IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

Tier 1 ISPs: What They Are and Why They Are Important

Sponsored by: NTT Communications

Mark Winther

May 2006

EXECUTIVE SUMMARY

The global Internet is an amalgam of separate, but semiautonomous networks. The binding element of the global Internet is that networks share a common IP addressing and global BGP routing framework. Generally, smaller tier 2 and 3 Internet service providers (ISPs) connect to larger tier 1 networks for delivery of their customers' packets to destinations outside the smaller providers' footprints.

The choice of which network to connect to for upstream bandwidth is important because it affects the experience of end-user customers. Many networks refer to themselves as tier 1 ISPs, but their specific capabilities often vary widely and attributes may vary.

The simple approach has been to label a network tier 1 if it has large traffic volumes, large capacities, large customer bases, and large numbers of routes and if it supports many autonomous systems (ASs) inside the network.

However, size and scale are not the only dimensions of tier 1 ISPs. The key attributes of tier 1 ISPs are as follows:

- They have access to the entire Internet routing table through their peering relationships.
- They have one or two AS numbers per continent or, ideally, one AS worldwide.
- They own or lease international fiberoptic transport.
- They deliver packets to and from customers and to and from peers around the world.

Global tier 1 ISPs have two additional characteristics:

- They peer on more than one continent.
- They own or lease transoceanic fiberoptic transport to facilitate the best possible customer access experience in diverse markets on more than one continent.

These criteria matter because the Internet interconnect model is under pressure from:

- Technology developments, such as broadband, VoIP, and multiservice networks, that magnify unpredictable swings in traffic patterns and usage
- Business factors, such as industry mergers, the possible end of network neutrality, and when and how to adopt IPv6, that add uncertainty to the business of network interconnect

Global tier 1 ISPs are best positioned to support these new challenges. They have the special engineering resources and architecture required to minimize the number of router/switch hops and to manage routing policies based on class-of-service (CoS) parameters as well as participation in industry working groups to set these standards. Because of the volume of traffic they carry and exchange with other networks, they are least likely to be affected by imbalances in terms of trade resulting from merged megapeers, network toll gating, or IPv6 timing triggers.

IN THIS WHITE PAPER

There is a lot of confusion about the structure of the Internet and the classification of the networks that make up the Internet. This confusion stems from the combination of technology and business pressures that dynamically evolve the Internet interconnect model and from the sales and marketing initiatives that misrepresent the nature of Internet connectivity and confuse the differences between tiers. This white paper defines a taxonomy of Internet networks according to tier 1, 2, and 3 networks; identifies the key criteria of and measurements for tier definitions; and describes why the differences between tiers are important.

INTRODUCTION: THE TRADITIONAL ISP HIERARCHY

The global Internet is an amalgam of separate, but semiautonomous networks. Each network and server provider is an independent entity with its own policies, services, and customer targets. The binding element of the global Internet is that networks share a common IP addressing and global BGP routing framework that allows all networks to interconnect with each other directly or indirectly. There is little regulation, and network operators are free to decide where, how, and with whom to connect. Generally, smaller ISPs connect to larger networks for delivery of their customers' packets to destinations outside the smaller providers' footprints. Networks are classified by tiers based on the nature of their connection to other networks (see Figure 1). The three tiers range from smallest (tier 3) to largest (tier 1). Tier 1 is further divided into global and regional tier 1 ISPs (see Figure 2).

FIGURE 1

Hierarchy of the Global Internet

Peering is done between equivalent-sized partners (tier 3 to tier 3).

Transit or fee-based peering is done where there are unequal traffic flows (tier 2 to tier 1).

Peering and transit arrangements may be established directly or at third-party exchange points.

Many ISPs have multihoming where they connect to more than one upstream provider for diversity and reach.

Customers can connect to any ISP.

