The Internet Protocol Journal

June 2013 Volume 16, Number 2

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

You can download IPJ back issues and find subscription information at:
www.cisco.com/ipj

ISSN 1944-1134

From The Editor

Fifteen years ago we published the first edition of *The Internet Protocol Journal* (IPJ). This seems like a good time to reflect on where the Internet is today and where it might be going in the future, instead of looking back at earlier developments the way we did in the tenth anniversary issue of IPJ.

In our first article, Geoff Huston discusses network service models, comparing the Internet to the traditional *Public Switched Telephone Network* (PSTN) in both technical and business terms, and asks if the fundamental architectural differences between these networks might explain the rather slow deployment of IPv6. Although the number of IPv6-connected users has doubled in the last year (see page 35), IPv6 still represents a small percentage of total Internet traffic.

The mobile device dominates today's Internet landscape. Smartphones and tablets are starting to replace more traditional computers for Internet access. Many technical developments have made this possible, including high-resolution screens; powerful processors; and compact, long-lasting batteries. Combine such developments with numerous radio-based technologies (GPS, cellular, Wi-Fi, and Bluetooth) and you end up with a handheld device that is always connected to the network and can perform almost any task, using an appropriate "app." Improvements to communications technologies such as the deployment of *Long-Term Evolution* (LTE) cellular data networks and *Gigabit Wi-Fi* (IEEE 802.11ac) are already underway.

We asked Vint Cerf, known to many as one of the "Fathers of the Internet," to look beyond what is possible with today's Internet and today's devices and predict what the future might look like in a world where every imaginable appliance is "smart," connected to the network, and location-aware. His article takes us through some history and current trends, and then describes how the future Internet might shape many aspects of society such as business, science, and education.

According to Wikipedia, a *Data Center* is "a facility used to house computer systems and associated components, such as telecommunications and storage systems." In our final article, Alvaro Retana and Russ White discuss how developments in link-state protocols, usually associated with wide-area networks, can be applied to data center networks.

—Ole J. Jacobsen, Editor and Publisher

ole@cisco.com

Network Service Models and the Internet

by Geoff Huston, APNIC

In recent times we've covered a lot of ground in terms of the evolution of telecommunications services, riding on the back of the runaway success of the Internet. We have taken the computer and applied a series of transformational changes in computing power and size, battery technology, and added radio capabilities to create a surprising result. We've managed to put advanced computation power in a form factor that fits in the palms of our hands, and have coupled it with a communications capability that can manage data flows of tens if not hundreds of megabits per second—all in devices that have as few as two physical buttons! And we have created these devices at such scale that their manufacturing cost is now down to just tens of dollars per unit. The Internet is not just at the center of today's mass market consumer service enterprise, it is now at the heart of many aspects of our lives. It's not just the current fads of the social networking tools, but so much more. How we work; how we buy and sell, even what we buy and sell; how we are entertained; how democracies function, even how our societies are structured; and so much more—all of these activities are mediated by the Internet.

But a few clouds have strayed into this otherwise sunny story of technological wonder. Perhaps the largest of these clouds is that the underlying fabric of the Internet, the numbering plan of the network, is now fracturing. We have run out of IP addresses in the Asia Pacific region, Europe, and the Middle East. At the same time, the intended solution, namely the transition to a version of the IP protocol with a massively larger number space, IPv6, is still progressing at an uncomfortably slow pace. Although the numbers look like a typical "up and to the right" Internet data series, the vertical axis tells a somewhat different story. The overall deployment of IPv6 in today's Internet currently encompasses around 1.3 percent[0] of the total user base of the Internet, and it is possible that the actions of the open competitive market in Internet-based service provision will not necessarily add any significant further impetus to this necessary transition.

