.<sup>4</sup>

Today the Internet arrives at a new threshold, driven by consumer expectations that an Internet connection should be able deliver video similar to what a television can do. However, currently deployed Internet backbone and access networks have not been engineered to deliver real-time, premium video services to large numbers of simultaneous viewers. Low resolution, short duration video intended for relatively small computer screens can work very well (e.g., YouTube). But this level of quality must evolve in order to match the viewing experience available over cable, satellite and other multichannel delivery platforms such as IPTV, which can provide programming in high definition on large screens. For the Internet to be capable of providing video at a similar level of quality and ubiquity, significant and sustained investments in both backbone and broadband service provider networks will be required.

---

<sup>3</sup> See "Global Internet Geography," TeleGeography Research

<sup>4</sup> See "The (un)Economic Internet," Claffy, Meinrath, Bradner in IEEE Spectrum, June 2007



**Exhibit 2:** *International Internet Traffic and Bandwidth Growth, 2003-2006* [Source: TeleGeography Research]

## 1.2 Founding Principles of Internet Architecture

Depending on how you measure it, the Internet is between three and five decades old. Initially spawned from a publicly-funded US government project, the Internet first evolved from a regional networking experiment in the 1960s to a national apparatus for sharing computing resources and email. Eventually the network began to add capacity and links that interconnected universities as well as private corporations. In the early 1990s a new method for navigating Internet resources began to gather momentum: the world wide web. At this point the core US Internet backbone was still run by NSFNET, under the auspices of the US National Science Foundation.<sup>5</sup>

<sup>5</sup> See "A Brief History of NSF and the Internet," National Science Foundation, [http://www.nsf.gov/news/special\\_reports/cyber/internet.jsp](http://www.nsf.gov/news/special_reports/cyber/internet.jsp)

When the decision was made in 1993 to privatize the backbone function performed by NSFNET, private companies were left to establish commercial arrangements for interconnection largely on their own. Although they started off with very simple arrangements, ultimately two primary ISP interconnection models emerged: peering and transit. Peering happens between two networks (usually exchanging equivalent volumes of traffic) which terminate traffic for the other without cash payments being made for the use of the network. Transit service, on the other hand, is a purchased service that provides the buying party with access to Internet addresses on the seller's network as well as delivery of packets to and receipt of packets from the rest of the Internet. There are many variations on these models which have evolved over time, but the basic principles have worked well without much government intervention for over a decade.<sup>6</sup>

### 1.3 How Networks Handle Traffic Today

Today the Internet functions primarily based on the technological and financial cooperation of thousands of large and small ISPs around the world, with very little government involvement. A few hundred large ISPs operate international backbones; thousands more operate domestic infrastructure only. Only a handful run large multi-country networks.

From the average user's perspective, email and web traffic moves transparently across the planet. This apparent simplicity masks an incredibly complex set of interdependencies among ISPs, content creators and other agencies around the world. One mouse click on a web page can generate a data connection that traverses 10 to 20 routing computers ("routers") which link together different ISP networks and ultimately deliver the content requested by the user. Some of the routers are likely to be managed by the same ISP, but usually multiple networks are involved in the handoff of packets between content server and end-user. The Internet only appears simple to users because ISPs are continuously engaged in complex negotiations on how they will transmit each other's traffic.

In addition to the core backbones which haul Internet traffic, a separate layer of Content Delivery Networks (CDNs) and other caching systems has sprung up to help web pages

---

<sup>6</sup> For more information on the evolution of peering and transit, see "The Digital Handshake: Connecting Internet Backbones," by Michael Kende, formerly Director of Internet Policy Analysis at the FCC and now Principal Consultant at Analysys

and other files get to their destination more rapidly and with more predictable performance. Companies such as Akamai and Limelight, which specialize in pre-distributing and storing copies of content closer to end users, sell directly to content producers like Apple and Disney that want to ensure that their web pages, music files and videos load quickly. CDNs emerged years ago to overcome the variable end-to-end performance of the Internet: if speed and quality are absolutely critical, the Internet's best-effort routing system often isn't good enough.

