How the Net Works: A Brief History of Internet Interconnection

BRET SWANSON > February 21, 2014

The Internet is an historic technological, social, and commercial success. It is also a success of self-organization and self-governance. Building something so complex requires exquisite planning by individuals and teams creating the hardware and software to power such a sprawling system. It also requires a conceptual framework that provides just enough commonality to make the pieces work together, but not so much top-down instruction that the system cannot adapt, grow, evolve, and innovate.[1]

We celebrate the Internet's dynamism – most apparent in the ever expanding choices of content, services, and devices that attach to it. Less heralded, but no less important, however, are the networks that power the whole system and the increasingly complex and creative ways all our networks connect to one another.

As the Internet grows in complexity and commercial importance, new network players, new network economics, and new interconnection practices can cause friction among the participants. Some argue we need new laws or regulations to govern the Internet from on high. But with all the industry's positive momentum, abandoning self-governance and commercial give-and-take would be a mistake. The market has proven it will adapt as circumstances change.

We have not reached the end of the line in network innovation. Cloud computing, mobile, real-time telepresence, and other network intensive services will require more bandwidth, more coverage, more connectivity, more up-time, and lower latency, all functions that will require more hyper-connected network capacity. The existing organic process, where engineers and businesses make pragmatic technical and financial decisions, is, in this dynamic environment, far more likely than government mandates to drive growth and accommodate unpredictable innovations.

The Early Internet

A brief history of the Internet helps make the point. In 1969, engineers working on a Department of Defense contract connected the

Interconnection circa 1998

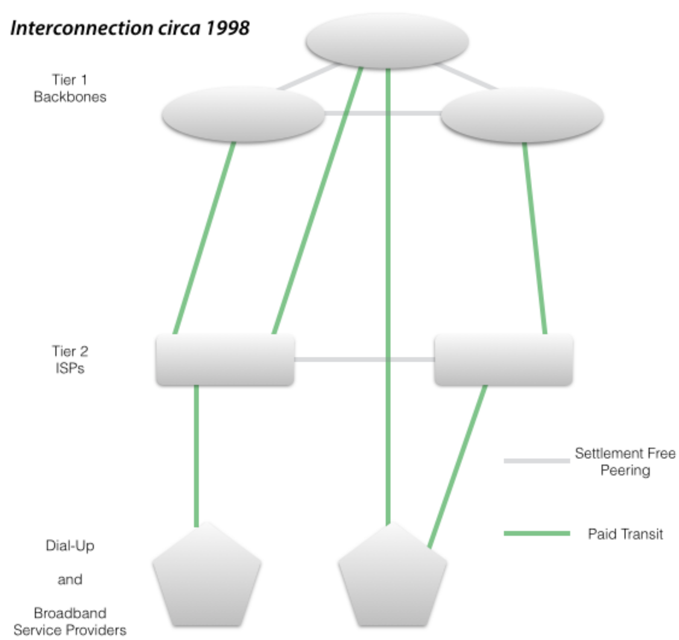

campus computer networks of UCLA, UC Santa Barbara, Stanford Research Institute, and the University of Utah. Arpanet, the seed of the Internet, was born.

Through the 1970s, more universities and government researchers joined Arpanet, and distinct teams built other experimental networks. Engineers created some of our well known languages and protocols, such as TCP/IP and Ethernet, but they also tried others that did not survive. In the 1980s, the National Science Foundation helped upgrade the backbone network from its original 50 kilobit-per-second telephone lines to faster 1.44-megabit T1 lines, and later to 45-megabit T3s. Private entities, such as UUNET and PSInet, however, also began building backbone networks. We started calling these data networks, collectively, "the Internet."

Getting all these systems to work together was a highly collaborative process. The Internet's early "stakeholders" circulated some one thousand Request for Comment (RFC) memos on protocols and interconnection schemes. In 1984, the domain naming system (.com, .edu, .gov) went into effect, and soon after practitioners from across the globe created two key groups – the Internet Society and the Internet Engineering Task Force – that would help develop the standards and customs that drove the next wave of growth. Between 1985 and 1987 the number of Internet hosts jumped from 2,000 to 30,000, then to 160,000 in 1989, and to one million by 1993.

By the early 1990s, the World Wide Web and Netscape browser shifted the Internet into an even higher gear. In 1990, NSF had lifted commercial restrictions on the NSFNET, and in 1995, NSF privatized it.