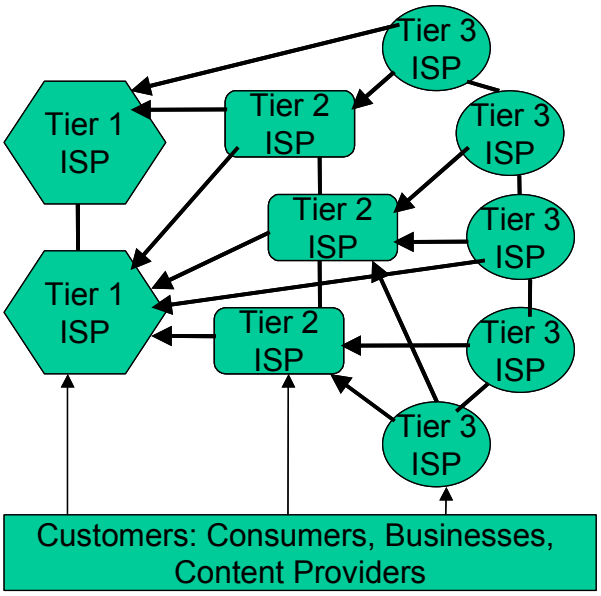

Source: IDC, 2006

FIGURE 2

Categories of ISPs

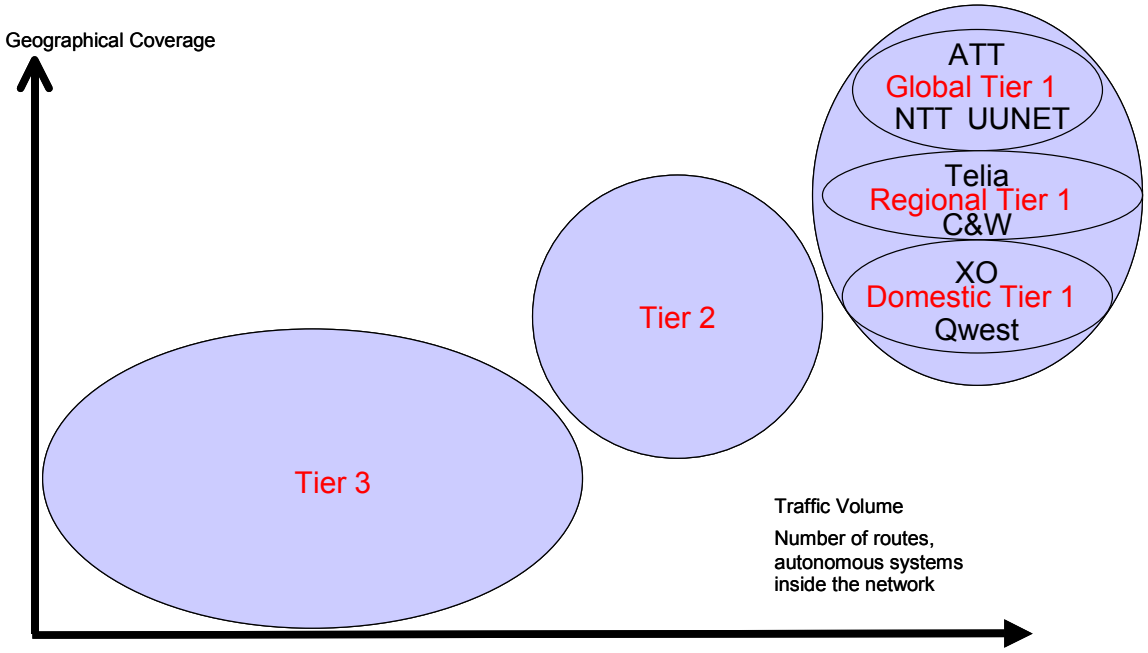

Note: The figure provides only a representative list of ISPs.

Source: IDC, 2006

ISP Tiers

Tier 1 ISPs own the operating infrastructure, including the routers and other intermediate devices (e.g., switches) that form the backbone, which is interconnected with other tier 1 ISPs via private peering in a "settlement-free" interconnection. This is also called free peering. They also interconnect at Internet Exchange (IX) points. Because a significant amount of today's Internet traffic is exchanged via private peering (discussed in the next section), tier 1 ISPs deliver the best network quality and throughput because they have the most direct control over the traffic that flows through these private peering connections. Other ISPs are completely dependent on tier 1 ISPs and their capabilities to properly manage the private peering infrastructure.

Tier 1 ISPs make use of self-owned telecommunications circuits for those parts of their networks in which they have such an infrastructure. However, this may not be the case in every market in which tier 1 ISPs operate. Tier 1 ISPs may choose to make use of circuits provided by alternative carriers because of a number of factors, including lack of self-owned circuits, contractual arrangements (e.g., reciprocity), facility availability (OC192 or DWDM might be available from another carrier in a specific market before it is available from the carrier that owns the tier 1 ISP), or a desire to maintain some level of carrier diversity to ensure more stability in the network.