Another method content creators use to shorten the path from user to content is multi-homing" on multiple networks. By purchasing direct connections to multiple ISPs in different locations, routers can redirect traffic to the less congested network so that the end-user's requested data will be travelling with less likelihood of delay. Some very large content creators even run their own private backbone networks for carrying only their content to nodes close to the end user – in effect running their own CDNs.<sup>7</sup>

Although adding network capacity is an essential part of dealing with traffic growth, networks also have to make the most of what they have in between upgrade cycles. No network can be engineered to be completely free of congestion on every link all the time. While the congestion that occurs may be for brief instances, it nevertheless causes variation in the delivery rate and even the order of packet arrival at the destination. This situation is not problematic for non-real-time low intensity applications such as e-mail and web-browsing. The software in end-user devices can adapt for these small variations. Increasingly, internet applications demand much more capacity and are less tolerant, if not intolerant, of variable network performance. Voice over Internet Protocol (VoIP), networked multi-user gaming and video teleconferencing are examples of very demanding emerging internet applications.

From the earliest days of the Internet, network managers have used various tools to manage the flow of traffic within the backbone and in the last mile. This has been essential for avoiding traffic jams at peak periods of usage and also important for giving priority to time sensitive applications (see Section 3.2 below for more on traffic management). These emerging applications require even more sophisticated network management techniques.

---

<sup>7</sup>

See "The Next Wave of Massive Disruption to the U.S. Peering System," Bill Norton, Equinix

## 2 The State of Internet Traffic Today

### 2.1 Internet Traffic Drivers

Both in the US and around the world, Internet traffic has grown for a sustained period due to a combination of factors:

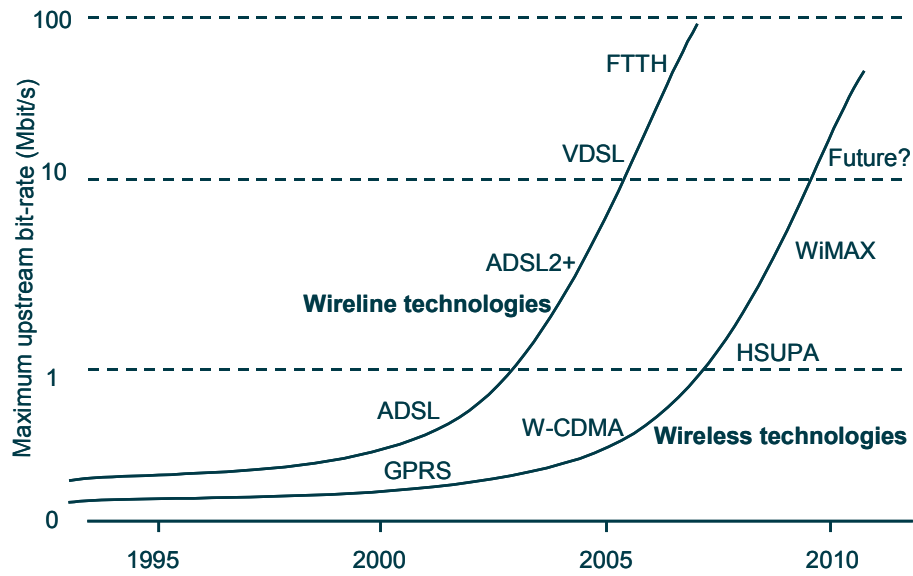
- Broadband subscriber penetration and usage has increased<sup>8 9</sup>
- End-user access speeds have increased (see Exhibit 3 below)
- Backbone network speed and router capacities have increased
- IP access and transit prices have consistently decreased on a per-Mbps basis
- PC processor and storage capabilities have increased
- The proliferation of Internet content
- Generally laissez-faire regulatory treatment
- Access to capital markets at all levels of the value chain

These drivers are all part of a complex supply-demand relationship which plays out at different rates on each network operator's backbone, depending on user demand, capacity deployments, and regional factors. On a global scale, international traffic has picked up pace over the last two years following a brief slowdown between 2004 and 2005 (see again Exhibit 2). Over the last few years, international traffic grew most rapidly on routes into Southern Asia, Eastern Europe, Latin America, and the Middle East, where markets are opening to competition and large populations are seeing broadband for the first time.