We have gone through numerous phases of explanation for this apparently anomalous success-disaster situation for the Internet. Initially, we formed the idea that the slow adoption of IPv6 was due to a lack of widely appreciated knowledge about the imminent demise of IPv4 and the need to transition the network to IPv6. We thought that the appropriate response would be a concerted effort at information dissemination and awareness rising across the industry, and that is exactly what we did. But the response, as measured in terms of additional impetus for the uptake of IPv6 in the Internet, was not exactly overwhelming.

We then searched for a different reason as to why this IPv6 transition appeared to be stalling. There was the thought that this problem was not so much a technical one as a business or a market-based one, and there was the idea that a better understanding of the operation of markets and the interplay between markets and various forms of public sector initiatives could assist in creating a stronger impetus for IPv6 in the service market. The efforts at stimulation of the market to supply IPv6 goods and services through public sector IPv6 purchase programs have not managed to create a "tipping point" for adoption of IPv6.

Some have offered the idea that the realization of IPv4 exhaustion would focus our thinking and bring some collective urgency to our actions. But although IPv4 address exhaustion in the Asia Pacific region in 2011 has created some immediate interest in IPv4 address extension mechanisms, the overall numbers on IPv6 adoption have stubbornly remained under 1.5 percent of the 2 billion user base of the Internet.

Why has this situation occurred? How can we deliberately lead this prodigious network into the somewhat perverse outcomes that break to basic end-to-end IP architecture by attempting to continue to overload the IPv4 network with more and more connected devices? What strange perversity allows us to refuse to embrace a transition to a technology than can easily sustain the connection needs of the entire silicon industry for many decades to come and instead choose a path that represents the general imposition of additional cost and inefficiency?

Perhaps something more fundamental is going on here that reaches into the architectural foundations of the Internet and may explain, to some extent, this evident reluctance of critical parts of this industry to truly engage with this IPv6 transition and move forward.

Telephony Network Intelligence

Compared to today's "smart" phone, a basic telephone handset was a remarkably basic instrument. The entire telephone service was constructed with a model of a generic interface device that was little more than a speaker, a microphone, a bell, and a pulse generator. The service model of the telephone, including the call-initiation function of dialing and ringing, the real-time synchronous channel provision to support bidirectional speech, all forms of digital and analogue conversion, and of course the call-accounting function, were essentially all functions of the network itself, not the handset. Although the network was constructed as a real-time switching network, essentially supporting a model of switching time slots within each of the network switching elements, the service model of the network was a "full-service" model.

The capital investment in the telecommunications service was therefore an investment in the network—in the transmission, switching, and accounting functions.

Building these networks was an expensive undertaking in terms of the magnitude of capital required. By the end of the 20th century the equipment required to support synchronous time switching included high-precision atomic time sources, a hierarchy of time-division switches to support the dynamic creation of edge-to-edge synchronous virtual circuits, and a network of transmission resources that supported synchronous digital signaling. Of course although these switching units were highly sophisticated items of technology, most of this investment capital in the telephone network was absorbed by the last mile of the network, or the so-called "local loop."

Although the financial models to operate these networks varied from operator to operator, it could be argued that there was little in the way of direct incremental cost in supporting a "call" across such a network, but there is a significant opportunity or displacement cost. These networks have a fixed capacity, and the requirements for supporting a "call" are inelastic. When a time slot is being used by one call, this slot is unavailable for use by any other call.

Telephony Tariffs

Numerous models were used when a retail tariff structure for telephony was constructed. One model was a "subscription model," where, for a fixed fee, a subscriber could make an unlimited number of calls. In other words the operator's costs in constructing and operating the network were recouped equally from all the subscribers to the network, and no transaction-based charges were levied upon the subscriber. This model works exceptionally well where the capacity of the network to service calls is of the same order as the peak call demand that is placed on the network. In other words, where the capacity of the network is such that the marginal opportunity or displacement cost to support each call is negligible, there is no efficiency gain in imposing a transactional tariff on the user. In the United States' telephone network, for example, a common tariff structure was that the monthly telephone service charge also allowed the subscriber to make an unlimited number of local calls.

