DISCOVERY INSTITUTE

Estimating the Exaflood

—

The Impact of Video and Rich Media on the Internet

—

A "zettabyte" by 2015?

Bret Swanson[1] & George Gilder[2]

January 2008

Discovery Institute

Seattle, Washington

Estimating the Exaflood

EXECUTIVE SUMMARY

An upsurge of technological change and a rising tide of new forms of data are working a deep transformation of the Internet's capabilities and uses. In this third phase of Net evolution, network architectures and commercial business plans reflect the dominance of rich video and media traffic.

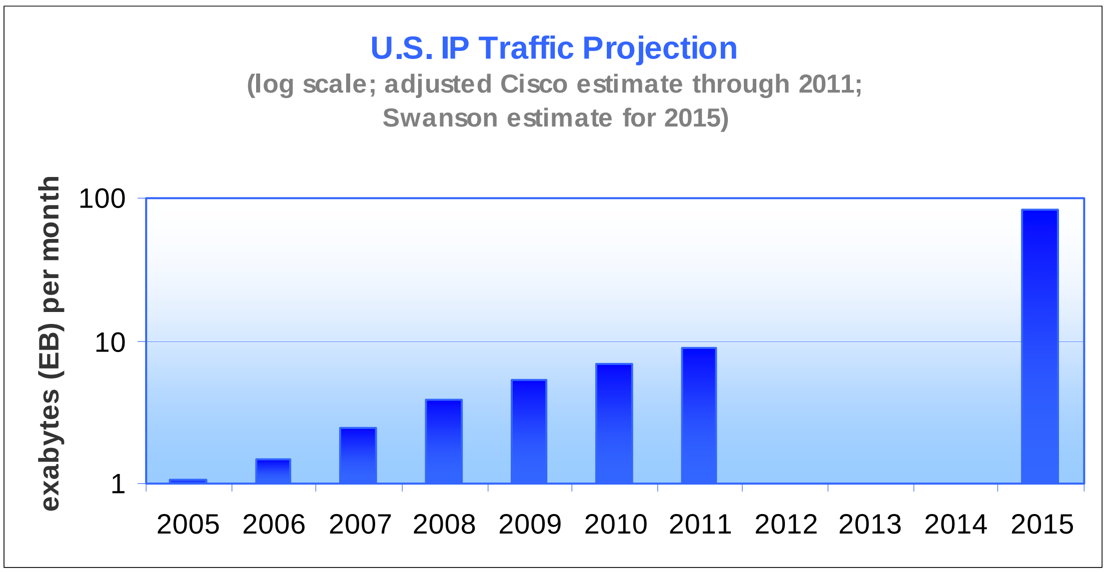

From YouTube, IPTV, and high-definition images, to "cloud computing" and ubiquitous mobile cameras—to 3D games, virtual worlds, and photorealistic telepresence—the new wave is swelling into an *exaflood* of Internet and IP traffic. An exabyte is 10 to the 18^{th}. We estimate that by 2015, U.S. IP traffic could reach an annual total of one zettabyte (10^{21} bytes), or one million million billion bytes.

We began using the term "exaflood" in 2001 to convey the vast gulf between the total traffic on the nation's local area networks, then 15 exabytes a month, and the thousandfold smaller flows across the Internet. We predicted then that the deployment of broadband networks would bring exafloods of data to the Net.

Today it is happening. We estimate that in the U.S. by 2015:

- movie downloads and P2P file sharing could be 100 exabytes
- video calling and virtual windows could generate 400 exabytes
- "cloud" computing and remote backup could total 50 exabytes
- Internet video, gaming, and virtual worlds could produce 200 exabytes
- non-Internet "IPTV" could reach 100 exabytes, and possibly much more
- business IP traffic will generate some 100 exabytes
- other applications (phone, Web, e-mail, photos, music) could be 50 exabytes

The U.S. Internet of 2015 will be at least 50 times larger than it was in 2006. Internet growth at these levels will require a dramatic expansion of bandwidth, storage, and traffic management capabilities in core, edge, metro, and access networks. A recent Nemertes Research study estimates that these changes will entail a total new investment of some $137 billion in the worldwide Internet infrastructure by 2010. In the U.S., currently lagging Asia, the total new network investments will exceed $100 billion by 2012.

Technology remains the key engine of U.S. economic growth and its competitive edge. Policies that encourage investment and innovation in our digital and communications sectors should be among America's highest national priorities.

Estimating the Exaflood

The Internet's Three Phases

New technologies are driving a deep transformation of the Internet's capabilities and uses. We are entering a new phase. The first phase of the Internet, starting with Arpanet in 1969, was a small research project that linked together a few, and then a few thousand, scientists. They exchanged rudimentary messages and data. In the mid-1990s the second broad phase delivered the Internet to the masses with e-mail, graphical browsers, and the World Wide Web.

Today, the third phase is underway. Video over the Net portends innumerable consumer and commercial possibilities. This new medium will change every realm of communication and content. The broadcast petabyte flows of radio and television will branch out into narrowcast, multicast, mobilecast, and everycast streams. With real-time transactions and collaborations, rich images, video, and interactive virtual worlds, the Net's current content of static text and pictures will swell to form exabyte rivers. We call this third phase of rich broadband content the Exaflood.[3]

At the time we began using this word in the *Gilder Technology Report* in 2001, we had only a vague notion of what it signifies. Ethernet inventor Robert Metcalfe had estimated that some 15 exabytes of data might be circulating in the nation's local area networks (LANs), waiting to be liberated. But what was an "exabyte"? It is 10 to the 18th bytes, an immensity best measured in LOCs. At 20 terabytes or 20 million megabytes, the LOC has found favor as a unit of measurement designating roughly the contents of the Library of Congress translated into digital form. Since a megabyte can hold the information in a book of some 400 pages, a LOC comprises about 20 million big books. An exabyte is 50 thousand LOCs, which comes to a trillion big books. The perhaps 15 exabytes unleashed onto the Net by broadband connections to homes and offices would add up to more than 15 trillion big books. Imagine a tower of tomes 200 million miles high, reaching twice as far as the sun.