---

<sup>8</sup> At the end of March 2007, PointTopic reported that global broadband subscribers reached 298 million, with over 60 million in the United States.

<sup>9</sup> According to the Pew Internet & Life Project, at March 2006 42% of Americans had broadband access, 40% more than the prior year and double the prior year growth; penetration has reached 73%.



**Exhibit 3:** Upstream Rates for Fixed and Mobile Technologies [Source: Analysys]

There is more to the traffic equation than end-user access speeds, however. The rate of traffic growth is both a function of end-user demand as well as content supply and backbone capacity. Although an end-user may have a brand new PC and a broadband Internet connection, the rate of network utilization (another word for traffic) is ultimately determined by the nature of the content placed on the Internet and the combined capacity of different networks and servers to handle the movement of traffic between the source and destination. This is why YouTube didn't exist five years ago. Now, with more broadband connections and faster more intelligent networks, low resolution online video can be reliably streamed to most devices.

## 2.2 New Applications Taxing the Internet

Online video, whether it be streaming on-demand, downloaded in the background or live video chat, is nothing new. However the combined quantity and quality demands of video content has reached a critical mass, passing a quality threshold necessary for mass adoption. According to a March 2007 comScore study, 123 million people in the US –

70% of Internet users – viewed 123 billion videos online, with an average length of 2.6 minutes.<sup>10</sup>

The volume of traffic alone presents a challenge, but so does the unpredictable *peak* of traffic. Known as the “flash crowd” effect,<sup>11</sup> popular video content can receive massive attention from a worldwide audience, often with no notice. As a result, enormous volumes of traffic may be sent from a single point (the server storing the video content) rather than through a distributed content delivery network. There is little likelihood that the network operator has notification that a particular video clip on a server connected to its network will capture the attention of worldwide audiences. The traffic pattern may be so unusual that intermediate network capacity is temporarily strained. Thus video, in particular, presents new challenges for network operators.

Some examples of the most popular new online video services include:

- Streaming User-Generated Video
  - YouTube – Serving 100 million videos per day, with viewers spending an average of 17 minutes per visit on the site.<sup>12</sup> According to Nielsen/NetRatings, YouTube serves 20 million unique visitors per month.
  - VideoEgg – Supplies wholesale video services to 70 social network sites, serving up 15 million videos per day. VideoEgg’s CEO, Matt Sanchez, estimates that based on current demand levels, video requests could triple by year end 2007.<sup>13</sup>

---

<sup>10</sup> See “Primetime U.S. Video Streaming Activity Occurs on Weekdays Between 5-8 P.M.,” comScore, <http://www.comscore.com/press/release.asp?press=1264>

<sup>11</sup> See “The Next Wave of Massive Disruption to the U.S. Peering System,” Bill Norton, Equinix

<sup>12</sup> See “YouTube hits 100m videos per day,” BBC, July 17, 2006

<sup>13</sup> See “Telecom: Back from the Dead,” Business Week, June 25, 2007

- Streaming Television
  - Disney/ABC – In 2005 Disney struck a deal with Apple to distribute downloadable episodes of a few popular shows. Just over a year later Disney began streaming episodes from its web sites, along with commercials. Each episode of “Desperate Housewives” occupies an estimated 210 megabytes.<sup>14</sup>
  - Joost – The founders of KaZaA and Skype have launched a new TV-over-Internet service using P2P “swarmcasting” to deliver streaming video at a rate of approximately 350 megabytes per hour. Although the service was still in beta testing at this writing, major entertainment firms such as Sony, Viacom, and Turner Broadcasting have signed deals to provide content.<sup>15</sup>
  