Another model in widespread use in telephony was of a smaller, fixed service charge and a per-transaction charge for each call made. Here a subscriber was charged a fee for each call (or "transaction") that the subscriber initiated. The components to determine the charge for an individual transaction included the duration of the call, the distance between the two end parties of the call, the time of day, and the day of the week. This model allowed a network operator to create an economically efficient model of exploitation of an underlying common resource of fixed capacity. This model of per-call accounting was widespread, used by some operators in local call zones, and more widely by telephone service operators in long distance and international calls.

This model allowed the operator to generate revenue and recoup its costs from those subscribers who used the service, and, by using the pricing function, the network operator could moderate peak demand for the resource to match available capacity.

This per-transaction service model of telephony was available to the operator of the telephone service simply because the entire function of providing the telephone service was a network-based service. The network was aware of who initiated the transaction, who "terminated" the transaction, how long the transaction lasted, and what carriers were involved in supporting it. Initially this transactional service model was seen as a fair way to allocate the not inconsiderable costs of the construction and operation of the network to those who actually used it, and allocate these costs in proportion to the relative level of use. I suspect, though, that this fair cost allocation model disappeared many decades ago because these per-transaction service tariffs became less cost-based and more based on monopoly rentals.

IP Network Minimalism

The Internet is different. Indeed, the Internet is about as different from telephony as one could possibly imagine. The architecture of the Internet assumes that a network transaction is a transaction between computers. In this architecture the computers are highly capable signal processors and the network is essentially a simple packet conduit. The network is handed "datagrams," which the network is expected to deliver most of the time. However, within this architecture the network may fail to deliver the packets, may reorder the packets, or may even corrupt the content of the packets. The network is under no constraint as to the amount of time it takes to deliver the packet. In essence, the expectations that the architecture imposes on the network are about as minimal as possible. Similarly, the information that the edge-connected computers now expose to the network is also very limited. To illustrate this concept, it is useful to look at the fields that the Internet Protocol exposes to the network.

In IPv4 the fields of the Internet Protocol header are a small set, as shown in Figure 1. An IP packet header exposes the protocol *Version, Header Length* (IHL), *Total Length* of the IP packet, packet *Fragmentation Offset,* and *Type of Service* fields, a hop counter (*Time To Live* field), a *Header Checksum* field, and the *Source and Destination Address* fields. In practice, the Type of Service field is unused, and the Length and Checksum fields have information that is also contained in the data link frame header. What is left is the protocol Version field, packet length (Total Length field), the Fragmentation Offset field, a hop counter, and the Source and Destination Address fields. Of these fields, the Packet Length, Fragmentation Offset, hop counter, and Destination Address are the fields used by the network to forward the packet to its ultimate destination.