This Exaflood is coming. However, it will only be possible with a vast new infrastructure to match the vastness of the coming digital deluge. Building this new infrastructure will be very expensive, likely requiring some $137 billion in global new investment over the next two years alone and at least $50 billion in the U.S. Technology remains the key engine of U.S. economic growth and its competitive edge. Consummating a true broadband Internet will depend on smart communications law and an investment friendly economy.

The second phase of the Net was chiefly enabled by two broad technical advances: (1) new computer modems at the edge of the network and dramatic advances in fiber-optic communications in the core of the network, both of which supplied unprecedented *physical* connectivity; and (2) new *logical* concepts like the HTTP-based World Wide Web and software applications like the browser, which made the Net accessible to the masses.

Today, the third phase is likewise being driven by a combination of advances in physical connectivity and software innovation. Today's residential cable modems now average more than 5 megabits per second, or 100 times faster than the 56-kilobit modems that mostly reigned at the outset of phase two. Many cable MSOs now offer 10- or even 15-megabit services. Meanwhile, the nation's telecom companies are building a new generation of fiber-optic networks to neighborhoods and homes that will reach tens of millions of consumers in the next few years. These networks will offer an *additional* factor of 20 capacity increase initially and are massively scalable for the future. On the software side, user-friendly self-publishing applications have given rise to millions of blogs and myriad social networking communities. Media players and Flash applications enable the easy creation and dissemination of rich visual content.

Prefixes —
Kilo – 10^{3}
Mega – 10^{6}
Giga – 10^{9}
Tera – 10^{12}
Peta – 10^{15}
Exa – 10^{18}
Zetta – 10^{21}

Three additional digital technologies are feeding the Exaflood—digital cameras, digital storage, and graphics processors (GPUs).

Digital cameras are now everywhere, but we forget that they only hit the mass market around the year 2000. Today, professional photographers have gone digital, and 500 million people take photos—and videos—with their phones. Once the domain of Hollywood studios and other elite professionals, the ability to create rich content and distribute it across the globe now devolves to every individual with a camera, computer, and communications link.

This dramatic increase in digital content creation requires a corresponding increase in **digital storage**. In 1991, $500 bought 100 megabytes worth of hard drive storage. Today, some $300 buys a terabyte drive. In a decade and a half, digital storage technology has thus advanced by a factor of more than 10,000. The pace of digital storage innovation is expected to continue or even slightly accelerate over the next decade.

Graphics processors from the likes of Nvidia and ATI (now part of AMD) are growing in power faster than traditional microprocessors, which were designed for the *serial* needs of the static computer. GPUs are inherently *parallel* and thus well suited to the video feeds and fiber speeds of the network. The new GPUs are taking 3D video and computer generated imaging (CGI) beyond the Hollywood studio and delivering it to the masses, first through game machines and soon through the Web and mobile devices. AMD’s recent purchase of ATI and Intel’s own internal graphics project called Larrabee show that the major microprocessor companies believe the *GPU* may soon rival the importance of the *CPU*.

Ubiquitous digital cameras, cheap digital storage, and powerful graphics processors are changing the *nature* of network traffic—moving from text and low-resolution graphics to visually rich, interactive, high-resolution images. The consultancy IDC estimates that in 2006 total worldwide digital information created was 161 exabytes. (This was the sum of new information generated, captured, or replicated but not necessarily stored or transmitted.) IDC predicts new digital content generated *annually* by 2010 will reach 988 exabytes.[4] Fast networks and new user-friendly software will unleash these huge troves of data and thus greatly increase the *volume* of network traffic.

A brief history of Internet traffic

In November 1993, John Markoff of *The New York Times* offered some early metrics on the nascent Net. "Call it a cautionary tale for the information age," wrote Markoff.

> The nation's increasingly popular data highway is beginning to groan under the load of rush-hour traffic.
>
> The problem, computer and telecommunications experts say, is that the computer network called the Internet, once a cozy community of a few thousand computer scientists, engineers and programmers who quietly and freely shared their on-line data, has suddenly been besieged by millions of newcomers. Anyone with a personal computer and a modem can easily and cheaply gain access to the Internet, a global web of thousands of computer networks.
>
> The Internet has been widely hailed as a here-and-now precursor to the interactive 500-channel cable television systems of the future. But the information highway metaphor is proving all too literal, as traffic jams become increasingly common.[5]

Markoff noted that the Net had grown to some 15 million users and that "the National Center for Supercomputer Applications in Champaign, Ill., a new service that answers requests to an electronic library called the World Wide Web, has seen the number of daily queries explode from almost 100,000 requests in June to almost 400,000 in October." Markoff even commented on the Net's most popular search tool, the Google of its day: "This year, information retrieved using a popular searching program called Gopher increased more than 400 percent, to almost 200 billion bytes a month…." Two hundred billion bytes may have seemed a large number at the time, but 200 gigabytes (GB) is about the size of today's average PC hard drive.

With the advent of e-mail and the Web browser, which marked the transition from phase one to phase two, we saw an almost 100-fold rise in U.S. Internet traffic between 1994 and 1996. After this enormous spike, when Internet activity grew near 10-fold per year for two years, traffic growth reverted to a level of about 100%, or a doubling each year, for the next several years. In the early 2000s, it appears the rate of U.S. traffic growth slowed somewhat to around 50-70% per year.