Global tier 1 ISPs have their own large Internet backbones with international coverage. They have large traffic volumes, large customer bases, and large numbers of routers and support many ASs inside the network. In addition to size and scale, the key attributes of global tier 1 ISPs are as follows:

- They don't pay to have their traffic delivered though similar-sized networks.
- They have access to the entire Internet routing table, solely through their peering relationships.
- They peer on more than one continent.
- They own or lease transoceanic fiberoptic transport.
- They deliver packets to and from customers and to and from peers around the world.

Global tier 1 ISPs include:

- AboveNet (AS 6461)
- AT&T (AS 7018, AS 2686, AS 5623, and others)
- Global Crossing (AS 3549)
- Level 3 (AS 3356)
- MCI EMEA (AS 702) and MCI UUNET (AS 701/703), now part of Verizon Business

- NTT Communications (AS 2914)
- SAVVIS (AS 3561)
- Sprint (AS 1239)
- Teleglobe (AS 6453), now part of VSNL

Only a few tier 2 ISPs are able to provide service to customers on more than two continents, thus extending the characteristics. Often, they have lower-quality networks and slower access speeds than tier 1 ISPs. They are at least one router hop away from the core of the Internet. Reach and Singapore Telecom/STIX are examples of tier 2 ISPs.

Tier 3 ISPs focus on local retail and consumer markets. They provide the "on ramp" or local access to the Internet for end customers. They have many end users, but no or very few destinations (i.e., servers) on their networks. Their coverage is limited to a specific country or to subregions, such as a metropolitan area, within a country. Tier 3 ISPs are customers of higher-tier ISPs for access to the rest of the Internet. Because tier 3 ISPs' traffic requires several router hops to get out to a URL, these ISPs tend to have relatively low network quality and access speeds.

The disadvantages of tier 2 and tier 3 ISPs are the number of router hops required to get to the Internet and the oversubscription of that bandwidth. Users of lower-tier ISPs share a common gateway to higher-tier ISPs, and the gateway bandwidth may degrade the access bandwidth.

Public and Private Peering

Because the Internet is made up of a complex hierarchy of separate networks, rules are applied to define the technical and business aspects of interconnection. There are three basic forms of interconnection rules: public peering, private peering, and transit.

Peering involves the interconnection of networks for the exchange of traffic for the mutual benefit of both parties. There are two types of peering. Private peering involves the direct connection over a Layer 1 or Layer 2 link between two ISPs with similar network capacity and traffic levels. This is a bilateral agreement with traffic exchange dedicated exclusively to the two ISPs connecting. The two parties involved in private peering are committed to maintaining adequate bandwidth to keep packet loss as close to zero as possible. Proper management of private peering by a provider is a key measure of the overall quality of the performance of the Internet for the provider's customers.

Historically, ISPs peered with each other for the express purpose of reducing transit costs and expanding their coverage of packet transfers. A significant percentage of the Internet's traffic is exchanged via private tier 1 peering agreements. However, the costs involved with direct peering make it a very exclusive model; typically, the highest-tier ISPs participate in this model.

Public peering is the predominant model for exchanging traffic between ISPs. It was originally accomplished at network access points (NAPs) primarily in the United States. This role is now complemented by about 250 IX points around the world. NAPs and IXs facilitate the exchange of traffic over Layer 2 fabrics such as ATM and Ethernet. A large exchange point may bring together hundreds of tier 1, 2, and 3 ISPs in a central physical location for access to multiple networks over a shared connection. Unlike the bidirectional private arrangements, public peering enables multiple streams of traffic exchange.