- Premium Video Downloads
  - iTunes Store – In September 2006 Apple released iTunes 7 along with the capability to download feature length movies (for a fee) and NFL game highlights (for free). Disney sold 125,000 videos in the first week. According to Sandvine, which sells traffic management systems, peak traffic increased 140% among iTunes users in North America.<sup>16</sup> This traffic came on top of demand for the 220 TV shows available for download through iTunes. Both the movie and TV content available through iTunes is relatively low quality, designed for playback on small computer screens or iPods.
  - Wal-Mart – In 2007 Wal-Mart, the largest distributor of DVDs in the US, began beta testing a video download service for “PC Format” movies.

---

<sup>14</sup> See “Peering into the Future,” I, Cringley, PBS, March 2, 2006

<sup>15</sup> See “Overdoing it?” Economist, June 7, 2007

<sup>16</sup> See “iTunes: Videos Drive Downloads,” Sandvine Inc., [http://www.sandvine.com/solutions/itunesreport\\_1.asp](http://www.sandvine.com/solutions/itunesreport_1.asp)

Wal-Mart estimates download times for these 1.5 gigabyte files of 90 minutes over DSL and 30 minutes over cable modem.<sup>17</sup>

- P2P Video Downloads
  - BitTorrent – According to reports from traffic management system vendors CacheLogic, Sandvine, and Ellacoya networks, somewhere between 40% and 60% of consumer Internet traffic originates with P2P software, which allows users to share audio and video files via their local hard drives (see Exhibit 4). As of October 2006, 6 million US Households had downloaded at least one digital video file (10Mb or larger) from a P2P service.<sup>18</sup> Despite its common association with illicit content distribution, the most popular P2P software, BitTorrent, has now become a platform for legitimate paid “PC Format” movies.

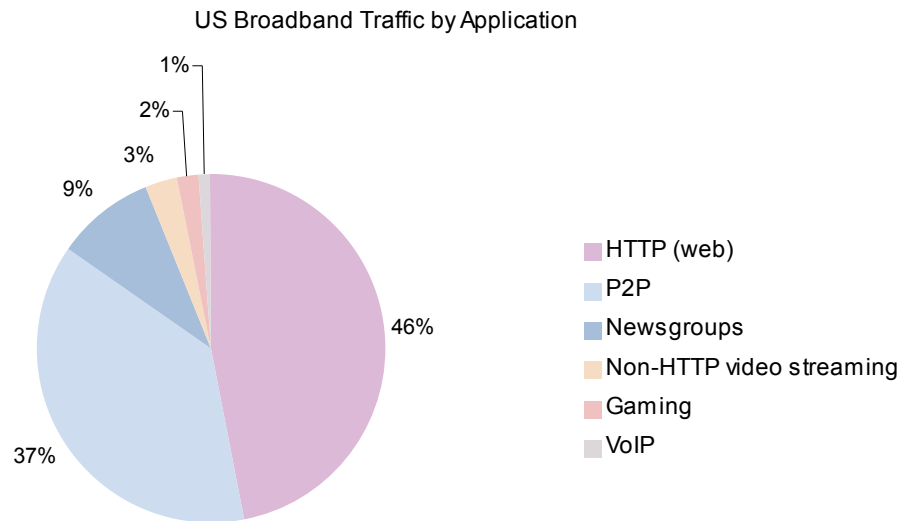
In addition to the easily understood impact of large video file downloads, the popularity of social networking sites has created unexpected problems for network infrastructure. Called the “MySpace Effect,” popular social networking sites are placing an unanticipated stress on the world’s DNS servers by linking to dozens of files outside the MySpace realm. According to Keith Oborn, network systems product architect at Virgin Media, “Because of the way MySpace pages are structured, a single page can generate hundreds of DNS queries.”<sup>19</sup> This increasing load has been augmented by anti-spam systems which check the validity of email domains on the never-ending flow of unsolicited email.