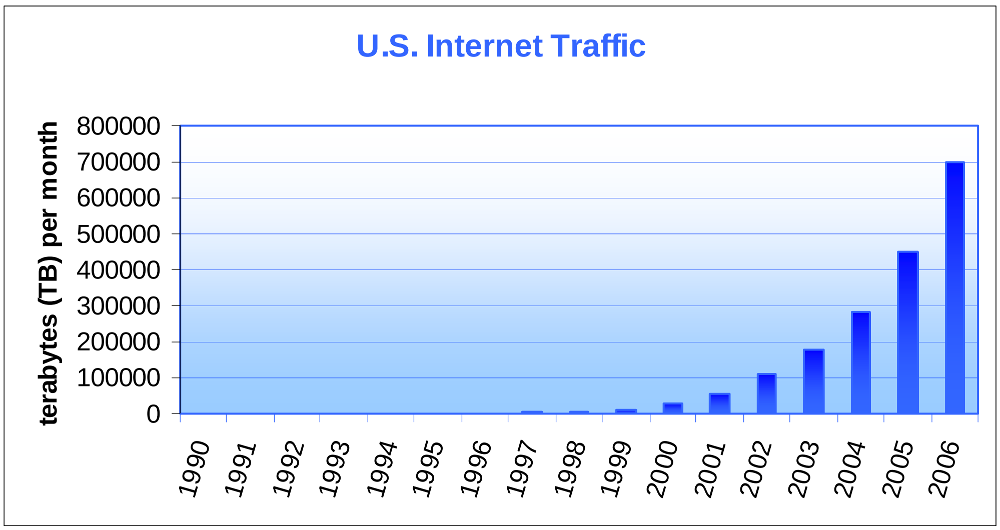

Figure 1

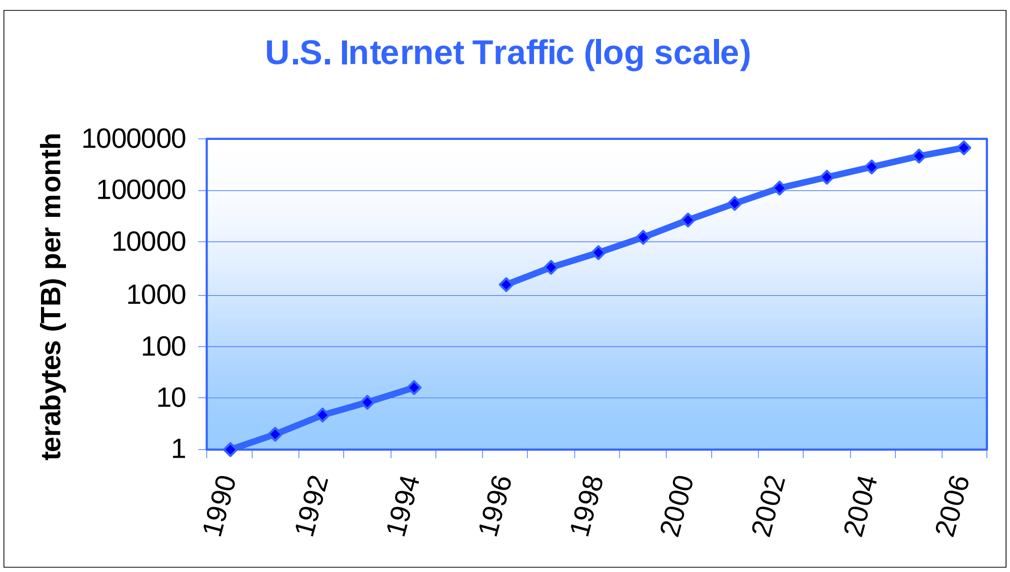

Figure 2

Over a decade, these growth rates produce large changes. U.S. Internet traffic in 1996 was about 1.5 petabytes (PB) per month, or about 18 PB per year. By the end of 2006 Internet traffic was probably 700 PB per month, yielding a rough run rate of 8.4 exabytes per year. Thus starting *after* the large spike of the mid-90s, U.S. traffic still grew almost 500-fold in the next 10 years.

To meet the traffic challenge of the Internet's second phase, we invented a host of ingenious new technologies, from wave division multiplexing (WDM) and the erbium doped fiber amplifier (EDFA) in the optical domain to copper DSL and coaxial cable modem technologies in the RF access domain. Without these miraculous optical and RF innovations and many tens of billions of dollars of network investment, we could never have accommodated the Web-e-mail-browser surge.

The future of Internet traffic

Professor Andrew Odlyzko of the University of Minnesota is among the world's leading experts on Internet traffic. Some have pointed to his scholarship—especially recent estimates of 50-60% per year U.S. traffic growth—in an effort to "debunk" the idea of video's large impact on the Net. Odlyzko's estimates of 50-60% U.S. traffic growth over the past few years appear to be correct.[6] But these growth rates (1) are still a remarkable phenomenon considering the ever larger base of traffic and (2) do not tell us what growth rates will be over the coming years, when the impact of the themes discussed in this paper come into full effect. Odlyzko acknowledges it is an open question. "There are huge potential sources of Internet traffic that already exist," Odlyzko writes, and "just a slight change in the velocity with which information circulates can have a large impact on Internet traffic."[7]

Cisco Systems made its own detailed projections in August of 2007. In "The Exabyte Era," Cisco concluded that "YouTube is just the beginning. Online video will experience three waves of growth."[8] The first wave is driven by YouTube and similar applications, which Cisco generalizes as Internet Video-to-PC. The second wave is Internet Video-to-

TV. Cisco believes this category, which includes IPTV, will overtake the first wave in traffic volume by 2009. The third (and potentially largest) wave is Internet Video Communication, or video calling and conferencing. Cisco's detailed projections, which are broken down by global region and by application, appear if anything to be conservative. For example, Cisco believes that the third wave of Video Communication will not begin to achieve popularity until 2015. Nevertheless, Cisco's estimates yield annual U.S. traffic in 2011 of more than 100 exabytes.