The IX role has expanded over the past few years to incorporate value-added services such as collocation and billing services for paid peering. IXs have also become cost-effective exchanges for private peering. Some IXs are nonprofit companies that are owned by multiple ISPs and other entities; others are for-profit specialist companies such as Equinix. Among the largest NAP/IXs are:

- London LINX and Amsterdam AMS-IX in Europe
- Palo Alto PAIX, San Jose MAE-West in North America
- Johannesburg JINX in Africa
- Sao Paulo in South America
- JPNAP and JPIX in Japan, HKIX in Hong Kong, and Seoul KIX in Asia

While public peering permits many networks to interconnect via a more cost-effective shared connection, it often does so at the expense of performance, service-level agreements (SLAs), and rules/considerations by other participants. Many public peering points are overloaded and create sources of packet loss, which results in the current standard of "best-effort" level of service. Commercially operated exchange points (such as PAIX and JPNAP) are always striving to equal the performance of private peering connectivity, and these exchanges are moving incredible amounts of traffic without packet loss today. For example, JPNAP carries over 80Gb of traffic. Some providers are even using these exchange points to support their private peering activities. However, it is hard to beat a tier 1 ISP that relies on well-managed private peering relationships and backs them up with SLAs.

The nature and purpose of peering vary around the world. In the United States, most public peering or IX points are operated by telephone companies such as AT&T and Verizon or by neutral collocation providers that are for-profit companies such as Equinix. Outside the United States, many public peering points are operated by academic/government research networks or by nonprofit membership organizations, but a few (e.g., JPNAP) are operated by commercial organizations.

Transit is a simpler form of interconnection. Most tier 1 and some large tier 2 ISPs are willing to sell dedicated access to their networks via private leased-line telecommunications circuits. Transit costs include the circuit required for the ISP and the variable cost associated with the traffic carried upstream to the Internet. This form of interconnection is attractive for smaller tier 2 and 3 ISPs that may not be located near a public IX. Also, many lower-tier ISPs have neither the technical resources nor the traffic volume to justify a private peering relationship with a higher-tier ISP.

Regional Differences in Internet Structure

There are regional differences in Internet connectivity. Key factors include the nature of international and interregional traffic flows, the depth of broadband access adoption, and the physical geography.

Many of the top Internet destination sites are in North America, so a lot of traffic originating in Europe and Asia is bound for the United States. As a result, significant investment in international connectivity is required to deliver Internet access and the ability to manage latency for quality customer experience. Also, Asian connectivity to European sites is burdened with delay if routed via the United States. Asian ISPs need an international backbone provider that has efficient connectivity to the United States and to Europe.

Asia has seen dramatic growth in the number of Internet users and Internet access hours. According to IDC forecasts, Asia/Pacific (excluding Japan) broadband access subscribers will rise from 57 million in 2005 to 104 million in 2009. IDC projects that broadband connections in Japan will grow from 23 million in 2005 to 34 million in 2009. Moreover, IDC estimates that fiber-to-the-home connections in Japan will increase from 4 million in 2005 to 13 million in 2009.

Mainland China is one of the fastest-growing online markets in the world, and the percentage of total Internet users will increasingly be weighted toward Asia, especially China. The surge in demand for data and voice traffic in the region is creating new challenges for Internet congestion controls. To grow their own transit business and develop into tier 1 ISPs, Chinese ISPs are controlling traffic flows in and out of China by constraining the existing peering relationships and making it difficult for other ISPs to gain Chinese peering.

The growing broadband Internet momentum in the Asia/Pacific region will increase the need for communications in that market as well as in North America and Europe. An international backbone provider with strong regional connectivity delivers the best solution for Asian ISPs. The Asian markets — especially Japan, South Korea, Hong Kong, and Taiwan — are leading the transition to higher-speed broadband. As broadband adoption in Asia outpaces that in North America and Europe, the new network challenges are felt first in Asia. There are several dimensions to this scenario. Strong intraregional and international connectivity is needed to support online gaming applications flowing across Korea, Japan, and Taiwan.

The diversity of the Asian geography means greater dependency on submarine cables for intraregional connectivity (e.g., APC, APCN2, C2C, i2i, TIS) and for connectivity to Europe, the United States, and the Middle East (e.g., Japan-US, RNAL, Sea-Me-We 3, Sea-Me-We 4). A backbone provider that can activate bandwidth capacity in the timeliest manner (via cable ownership or leasing) will have the best cost structure to compete effectively for intraregional and international transit business.

THE EVOLVING INTERNET INTERCONNECT MODEL

The Internet interconnect model is under pressure from technology developments and business factors.

