Internet Peering and Transit

Anna-Maria Kovacs Ph.D., CFA[1]

April 4, 2012

Executive Summary

The Internet has evolved from four university mainframes connected by 64-kilobit lines to a network of networks that spans the globe, connecting hundreds of millions of devices and carrying hundreds of exabytes of traffic per year. Its key characteristics have been adaptability and responsiveness. It has accommodated changes in its own infrastructure and topology, changes in technology at its edge, changes in type and volume of traffic, changes in types of market participants and their relationships, and changes in methods of compensation for the exchange of traffic.

The Internet has had very little formal governance and no regulation to slow its responses or divert its course. The absence of rate regulation has been particularly helpful, providing the flexibility that allows the Internet to accommodate seamlessly ever-larger volumes, suddenly-emerging new types of traffic, and new players both at the core and the edge. The Internet's traffic moves along on commercial agreements, many of them hand-shake agreements. The system has evolved from peering among academic and government networks to a mixture of paid and unpaid peering and transit among various types of commercial networks. Competition is rife, with the five largest bandwidth providers carrying only about a third of the international capacity of the U.S., and new entrants like Google and Comcast smoothly displacing incumbents in the top 10. Because of its informality, this unregulated system has facilitated the necessary investment in infrastructure to support the explosive growth of traffic that the Internet carries, and carries at ever-lower prices. That, in turn, has stimulated economic growth and job creation in the U.S. and around the world.

---

[1] Anna-Maria Kovacs is a Visiting Senior Policy Scholar at the Center for Business and Public Policy at the McDonough School of Business at Georgetown University. She has covered the communications industry for more than three decades as a financial analyst and consultant.

The author gratefully acknowledges Broadband for America for providing financial support for this paper and for obtaining all the required permissions for citations.

Internet Peering and Transit

From its inception as four university mainframes connected by experimental packet switches and 64-kilobit lines to its present incarnation as a commercial network of networks connecting hundreds of millions of devices via routers using Internet Protocol over multi-gigabit links, the Internet has constantly evolved. New technologies and demand for new services have grown in tandem, driving and reinforcing each other. Adaptability and responsiveness have been the distinguishing characteristics of the Internet. With very little formal governance and no regulation to slow its responses, it has accommodated:

- Changes in the infrastructure over which the Internet itself rides
- Changes in the topology of the Internet
- Changes in technology at the edge
- Changes in type of traffic
- Changes in volume of traffic
- Changes in types of market participants and their relationships
- Changes in methods of compensation for exchange of traffic

The Internet responds swiftly and automatically to changes in the most crucial factors that affect it. The Internet's ability to handle vast increases in the volume of traffic and accommodate new types of traffic, to be indifferent to the types of players who own and manage its physical infrastructure and to the evolving relationships among them, to move seamlessly from generation to generation of innovative hardware and software both at the core and the edge—this is the greatest gift of the Internet.

The Internet has been unregulated, and that lack of regulation is a key factor in its success. Any particular regulation may be good or bad, but the process is rarely swift, and by its nature tends to look backward rather than forward. Laws are built upon precedent and created with deliberation—that is their strength and their weakness. Regulation of a global network that crosses myriad jurisdictions is exponentially vulnerable to delay. In dealing with forces at the edge whose constant shape-shifting cannot be predicted, unintended consequences are inevitable. Because it is not regulated, the Internet is free to respond quickly at its core to the changes innovation brings at its edge.

What has worked particularly well for the Internet has been the absence of rate regulation. Much of the Internet's traffic moves along on hand-shake agreements, and the rest on commercial agreements. Despite—or because of—its informality, this system has facilitated the necessary investment in infrastructure to support the explosive growth in the volume of traffic that the Internet carries, and carries at ever-lower prices.

Without rigid boundaries to limit its growth and strait-jacket its development, the Internet has not only grown rapidly itself, but has fostered the growth of the vast and ever-changing ecosystem that is the foundation of today's global economy. It is a truly remarkable example of Adam Smith's "invisible hand" in action.

Early history of the Internet--Inception to commercialization[2]:

In 1969, Bolt Beranek and Newman (BBN) installed Interface Message Processors (IMPs), which were experimental packet switches, at four universities across the U.S. and connected them with phone lines. This was the origin of ARPANET, funded by DARPA (Defense Advanced Research Projects Agency), which linked mainframes via remote login for timesharing.

- In 1970, Network Control Protocol (NCP) was developed as ARPANET's first host-to-host protocol. It was implemented in 1971-1972.
- In 1972, BBN introduced the first email system.
- In 1972, Metcalfe invented Ethernet. It was established as a standard protocol and was commercialized in the 1980's.
- In 1973, the first written version of Transmission Control Protocol/Internet Protocol (TCP/IP) was developed. It was designed for nationwide networks, to be the general infrastructure on which applications could be layered.
- In 1980, TCP/IP was adopted as a defense standard, and in 1983, the shift from NCP to TCP/IP was implemented. TCP/IP remains the standard for the global Internet.
- In 1986, NSFNET was funded by NSF to connect academic supercomputing sites. ARPANET continued to serve and be funded by the military.
- In 1989, Tim Berners-Lee, then at CERN, invented the World Wide Web.
- In 1989, Border Gateway Protocol (BGP) became an Internet standard, with the current version BGP4 adopted in 1995.
- In 1990, ARPANET was decommissioned.
- In 1991, a Commercial Internet Exchange was established in Santa Clara, CA to connect burgeoning private Internet networks. In 1992, arrangements were made for it to exchange traffic with NFSNET.
- In 1993, saw the first commercial web browser.
- In 1995, NFS privatized the Internet. Government funding of the Internet ceased and the Internet became a commercial enterprise.
- Since 1995, the Internet's infrastructure has been privately funded.

What is the Internet now?

The Internet is a global network of networks that transport communications in Internet Protocol (IP).[3] It is, simply, an ever-expanding set of networks that exchange traffic, whatever its origin or destination, as long as it is expressed in IP.

[2] This section draws most heavily on Leiner, Cerf, Clark et al, "Brief History of the Internet", but also on Kende, "IP Interconnection Ecosystem"; Oxman, "Unregulation of the Internet"; Yoo, "Internet's Architecture".
[3] TeleGeography, "Market Structure," p. 1. For discussions of the Internet as a network of networks, see also: Norton, Peering Playbook; Oxman, "Unregulation of the Internet"; ENISA, "Inter-X"; Yoo, "Internet's Architecture".

Another way to look at it is that the Internet has become the primary medium for the carriage of traffic globally. Figure 1 below[4], from TeleGeography, shows that by mid-2002, seven years after the Internet became privately owned and managed, it accounted for 70% of all international traffic, with switched voice accounting for 5% and private networks accounting for the remaining 25%. By mid-2011, the Internet's role became even more significant, with 80% of international traffic carried on the Internet. Switched voice has all but disappeared at 0.3% of traffic, and private networks now carry less than 20% of traffic.

Figure 1:

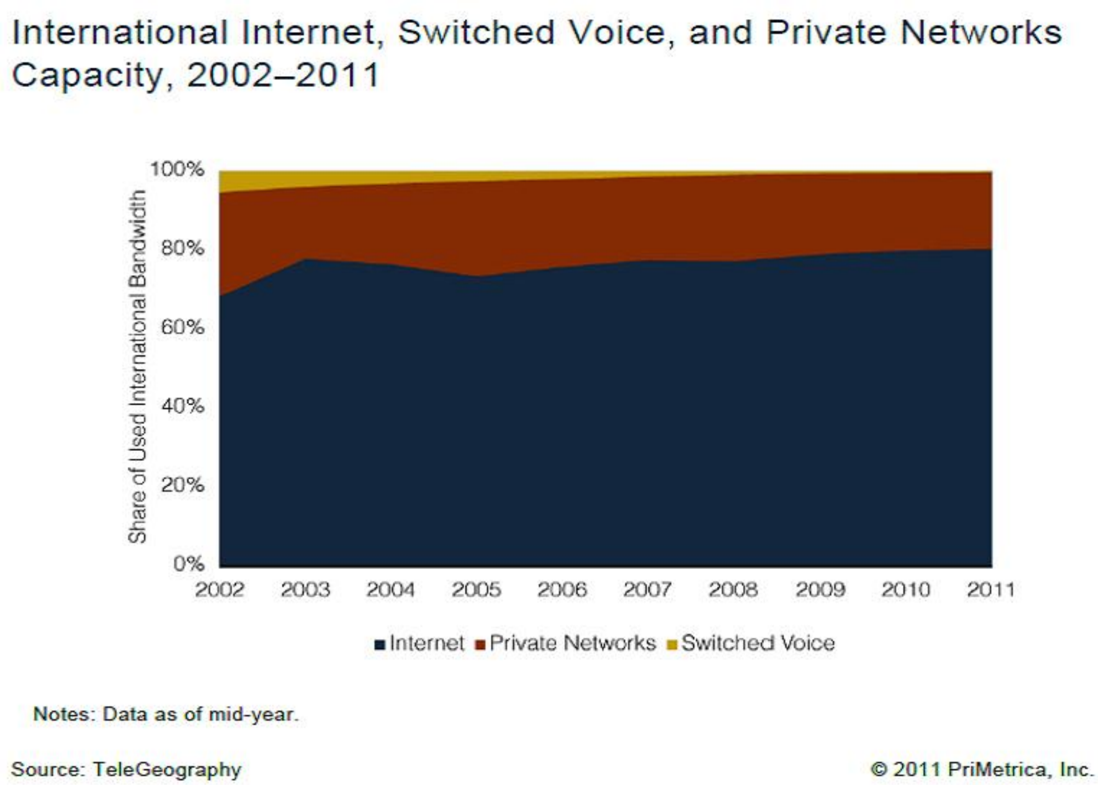

Although the Internet quickly assumed the role of primary global communications medium, it continues to grow at a phenomenal rate. At the time the Internet was privatized in 1995, it connected fewer than 10 million hosts. By January 2002, it connected roughly 147 million hosts and by January 2012, it connected roughly 888 million hosts.[5]

---

[4] Figure 1: Reproduced with permission from TeleGeography, "Capacity", Figure 4, p. 6.
[5] Internet Systems Consortium, "Host Count".