Figure 1: The IPv4 Packet Header

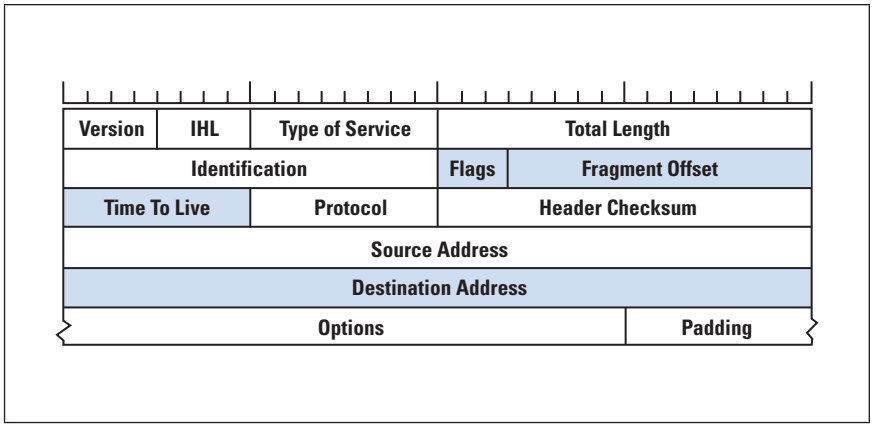

In IPv6 this minimal approach was further exercised with the removal of the Fragmentation Control fields and the Checksum fields (Figure 2). Arguably, the *Traffic Class* and *Flow Label* are unused, leaving only the *Protocol Version, Payload Length,* a *Hop Counter,* and the source and destination addresses exposed to the network. In IPv6 the minimal network-level information is now reduced to the packet length, the hop counter, and the destination address.

Figure 2: The IPv6 Packet Header

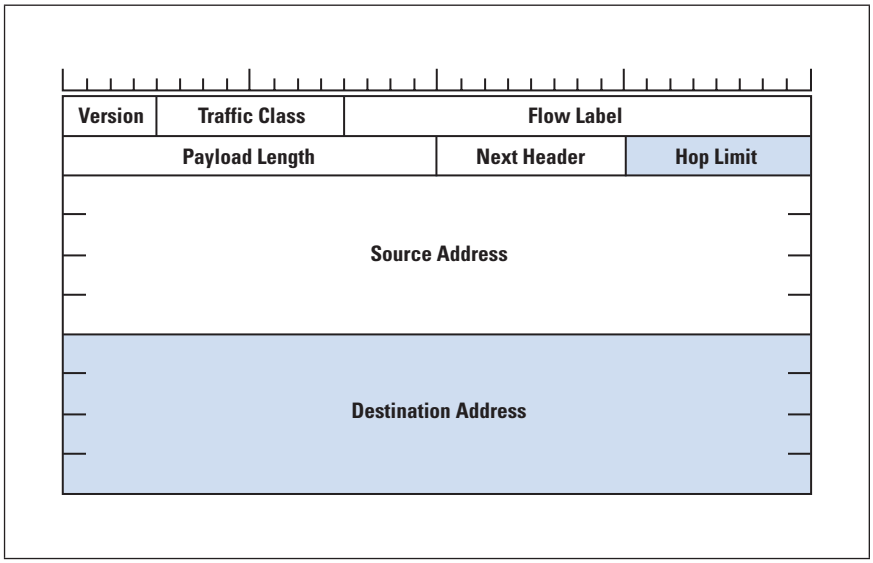

These fields represent the totality of the amount of information that the Internet Protocol intentionally exposes to the network. There are no transaction identifiers, no call initiation or call teardown signals, or even any reliable indication of relative priority of the packets. All the network needs to "see" in each carried packet is a hop counter, a packet length, and a destination address.

Within this model the actions of each of the network's switching elements are extremely simple, as shown in Figure 3.

Figure 3: IPv4 and IPv6 Packet Processing

```
for each received packet:
      decrement the hop counter
          if the counter value is zero then discard the packet, otherwise…
      look up the packet's destination address in a local table
          if the lookup fails then discard the packet, otherwise…
      look up the output queue from the located table entry
          if the queue is full discard the packet, otherwise…
          if the packet is too large for the outbound interface then
                 fragment the packet to fit, if permitted (IPv4)
                 or discard the packet (IPv6), otherwise…
      queue the packet for onward transmission
```

The Internet Service Model

What happened to "transactions" in this service model? What happened to network state? What happened to resource management within the network? What happened to all the elements of network-based communications services? The simple answer is that within the architecture of the Internet it is not necessary to expose such a detailed view of transactional state to the underlying network just to have the network deliver a packet. From a network perspective, IP has thrown all of that network level function away!

In the context of the Internet service architecture, a "transaction" is now merely an attribute of the application that is run on the end systems, and the underlying network is simply unaware of these transactions. All the network "sees" is IP packets, and each packet does not identify to the network any form of compound or multi-packet transaction.