Technology Developments

Proliferating Broadband

The proliferation of broadband access increases traffic loads, but more significantly for network architectures, it changes users' behavior. At the end of 2005, IDC estimated the worldwide broadband market at 185 million subscriptions, including DSL, cable modem, and fiber-based facilities. This total represents nearly 50% of worldwide Internet households. IDC projects that worldwide broadband subscriptions will exceed 317 million by 2009. Approximately 80% of the world's broadband connections in 2005 are DSL with average speeds of 3MB downstream. A small, but fast-growing number of households are getting access to direct fiber connections (e.g., Japan with more than 4 million fiber-connected households), which provide exponential speed increases and generate a significant impact to backbones. ADSL2+, VDSL, and fiber initiatives are moving ahead in North America and Europe, with most carriers planning for gradual expansion to 50MB services to cover 40% of households and 10–20MB to cover most of the balance.

In contrast to narrowband access, which constrains users' behavior both in time and across applications, broadband opens a much wider range of activities and explorations. This range adds increased variability between users and magnifies the difference between peak and average use. Content providers and application developers have taken advantage of new broadband behavior to proliferate new applications — VoIP, collaboration, video streaming, multiplayer online gaming, IP TV — that drive higher broadband usage. All of this adds unpredictability to traffic flows.

Broadband was originally designed and targeted as a consumer residential service; businesses did not consider it a reliable solution. However, broadband is gaining momentum across all business sectors from small single-site businesses to large, geographically dispersed enterprises. Because businesses are more sensitive to latency and packet loss than are consumers, traffic handoff policy is critical. Most DSL providers hand off traffic at the nearest peering point and no longer concern themselves with transiting traffic. This approach triggers several risks for business users. Traffic handoff with no end-to-end performance monitoring means local Web browsing and application transit times may result in unsatisfactory user experiences. Because several upstream providers may be involved, no one has end-to-end responsibility for trouble resolution or performance. Interprovider connectivity may introduce congestion and performance degradations. To deliver business-quality services, ISPs need to find ways to develop transit relationships with backbone providers that would allow SLA performance guarantees to be extended to the edge.

The diversifying behavior of high-speed broadband households and the heightened performance demands of business users put new pressures on the Internet interconnect model. Tier 2 and 3 ISPs need an upstream transit partner with robust, high-quality Internet connectivity to deliver the best Internet access to end customers.

Emergence of Multiservice Networks

The classic Internet relies on dynamic routing to determine optimum resource utilization to convey the transmission between source and destination nodes. The routing control searches for and selects a path, determined as a function of the network state. In this way, traffic naturally gravitates to the best connection.

The tendency to integrate services — voice, data, video — defines a new and more complex routing panorama. It requires a network to introduce a more personalized treatment for each type of service. Routing policies must now base their decisions not only on the topology of the network but also on the CoS parameters demanded by the traffic flows and on the actual ability of the network to guarantee this quality via SLAs.

CoS facilitates the provisioning of differentiated levels of service and resource assurance across a network. It enables application prioritization to guarantee performance of VoIP calls, VPN links, and mission-critical applications.

Differentiated services require SLA policy management facilities that define service levels for optimum price-performance packages. The multiservice network requires tools to define policies that control network behavior and intelligently shape network bandwidth to deliver systemwide service-level guarantees.

Extending CoS parameters and SLAs, as well as traffic controls and security protections, across the Internet hierarchy is a challenge for the traditional peering structures. Interdomain CoS is required to support integrated service network and premium applications. Most peering relationships do not involve SLAs that enable ISPs to guarantee their business end customers four-hour mean-time-to-repair conditions on end-to-end service. It is true that both parties in a peering relationship benefit from rapid outage restoration and may include language in their peering agreement to diligently correct problems. However, a peering agreement has little "contractual" power, so if an imbalance in CoS occurs, a customer relationship may work better than a peering relationship.

Peering on the Internet has been and continues to be based on the assumption of symmetry in traffic streams. However, traffic asymmetry and investment asymmetry mean that one party bears more of the cost as a result of peering. For example, content providers (e.g., CNN or Amazon) are a huge source of traffic, in contrast to consumers or businesses that do not provide but instead want access to content. End users use their Internet connections to send short queries for Web pages, while content providers send a large amount of data in response, such as the Web page. When an end user–heavy network (i.e., a network with few destination addresses) exchanges traffic with a content-heavy network (i.e., network with lots of servers and URL destinations), the content-heavy network has a high outgoing-to-incoming traffic flow.