Figure 2:

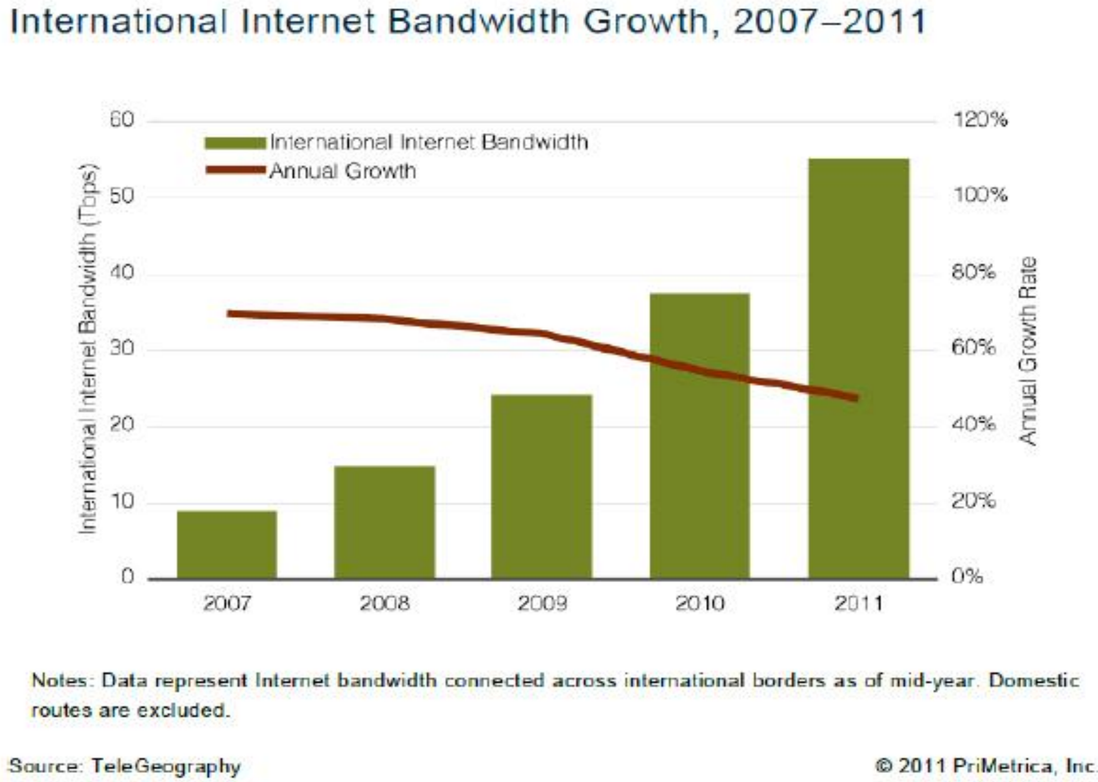

There are various ways to look at the Internet's growth. TeleGeography measures it in international Internet bandwidth—i.e., its capacity for carrying traffic between nations. Since 2002, that has grown from 0.9 Tbps to 54.9 Tbps, a compounded annual growth rate (CAGR) of 58%.[6] As Figure 2 above shows, even in more recent years, international Internet bandwidth has more than quintupled.[7] From mid-2007 to mid-2011 it grew from 8.7 Terabits per second (Tbps) to 54.9 Tbps. Even those statistics understate total global Internet bandwidth, because they measure only capacity for traffic that crosses international boundaries. For the U.S., in particular, that ignores roughly 80% of the bandwidth capacity (as measured at the ten largest U.S hubs), which has been growing at 37% CAGR over the last four years.[8]

And the growth is expected to continue. Cisco Visual Networking Index (VNI) publishes an annual forecast of Internet traffic. Unlike the TeleGeography graph which measures *capacity*—i.e., how much bandwidth the network could provide per second at any given point in time, Cisco's graph, shown below, measures *actual traffic* carried over a period of time. Looking back at roughly the same period covered by TeleGeography, Cisco translates that capacity to exabytes used per month. In Figure 3, Cisco shows that traffic grew from roughly 80 petabytes per month in 2001 to more than 20 thousand

[6] Calculation based on TeleGeography, "Capacity," Figure 3, pp. 4-5.
[7] Figure 2: Reproduced with permission from TeleGeography, "Capacity," Figure 1, p. 1.
[8] Calculation based on TeleGeography, "Region: United States," Figures 4 and 5, pp. 4 and 6.

petabytes--i.e., 20 exabytes--per month in 2010.[9] Given the already large base of traffic, growth can't continue at that torrid rate. However, Cisco still projects a CAGR of 32% which will result in a quadrupling of traffic from 2010 to 2015, with traffic in 2015 expected to reach over 80 exabytes per month. On an annual basis, that will equate to a zettabyte, hence the name of Cisco's mid-2011 VNI report—"Entering the Zettabyte Era".

Figure 3:

Five Traffic Milestones and Three Traffic Generator Milestones by 2015

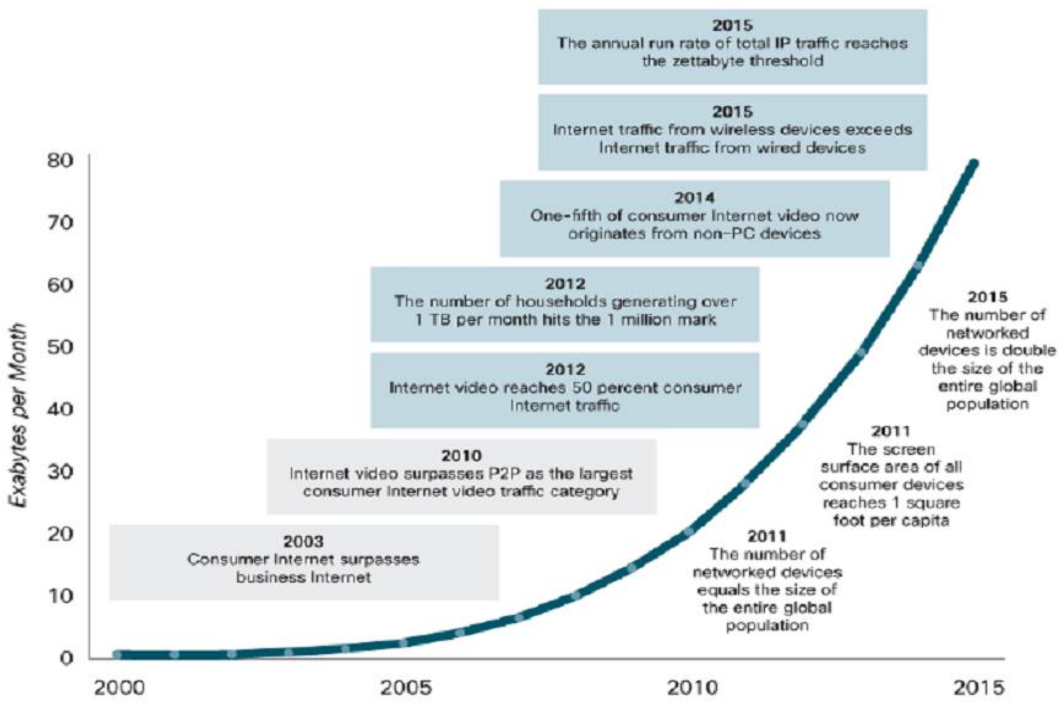

Source: Cisco VNI, 2011

While both TeleGeography's and Cisco's statistics show the Internet's remarkable ability to expand raw capacity to meet demand, Cisco VNI's graph shows in greater detail the variety in types of demand that the Internet has met with seamless supply. The transformational milestones Cisco points out include shifts in types of end-users, devices, and types of traffic. Consumers overtook businesses as users of the Internet in 2003. In 2010, Internet video displaced peer-to-peer (P2P) as the largest source of consumer Internet traffic, a distinction P2P had claimed for the prior decade, according to Cisco. One critical driver--the number of networked devices--is expected to double in the next four years, even though in 2011 the number of networked device matches the entire global population. The nature of the devices also is changing with non-PC devices becoming increasingly prevalent. The nature of the traffic is changing in another respect, as well, with wireless--both fixed and mobile--playing an increasingly important role as origin or destination of traffic.

[9] Figure 3: Reproduced with permission from Cisco, "Zettabyte Era," Figure 1, p. 5.

Capacity on the Internet is not only growing rapidly, it is dispersed among a large number of owners. TeleGeography points out in its "Capacity" report that on international routes connected to the U.S., the top five carriers have only about 35% of the bandwidth, the next fifteen have about 33%, and the rest control another 32%.[10] That distribution of capacity has been remarkably stable over the past five years.[11]

How does the Internet work?

An obvious question is, how does this wondrous creature work? What makes it so adaptable to changes in its ecosystem as well as so responsive to such large increases in the sheer volume of demand?

The Internet is actually a pretty simple construct that links independent networks that are responsible for maintaining themselves, and that interact via commercial agreements that may be as informal as "handshakes":

- Each network that transmits data on the Internet is an Autonomous System (AS) and is identified via an ASN (Autonomous System Number). Each AS is responsible for maintaining the IP addresses of its own customers. To interconnect, the networks direct their IP traffic via their routing tables.

- The relevant protocols are TCP/IP at the transmission and Internet layers and BGP4 at the application layer.

- Physically, the Internet consists of routers that direct streams of packets over fiber, copper, coax, wireless, and satellite links. Each network is responsible for its own facilities. They meet at interexchange points that may be owned by one of the networks or by a commercial vendor.

- The infrastructure of the Internet is privately owned, and private investment has upgraded it continuously since 1995 to meet increasing and changing demands.

- Traffic on the Internet is exchanged based on commercial agreements among the relevant networks, with or without compensation. The agreements may or may not be formalized via written contracts.

Those are the essential characteristics of the Internet, and they have not changed fundamentally since 1995. It is a very simple, flexible, and resilient system that works well even though those agreements are now made among thousands of networks operating in countries across the globe.

What has changed is the level and type of demand placed upon the Internet, as users and applications have proliferated. The Internet has responded by modifying its topology and, along with it, some

[10] TeleGeography, "Capacity," pp. 31-32.
[11] TeleGeography, "Capacity," Figure 33, p. 32.

aspects of its compensation system. It is an outstanding model of supply responding to demand, and thus enabling greater and more varied demand.

Who are the players and how have the relationships among them changed?

Network topology:

As Professor Yoo of the University of Pennsylvania explains in a paper about innovation in the Internet's architecture,[12] once the Internet was commercialized in 1995, several commercially-owned backbone providers took over from the prior NFSNET backbone, interconnecting with each other at Network Access Points (NAPs). These backbone Internet Service Providers (Tier-1 ISPs) also connected to regional ISPs (Tier-2 ISPs), whose networks in turn connected to local distribution facilities, i.e., central offices and cable head-ends. As Figure 4 shows,[13] the relationship of last-mile providers, regional ISPs, and

Figure 4:

The Architecture of the Early Internet

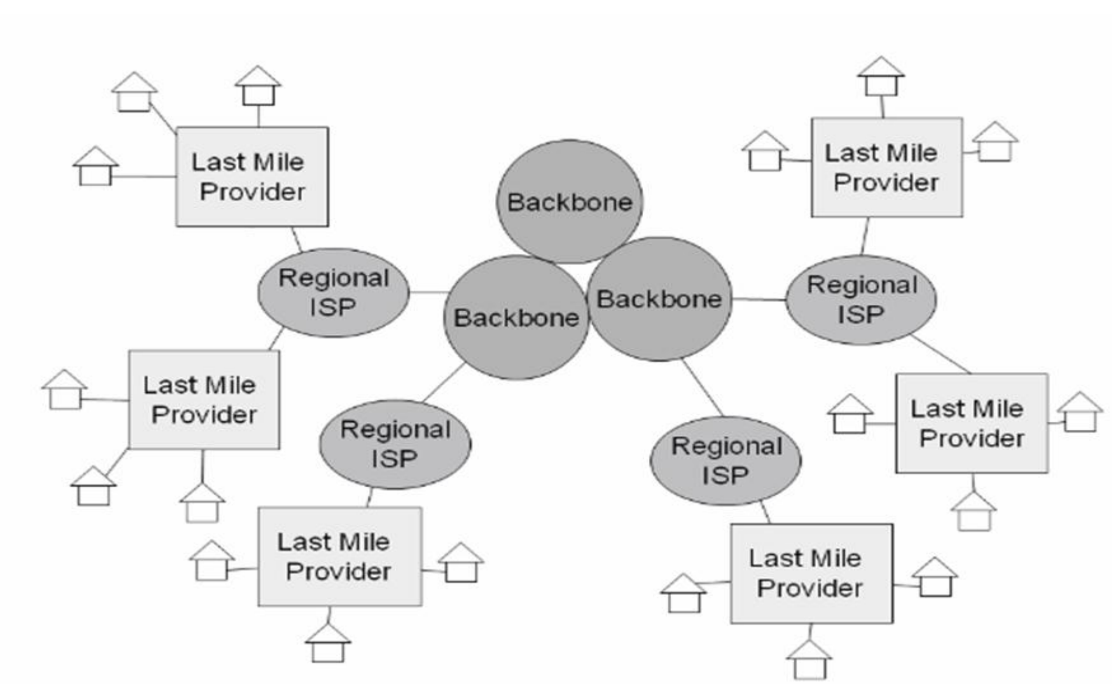

Source: Christopher S. Yoo, Journal of Telecommunications and High Tech Law, 2009.

backbone providers was hierarchical (as had been NFSNET), providing a unique path between each set of nodes.