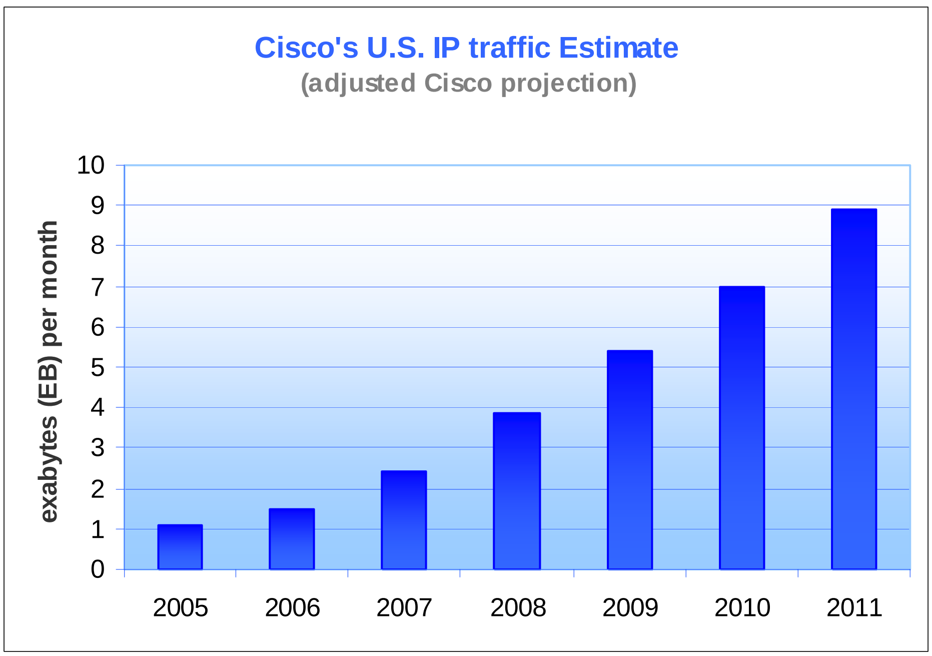

Figure 3

A new study by Nemertes Research approaches the Internet traffic question from a different angle.[9] Nemertes does not target particular consumer or business applications but employs estimates of network capacity and utilization rates to project data traffic flows. Nemertes believes that many new applications and innovations are beyond our ability to predict with any degree of precision. The one thing we know for sure, they say, is that we will find creative ways to use up a certain large portion of new bandwidth, just as we have learned to exploit or even "waste" abundant digital compute cycles and disk storage. Nemertes suggests that U.S. Internet traffic could reach some 50 exabytes per month by 2012, for an annual total of around 600 exabytes.

Any long-range forecast in the dynamic world of technology exposes itself to charges of "speculation." Yet Nemertes' approach is conceptually appealing for the very reason that most observers, unable to grasp the power of exponential change, have consistently underestimated the mid- to long-range pace of advance in the world of computation and communication. Professor Odlyzko believes Nemertes' projections of traffic (or what it calls "demand") are very aggressive but not beyond plausibility. We agree.

It is impossible to predict exactly when new applications that drive traffic will penetrate the market, how fast new broadband links will be built, and how fast they will be adopted by consumers and businesses. But a long list of Exaflood applications are about to break through across the globe and drive a new round of dramatic Internet growth. Some, like the three-year old YouTube, have already exploded on the Internet scene, and we have begun to measure their impact. Other essential innovations, like fiber-speed electronic nodes, are just now being deployed. And new network, hardware, software, and application ideas will continue to come online for decades to come.

The Exaflood

Here we list several major paradigms that comprise the Exaflood and describe the types of applications and phenomena that will drive new traffic growth. These trends are: Life After Television; the Global Sensorium; and LAN's End. After describing these trends, we then offer rough estimates of the new traffic that each might generate over the next eight to 10 years.

Life After Television

As first described in Gilder's 1990 book *Life After Television*, these applications all fit the broad paradigm of video moving out of the broadcast mode and into the narrowcast, multicast, and consumer-driven models of the Net. (Many of these video applications will be viewed on displays that look like "televisions," but in the Exaflood era displays will be multipurpose, networked devices that can accommodate a broad diversity of content types, sources, and distribution channels.)

- *YouTube*

By mid-2007, YouTube was streaming around 50 petabytes per month, or 600 petabytes (PB) per year. This was approximately ~7% of all U.S. Internet traffic. For another reference point, consider that all original broadcast and cable TV and radio content totals around 75 PB per year. YouTube streams that much data in 1.5 months. A Hi-Def YouTube would mean 12 exabytes per year, or almost as large as the entire U.S. Internet in 2007. YouTube and its many competitors are still in their infancy.

- *Video Conferencing*

Today we transmit close to 30 exabytes of telephone traffic globally each year. Gradually, a large portion of telephony will shift to video conversation: via high-end video conferencing in offices, mid-range video-calling from homes, and even on mobile phones. The Dick Tracy video watch is today a reality. By mid-2007, MSN Video Messenger was already generating 4 PB per month, or as much as the entire Internet in 1997. A move to video-phones would mean 300 exabytes—*at least*—or 30x the size of the existing U.S. Internet. Cisco's new HD Telepresence system requires a symmetrical 15 Mbps connection. A one-hour conference call would thus generate around 13.5 GB of data. Just 75 of these Telepresence calls would generate as much traffic as the entire Internet in 1990. In coming years, we can easily imagine high-definition always-on "Virtual Windows" that open permanent visual portals across the globe.

- *Amateur Video*

Amateurs capture some 10 exabytes of video each year. These videos will be stored on computers, shared with friends and family, and backed up in remote data storage facilities. Conversion of amateur video capture to HD would mean 100 exabytes per year, or 10 times today's annual U.S. Internet traffic.

- *Motion Pictures*

Amazon, Netflix, Blockbuster, and Apple are all in the movie download business, not to mention CinemaNow, MovieLink, and others. The cable TV operators are of course already the major player in residential movie watching—with basic, premium, and pay-per-view films. The telecom companies are now seeking to become cable's chief competitor in this entertainment space. One HD film (motion picture) is about 10 GB. With HD, NetFlix today would ship 5.8 exabytes of DVDs each year. HD movie downloads and streams could generate 100 exabytes per year, or around 10x today's U.S. Internet.