Multinational and multicontinent networks create congestion and traffic flow control challenges. Traditionally, large enterprises like their local offices to have local connections to the Internet because they provide good user experiences and accommodate local Web site interest areas. However, compliance challenges are forcing enterprises to centralize Internet access to a few gateway connections that aggregate traffic, providing more secure connections than is possible via local

connection, as they can be more effectively monitored and audited according to compliance guidelines. The trend toward defragmenting the IT environment for more centralized and controlled operations puts greater pressure on the network to make around-the-world look like next door.

VoIP and the Challenges of Jitter SLAs

There is growing pressure on ISPs to enter the VoIP market. ISPs have traditionally sold Internet access and other data services (e.g., email, Web hosting, security services). VoIP is an opportunity for ISPs to tap into their customers' spending on voice telephony and represents an important new revenue source.

When voice calls are delivered over the Internet, the quality of the calls is affected by three factors: packet loss, latency, and jitter.

- Packet loss is caused by packet buffer overflow and relates to bandwidth, congestion, and network function or dysfunction.
- Latency can be caused by end points, gateways, or the network and relates to propagation delay, serialization, and hop count.
- Jitter is caused by the IP network and occurs when there is variation in delay because of bandwidth congestion.

Underlying IP network performance is a prerequisite for acceptable voice quality. Standard guidelines for VoIP networks include the ability to consistently deliver less than 1% packet loss, 2 milliseconds of jitter or less, and less than 150 milliseconds of latency in one direction.

Unlike data communication, voice communication occurs in real time and does not tolerate packet delay on the network. Voice is symmetric and requires a network that can send and receive at the same bandwidth. Data communication is often asymmetric, like DSL and cable modem with different upstream and downstream bandwidth. Voice requires relatively little bandwidth — 64kbps can be sufficient for toll-quality voice — but it requires no fluctuations and no dropped packets. Data traffic tends to be bursty and is relatively insensitive to a small number of dropped packets.

The demands of voice traffic create challenges for the data network. Despite being a low-bandwidth application, the latency and sensitivity of voice traffic require special engineering. In the IP data network, routers are optimized for data streams. Engineering the network for VoIP is a nontrivial problem and requires understanding of IP network as well as voice services. ISPs that want to deliver VoIP services to their customers need a backbone network provider whose infrastructure is set for the minimal number of router/switch hops and arrange to have multiple queues on those devices so that voice traffic has its setting on highest priority queue. A single global tier 1 backbone has advantages over a regional tier 1 or tier 2 backbone that requires multiple router hops per region or per continent.

IPv6

Routing over the Internet is done via the IPv4 addressing scheme. IPv4 is a 32-bit address technology that was originally defined in 1981. Since then, the steady upward trajectory of network-connected devices points to a shortage of IP addresses. To extend the life of IPv4, network address translation (NAT) has been widely applied to conserve addresses. By providing a way to hide many internal addresses behind a single or small number of routable addresses, NAT has been effectively used. But workarounds create their own problems as maintenance costs become excessive and the end-to-end model is broken, so NAT is not sustainable.

In contrast to IPv4, IPv6 is a 128-bit protocol — four times larger than IPv4. The practical effect of a fourfold size increase in the number of bits per address is an exponential increase in the number of practical addresses. Because of its ability to deal with address space exhaustion and other key improvements, IPv6 delivers several benefits, including pervasive networking, lower facilities costs, and improved security.

From a technical point of view, transition to IPv6 is a good thing, but retrofitting networks with new addressing schemes is an expensive and disruptive process. In the United States and many parts of the world today, there are no clear rewards or triggers for IPv6 adoption. The proliferation of IP services and of IP-connected devices will eventually compel IPv6 adoption. Meanwhile, working with backbone ISPs that have expertise in IPv6 will benefit ISPs when the time comes. NTT Communications deployed the first worldwide IPv6 dual-stack backbone in 2003 and has the most experience providing commercial IPv6 service.

Business Factors

Industry Restructuring

The telecommunications industry is consolidating and integrating formerly diverse back-office operations. In the United States, the mergers of SBC (AS 7132) with AT&T (AS 7018 and others) and with BellSouth (AS 6197) (expected to be completed in late 2006) and of Verizon (AS 19262) with MCI (AS 701 MCI UUNET, AS 702 MCI EMEA, AS 703 MCI APAC) will affect the peering community. These are the largest companies, usually ranking among the top 10 carriers of Internet traffic.