[12] Yoo, "Internet's Architecture," pp. 81-85.
[13] Figure 4: Reproduced with permission from Yoo, "Internet's Architecture," Figure 3, p. 85.

Over time, that simple hierarchical topology has evolved into a more complex topology as it has responded to changes in volume and type of demand. In addition to backbone and regional ISPs, i.e., Tier-1 and Tier-2 ISPs, content delivery networks (CDNs)--some owned and/or dedicated to a single content provider and some serving many content providers--have been added to the mix. Rather than being limited to a single relationship, each of these may now exchange traffic with any of the others. Figure 5[14] below from ENISA's report describing the resilience of the Internet interconnection system shows how complex the exchange has become.

Figure 5:

The System of Connections

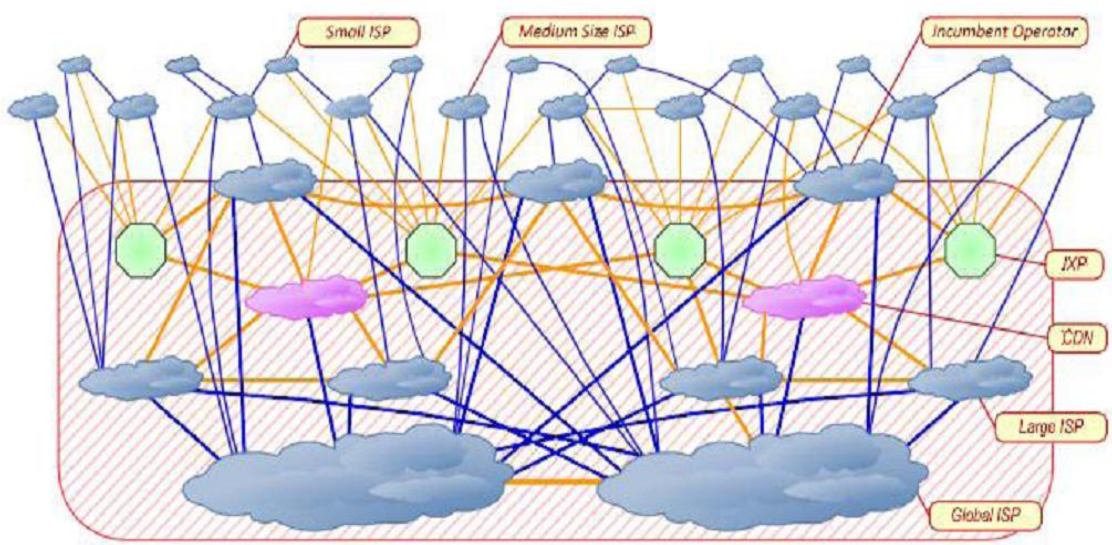

Source: ENISA, 2011

The players and the mesh of commercial arrangements among them—peering, Internet transit, direct Internet access:

To understand the differences between the early topology and the current version, it is important to understand the function of each player and the basis on which traffic is exchanged between different types of players.[15]

- Backbone providers, also known as Tier-1 ISPs, often maintain backbone networks as well as IP-address routing tables that include all IP addresses, so that they are able to route traffic anywhere on the Internet.

[14] Figure 5: Reproduced with permission from ENISA, "Inter-X," Figure 28, p. 89.

[15] See Norton, Peering Playbook; TeleGeography, "Market Structure" and "Pricing"; Woodcock and Adhikari, "Survey"; Yoo, "Internet's Architecture."

- The Regional ISPs, also known as Tier-2 ISPs, maintain smaller networks and IP-address routing tables that contain their own customers. They can also provide their customers with access to the whole Internet, but they do so by supplementing their own network with either unpaid or paid partners.

- Content delivery networks (CDNs) such as Akamai or Limelight enable content providers to cache their content at various points close to the edge of the Internet. Some CDNs serve multiple content providers, while others are dedicated to and owned by one specific content provider such as Google or Amazon. CDNs may rely heavily on ISPs for transport, but a few own facilities of their own that are as extensive as those of some backbone providers. In some cases, e.g. Level 3, an ISP may also take on the function of a CDN.

- IXPs are Internet exchange points. ISPs connect to one another at IXPs, which may host many ISPs or only a few. The IXPs may be owned by a particular ISP or may be owned by a third-party such as Equinix that is not an ISP itself, making that a 'neutral' IXP. CDNs may also connect to ISPs at IXPs.

All traffic among ISPs--or ISPs and CDNs--is exchanged on the basis of commercial agreements. There are no regulated rates for exchanging traffic over the Internet. As Figure 6 below shows, the agreements may be for paid or unpaid exchange, and may provide access to part or to all of the Internet.

- "Internet peering is the business relationship by which two companies reciprocally provide access to each other's customers," according to Bill Norton, one of the founders of Equinix and an expert on peering and Internet transit arrangements.[16]

- Peering provides two ISPs access to each other's customers, not to the entire Internet. Peering has generally been unpaid—i.e., settlement-free—but a trend toward paid peering has developed, as well.
- Internet transit is a paid arrangement. It gives the Tier-2 ISP or the CDN who pays Internet transit the right to reach the entire Internet. That reach comes not only via its partner's own customer base, but also through all the relationships that the partner has with other ISPs.

- Bill Norton defines Internet transit: "Internet transit is the business relationship whereby an Internet Service Provider provides (usually sells) access to the global Internet." He adds, "Customers connect their networks to their Transit Provider, and the Transit Provider does the rest."[17]

- Thus, peering stops at the border of the partner's own customer base. Transit, however, allows both parties to pass through and beyond the borders of the partner's customer base, to enjoy the

[16] Norton, Peering Playbook, p. 83.
[17] Norton, Peering Playbook, p. 41.

benefit of all the relationships that the partner has with other partners, and that those, in turn, have with yet more providers, so that ultimately the entire Internet becomes reachable.

Figure 6:

Internet peering v. transit

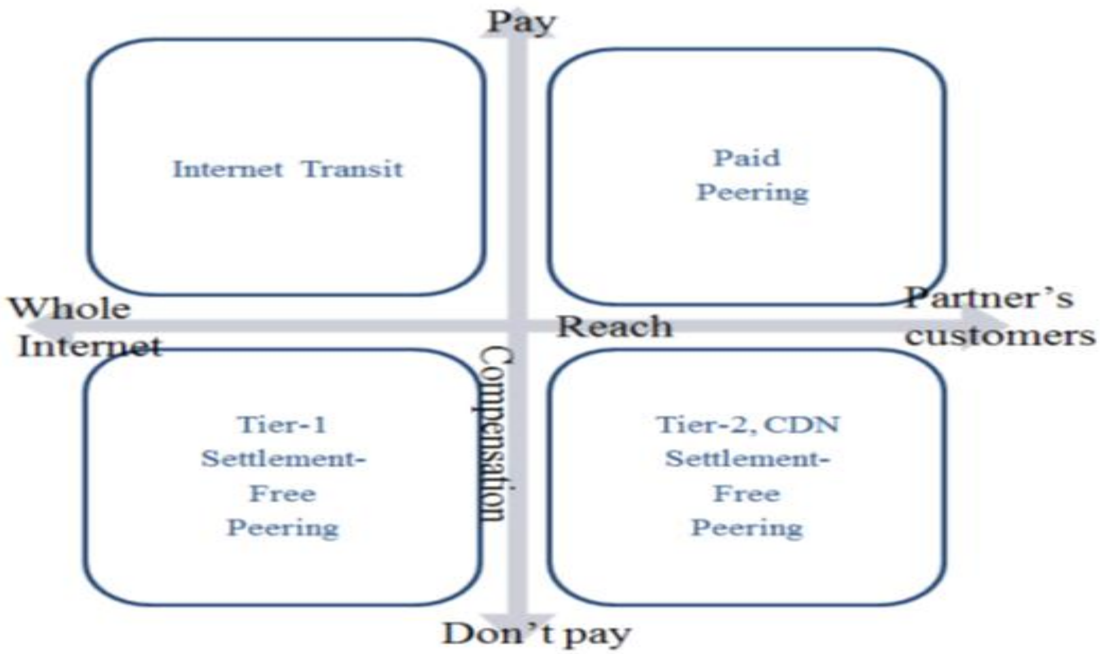

Source: Anna-Maria Kovacs , 2012

- Settlement-free peering:
    - As the name suggests, settlement-free peering is done between "equals". It does not involve payment—the traffic is freely exchanged. Each party provides the other partner with access to its own customers, not to the entire Internet.
    - When the Internet was first commercialized, in the timeframe represented by Figure 4 above, settlement-free peering was the only form of peering available and was done only among Tier-1 ISPs (and universities).
    - Settlement-free peering is still the most common method when Tier-1 ISPs exchange traffic with each other.
    - However, today, settlement-free peering may occur not only between Tier-1 ISPs, but also between Tier-2 ISPs exchanging balanced amounts of traffic among themselves, or between a combination of CDNs, Tier-1 ISPs and Tier-2 ISPs, again depending on the balance of traffic and benefit to the parties.

- As Bill Norton points out in his *Playbook*, that balance is not just based on volume, but mutual benefit. It may consist of the volume of traffic peered bi-directionally, or it may be proportional to the desirability and/or uniqueness of the routes, or it may represent the number of end-users reached.[18]

- Paid peering: May be used if there is a traffic imbalance, so that settlement-free peering is not appropriate, but the paying ISP only wants access to its larger partner's customer base, not to the whole Internet.

- Looking at it from the other angle, a payment-free relationship occurs between partners who are peers in the sense of bringing relatively equal traffic to the relationship. A paid relationship assumes that traffic will be out of balance in some way, and the form it takes—paid peering v. transit depends on the extent of the reach the payer wants.