Online Gaming / Virtual Worlds

Graphics chips (GPUs) from Nvidia and ATI (now a unit of AMD) make online 3D gaming and virtual worlds a possibility for the first time. GPU performance is advancing several times faster than CPU performance, and the GPU's inherent parallelism is well suited not only for rendering video pixels onto displays (obviously) but also for some high-speed networking functions, as in systems offering GP GPUs (general purpose GPUs). Reflecting the rise of video detailed in this paper, Intel and AMD are putting ever greater emphasis on massively parallel GPUs and multi-core CPUs.

At the same time, new software advances, such as Otoy, a software language that directly addresses the GPU instruction set, makes possible real-time 3D rendering and mass- and peer-to-peer distribution of rich video. GoogleEarth is another early example of our ability to render either real or virtual worlds within the computer and over the Net. New Internet software will enable images and video to be embedded in any object on the Net—and to be played back on any device, even thin-clients like mobile phones.

One massively parallel HD virtual world with 1 million players could generate 100 PB per month—more than an exabyte a year—or one-tenth today's U.S. Internet. The virtual world Second Life now says it has more than 11 million "residents" (although the relevance of this top-line number is disputed) and reports that, since the autumn of 2007, it has consistently reached levels of more than 50,000 concurrent users. World of Warcraft reports some 8 million users. Webkinz, Gaia, and numerous other virtual worlds are rapidly gaining popularity. These games today are relatively crude and clunky but are about to evolve into full-motion, real-life, networked worlds. Some of them, based on Otoy, will be rendered not by a gaming box in your home but by a data center or Googleplex out in "the cloud."[10]

IPTV / Internet TV

The telecom companies are moving to IPTV, with AT&T the first mover, and the cable TV operators are likely to follow. Cable has already inched in the direction of IPTV with its popular pay-per-view and on-demand services. Unlike the tree-and-branch broadcast technology long used by cable TV, IPTV is a switched, interactive, customized service. Although IPTV is not necessarily an Internet service, television and the Internet will over time merge into something entirely new. It remains to be seen just what portion of "IPTV" content will actually traverse the Internet backbone or be stored at local hubs and cable head-ends. Regardless, *metro* bandwidth and *last mile* bandwidth must expand between 10 and 100 times to meet new IPTV capacity challenges. Although targeted single-cast and multi-cast content could in a narrow sense actually reduce the amount of data flowing to a particular end-point—such as a set-top box that today receives a continuous stream of up to a gigabit per second in cable TV content—these new IPTV and Internet video architectures could generate many times the traffic of broadcast at other nodes and trunks in the network. Joost, led by former top Cisco executive Mike Volpi, and a number of similar services offer free, advertising-based TV over the Net. The networks have also begun parallel strategies focusing on the Internet: NBC, for example, streamed 50 million of its shows over the Net in October 2007. Now NBC Universal and News Corp. are set to launch Hulu, a premium online video site that combines their considerable TV and film resources.

Because of the diversity of the new paradigm—where many more content sources and types travel over many more distribution channels to numerous storage caches and a multitude of varied end-point devices—bandwidth across the network will have to expand to accommodate very high peak rates anywhere on the Net: core, edge, metro, or access.

- *3D Video / Home Theater*

Today's HD video requires between 8 and 18 Mbps, depending on the codec. Next generation 3D video could require between 50 and 100 Mbps. HP Labs' "IMAX at Home" generates 4,096 x 2,304 pixels for a 16 x 9 ft. image. A two-hour movie in this format would be 250 GB (uncompressed). Ultra High Definition Video (UHDV), under development by Japan's NHK, could appear circa 2016. The UHDV format generates:

- 7,680 x 4,320 pixels
- 33 megapixels @ 60 frames per second
- or 16-32 times the pixels/sec of HDTV
- uncompressed two-hour movie of ~ 25 terabytes
- or an MPEG4 two-hour movie of ~ 360 gigabytes

Mere HD systems do not have sufficient resolution for the coming generations of displays. We will need UHDV for the ever-larger flat panel and laser projector displays coming to the mass home theater market. A decade hence, with UHDV, all the HD numbers we consider in this paper thus get multiplied by another 5-10x, which is 50-100x more than standard definition video. One hundred exabytes of HD video becomes a zettabyte of 3D video.

Thus, even as video applications and content proliferate in number, type, and popularity, the increasing picture quality, or data-density, of the videos will compound the effect on traffic growth. On the following page, compare and contrast the coming generations of video. High definition video has six times as many pixels as standard TV. ***Ultra*** HD video has 16 times the pixels of HD video.