Connecting the First Networks

During this period of expanding usage and new, private networks, a number of "exchange points," or network meeting places, emerged. MAE-East, Commercial Internet eXchange (CIX), NSF's Network Access Points (NAPs), and, later, MAE-West and Palo Alto Internet Exchange (PAIX) connected the various networks to one another. These were physical locations where the cables of the various networks connected to allow data traffic to flow from one to another.

This was an unregulated arena, so unlike the world of telecom at the time, with its government-set tariffs, geographic boundaries, and access charges, the Internet players were making up the technical and commercial rules as they went along.

At the exchange points, some of the larger networks with roughly equal traffic flows agreed to trade data traffic at no cost. They called it "settlement free peering," and the choice of words was appropriate. "Peers" were networks similar in size and capability. Because most of the traffic was email, text, and Web pages, traffic tended to be roughly

Interconnection Terms

Tier One ISP — a large continental or global network that, through its own infrastructure and its peering relationships with other networks, can reach any point on the Internet. It does not pay others for transit.

Tier Two ISP — a network, often regional in nature, that connects broadband service providers, content providers and websites, and enterprises to larger Tier One networks. These entities pay Tier Two networks for transit to the Tier One networks, and Tier Two networks pay Tier One networks for transit to the rest of the Internet.

Content Delivery Network (CDN) — a network of computers and "caches" that stores data, webpages, and videos close to end users and optimizes routes across the Internet, both logically and geographically. Content providers and websites pay CDNs to speed their content to end users. Some large content providers like Google have their own CDNs.

Transit — a network access service in which, most often, a smaller entity or network pays a larger entity or network for access to the larger network. Consumers pay their broadband service provider for "transit" to the Internet. Broadband service providers, Tier Two ISPs, and CDNs pay Tier One ISPs for "transit" to the Internet.

Settlement Free Peering — an interconnection agreement in which two networks trade traffic with one another at no cost.

Paid Peering — an interconnection agreement in which networks trade traffic with one another but, because the traffic is "asymmetric" (one network is carrying far more data than the other, incurring higher costs), the party carrying less traffic pays the other a fee to make up the disparity.

symmetrical. Each network was likely to give and receive similar amounts of traffic to the other networks with whom it "peered." Why engage in extra financial transactions with one another if the payments would just cancel out?

Smaller networks and the early Internet access providers like Compuserve and AOL purchased "transit" connections to the larger Internet backbones. These "Tier 2" Internet service providers thus paid to gain access directly to a "Tier 1" Internet backbone and, because the large backbones peered with one another, all points across the Internet. Transit providers could thus be thought of as "ISPs for ISPs."

The First Web Boom

The Internet exploded in the mid- to late-1990s, and its architecture continued to change. Between 1994 and 1996, Internet traffic grew 100-fold, or 10-fold two years in a row. And commercial Tier 1 backbones struggled to keep up. The exchange points were no longer up to the task of establishing enough connectivity, in the right places, in a timely manner. So the backbone networks started to connect to one another in a wider number of large markets using metro area circuits.

Peering politics was sometimes fierce. Networks fought with each other over who was Tier 1 versus Tier 2 and bickered over interconnection terms. Each network carrier wanted, as much as possible, the other networks to connect with it at its preferred location on its preferred terms. (In many ways, this is happening again today.) And yet the market successfully adjusted to the changing environment.

By 2000, a new model was emerging — the large, carrier-neutral, data exchange center. A company called Equinix proposed this new model. It would build large, modern, secure data centers and allow all comers to connect inside its facilities on their own terms. Because it supplied only the meeting space, Equinix marketed itself as a neutral party, a sort-of open super hub for all types of network and content firms. It was a place where you knew all the other networks would have a presence and where, as peering expert Bill Norton described, "large-scale peering interconnections could be established within 24 hours rather than 24 months." [2]

At about the same time, in the late-1990s, two other significant dynamics were changing the interconnection market — broadband access networks for consumers and content delivery networks.

Broadband Access Providers

The cable TV firms grew up serving their customers video content, first via antennas on tops of hills and then via large satellite collectors at their "head-end" facilities in each town or market. The cable firms did not have connections to cross-country or global telecom networks. But the advent of the cable modem meant cable needed a path to the Internet. In the late-1990s, cable's chief links to the Internet were through paid transit arrangements from Tier 2 ISPs such as @Home and Roadrunner.