- One of the most important points made by Bill Norton is that even settlement-free peering is not free, in that each party has its own costs to cover—costs of both links and the facilities at the IXP. Thus, each provider must decide for each route what method to use—peering, transit, or self-provisioning of facilities. The decision has to be based on the breakeven between the cost of transit and the cost of peering, on the volumes involved, and possibly based on other factors such as acceptable latency.[19]

In the early years of the commercial Internet, as depicted in Figure 4 above, Tier-2 ISPs connected to the entire Internet by buying transit from one particular backbone provider at one of the relatively few available interconnection points (IXPs).

Today, as depicted in Figure 5 above, the mesh of connections is much richer and more complex. In Figure 5, the yellow lines represent peering and the blue lines transit. A Tier-2 ISP may connect at any of a large number of IXPs, and may do so not only with multiple Tier-1 providers but also with Tier-2 providers of various sizes as well as with CDNs, some of which are owned by and serve one specific content provider.

Some of the largest content providers have, in effect, created networks equivalent to those of Tier-1 ISPs. Google has the most extensive set of peering points, covering every continent, but Microsoft, Yahoo, Amazon and Facebook also have extensive peering. TeleGeography, citing Arbor Networks, indicates that "Google has boosted its share of traffic exchanged via peering links from 40 percent in 2008 to 60 percent in 2010."[20]

Not surprisingly, one of the best explanations of the range of choices available to a content provider who wants to distribute content over the Internet comes from Level 3, which is both a Tier-1 ISP and

[18] Norton, Peering Playbook, p. 83.
[19] Norton, Peering Playbook, p. 109.
[20] TeleGeography, "Market Structure," p. 6.

CDN. Figure 7 below[21] was submitted by Nicolas Pujet, SVP of Corporate Strategy for Level 3, as part of a Declaration he made to the FCC in July 2011 in the Level 3 – Global Crossing merger docket. Pujet explained in that Declaration : "In the early days of the Internet, to reach the Internet and each other, content providers and ISPs were in large part dependent on purchasing IP transit from Tier 1 backbone providers, which acted as middlemen, exchanging their customers' traffic between themselves through peering. Today's ISPs and content providers have much more choice."[22] He details some of the options available today to a content provider or regional-ISP:[23]

- Because of the low cost of dark fiber as well as of dedicated wavelengths and of the associated equipment, it is increasingly common for content providers and non-backbone ISPs to build or lease their own physical networks, and then exchange traffic either via "direct peering"[24] without using a 3d-party transit provider, or via public peering at interexchange points.

Figure 7:

Options for customers to move traffic

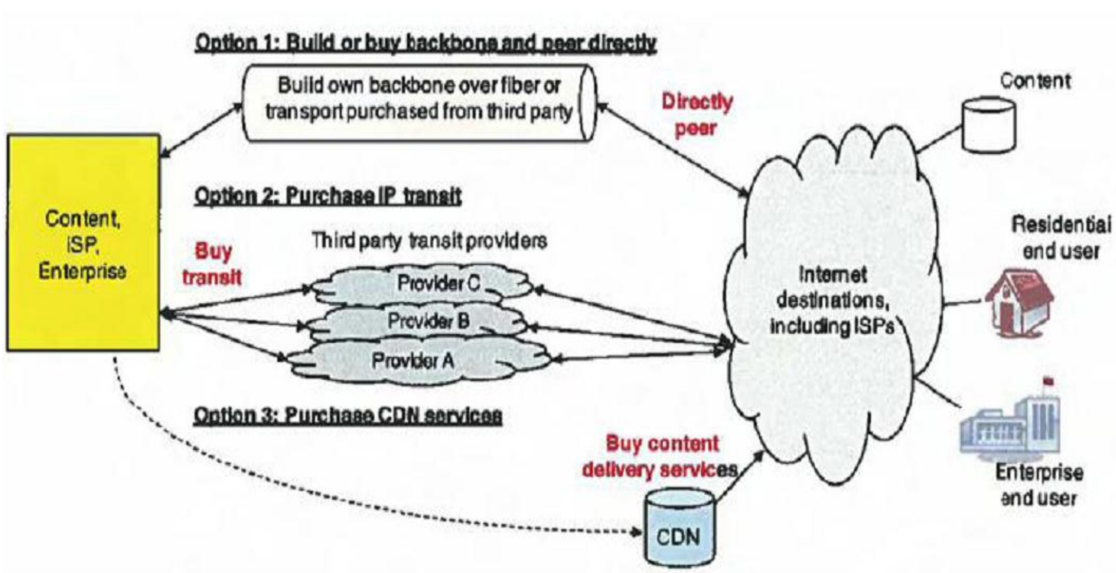

Source: Level 3 declaration in FCC docket IB 11-78

- Another option is to buy Internet transit from one or more providers who will then carry the IP traffic. Pujet notes that "Today there are dozens of providers with U.S. backbone networks

[21] Figure 7: Reproduced from Level 3, "Pujet Declaration," Figure 1, p. 6.
[22] Level 3, "Pujet Declaration," p. 5.
[23] Level 3, "Pujet Declaration," pp. 5-8.

[24] Level 3, "Pujet Declaration," p. 6.

providing transit service nationwide."[25] He includes a graphic (which we show as Figure 11) from DrPeering, i.e., William Norton, to show the radical decreases in Internet transit pricing that have occurred between 1998 and 2011.

- Yet another option is to pay a CDN to store content close to the ISP and end users, bypassing backbone transit providers completely. Specifically, Pujet notes that the Renesys dataset lists thousands of autonomous systems and says that Level 3's IP Product Management team estimates that 38 companies sell transit or offer peering on a national basis using an IP backbone.[26]

- He also points out that multi-homing, i.e., sharing traffic over at least two providers, is a common practice, and that only about 10% of Level 3's and Global Crossing's customers are single-homed to one or the other.[27]

The FCC, in the order approving the Level 3 – Global Crossing merger without any conditions, discussed the various choices available for the exchange of Internet traffic, including the various forms of Internet peering and transit, and the use of CDNs. In its rationale for allowing the merger, the FCC noted that the number of Tier-1 ISPs has increased since 2005. It cited Level 3's assertion that there may be as many as 38 providers that sell transit or offer peering on a nationwide basis. And it accepted Level 3's arguments regarding multi-homing. The FCC noted that the relationships between the entities carrying Internet traffic are governed by privately negotiated commercial agreements. In allowing the merger, it described a market that offers numerous choices to customers large and small.[28]

Peering continues evolving, but remains highly informal:

What is particularly interesting is that this complex mesh of facilities and relationships works seamlessly, based on commercial relationships that in many cases are highly informal.

Packet Clearing House (PCH) published the results of a survey on May 2, 2011 in which it focused on peering relationships.[29] PCH points out that there are 5039 ISPs that interconnect to form the mesh that is the Internet. PCH distributed a voluntary survey between October 2010 and March 2011. It received responses from roughly 86% of the ISPs, representing 96 countries, including all OECD member countries. As a result, PCH was able to study 142,210 peering agreements among 4331 different ISP networks. In 1032 cases, both parties to the same agreement responded, and PCH found that in 99.52% of the cases, their responses to the survey were identical, giving PCH comfort that survey respondents understood the questions clearly. PCH notes the limitations of its methodology, in that the survey was voluntary and does not represent all interconnections agreements that exist. Nevertheless, the sheer size of its sample makes its observations important.

[25] Level 3, "Pujet Declaration," p. 7.
[26] Level 3, "Pujet Declaration," p. 5.
[27] Level 3, "Pujet Declaration," p. 7.
[28] FCC, "Global Crossing and Level 3 Consent to Transfer Control, pp. 9-16."
[29] Woodcock and Adhikari, "Survey."

PCH found that:

- 99.51% of the peering agreements—i.e., 141,512 of the 142,210--were informal. They "were 'handshake' agreements in which the parties agreed to informal or commonly understood terms without creating a written document. The common understanding is that only routes to customer networks are exchanged, that BGP version 4 is used to communicate those routes, and that each network will exercise a reasonable duty of care in cooperating to prevent abusive or criminal misuse of the network."[30]

- 99.73% of the agreements were symmetric. All but 374 or the 142,210 agreements "had symmetric terms, in which each party gave and received the same conditions as the other." In the few asymmetric cases, one party compensates the other via paid peering to receive routes it might not receive, or has to meet minimum peering requirements. But in 141,836 cases, parties simply exchanged routes without settlements or other requirements.[31]

- Most of the networks have small numbers of interconnection partners, with 62% having ten or fewer. Only 12 networks had more than 700. However, there are some spikes which are the result of large multilateral peering agreements (MLPAs), at the Hong Kong, Warsaw, and Frankfurt Internet exchange points. The multilateral peering agreements tend to follow the same form and terms as bilateral, except that they have more than two parties.[32]

- Multilateral agreements appear to be gaining ground, especially outside the U.S. While most networks have few interconnection partners, those that do have them account for a large portion of the Autonomous System pairs PCH observed. PCH notes that "As an example, the 144 participants in the Hong Kong Internet Exchange multilateral peering agreement represent 10,296 AS-pair adjacencies, and *each one* of those participants individually exceeds the average "tier-1" carrier in degree of interconnection."[33] PCH speculates that "just as 'donut peering' overtook 'tier-1' peering in the late 1990's, multilateral peering may now be overtaking bilateral peering, at east in sheer numbers, if not necessarily in volume of traffic."[34]

- PCH finds that Tier-1 carriers tend to have many IPv4 addresses, but very few interconnection partners. On the other hand, large CDNs with similar scale and degree of infrastructure investment tend to have very broad interconnection, both in absolute numbers and in geographic diversity.[35]

[30] Woodcock and Adhikari, "Survey," p. 2.
[31] Woodcock and Adhikari, "Survey," p. 2.
[32] Woodcock and Adhikari, "Survey," p. 5.
[33] Woodcock and Adhikari, "Survey," p. 6.
[34] Woodcock and Adhikari, "Survey," p. 6.
[35] Woodcock and Adhikari, "Survey," p. 7.

Thus, whatever advantage Tier-1 ISPs may have had in the past as a result of their extensive routing tables, new entrants can now achieve the same result by interconnecting with many partners, both via bilateral and multilateral agreements.