HDTV [11] vs. Standard TV

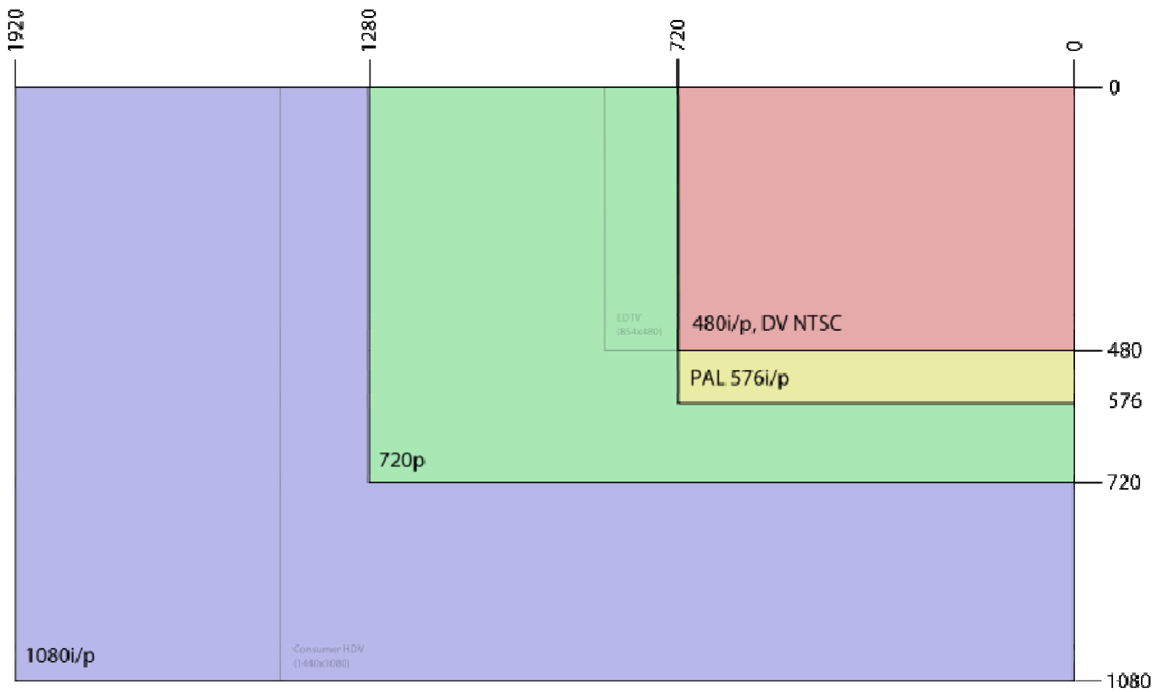

Figure 4

UltraHDV [12] vs. HDTV

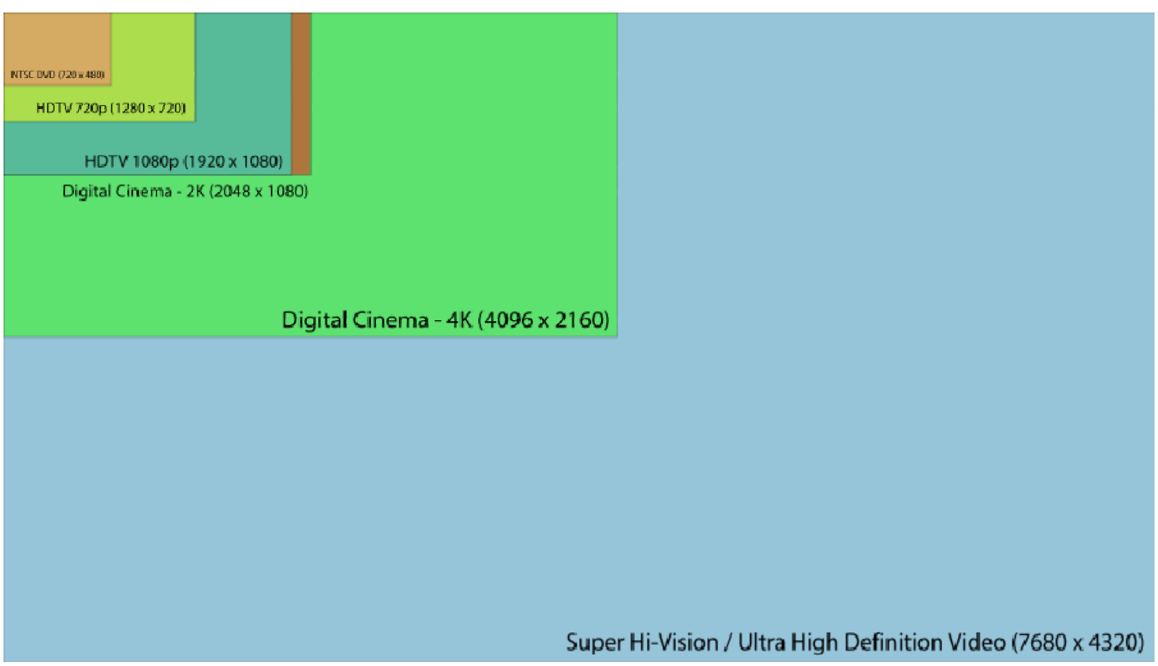

Figure 5

Global Sensorium

Again, we examine the technologies that will form a global network of data gathering sensors and transmitters. Cameras probably will be the most important factor in generating new Net traffic, but eventually other types of automated sensors in this category will get into the act and constantly and remotely feed the Net with exabytes of real-time information.

- *Imaging*

Consumers purchased more than 1 billion mobile phones around the globe in 2006, and they bought even more in 2007. Of these billion phones sold in 2006, more than 400 million were *camera* phones. And an even larger portion had cameras in 2007. At least 700 million *camera* phones will be sold in 2009. Mobile cameras are moving from the poor quality of the original 1-megapixel sensors to much higher resolutions—up to 8 megapixels and 30 frames-per-second VGA high quality video. By 2009, consumers will buy more than 100 million compact digital cameras—generating files of around 3 MB per photo—and 6 million high-end DSLR cameras—generating photos of 20 MB or more. The mass volumes, and corresponding price drops, of the phone and consumer camera markets will translate into mass deployment of image sensors everywhere. For surveillance, in medicine, in automobiles, and in every PC, just to name the obvious. Can we imagine a *dozen* cameras on each city block or building entrance? A *dozen* digital cameras in every automobile? A camera in every PC at more than 100 million per year? The world's consumers and businesses are about to deploy tens of billions of digital cameras across the globe.

- *Mobile revolution*

The mobile revolution means more people will be connected to the Net more of the time. iPhones, Treos, and Blackberries are not phones or PDAs—they are network computers. Rich content is going mobile, and for the first time, with the Apple-AT&T iPhone, mobile video is for real. With high resolution mobile displays and cameras, we can now consume *and produce* rich content anywhere, anytime, not just at our desktop or laptop. Apple sold four million iPhones in its first 200 days, and Google reports that on Christmas Day 2007 new iPhones generated a "surge" in traffic, consuming more network data than any other wireless device.[13] In just the last few months, the idea of consuming rich visual data like video clips and interactive satellite maps on mobile devices has gone from clunky fantasy to elegant and sensational reality. More people connected more of the time means more data and new traffic patterns.

- *RFID – Radio Frequency Identification*

Over the coming decade, we will tag just about every item everywhere with tiny RFID sensors. How many hundreds of billions of items? No one knows. How much information will these tags collect? How will we transmit and store this information? These are difficult questions to answer. John Chambers says by 2000 we had connected 100 million devices to the Net. By 2010 we expect to connect 14 billion devices.