Because a transaction is not directly visible to the IP network operator, the implication is that any effort for an IP service provider to use a transactional service tariff model becomes an exercise in frustration, given that there are no such network-visible interactions that could be used to create a transactional service model. In the absence of a network-based transactional service model, the *Internet Service Provider* (ISP) has typically used an access-based model as the basis of the IP tariff. Rather than paying a tariff per "call" the ISP typically charges a single flat fee independent of the number or nature of individual service transactions. Some basic differentiation is provided by the ability to apply price differentials to different access bandwidths or different volume caps, but this form of market segmentation is a relatively coarse one. Finer levels of transactional-based prices, such as pricing each individual video stream—or even pricing every individual webpage fetch—are not an inherent feature of such an access-based tariff structure.

The consequence for ISPs here is that within a single network access bandwidth class, this service model does not differentiate between heavy and light users, and is insensitive to the services operated across the network and to the average and peak loads imposed by these services. Like the flat-rate local telephone access model, the Internet pricing model is typically a flat-rate model that takes no account of individual network transactions. The ISP's service-delivery costs are, in effect, equally apportioned across the ISP's user base.

Interestingly, this feature has been a positive one for the Internet. With no marginal incremental costs for network usage, users are basically incented to use the Internet. In the same vein suppliers are also incented to use the Internet, because they can deliver goods and services to their customer base without imposing additional transaction costs to either themselves or their customers. For example, we have seen Microsoft and Apple move toward a software distribution model that is retiring the use of physical media, and moving to an all-digital Internet-based service model to support their user base. We have also seen other forms of service provision where the access-based tariff model has enabled services that would otherwise not be viable—here Netflix is a good example of such services that have been enabled by this flat-rate tariff structure. The attraction of cloud-based services in today's online world is another outcome of this form of incentive.

The other side effect of this shift in the architecture of the Internet is that it has placed the carriage provider—the network operator—into the role of a commodity utility. Without any ability to distinguish between various transactions, because the packets themselves give away little in terms of reliable information about the nature of the end-to-end service transaction, the carriage role is an undistinguished commodity utility function. The consequent set of competitive pressures in a market that is not strongly differentiated ultimately weans out all but the most efficient of providers from the service provider market—as long as competitive interests can be brought to bear on these market segments.

Invariably, consumers value the services that a network enables, rather than the network itself. In pushing the transaction out of the network and into the application, the architecture of the Internet also pushed value out of the network. Given that a service in the Internet model is an interaction between applications running on a content service provider's platform and on their clients' systems, it is clear that the network operator is not a direct party to the service transaction. An ISP may also provide services to users, but it is by no means an exclusive role, and others are also able to interact directly with customers and generate value through the provision of goods and services, without the involvement of the underlying network operators. It is not necessary to operate a network in order to offer a service on the Internet. Indeed, such a confusion of roles could well be a liability for such a carriage and content service provider.

The Content Business Model of the Internet

This unbundling of the service provision function from the network has had some rather unexpected outcomes. Those who made the initial forays of providing content to users believed that this function was no different from that of many retail models, where the content provider formed a set of relationships with a set of users. The direct translation of this model encountered numerous problems, not the least of which was reluctance on the part of individual users to enter into a panoply of service and content relationships. When coupled with considerations of control of secondary redistribution of the original service, this situation created some formidable barriers to the emergence of a highly valuable market for content and services on the Internet.

However, as with many forms of mass market media, the advertising market provides some strong motivation. With a traditional print newspaper, the full cost of the production of the newspaper is often borne largely by advertisers rather than by the newspaper readers. But newspaper advertising is a relatively crude exercise, in that the advertisement is visible to all readers, but it is of interest to a much smaller subset. The Internet provided the potential to customize the advertisement.