During the technology crash of 2000, however, @Home failed, and the cable firms began buying transit directly from the Tier 1 backbone providers. The cable firms noticed something else. Much of their traffic was being sent to and from other cable providers. Instead of employing a Tier 2 ISP to reach the Tier 1 backbone, who would then connect to yet another Tier 2 ISP, and then down to the cable firm, why not just establish direct connections with other cable firms?

The broadband service providers — the cable firms and telecom DSL networks — thus began directly exchanging traffic with one another, often inside the new neutral exchange point data centers. Because they were carrying so much traffic within their own customer bases, the larger cable companies, such as Comcast, also began building larger nationwide backbones of their own.

Content Delivery Networks

As the visual Web grew in the late-1990s, content firms, including big dot-coms, news sites, and ecommerce providers, needed to get closer to end users. If an Internet user in New York clicked on a webpage hosted on a server in San Francisco, the content of that webpage would have to traverse the country, often taking indirect routes through as many as 17 router and switching "hops." (A hop is a physical node on the network — a router or a switch — that data packets touch on the way from origin to destination. More hops mean a less direct transmission, more electronic processing of packets, and ultimately slower and less reliable delivery of packets.) The physical distance and high hop-counts delayed the delivery of packets to the end user and eroded the experience, especially for photos, artwork, banner ads, and other multimedia content. Content providers, who purchased transit through Tier 2 and even Tier 1 ISPs, were dissatisfied.

Akamai, one of the first content delivery networks (CDNs), offered a solution. Replicate and store the most popular webpages and other content in multiple servers, strategically placed geographically and with more closely-coupled connections to broadband access networks. This would reduce both the light speed delay and the hop delay and might even reduce a content provider's transit bill.

Content firms and websites paid CDNs to get their content closer to end users. CDNs, which consist of tens of thousands of geographically dispersed servers running specialized software that optimizes routes across the Internet, would often pay for multiple high-throughput connections to the broadband providers at strategic points around the country, and around the world.

Few of the early Internet pioneers could have imagined these creative network innovations happening within their conceptual framework, but there were even bigger changes on the way.

Web Video and the Hyper Giants

Launched in 1998, Google, by 2003-04, was growing so fast that it was rapidly taking over entire data centers where it rented space. In 2006, Google acquired YouTube, and with broadband access networks now delivering multi-megabit speeds, Web video exploded. Google needed not just its own data centers but its own content delivery networks and global fiber network. It built them all.

Soon, Microsoft, Facebook, Amazon, Apple, and other content and software firms would do the same. The largest content firms (later dubbed "Hyper Giants" by network scientist Craig Labovitz) had suddenly become some of the world's largest network firms. This was a silent revolution.

Netflix, the DVD-by-mail company, meanwhile, launched its Web streaming service, and seemingly overnight became one of the biggest bandwidth users on the planet.

The rise of Web video did something else. It substantially altered the mix of downstream and upstream data traffic. Video is thousands of more times bandwidth-intensive than text or webpages, and for movies, sports, and video clips, it is nearly all downstream. That is, end users consume vastly more traffic than they put back into the network.

Transit payments had always been used by smaller networks or content providers seeking connectivity with more end points (that is, seeking to reach a larger audience). And settlement free peering often made sense between similarly situated networks — for example, between two Tier 1 ISPs. But in the past, the traffic and payment flows were simpler and more hierarchical (see network maps on page 5). In general, end users paid broadband service providers and content providers, who paid Tier 2 ISPs, who paid Tier 1 ISPs.

In the new world of YouTube, Netflix, and CDNs, however, an even larger share of the traffic is one-way, at least on many portions