---

<sup>17</sup> See <http://mediadownloads.walmart.com>

<sup>18</sup> See “Peer-to-Peer Digital Video Downloading Outpacing Legal Alternatives Five to One,” NPD Group, December 20, 2006

<sup>19</sup> See “How MySpace Is Hurting Your Network,” PC World, June 24, 2007



**Exhibit 4:** US Broadband Traffic by Application, 2007 [Source: Ellacoya Networks]

### 3 Keeping the Traffic Flowing

#### 3.1 Investment in Capacity is the First Step

Because traffic has grown so quickly, a bystander might wonder how the Internet avoids traffic jams. The most direct way for an ISP to handle growing traffic volumes is to add capacity where it is most required. For a large ISP this typically involves:

- Activating already deployed dark fiber, upgrading the transmission electronics on existing facilities to higher line rates, or laying new fiber where needed,
- Connecting the circuits to very high speed routers which may also need upgrading to accommodate higher port speeds,
- Dedicating the circuits to the task of carrying Internet traffic, and
- Expanding the capacity of end user access lines where needed.

For smaller ISPs the process is similar; existing network connections must be upgraded and, more than likely, additional IP transit service must be obtained.

Although Internet capacity is upgraded to meet new traffic requirements, the increase in capacity and capabilities also stimulates new traffic growth as well. As described above, aggregate increases in available network capacity supply can also create demand in the network by making new types of content delivery possible. This occurs because there is less contention for resources on end-to-end connection. As a result, Internet traffic growth typically moves in relation to the growth of capacity. However the relationship does typically move in lock-step: in some years traffic grows faster than new capacity is deployed. According to TeleGeography, in 2006 the growth in average traffic level (75%) outpaced the growth of capacity (47%) on the world's international Internet backbones for the third consecutive year.<sup>20</sup> Significantly, this data only captures the early stages of surging video traffic, which is putting new strains on conventional capacity upgrade processes.

### 3.2 Smarter Networks

Going back to the foundations of the modern Internet, John Postel's original 1981 design for Internet Protocol packets contained the ability to specify quality of service (QoS). This later evolved into a mechanism called Differentiated Services (DiffServ), which allows additional categories of service to be handled (i.e., up to 64 rather than 8). Although DiffServ has been deployed on many corporate and ISP networks, it has never been widely applied as a method for achieving QoS across the public Internet due to the complexities of consistently defining classes and "allocating" performance margins to the potentially numerous networks employed in an end-to-end connection.<sup>21</sup>

Some network engineers believe that the cost of building more elaborate quality management systems could be well worth it: A recent study found that if ISPs deployed differentiated class of service for delay-sensitive traffic, the network could require 60%

---

<sup>20</sup> See "Global Internet Geography," TeleGeography Research

<sup>21</sup> See "The (un)Economic Internet," Claffy, Meinrath, Bradner in IEEE Spectrum, June 2007

less capacity at relatively moderate traffic levels than would be required to deliver the same Class of Service (CoS) with the typical “best effort” traffic delivery system. At higher traffic levels, there could be 100 percent or more difference in capacity requirements.<sup>22</sup> In other words, the differentiated services model in this study would deliver substantially more bang for the buck.

Despite the long history of development in the technology enabling QoS, operational issues remain that must be resolved for QoS to work well across multiple networks. Primary among these are:

- The CoS “marking” is created by the software at the ends of the network, requiring the network operators in the middle to interpret the mark properly. For example, a content provider applying a CoS mark of “4” may have entirely different handling expectations compared to a different content provider using the same mark.
- The end-to-end performance expectation logically dictates the CoS mark applied, but such service level assurance must be delivered across multiple networks. Until there is a means to allocate allowances for less than perfect performance to each of the participating networks in end-to-end service assurance, CoS is not realistic (except on single network connections).

While these and other issues are solvable with industry collaboration and additional investment, such resolution may not occur until regulatory uncertainty with respect to QoS systems is removed.