***Internet Transit—prices are falling*:**

When an ISP (or CDN) buys Internet transit, it is paying its partner to transmit its traffic to the entire Internet, not only on the partner's own network but on the entire mesh of networks with which the partner and the partner's partners and their partners *ad infinitum* have Internet transit or peering agreements. The ISP pays for transit by port capacity, possibly but not necessarily with a minimum commitment. Unlike leased lines for which payment is a function of distance, Internet transit is not paid on a distance-basis, but on a capacity-basis, e.g. for full port size of 100 Mbps or 1 Gbps or 10 Gbps, or for a commitment for a portion of a port, i.e., for a specified data rate on a particular type of port. Alternately, payment may be for the data-rate that is actually experienced in a given month, usually measured at the 95th percentile (thus disregarding occasional peaks). Thus, the ISP is paying its partner to pass a certain amount of traffic through its port, from which point it can go anywhere on the Internet, by whatever route is available through that partner's mesh of partners. Depending on the routes sought, an ISP (or CDN) might buy transit from a Tier-1 ISP or from a Tier-2 ISP.[36]

A general but not invariable rule is that rates are lower per Mbps:

- if a minimum capacity is committed—the vendor can plan for the capacity
- for greater capacity—e.g. 900 Mbps costs less than 200 Mbps on a GigE port <u>*measured per Mbps*</u>
- if it is on a higher capacity port (which to some extent also presumes greater volume, as well as possibly a newer and more cost-effective technology)

The TeleGeography database provides extensive detail about Internet Transit pricing: by city--both international and U.S. domestic, by type (SDH/SONET or Ethernet) and bandwidth of connection, and by carrier. These are summarized in various ways in TeleGeography's Global Internet Geography Reports each year. A pair of tables from the 2011 pricing chapter, showing prices for some U.S. and Canadian cities, helps illustrate the point. The first table, Figure 8 below,[37] shows median prices for the monthly recurring cost of a full-commit 1 GigE port and the second, Figure 9 below, for a 10 GigE port, for the years 2008-2011 (measured at mid-year).

- The tables show that prices have been falling steadily over the 2008-2011 period, even if one looks within a single technology and port size: down 12-18% compounded per year in the U.S. for 1 GigE port and down 17% to 25% compounded per year for 10 GigE ports.

- The tables show that there is some variation in median rates in the U.S. from city to city, but not much. The median across the country is fairly uniform, particularly at major hubs.

[36] See: Norton, Peering Playbook; TeleGeography, "Market Structure" and "Pricing."
[37] Figures 8 and 9: Reproduced with permission from TeleGeography, "Pricing," Figures 8 and 9, pp. 9 and 10.

- However, while the median price is fairly uniform across the U.S., TeleGeography also shows that there is quite a range around the median. For example, in New York, the price per megabit for a full GigE port at the end of second quarter 2011 ranged from $3.00 to $10.90, around the median of $6.70.[38]

- The price per Mbps in a given city, on a given port size, also ranges considerably by data-rate commitment, at least for some carriers. Again in New York in 2011, over a 1 GigE port, one carrier charges roughly $18 per Mbps when selling 100 Mbps but only about $6 when selling 1000 Mbps on a 1 GigE port. In the same city, for the same port type, another carrier's price declines from $8 to $4 over the same range, while yet another's price is fairly steady over the whole range at about $7 to $6.[39]

Figure 8:

Median North American IP Transit Prices per Mbps, Gigabit Ethernet (1,000 Mbps), Q2 2008-Q2 2011

	2008	2009	2010	2011	2010-11	CAGR 08-11
Chicago	$10.95	$9.00	$7.00	$6.02	-14%	-18%
Houston	$10.98	$9.00	$7.30	$6.31	-14%	-17%
Los Angeles	$10.98	$10.74	$8.20	$7.50	-8%	-12%
Miami	$10.98	$9.00	$7.00	$6.64	-5%	-15%
New York	$11.00	$9.49	$7.61	$7.00	-8%	-14%
San Francisco	$11.00	$9.00	$8.00	$7.00	-12%	-14%
Toronto	$9.90	$9.00	$7.00	$6.75	-4%	-12%

Notes: Data shown are monthly USD per Mbps prices for a full-port commit, excluding local access and installation fees. Data as of Q2 of each year.

Source: TeleGeography Research

[38] TeleGeography, "Pricing," p. 7.
[39] TeleGeography, "Pricing," p. 9.

Figure 9:

Median North American IP Transit Prices per Mbps, 10 Gigabit Ethernet (10,000 Mbps), Q2 2008-Q2 2011

	2008	2009	2010	2011	2010-11	CAGR 08-11
Chicago	$8.00	$6.00	$4.42	$3.40	-23%	-25%
Houston	$7.11	$6.25	$4.25	$4.00	-6%	-17%
Los Angeles	$8.00	$7.25	$4.68	$4.00	-15%	-21%
Miami	$8.00	$6.00	$4.52	$4.00	-12%	-21%
New York	$8.00	$6.25	$4.31	$4.00	-7%	-21%
San Francisco	$7.11	$6.00	$5.00	$3.40	-32%	-22%
Toronto	$7.11	$6.00	$4.25	$3.33	-22%	-22%

Notes: Data shown are monthly USD per Mbps prices for a full-port commit, excluding local access and installation fees. Data as of Q2 of each year.

Source: TeleGeography Research

Figure 10:

Median GigE IP Transit Prices in Major Global Cities, Q2 2005-Q2 2011

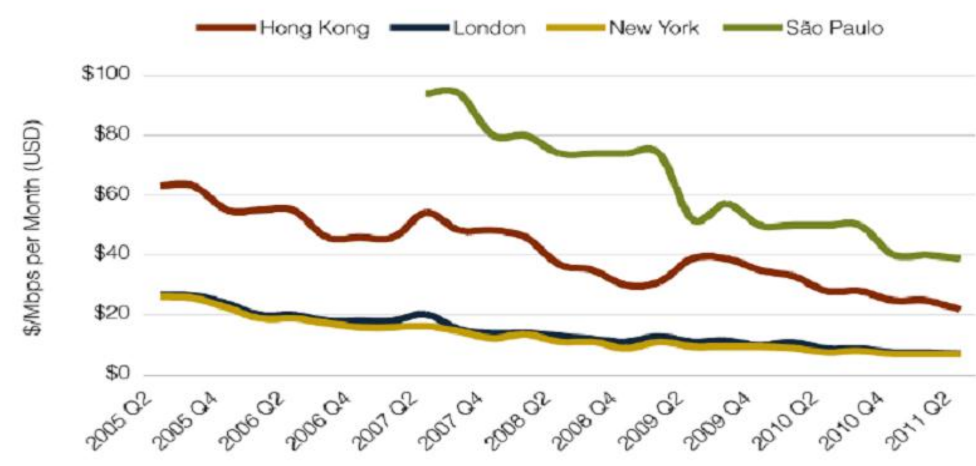

Notes: Each line is the median monthly price per Mbps in that city. Data exclude installation and local access fees. Gigabit Ethernet (GigE) = 1,000 Mbps.

Source: TeleGeography

It is also important to note that while median-price variations within the U.S. for particular types and volumes are small, they are significant if one looks across the globe. Prices in the U.S. and Europe are roughly a third of what they are in Asia which in turn is about half the price in Latin America (disregarding Sydney and Manila, which appears to be the most expensive markets by far).[40] But, as Figure 10, from TeleGeography,[41] shows, prices are converging over time. As Norton explains extensively in his *Playbook*, carriers from high-priced regions, e.g. Latin America, Africa or Asia, have the option of using a direct connection to a low-priced region such as the U.S. and then using Internet transit from there. The convergence of transit prices in the international graph above may be an indicator that ISPs are making use of that option.

Focusing again on the U.S., the decline in Internet transit prices over the last 13 years, and the expectation that it will continue, are best summed up by the graphic below from Norton's DrPeering website (and *Playbook*), Figure 11 below.[42] Norton shows that the price of transit has been falling steadily and sharply. This chart tracks the price for Internet transit in the U.S. since 1998 and projects ahead to 2015. The price is expressed on a megabit/second basis for committed capacity. As Norton explains, it is a rough indicator, but it is a useful one in that it provides an extensive time series. It is also confirmed to some extent by various TeleGeography reports. For example, the February 2005 TeleGeography Bandwidth Pricing Report shows the price per Mbps in New York on an OC-3 port declining from 2003-2005 from a range of about $80 to $120 in 2005 to about $30 to $50 in 2005. Recent prices in the DrPeering graph correspond to the price for 10GigE ports.[43]

Norton explains that he polls 30-50 people three to five times a year at operations conferences that he attends each year.[44] While the TeleGeography data we cite above is specific to particular cities and port types and sizes, our impression is that Norton's data cuts across those variations. In other words, what he is gathering is what ISPs and CDNs are paying for the best Internet transit deal they can get, when they have done all their shopping across carriers, port size and types, potential routing, required quality, etc.

A slightly earlier version of this graph was cited by Level 3 in its FCC docket to show the rapid decrease in Internet transit prices. But Internet traffic exchange can be optimized even further. As Level 3's Pujet pointed out to the FCC, content providers, CDNs and ISPs choose between directly-owned/leased facilities, Internet transit, and settlement-free or paid peering. Thus, even the sharply declining Internet transit prices in the DrPeering graph above do not fully reflect the lowest price available to those who wish to exchange traffic over the Internet.

[40] TeleGeography, "Pricing," p. 2 and "Regional Analysis Asia", p. 16.
[41] Figure 10: Reproduced with permission from TeleGeography, "Pricing," Figure 1, p. 2.
[42] Figure 11: Reproduced with permission from Norton, "Internet Transit Pricing." Also in Norton, Peering Playbook, Figure 2-4, p. 48.
[43] TeleGeography, "Pricing," p. 10 and "Bandwidth Pricing Report" p. 2.
[44] Norton, Peering Playbook, p. 48.

Figure 11:

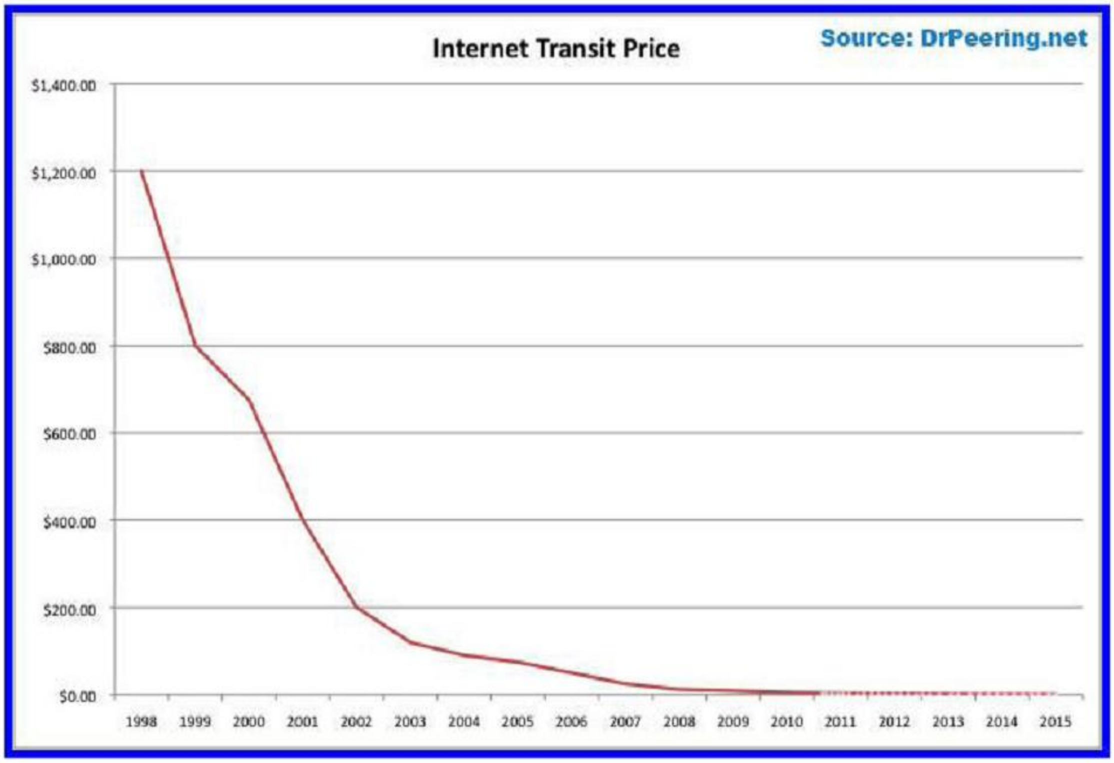

Even within a region, there is now usually a trade-off between peering vs. Internet transit for Tier-2 ISPs, large CDNs and content providers. One of the most interesting graphs in Norton's *Playbook*, which we show as Figure 12 below,[45] depicts the break-even between peering and transit. Even when peering is settlement-free it does have costs—facility rent, utilities, router, etc. Norton explains how an operations manager or peering coordinator would go about deciding over what volume of traffic to peer rather than use Internet transit.