Interconnection, Then and Now — *These figures show simplified network maps, circa 1998 and 2014. Notice the big changes over a mere decade and a half — more players, new connection types, the rise of the "hyper giants," and greater overall complexity. Also notice that the Internet is composed of a mix of paid transit, paid peering, and settlement free peering relationships, among others. (Lines connecting specific firms do not necessarily represent actual network or business relationships. Rather, they show typical connections and business transactions between firms of the type shown — i.e., broadband service provider, Tier 1 backbone, CDN, content firm, etc.)*

c. 1998

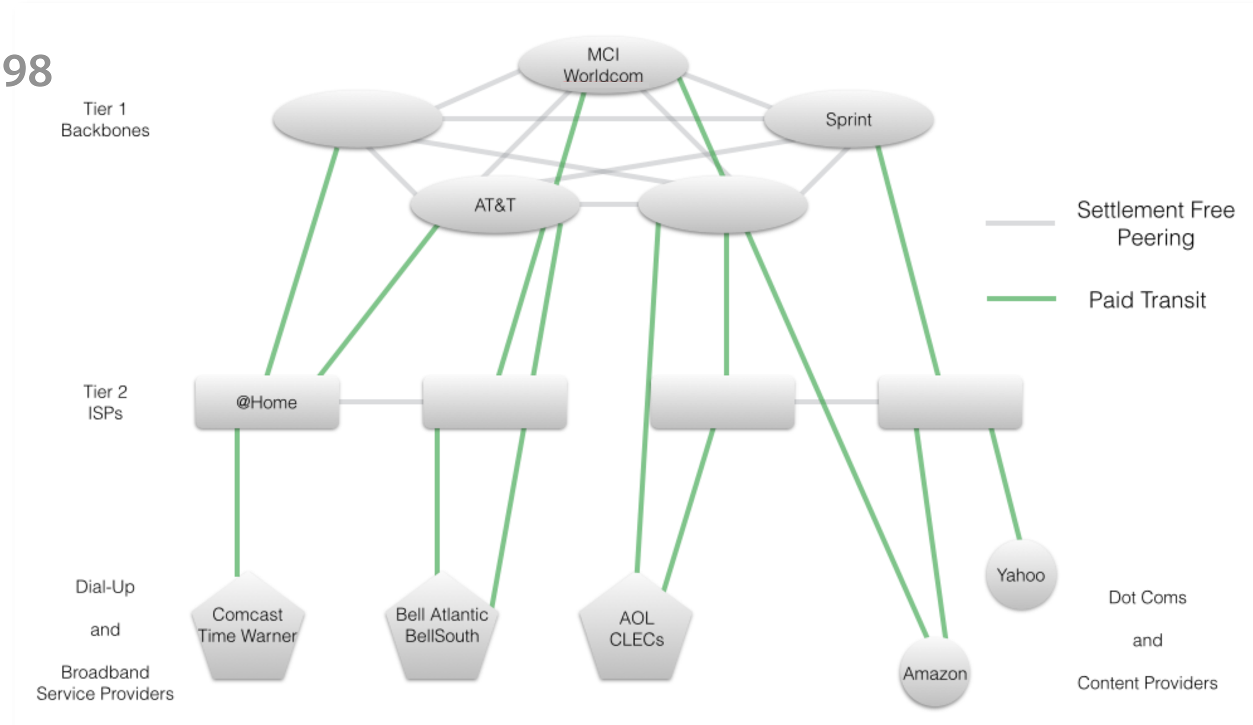

c. 2014

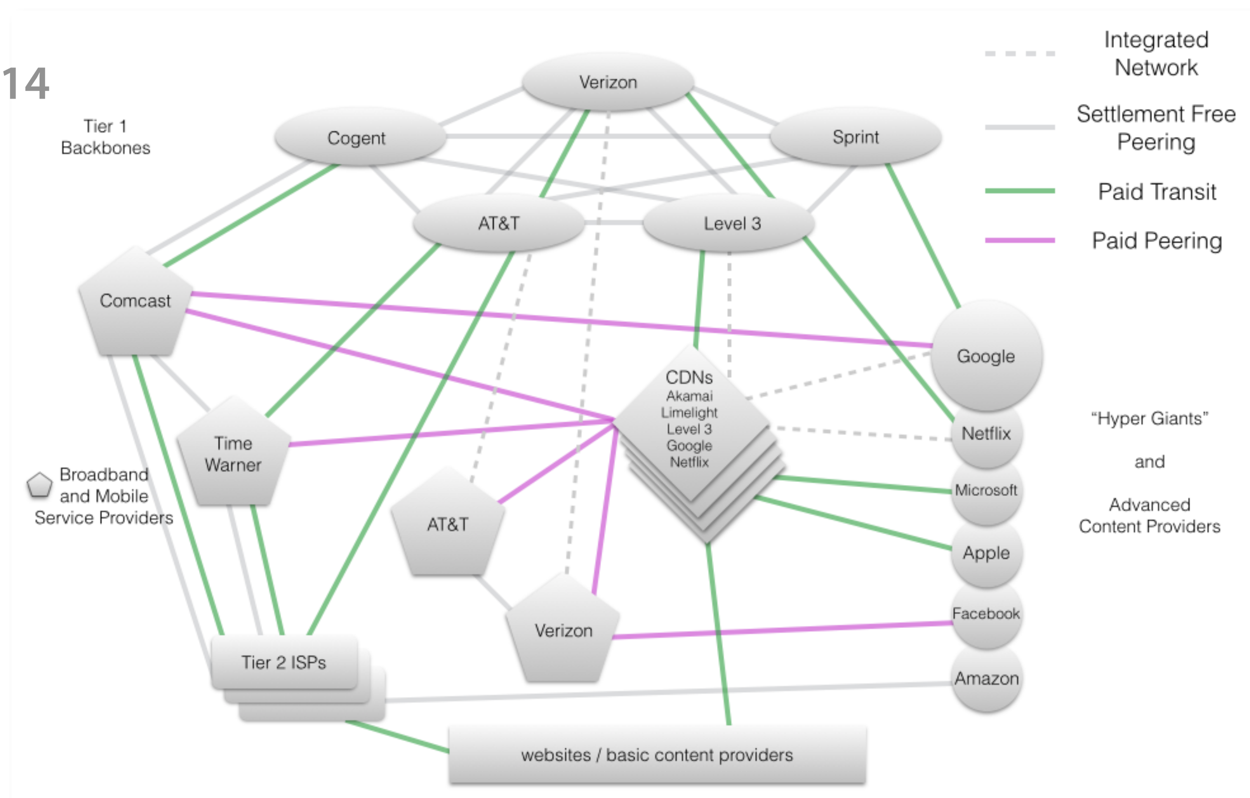

of the network, at many times of the day. And the traffic does not necessarily simply flow “vertically” up to Tier 1 backbone networks and back down. More networks and content providers often connect to one another more directly — or “horizontally” — and in more places (again, see page 5). More networks and content providers thus use more varied and more sophisticated paid transit arrangements and even “paid peering” to account for these highly “asymmetric” traffic flows.

A Rare Public Battle

In 2010, Comcast, Level 3, and Netflix engaged in a high profile battle over the ways Netflix’s traffic would reach customers on Comcast’s network. Level 3 and Comcast had both transit and peering relationships. And Netflix, through CDNs, had paid Comcast for access. But Netflix and Level 3 had an idea. If Netflix housed its content within Level 3, it could deliver its video to Comcast for free as if it were a peer. Level 3 would enter the CDN business and host the Netflix content for a lower price than other CDNs were charging Netflix to connect to Comcast. Level 3 would get a little extra revenue, and Netflix would cut costs by by routing this traffic over Level 3’s settlement free peering links. Comcast would get the downside. Firms reorganize their network operations and business relationships often, and there is nothing wrong with seeking more efficient architectures.