- *IPv6 (Internet Protocol version six)*

The new version of Internet Protocol (IPv6) means every person, device, product, object, place, and virtual space can have thousands of addresses. The previous and reigning version of IP, IPv4, is a 32-bit protocol. IPv4 had 4 billion unique addresses and payloads up to 64 KB. Address shortages and awkwardly disseminated blocks of addresses required Network Address Translation and other work-arounds. We are now moving to 128-bit IPv6, with 3.4×10^{38} unique addresses (340 billion billion billion billion) and payloads up to 4 GB "jumbograms." In the coming decade, we can therefore connect everything to the Net.

LAN's End

With the consummation of secure, broadband links to homes and businesses, Local Area Networks will dissolve into the global Internet. We have long predicted this phenomenon as "LAN's End."

Trusted Computing

More than 100 million PCs and laptops now contain Trusted Platform Modules; almost every computer, and most hard-drives, built going forward will have TPMs. The Trusted Computing paradigm moves security to the very edge of the Net. Over time it could eliminate not only most firewalled LANs but also most infuriating virus software and passwords, which none of us can remember. Your computer and hard-drive will be secure and smart—they will know who you are. As important, you will run online applications with the security and customization as if they were running on your PC. Trusted Computing could allow all local traffic, storage, and apps seamlessly to flood the Internet.

Network Computing / Web Services

Twenty years ago, Sun Microsystems told us, "The network is the computer." Today, this aspiration is reality. It is the paradigm of centralized "cloud" computing most obviously manifest in the Googleplex—Google's intelligent global network of bandwidth, storage, and computer power. Computing and storage functions will forever be shifting from the edge to the core and back to the edge, depending on the relative prices and capabilities of bandwidth, storage, and processing. The PC revolution that began in the late 1970s and early 1980s moved computing and storage to the edge, away from the centralized mainframes and time-share computers. Today, we are in the midst of a general reversal of that broad trend, as bandwidth makes the "cloud" more capable, adaptable, and attractive than your PC. Network applications will continue to proliferate, and many local applications will move to the cloud, too. Some hosted real-time business applications will require robust connections of 25 Mbps or more.

Peer to Peer

File sharing of music, TV, amateur video, and games has been a monstrous phenomenon since Napster in the late 1990s. Despite copyright issues and legal and structural changes, it continues today. Some have claimed the peer-to-peer file-sharing system BitTorrent accounts for one-third of Internet traffic. This estimate is far too high, but BitTorrent and its competitors are very substantial bandwidth consumers. Looking further ahead, as legal or structural challenges possibly reduce the growth rate of some peer-to-peer systems, other applications—such Microsoft's experimental Photosynth, which reach out, gather, and synthesize data and content from peers across the Web—will emerge.

- *All Things Digital*

According to IDC and EMC, total digital information created, captured, or replicated (though not necessarily stored or transmitted) in 2006 was 161 exabytes. By 2010, IDC estimates we will generate digital information totaling 988 exabytes. John Chambers of Cisco writes that the unconnected, firewalled, quarantined data of the "Dark Web" could be 500 times the size of the existing Internet. What happens when storage is truly virtualized across the Net?

- *Remote Backup*

If, in the near future, we assume 2 billion PCs worldwide with each storing 25 GB per PC, and if we assume that remote backup will be a common feature provisioned by the popular Internet content companies (Google, Yahoo!, Microsoft, AOL) and broadband providers (Verizon, AT&T, Comcast, Cox), we could see 50 exabytes of worldwide remote backup in the coming years. In the U.S. by 2015, we could have 500 million PCs (and other computing/storage devices), with 100 GB per device, for a total of 50 exabytes of remote backup in U.S.

Acceleration? Fast, or Faster

These trends are all swelling tributaries funneling into the Exaflood. But will Internet traffic accelerate from its already fast growth rate? Or will the Exaflood merely mean a continuation of fast, but not accelerating, growth on top of our high base of traffic?

No one can know just how fast Internet traffic will grow, making capacity planning for network *backbone* engineers a difficult task. In fact, like network data itself, the macro trends discussed in this paper are "bursty"—they are unpredictable and potentially explosive. Networks will thus require much more raw bandwidth, yes, but also sophisticated new techniques to finely manage traffic flows. We do know for certain that if consumers and innovators are to enjoy the rising tide of the Exaflood, capacity in *access* networks to homes and businesses will have to expand by a factor of between 10 and 100 in just the next few years. One- and 10-megabit links will have to become 25- and 100-megabit links.

We also know that in an astounding feat South Korea, with just one-sixth the population of the U.S., now approaches the U.S. in Internet traffic.[14] Partly due to a more friendly regulatory environment, South Korea deployed fiber optic networks sooner than the U.S. did. South Korea also was an aggressive first builder of 3G mobile networks.

Summing the trends outlined in this paper over the next 8-10 years, we find that annual U.S. IP traffic circa 2015 could reach 1,000 exabytes, or 1 zettabyte.

Rough estimate of annual U.S. IP traffic, by application, circa 2015

Movie downloads and P2P	100 exabytes
Video calling and virtual windows	400 exabytes
“Cloud” computing / remote backup	50 exabytes
Internet video, gaming, virtual worlds	200 exabytes
Non-Internet “IPTV”	100 exabytes, or more
Business IP traffic	100 exabytes
Other (phone, Web, e-mail, photos, music)	50 exabytes
total	1,000 exabytes = **1 zettabyte**

Table 1

The U.S. Internet of 2015 will thus be some 50 times larger than it was in 2006. Reaching this projected zettabyte level by 2015 would require an acceleration of traffic from the modest levels projected by Cisco. It would certainly require video-calling and conferencing to take-off prior to Cisco's conservative projection of 2015. But it would only require traffic increases of around 90% per year using Cisco's 2011 estimate as a starting point. If Cisco's estimates over the coming few years prove conservative, we could reach the zettabyte mark in 2015 with far less than 90% compound growth.