Similar sorts of calculations could be—and no doubt are--made for the breakeven between owning/leasing facilities, peering, and Internet transit.

[45] Figure 12: Reproduced with permission from Norton, "Internet Transit Pricing." Also in Norton, Peering Playbook, Figure 5-2, p. 109.

Figure 12:

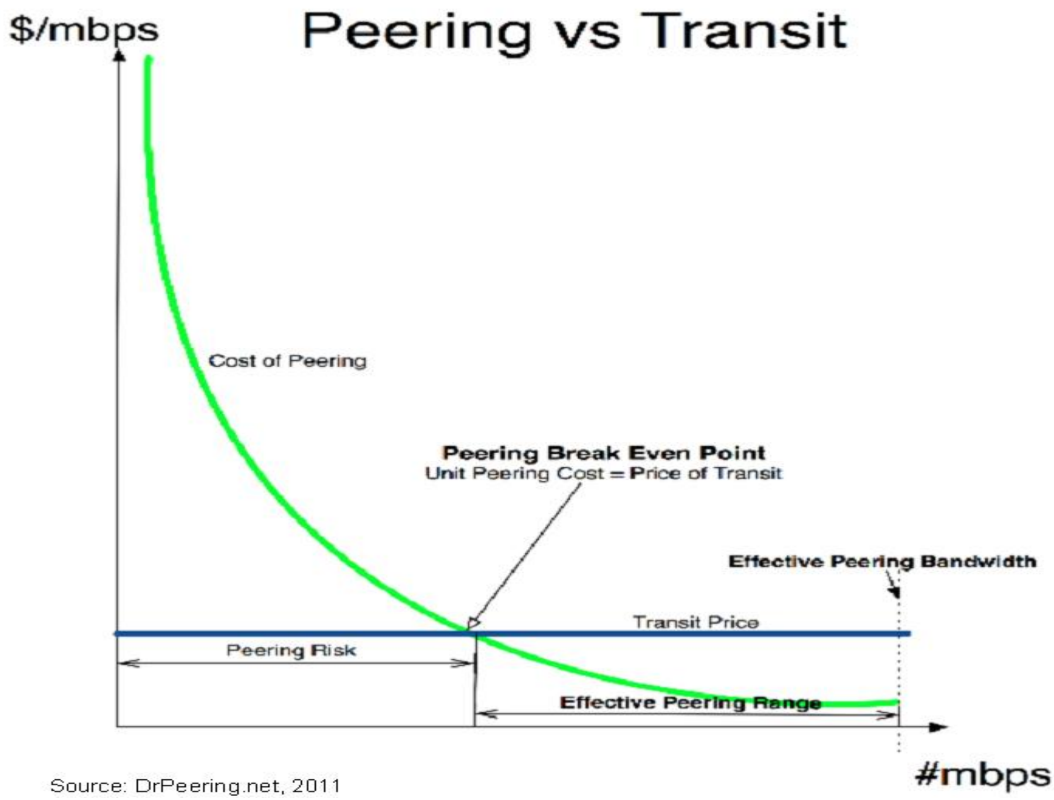

As Packet Clearing House, Level 3, and others point out, there is an enormous universe of providers of peering and transit among whom one can choose when seeking to exchange traffic. The simple universe of bilateral relationships that prevailed at the time the Internet was privatized has been replaced with a complex web of players, who can optimize for both price and quality.

<u>How is the Internet ecosystem evolving?</u>

The ability of the Internet to adapt to change has been critical to the success of the ecosystem that relies on it. New players have entered (and in some cases left) the market at various levels—backbone, regional ISP, CDNs, and content providers—some providing their own facilities, while others rely on the open market.

The networks they run are very different from the 64 kilobit lines that attached four mainframes in 1969--capacity is now measured in 100 Gigabit increments and the networks connect to and through myriad other networks. Not only are the networks that constitute the Internet itself changing, so are the things it connects. The Internet is simultaneously meeting the very different demands of ever-more-powerful mainframes, PCs in various client-server configurations and the emergence of multi-core PCs, as well as hundreds of millions of individual devices—from PCs to tablets to smartphone to game boxes--each making its own separate demands over a variety of different broadband technologies. Upgrades to access technologies—wired and wireless—are expect to result in a quadrupling of global average residential Internet connection download speed from 7 mbps to 28 mbps from 2010 to 2015.[46]

The Internet manages to respond to it all, without any 'central planning.'

As Figure 13 below[47], from TeleGeography, shows, the Internet keeps responding to enormous increases in demand for throughput with corresponding increases in supply. As peak utilization of a route gets to the 50%-60% level, the necessary investments are made to add capacity, to keep utilization within that margin of safety.

Figure 13:

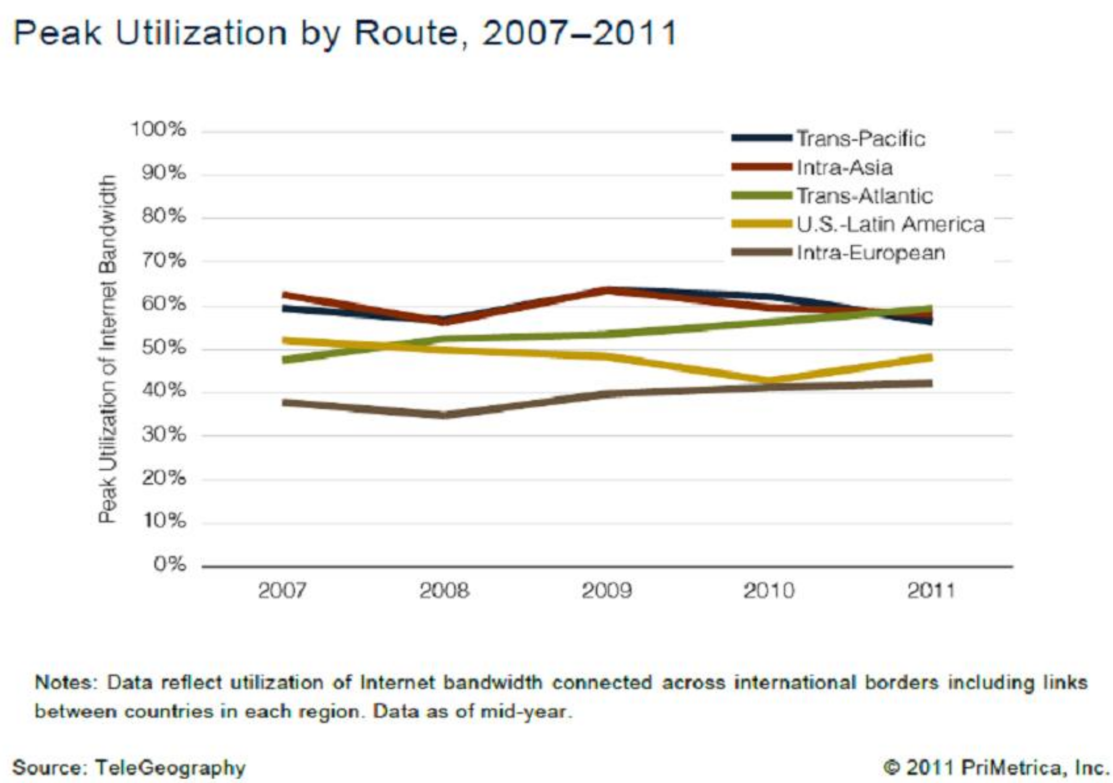

One aspect of the Internet's adaptability is its accommodation of sheer volume growth. McKinsey's recent study of the Internet's impact on the economy notes that 2 billion people are now connected to the Internet.[48] As Figure 14 below,[49] from Cisco VNI's "Entering the Zettabyte Era" white paper shows, consumption of bandwidth has grown exponentially. Global Internet traffic per capita increased from 1 megabyte per month in 1998 to 1 gigabyte per month in 2008 and is expected to grow to 9 gigabytes per month by 2015. Cisco expects total Internet traffic to grow from 1 exabyte per *year* (1000 gigabytes) in

[46] Cisco, "Zettabyte Era," p. 11.
[47] Figure 13: Reproduced with permission from TeleGeography, "Traffic," Figure 5, p. 6.
[48] Manyika and Roxburgh, "Great Transformer," p. 1.
[49] Figure 14: Reproduced with permission from Cisco, "Zettabyte Era", Figure 3, p. 7.

2001 to 1 exabyte per *day* in 2013.[50] Global IP traffic has increased eight-fold from 2005 to 2010 and is expected to quadruple by 2015.[51] In other words, total global Internet traffic is forecast to grow at a compounded annual rate of 32%, between 2010 and 2015, reaching roughly 1 zettabyte during 2015.[52]

Figure 14:

The VNI Forecast Within Historical Context

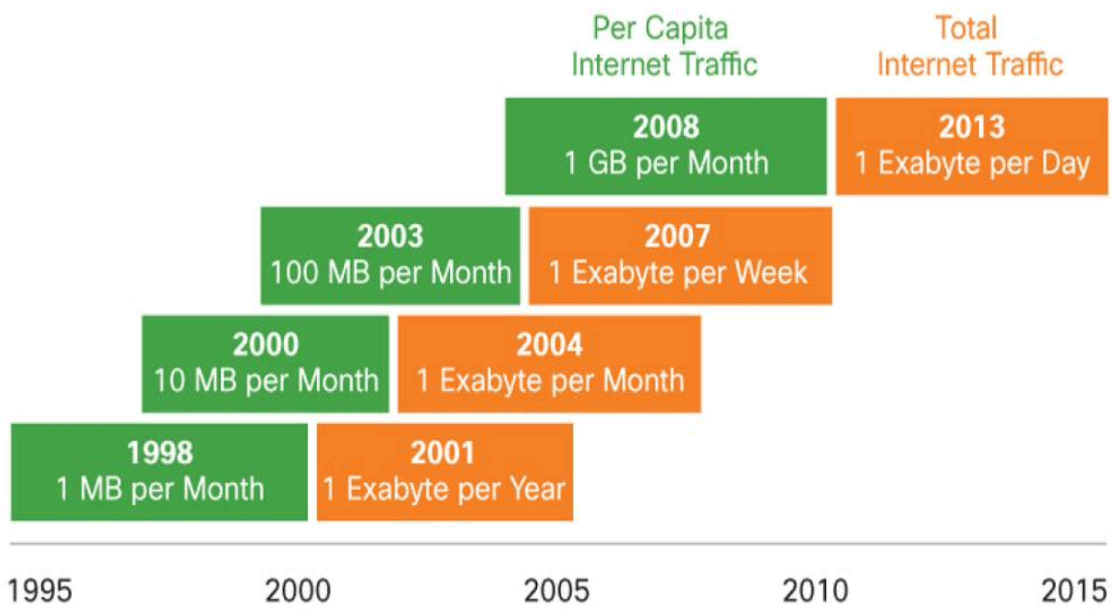

Source: Cisco VNI, 2011

But what also matters are the abrupt changes in the type of traffic. Steady flows of data were replaced around 2000 by vast amounts of P2P (peer-to-peer), which grew to two thirds of consumer Internet traffic by 2005. That, in turn, has been supplanted by video downloading, which exceeded P2P in 2010. By 2015, Cisco expects video to constitute 61% of consumer traffic and file sharing 24%, online gaming 15% and VOIP only 1%.[53] Traffic from wireless devices—using WIFI as well as mobile--will exceed traffic from wired devices by 2015, at 54% up from 37% in 2010. Mobile data Internet traffic will grow at a compounded rate of 92%, to comprise 8% of total IP traffic by 2015, with fixed Wi-Fi comprising nearly half of total IP traffic in 2015[54]. Other categories that could grow quickly include cloud-gaming, unicast rather than multicast TV broadcasting over the Internet, and adoption of 3DTV.[55]

---

[50] Cisco, "Zettabyte Era", pp. 6-7.
[51] Cisco, "Forecast," p. 1.
[52] Cisco, "Zettabyte Era", pp. 1 and 5.
[53] Cisco, "Zettabyte Era," pp. 2, 4, and 8.
[54] Cisco, "Zettabyte Era," pp. 2, 3, and 10.
[55] Cisco, "Zettabyte Era," pp. 13-14.

Cisco is far from alone in showing how dynamic the Internet is. Network upgrades—and thus investment--are a function of both the type of traffic and peak levels. Sandvine, which provides information on Internet traffic flows, shows the rapid change in the composition of Internet traffic. As Figure 15 from Sandvine shows,[56] between 2009 and 2011, over North American fixed access networks, real-time entertainment increased from 29.5% of the peak-time traffic to 53.6%. In the same timeframe, web-browsing fell from 38.7% of the traffic to 16.6%.

Figure 15:

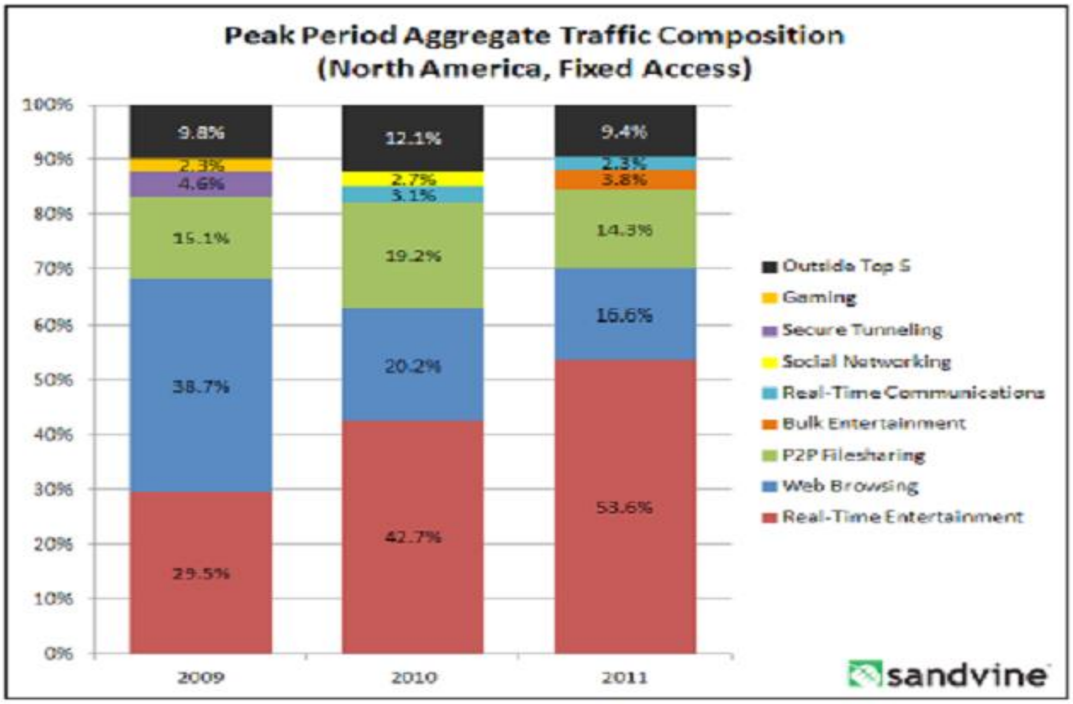

Source: Sandvine Global Internet Phenomena Report, 2011

As Figures 16 and 17 below, both from Sandvine,[57] shows that very different applications constitute the traffic on fixed access v. mobile access, and for upstream v. downstream. Netflix makes up 29% share of aggregate fixed-access traffic, all of it on the downstream. BitTorrent (P2P) still makes up 47.6% of the upstream traffic on fixed networks. On mobile networks, however, BitTorrent only makes up 3.8% of the upstream, and Netflix makes no appearance in the top-10 applications at all. For mobile, Facebook, HTTP, and YouTube are the applications that account for the greatest traffic.

[56] Figure 15: Reproduced with permission from Sandvine, "Global Internet Phenomena," Figure 2, p. 6.
[57] Figures 16 and 17: Reproduced with permission from Sandvine, "Global Internet Phenomena," Tables 1 and 2, pp. 8 and 12.

Figure 16:

Top Peak Period Applications by Bytes for North American Fixed Access

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	47.55%	Netflix	32.69%	Netflix	29.03%
2	HTTP	11.45%	HTTP	17.48%	HTTP	16.59%
3	Netflix	7.69%	YouTube	11.32%	BitTorrent	13.47%
4	Skype	4.27%	BitTorrent	7.62%	YouTube	9.90%
5	SSL	3.57%	Flash Video	3.41%	Flash Video	3.04%
6	Facebook	2.19%	RTMP	3.12%	RTMP	2.81%
7	PPStream	1.73%	iTunes	3.05%	iTunes	2.69%
8	YouTube	1.64%	Facebook	1.78%	SSL	1.96%
9	Xbox Live	1.31%	MPEG	1.72%	Facebook	1.84%
10	Teredo	1.25%	SSL	1.69%	MPEG	1.49%
	Top 10	82.63%	Top 10	83.88%	Top 10	82.83%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Source: Sandvine Global Internet Phenomena Report, 2011

Figure 17:

Top Peak Period Applications by Bytes for North American Mobile Access

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	Facebook	30.85%	HTTP	27.46%	HTTP	27.31%
2	HTTP	26.24%	YouTube	19.99%	Facebook	19.29%
3	SSL	6.05%	Facebook	17.62%	YouTube	18.23%
4	YouTube	6.01%	Windows Update	5.17%	Windows Update	4.70%
5	BitTorrent	3.83%	Android Market	4.09%	Android Market	3.75%
6	Ares	3.45%	Flash Video	2.96%	Flash Video	2.66%
7	Oovoo	2.57%	SSL	1.97%	SSL	2.48%
8	Skype	1.81%	RTSP	1.89%	RTSP	1.67%
9	Gmail	1.49%	Shockwave Flash	1.75%	Shockwave Flash	1.63%
10	Windows Update	1.48%	MPEG	1.67%	MPEG	1.53%
	Top 10	83.77%	Top 10	84.57%	Top 10	83.26%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Source: Sandvine Global Internet Phenomena Report, 2011

The difference in the type of application that users access over fixed v. mobile broadband makes the prediction of explosive growth in mobile broadband over the next five years particularly significant. Ericsson's *Traffic and Market Data Report: on the Pulse of the Networked Society,* which was published in November 2011, forecasts that mobile broadband subscriptions will rise from about 1 billion in 2011 to nearly 5 billion in 2016, with roughly half a billion PCs and tablets and the vast majority handheld devices. Ericsson expects fixed broadband subscriptions to remain well under 1 billion in 2016, but points out that while mobile subscriptions generally mean one user per subscription, fixed means one per household, with as many as three users per fixed subscription.[58]

The change in the type of applications that account for Internet traffic has been reflected in changes among those who carry the traffic. ATLAS Observatory analyzes traffic based on anonymous Autonomous System Number (ASN) (origin/transit) data. In 2007, the top-10 networks that accounted for 29% of Internet traffic were all traditional Tier-1 ISPs. By 2009, Google and Comcast had entered the top 10. Google, in fact, took over the #3 spot with 5.2% of traffic, more than any ISP except Level 3 and Global Crossing, who have since merged and whose combined share increased from 10.3% in 2007 to 15.1% in 2009.[59]

<u>The Internet and the global economy</u>:

The OECD did an extensive study called the "Future of the Internet Economy" on the contribution of the Internet to the global economy. As the FCC's "National Broadband Plan" had found for the U.S., the OECD found that the Internet can improve people's lives via e-commerce, education and job training, e-health, and access to employment. The study also found that from 1995 to 2008, the ICT-services sector grew gross-value-added much faster than the OECD's total businesses and total services, and that ICT-services employment growth over that period much exceeded that of business as a whole, with ICT employment growing at 18% v. business employment at 12%. The OECD study also found that ICT investment contributed more to GDP growth in most OECD countries than non-ICT investment in the 2000-2009 period. The OECD study also noted particularly the contributions ICT makes to labor productivity and R&D investment.[60]

McKinsey Global Institute's October 2011 paper "The great transformer: The impact of the Internet on economic growth and prosperity" points out that the Internet accounted for 21% of the GDP growth for mature economies over the past five years. McKinsey found that the Internet accounts, on average, for 3.4% of GDP in the large economics that account for 70% of global GDP.[61] McKinsey also highlights that maturity of the Internet ecosystem has a positive impact on living standards, via an increase in real per capita GDP of $500. The Internet rejuvenates traditional businesses, as well as creating new online businesses, and is a catalyst for job creation.[62] Summing up, McKinsey highlights that the Internet is a

[58] Ericsson, "Traffic," p. 6.
[59] McPherson, "Bandwidth Bandwagon," p. 5.
[60] OECD, "Future of the Internet Economy," pp. 26-27, 29.
[61] Manyika and Roxburgh, "Great Transformer," p. 1.
[62] Manyika and Roxburgh, "Great Transformer," p. 3.