## 4 Policy Considerations

Throughout the history of the Internet, dealing with very rapid growth has become standard operating procedure for network managers, engineers and planners around the world. The companies which operate the Internet’s backbones and access networks have managed consistently robust and sometimes incredible traffic growth by carefully managing

---

<sup>22</sup> See “Value of Supporting Class-of-Service in IP Backbones,” M. Yuksel, K.K. Ramakrishnan, S. Kalyanaraman, J.D. Houle, R. Sathvani, June 2007

available resources and augmenting network capacity when needed. Although this process is now accepted as the norm, it should not be taken for granted: to maintain their networks' health, service providers must continually invest in new capacity, hardware, and software while applying the latest technological innovations to manage traffic intelligently and efficiently.

Given the continuous need for capacity expansion and upgrades in all aspects of the Internet ecosystem and the recent explosion of exceptionally demanding applications such as online video, there is a role for policymakers in stimulating the next wave of Internet evolution:

- First, careful attention should be paid to any new regulation which might adversely impact the business case for internet investment or set preference for one business model over another.
- Second, policymakers should consider methods to encourage network investment, particularly for the last mile of underserved locations.
- Third, regulators should not inhibit Internet service providers' flexibility to experiment with new traffic management technologies and strategies in order to deliver new services with the quality and security options that Internet users desire.

By following these simple principles, policymakers can help provide a framework in which network engineers can keep the Internet running reliably and create the additional network capabilities required to meet the needs of future applications.

---

*“The Internet represents one of the most successful examples of the benefits of sustained investment and commitment to research and development of information infrastructure... [I]t started as the creation of a small band of dedicated researchers, and has grown to be a commercial success with billions of dollars of annual investment... At the same time, the industry struggles to find the economic rationale for the large investment needed for the future growth, for example to upgrade residential access to a more suitable technology. If the Internet stumbles, it will not be because we lack for*

*technology, vision, or motivation. It will be because we cannot set a direction and march collectively into the future.”*

From “A Brief History of the Internet,” by Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff, Internet Society, last updated April 5, 2007 <http://www.isoc.org/internet/history/brief.shtml>

---

### **About the Author**

Based in Washington, DC, Jason Kowal is responsible for Analysys Research operations in North America, including syndicated and custom research projects covering consumer, business, and wholesale communications markets. Prior to joining Analysys, Jason was President of TeleGeography, a US-based firm specializing in primary research on communications traffic and networks. During his 11 years leading the company, TeleGeography became the authoritative source of information and analysis on Internet traffic flows, backbone networks, and Internet access pricing.

# About Analysys

Analysys provides strategy and management consultancy, information services and start-up support throughout the telecommunications, IT and media sector. We help clients to understand and manage business risk, launch new market propositions and develop and implement ongoing and sustainable business improvement programmes. We have developed unparalleled industry knowledge, techniques and data sources over the last 20 years, giving us – and our clients – a head start in the search for competitive advantage.

Our services include strategy and business planning; investment appraisal, profitability analysis, corporate finance and venturing; public sector policy definition, procurement support and major project implementation; market sizing and forecasting; and litigation support. These are complemented by wide-ranging market intelligence which informs customers worldwide on the issues, the companies and the trends affecting the industry.

Major players at the forefront of the sector confidently base their strategic decisions on the advice and information that we provide. Our international client base includes: network operators and media groups; equipment manufacturers; regulators and policy makers; banks, investment and legal institutions; corporate users; regional development organisations and local government bodies.

Our consultants and analysts share a passion and enthusiasm for the sector and are committed to securing tangible results and performance enhancement. We have over 160 staff in offices in Cambridge, London, Madrid, Milan, Paris, Singapore and Washington DC.

Analysys Research is part of Analysys Mason Group Limited, owned by the international networking and IT services company Datatec Limited.

For further information about Analysys please visit our Web site at <http://www.analysys.com> or email [enquiries@analysys.com](mailto:enquiries@analysys.com).