The greatest market value for advertisements is generated by those operations that gain the most information about their customers. These days it has a lot to do with knowledge of the consumer. It could be argued that Facebook's $1B purchase of Instagram was based on the observation that the combination of an individual's pictures and updates forms an amazingly rich set of real-time information about the behavior and preferences of individual consumers. It could also be argued that Google's business model is similarly based on forming a comprehensive and accurate picture of individual users' preferences, which is then sold to advertisers at a significant premium simply because of its tailored accuracy. And the mobile services are trying to merge users' current locations with the knowledge of their preferences to gain even greater value.

These developments are heading in the direction of a multiparty service model, where the relationship between a content provider and a set of users allows the content provider to resell names of these users to third parties through advertising. This on-selling of users' profiles and preferences is now a very sophisticated and significant market. As reported in [1], some 90 percent of Google's $37.9B income was derived from advertising revenue. The cost per click for "cheap car insurance" is reported in the same source to be $33.97!

The Plight of the Carrier

Although the content market with its associated service plane is now an extraordinarily valuable activity, the same is not true for the network operator—whose carriage function has been reduced from complete service-delivery management to a simple packet carrier without any residual visibility into the service plane of the network.

Obviously, network carriers look at these developments with dismay. Their own traditional value-added market has been destroyed, and the former model where the telcos owned everything from the handset onward has now been replaced by a new model that relegates them to a role similar to electricity or water reticulation—with no prospect of adding unique value to the content and service market. The highly valuable service-level transactions are effectively invisible to the carriage service providers of the Internet.

There is an evident line of thought in the carriage industry that appears to say: "If we could capture the notion of a service-level transaction in IP we could recast our service profile into a per-transaction profile, and if we can do that, then we could have the opportunity to capture some proportion of the value of each transaction."

Short of traffic interception, could the network operators working at the internet level of the network protocol stack have a means to identify these service-level transactions? The generic answer is "no," as we have already seen, but there are some other possibilities that could expose service-level transactions to the network operator.

QoS to the Rescue?

The recent calls by the The *European Telecommunications Network Operators' Association* (ETNO) advocating the widespread adoption of IP *Quality of Service* (QoS) appear to have some context from this perspective of restoring transaction visibility to the IP carriage provider. In the QoS model an application undertakes a QoS "reservation" with the network. The network is supposed to respond with a commitment to reserve the necessary resources for use by this transaction. The application then uses this QoS channel for its transaction, and releases the reservation when the transaction is complete.

From the network operator's perspective, the QoS-enabled network is now being informed of individual transactions, identifying the end parties for the transaction, the nature of the transaction and its duration, as well as the resource consumption associated with the transaction. From this information comes the possibility for the QoS IP network operator to move away from a now commonplace one-sided flat access tariff structure for IP services, and instead use a transactional service model that enables the network operator to impose transaction-based service fees on both parties to a network service if it so chooses. It also interposes the network operator between the content provider and the consumer, permitting the network operator to mediate the content service and potentially convert this gateway role into a revenue stream.

Of course the major problem in this QoS model is that it is based on a critical item of Internet mythology—the myth that inter-provider QoS exists on the Internet. QoS is not part of today's Internet, and there is no visible prospect that it will be part of tomorrow's Internet either!

Knotting up NATs

But QoS is not the only possible approach to exposing service-level transactions to the carriage-level IP network operator. Interestingly, the twin factors of the exhaustion of IPv4 addresses and the lack of uptake of IPv6 offers the IP network operator another window into what the user is doing, and, potentially, another means of controlling the quality of the user's experience by isolating individual user-level transactions at the network level.

When there are not enough addresses to assign each customer a unique IP address, the ISP is forced to use private addresses and operate a *Network Address Translator* (NAT)[2] within the carriage network.