"positive disruptor": "we should not lose sight of the enormous value the Internet economy has already brought to rich and poor nations alike and its potential to boost growth across the globe."[63]

Finally, a recent study by Ericsson, Arthur D. Little, and Chalmers University of Technology found that doubling the broadband speed for an economy increases GDP by 0.3%, and that additional doublings of speed can yield an additional GDP stimulation of 0.3%, so that a quadrupling of broadband speed would result in 0.6% GDP stimulation.[64] That makes Cisco's expectation of quadrupling of download speed during the 2010-2015 period particularly encouraging.[65]

Summary and Implications:

What does all of this say about the Internet? It is an extraordinarily successful demonstration of Adam Smith's "invisible hand." The Internet evolved, naturally, in tandem with the evolution of its ecosystem, both depending upon and enabling changes in computing, software, and network technologies. Both the building blocks of the Internet itself and all those things the Internet connects have been transformed repeatedly since the 1960's. With absolutely no regulatory intervention, with a compensation system that is entirely unregulated and to a great extent informal, the Internet has responded to demand that is not just explosive in volume but unpredictable in type. Supply has unfailingly met demand, at ever-lower prices.

The Internet's responsiveness and adaptability stands in stark contrast to the rigidity created by the regulatory compensation regimes that have stifled conventional telephony. An April 2009 OECD paper points out that the accounting-rate system constrained the use of international telecommunications services and did not tend to result in network investment.[66] The problematic effects of the U.S. intercarrier compensation system have been thoroughly aired in the recent FCC reform docket.[67]

What is worth highlighting is that the reform the FCC just ordered involves a transition that will take seven more years to complete, having been ordered fifteen years ago in the Telecom Act of 1996, having been attempted several times in the years since then. In the meantime, the U.S. switched-communications market has withered.

The Internet, on the other hand, was commercialized in 1995 and in the sixteen years since then has added thousands of networks serving hundreds of millions of people and their devices around the globe. It responds to new applications, such as video-downloading, by expanding bandwidth at a rate of nearly 60% per year, spawning a new layer of network architecture around CDNs, and welcoming Google into the top-10 network providers in the number three spot as soon as Google chooses to provide its own network.

---

[63] Manyika and Roxburgh, "Great Transformer," p. 9.
[64] Chalmers University, press release.
[65] Cisco, "Zettabyte Era," p. 11.
[66] OECD, "Network Externality," pp. 11-12.
[67] FCC, Report and Order.

What makes the Internet so effective is not just its own simplicity and adaptability, but the absence of externally imposed rigidity. In 1999, Jason Oxman wrote in an FCC working paper: "The success of the Internet has not been an accidental development. Market forces have driven the Internet's growth, and the FCC has had an important role to play in creating a deregulatory environment in which the Internet could flourish."[68] In June of 2011, an OECD "Communiqé on Principles for Internet Policy-Making" makes the same point: "As a decentralized network of networks, the Internet has achieved global interconnection without the development of any international regulatory regime. The development of such a formal regulatory regime could risk undermining its growth."[69]

Indeed, any attempt to impose economic regulation on the Internet is likely to slow not only its own evolution but the innovation at the edge that depends on the internet's core. The multi-year delays and the rigidity that accompany regulatory proceedings are not designed for either the growth rate or shape-shifting of the Internet. Were the Internet subjected to economic regulation, investors would expect slower growth and less responsiveness not only in the market for infrastructure, but in the edge markets for services, applications and devices that rely on it. Funding the Internet's infrastructure would become more difficult.

Vinton Cerf, David Clark, and several others who contributed to the genesis of the Internet recently wrote a brief history of the Internet, from which we cited extensively at the beginning of this paper. Their summation was: "One should not conclude that the Internet has now finished changing….[It] is making possible a new paradigm of nomadic computing and communications."[70] That paradigm will bring its own new challenges, which can best be met within a commercial, unregulated framework.

---

[68] Oxman, "Unregulation of the Internet," p. 3.
[69] OECD, "Communiqué," p. 3.
[70] Leiner, Cerf, and Clark, "Brief History of the Internet."

Internet Peering and Transit bibliography:

Cisco. "Cisco Visual Networking Index: Forecast and Methodology 2010-2015." June 1, 2011. Last accessed on April 1, 2012 at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html

Cisco. "Cisco Visual Networking Index: Entering the Zettabyte Era", June 1, 2011. Last accessed on 12/6/11 on the Cisco website under White Papers. Last accessed on April 1, 2012 at http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/VNI_Hyperconnectivity_WP.html

Chalmers University of Technology, Ericsson, and Arthur D. Little. Press release. September 27, 2011. Last accessed on April 2, 2012 at http://www.chalmers.se/en/news/Pages/New-study-quantifies-the-impact-of-broadband-speed-on-GDP.aspx

Elliott, Jim. "IBM Mainframes – 45+ Years of Evolution." IBM Corporation, 2010. On IBM website, last accessed on April 2, 2012 at http://www.vm.ibm.com/devpages/JELLIOTT/pdfs/zhistory.pdf

European Network and Information Security Agency (ENISA) "Inter-X: Resilience of the Internet Interconnection Ecosystem." Ed. Panagiotis Trimintzios. Full Report, April 2011. Last accessed on April 2, 2012 at http://www.enisa.europa.eu/activities/Resilience-and-CIIP/networks-and-services-resilience/inter-x/interx/report/interx-report

Ericsson. "Traffic and Market Data Report – on the Pulse of the Networked Society." November, 2011. Online. Last accessed on April 2, 2012 at http://hugin.info/1061/R/1561267/483187.pdf

Federal Communications Commission. "Memorandum Opinion and Order and Declaratory Ruling: In Matter of Applications filed by Global Crossing Limited and Level 3 Communications, Inc. for Consent to Transfer Control." IB Docket No. 11-78. Adopted September 29, 2011. Last accessed on April 2, 2012 at http://transition.fcc.gov/wcb/cpd/214Transfer/214completed2011.html

Federal Communications Commission. "National Broadband Plan." 2010. Last accessed on April 2, 2012 at http://www.broadband.gov/plan/

Federal Communications Commission. "Report and Order and Further Notice of Proposed Rulemaking." WC dockets 10-90, 01-92, 96-45. Released November 18, 2011. Last accessed on April 2, 2012 at http://transition.fcc.gov/Daily_Releases/Daily_Business/2012/db0206/FCC-11-161A1.pdf

Internet Systems Consortium. "Internet Host Count History." 2011. Last accessed April 2, 2012 at http://www.isc.org/solutions/survey/history

Kende, Michael. "The Digital Handshake: Connecting Internet Backbones." OPP Working Paper No. 32, Federal Communications Commission, September 2000. Last accessed on April 2, 2012 at http://www.fcc.gov/working-papers/digital-handshake-connecting-internet-backbones

Kende, Michael. "Overview of recent changes in the IP interconnection ecosystem." Analysys Mason, January 23, 2011.

Leiner, Barry M., Vinton G. Cerf, David D. Clark, et al. "A Brief History of the Internet," Internet Society, 2011. Last accessed on April 2, 2012 at http://internetsociety.org/internet/internet-51/history-internet/brief-history-internet

Level 3 Communications Inc. "Declaration of Nicolas Pujet." In Federal Communications Commission IB Docket 11-78, July 21, 2011. Last accessed on April 2, 2012 at http://apps.fcc.gov/ecfs/document/view?id=7021695741

Level 3 Communications Inc. and Global Crossing Limited. "Joint Opposition and Reply Comments." In Federal Communications Commission IB Docket 11-78, July 21, 2011. Last accessed on April 2, 2012 at http://apps.fcc.gov/ecfs/document/view?id=7021695740

Labovitz, Craig, Scott lekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. "Internet Inter-Domain Traffic." New Delhi, India: SIGCOMM'10, August 30-September 3, 2010.

Labovitz, C., lekel-Johnson, S., McPherson, Danny, et al. "ATLAS Internet Observatory 2009 Annual Report." Last accessed on April 2, 2012 at http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf

Manyika, James and Charles Roxburgh. "The great transformer: The impact of the Internet on economic growth and prosperity." McKinsey Global Institute, October 2011. Last accessed on April 1, 2012 at http://www.mckinsey.com/Insights/MGI/Research/Technology_and_Innovation/The_great_transformer

McPherson, Danny. "Bandwidth Bandwagon: An ISOC Briefing Panel." ATLAS Internet Observatory. November 11, 2009. Last accessed on April 2, 2012 at http://www.isoc.org/isoc/conferences/bwpanel/docs/20091111_bandwidth_mcpherson.pdf

Norton, William B. The Internet Peering Playbook: Connecting to the Core of the Internet. Palo Alto: DrPeering Press, 2011. Paper edition.

Norton, William B. "Internet Transit Pricing – Historical and Projected." DrPeering.net, 2011. Last accessed on April 2, 2012 at http://drpeering.net/white-papers/Internet-Transit-Pricing-Historical-And-Projected.php

OECD. " Communiqué on Principles for Internet Policy-Making: OECD High-Level Meeting on the Internet Economy". Paris: June 28-29, 2011. Last accessed on April 2, 2012 at http://www.oecd.org/dataoecd/33/12/48387430.pdf

OECD. "The Future of the Internet Economy – A Statistical Profile." June 2011 update. Last accessed on April 2, 2012 at http://www.oecd.org/dataoecd/24/5/48255770.pdf

OECD. "Network Externality Premiums and International Telecommunication Traffic Exchange." OECD Digital Economy Paper 152. OECD Publishing, April 15, 2009. http://dx.doi.org/10.1787/222470556326

Oxman, Jason. "The FCC and the Unregulation of the Internet." OPP Working Paper No. 31, Federal Communications Commission, July 1999. Last accessed on April 2, 2012 at http://www.fcc.gov/working-papers/fcc-and-unregulation-internet

Sandvine Intelligent Broadband Networks. "Global Internet Phenomena Report." Fall, 2011. Last accessed on April 2, 2012 at http://www.sandvine.com/downloads/documents/10-26-2011_phenomena/Sandvine%20Global%20Internet%20Phenomena%20Report%20-%20Fall%202011.pdf

TeleGeography. "Global Internet Geography: Capacity Trends." PriMetrica Inc., 2011.

TeleGeography. "Global Internet Geography: Bandwidth by Region." PriMetrica Inc., 2011

TeleGeography. "Global Internet Geography: Market Structure." PriMetrica Inc., 2011.

TeleGeography. "Global Internet Geography: Pricing." PriMetrica Inc., 2011.

TeleGeography. "Global Internet Geography: Regional Analysis--Asia." PriMetrica Inc., 2011.

TeleGeography. "Global Internet Geography: Traffic." PriMetrica Inc., 2011.

TeleGeography. "TeleGeography's Bandwidth Pricing Report." PriMetrica Inc., February 2005 and July 2005.

Telephone World, "Overview & Background on Electronic & Digital Switching Systems." December 17, 2010.

Woodcock, Bill and Vijay Adhikari. "Survey of Characteristics of Internet Carrier Interconnection Agreements." San Francisco: Packet Clearing House, May 2, 2011. Last accessed on April 2, 2012 at http://www.pch.net/resources/papers//peering-survey/PCH-Peering-Survey-2011.pdf

Yoo, Christopher S. "Innovations in the Internet's Architecture that Challenge the Status Quo." Journal on Telecommunications and High Tech Law 8 (2009): 79-100.