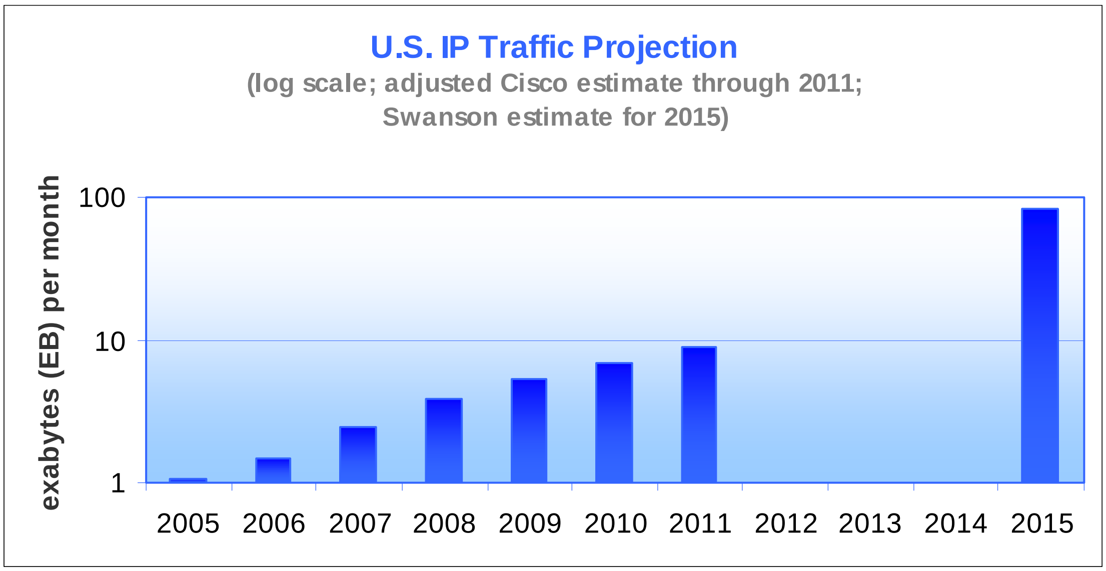

Figure 6

With U.S. traffic now around 30% of world traffic, global IP traffic will probably pass the zettabyte mark several years earlier, around 2012 or 2013.

Although U.S. adoption and usage patterns won't precisely track the Asian example, Korea's booming Internet traffic suggests that building the infrastructure first is the key to a robust digital economy. Time and time again, we see that consumers, businesses, and entrepreneurs find both serious applications and also whimsical ways to fill the pipes. None of the enticing Exaflood trends above can be achieved without a new network. The new optical network, with mobile and wireless capillaries, is the prerequisite—the very foundation—of the Exaflood.

[1] Bret Swanson is an adjunct fellow at the Discovery Institute and a senior fellow and director of the Center for Global Innovation at The Progress & Freedom Foundation.

[2] George Gilder is a senior fellow at the Discovery Institute and the founder of Discovery's Technology & Democracy Project. He is also chairman of George Gilder Fund Management, LLC and moderator of the Gilder Telecosm Forum.

[3] We are grateful to the Discovery Institute, where we first developed and named many of the concepts and themes outlined in this paper. We first used the term "exaflood" in a joint article for the June 2001 issue of the *Gilder Technology Report* entitled "Metcalfe's Exaflood." Gilder then expanded on the exaflood theme in an afterward to the paperback edition of his book, *Telecosm*. This paper is a further expansion of "The Coming Exaflood," Swanson's article that appeared in *The Wall Street Journal* on January 20, 2007. We are also grateful for the helpful input of Hance Haney, director of the Technology & Democracy Project at the Discovery Institute; Professor Andrew Odlyzko at the University of Minnesota; and Adam Thierer, senior fellow at the Progress & Freedom Foundation.

[4] Gantz, John F., et al. "The Expanding Digital Universe: A Forecast of Worldwide Information Growth Through 2010." IDC. March 2007. http://www.emc.com/about/destination/digital_universe/

[5] Markoff, John. "Jams Already on Data Highway." *The New York Times*. November 3, 1993.

[6] Odlyzko, Andrew. MINTS—Minnesota Internet Traffic Studies. http://www.dtc.umn.edu/mints/home.html

[7] Odlyzko, Andrew. MINTS—Minnesota Internet Traffic Studies. Internet Growth Trends and Moore's Law. http://www.dtc.umn.edu/mints/igrowth.html

[8] "The Exabyte Era." Cisco Systems, Inc. August 2007. http://www.cisco.com/application/pdf/en/us/guest/netsol/ns537/c654/cdccont_0900aecd806a81a7.pdf "Global IP Traffic Forecast and Methodology, 2006-2011." Cisco Systems, Inc. August 2007. http://www.cisco.com/application/pdf/en/us/guest/netsol/ns537/c654/cdccont_0900aecd806a81aa.pdf

[9] "The Internet Singularity, Delayed: Why Limits in Internet Capacity Will Stifle Innovation on the Web." Nemertes Research. November 2007. http://www.nemertes.com/system/files/Internet+Singularity+Delayed+Fall+2007.pdf

[10] A year ago we believed that almost all gaming and newer interactive virtual worlds would remain "local," or attached to and rendered by a gaming box or PC at the user's location. Perhaps control data would be transmitted over the Net, we surmised, but the images would be rendered locally. Conceptual input from Adam Thierer at the Progress & Freedom Foundation and technical observations from Jules Urbach of JulesWorld and Otoy, however, helped us understand how a much larger percentage of online games and virtual worlds can actually be rendered not locally but "in the cloud." This possible architecture could have a significant impact on Internet traffic.

[11] http://en.wikipedia.org/wiki/High-definition_television

[12] http://en.wikipedia.org/wiki/Ultra_High_Definition_Video

[13] Helft, Miguel. “Google Sees Surge in iPhone Traffic.” *The New York Times*. January 14, 2008. http://www.nytimes.com/2008/01/14/technology/14apple.html?_r=1&oref=slogin

[14] Odlyzko.