Comcast, however, noticed a significant spike in traffic coming from Level 3 (due to Netflix) and pointed out that this violated its peering agreement with Level 3. Settlement-free peering, remember, had long been limited to situations where networks exchange roughly similar amounts of traffic. Comcast believed Level 3 and Netflix were trying to game the system by exploiting the Comcast-Level 3 peering relationship to dump costs onto the Comcast network. (A network or content firm that mostly sends traffic to others, but does not carry much traffic in return, can impose large financial and network quality costs and upset the economics of the network value chain.) Comcast thus sought to adjust its agreement with Level 3 to reflect this traffic asymmetry. Level 3 and Netflix cried foul, using publicity and regulatory pressure to improve their negotiating leverage. In the end, however, the companies settled on a new agreement, the details of which were confidential — without regulatory intervention.

Considering the number of firms, the complexity of networks, and the pace of change, these episodes have been remarkably rare. The industry is highly competitive but, like most environments free from too much regulation, also highly cooperative.

Ever Changing Interconnection

None of the interconnection arrangements has totally displaced the others. Settlement free peering, Tier 1 and Tier 2 transit, paid peering, and CDNs, among other arrangements, exist side by side. Network relationships and commercial arrangements change according to the quickly advancing technological and financial realities of one of the world’s fastest moving industries.