The merging companies have stated that all peering relationships should maintain the status quo. Maintaining peering status quo has been imposed as a condition of approving the mergers in the United States and Europe. The FCC approved the AT&T-SBC and Verizon-MCI mergers on the condition that they maintain the same number of peers for three years (i.e., through 2008). However, concentration of the Internet backbone market combined with the sharp traffic inflation of a merged entity is likely to have a ripple effect through the peering structure. Because a larger traffic volume makes a peer a more powerful competitor, there is a strong disincentive to peer. By creating a bigger on-net footprint where traffic originates and terminates within one network, merger companies can become more restrictive in their peering terms.

The likelihood of unbalanced traffic flows due to the predominance of users versus content providers is magnified by the mergers. Both AT&T (formerly SBC) and Verizon have a disproportionately strong position in end-user Internet access, because of their legacies as local access DSL providers. The backbones of AT&T and Verizon (formerly MCI) tend to be dominated by business customers, which are not content providers. Consequently, the merged Verizon-MCI and AT&T-SBC have resulted in megapeers that primarily carry end-user traffic. This situation could affect the peering structure if the new AT&T or Verizon leverages its market position to:

- Degrade the quality of interconnection between peers such as slow transfer speeds
- Depeer other ISPs and refuse to accept their traffic
- Impose higher transit or paid-for peering charges

As China continues to grow its internal infrastructure and works to create a tier 1 capability through its control of the regulatory environment, it will also drive changes in how existing Internet suppliers will adapt to provide services to this emerging and important new market. It could affect the peering structure by creating more intra-Asia peering points, which could keep more and more intra-Asia traffic from leaving Asia for peering exchanges in the United States. This situation would weaken the strength of the United States–centric tier 1 carriers.

Network Neutrality

Network neutrality is a modern term for common carriage and specifies no discrimination on content or applications carried over the network, or on devices connected to the network. This has been the principle by which the Internet architecture has been designed to date, and many claim it is responsible for explosive innovations such as universal browser, instant messaging (IM), blogging, and peer-to-peer (P2P) downloads.

Network operators in the United States, Canada, Europe, and elsewhere are now calling the principle of network neutrality into question. It is understandable that networks want to maximize the return on their investments in broadband networks by charging a premium for special treatment of certain kinds of content. If adopted, a network toll-gating model has implications for the open network assumptions that most ISPs make for their future network deployments. Moreover, traffic imbalances could result from toll gating. Working with a large tier 1 ISP is the best way to protect from the unknown impact of network toll gating.

Bandwidth Brokers

Many ISPs go to bandwidth brokers (e.g., Arbinet, Interoute) for their Internet capacity needs. The bandwidth broker market works by posting buy or sell orders for Internet capacity to certain routes at specified prices and route quality. The bandwidth broker automates the buy-sell process, handles invoices and payment, and manages credit risk. The proposition is that it is easier and more efficient to go to a central market rather than individually negotiate and buy access to the network services of other communications services providers.

It is important that buyers understand what they are getting when purchasing at a bandwidth broker. Buyers can specify price, quality, and protocol, but they have little visibility into the architecture or routing practices of the seller. It is not known whether the bandwidth available is supplied by a tier 1, 2, or 3 ISP. Moreover, many of the technology and business factors driving change in the structure and performance of the Internet are not accounted for in the service descriptions available from a bandwidth broker.

CONCLUSION

There has always been a dynamic evolving structure in the nature of interconnects among the thousands of ISPs that make up the global Internet. Broadband connections are used by nearly 50% of worldwide Internet households. This widespread use of broadband accelerates application innovation and diversity on the Internet and magnifies variability in usage behaviors.

To survive the increasing pace of change and innovation, ISPs should work with large, well-funded backbone network providers. ISPs that want to unlock new revenue opportunities need to differentiate between backbone providers that are large and those that are large and that have the ability to support the performance needed for broadband applications, multiservice networks, VoIP, and so forth. The global tier 1 networks supported by companies with a tradition of financial stability and technical innovation will make the best partners when seeking the highest-quality, most stable Internet access.