Broadband service providers now even house within their own networks Google Global Cache (GGC) servers, which contain its most highly trafficked content. Netflix, likewise, within the last 18 months, moved most of its video content from third party CDN providers to its own OpenConnect CDN infrastructure. Netflix is also attempting to forge relationships with broadband providers where, like GGC, it would house its content directly within the broadband networks, close to end users.

By 2010, Google’s network had grown so large that, according to network scientist Craig Labovitz, it accounted for 6-7% of all Internet traffic. But by 2013, that number paled: Google, says Labovitz, now accounts for up to 25% of the Internet. Netflix, meanwhile, accounts for up to a third of the data flowing over U.S. broadband access networks in evening hours.

Despite the rapid change, tumult, and occasional friction, most of the interconnection world "just works." For example, according to a Packet Clearing House survey of the world's 5,000 ISPs, 99.51% of peering relationships in 2011 occurred without contract, or merely on a "handshake" agreement.

The industry over many decades developed these customs because networks, by their very nature, are highly interdependent. A network that does not have good connectivity to other networks plunges in value. Connectivity is king. The incentives motivate each network player to seek the best service for its customers. ISPs and broadband service providers want their customers to be able to reach as much content as possible, as reliably as possible.

Because of the dramatic changes in content, traffic flows, and the number and type of new network players (the Hyper Giants, for example), the types and terms of interconnection agreements have continued to evolve. Paid transit, paid peering, and other network arrangements will proliferate as the Internet evolves.

The Future

Networks will continue to grow, and interconnections will continue to grow in number and complexity.

Real-time multimedia streams for cloud-based gaming, desktops, and apps will replace many kinds of localized content. These data streams (such as ultra high definition 4K video) will need geographic proximity and, in some cases, interoperability of Quality of Service (or Quality of Experience) regimes that can prioritize content across multiple networks. The delivery of cloud-based apps, services, and content to mobile devices will especially benefit from closely coupled, low-latency links between data centers and mobile access points. (Because a mobile device relies so heavily on the cloud for its computer power and data storage needs — think Siri voice search, Google Docs, or cloud gaming — and because wireless is trickier and more capacity-constrained than is wired, optimizing the links between mobile devices, wireless nodes, and cloud resources can make a big difference in the user's experience.)

Software defined networks will also make new demands on and change the nature of interconnection. Moving network functionality like security, access control, QoS/QoE, remote peering, and network configuration to the cloud will yield large efficiencies and cost savings. Some firms are even considering the centralization and thus virtualization of individual wireless base station functions in remote cloud centers. But these cloud advances will also require big capacity, low latency, and high reliability, straining network performance.

Although asymmetric traffic flows dominated the last decade of Internet content, applications like high-resolution video chatting and conferencing may finally become widespread enough to reverse at least part of that trend, producing more symmetric content.

Whatever the case, all these technologies, products, traffic flows, and business relationships are difficult to predict. The numbers and types of networks will continue to grow, as will the interconnection relationships and overall complexity. Flexibility in network architecture and business relationships is thus crucial to accommodate these innovations.

Conclusion

The Internet is an ever expanding network of networks, where the whole and its constituent parts are ever changing. Where Arpanet linked four entities, each composed of a few end points (primitive computer terminals), today's Internet links thousands of large networks, millions of smaller networks, and billions of increasingly diverse end points (PCs, smartphones, web servers, cloud clusters, cars, and machines and sensors of all types).

To link billions of end points to one another, however, requires organization, cooperation, and trillions of dollars in infrastructure investment. It requires universal standards, like the Internet Protocol (IP), so all the parts work together. But it also requires enough flexibility – in technology, architecture, and commercial relationships – to allow for innovation in networks, content, and services.

From the beginning, our networks have never stopped changing. Nor have the ways networks connect to one another, or the terms. Interconnection disputes are not new, but they have been and remain rare. The size of the Internet economy dictates there will be more disputes (as in any industry), but the industry has and will continue to resolve these disputes in a dynamic, rapidly changing environment, without regulatory involvement. EE

---

[1] The author acknowledges and thanks Verizon for supporting the research in this report.

[2] Bill Norton's website drpeering.net and his books, including *The Internet Peering Playbook, 2013 Edition*, are excellent resources for both the lay reader and the industry insider.