However, NATs are not stateless passive devices. A NAT records every TCP and *User Datagram Protocol* (UDP) session from the user, as well as the port addresses the application uses when it creates a binding from an internal IP address and port to an external IP address and port. A new NAT binding is created for every user transaction: every conversation, every website, every streamed video, and literally everything else. If you were to look at the NAT logs that record this binding information, you would find a rich stream of real-time user data that shows precisely what each user is doing on the network. Every service transaction is now visible at the network level. How big is the temptation for the IP network operator to peek at this carrier-operated NAT log and analyze what it means?

Potentially, this transaction data could be monetized, because it forms a real-time data feed of every customer's use of the network. At the moment carriers think that they are being compelled to purchase and install this NAT function because of the IPv4 address situation. NATs offer a method for the carriage operator to obtain real-time feeds of customer behavior without actively intruding themselves into the packet stream. The NAT neatly segments the customer's traffic into distinct transactions that are directly visible to the NAT operator. I suspect that when they look at the business case for purchasing and deploying these *Carrier-Grade NAT* devices, they will notice a parallel business case that can be made to inspect the NAT logs and perhaps to either on-sell the data stream or analyze it themselves to learn about their customers' behavior.[3] And, as noted, there is already market evidence that such detailed real-time flows of information about individual users' activities can be worth significant sums.

But it need not necessarily be limited to a passive operation of stalking the user's online behavior. If the carriage provider were adventurous enough, it could bias the NAT port-binding function to even make some content work "better" than other content, by either slowing down the binding function for certain external sites or rationing available ports to certain less-preferred external sites. In effect, NATs provide many exploitable levers of control for the carriage operator, bundled with a convenient excuse of "we had no choice but to deploy these NATs!"

Where Now?

In contrast, what does an investment in IPv6 offer the carriage provider? An admittedly very bleak response from the limited perspective of the carriage service provider sector is that what is on offer with IPv6 is more of what has happened to the telecommunications carriage sector over the past 10 years, with not even the remote possibility of ever altering this situation. IPv6 certainly looks like forever, so if the carriers head down this path then the future looks awfully bleak for those who are entirely unused to, and uncomfortable with, a commodity utility provider role.

So should we just throw up our hands at this juncture and allow the carriage providers free rein? Are NATs inevitable? Should we view the introduction of transactional service models in the Internet as a necessary part of its evolution? I would like to think that these developments are not inevitable for the Internet, and that there are other paths that could be followed here. The true value for the end consumer is not in the carriage of bits through the network, but in the access to communication and services that such bit carriage enables. What does that reality imply for the future for the carriage role? I suspect that despite some evident misgivings, the carriage role is inexorably heading to that of a commodity utility operation.

This is not the first time an industry sector has transitioned from production of a small volume of highly valuable units to production of a massively larger volume of commodity goods, each of which has a far lower unit value, but generates an aggregate total that is much larger. The computing industry's transition from mainframe computers to mass market consumer electronics is a good example of such a transformation. As many IT sector enterprises have shown, it is possible to make such transitions. IBM is perhaps a classic example of an enterprise that has managed numerous successful transformations that have enabled it to maintain relevance and value in a rapidly changing environment.

The models for electricity distribution have seen a similar form of evolution in the last century. In the 1920s in the United Kingdom, electricity was a low-volume premium product. The prices for electricity were such that to keep just 5 light bulbs running for 1 day in a household cost the equivalent of an average week's wages. The consequent years saw public intervention in the form of nationalization of power generation and distribution that transformed electricity supply into a commonly available and generally affordable commodity.

The challenge the Internet has posed for the carriage sector is not all that different from these examples. The old carriage business models of relatively low-volume, high-value, transaction-based telecommunication services of telephony and faxes find no resonance within the service model of the Internet.

In the architecture of the Internet, it is the applications that define the services, while the demands from the underlying carriage network have been reduced to a simple stateless datagram-delivery service. Necessarily, the business models of carriage have to also change to adapt to this altered role, and one of the more fundamental changes is the dropping of the transaction-based model of the provision of telecommunications services for the carriage provider. What this situation implies for the carriage sector of the Internet is perhaps as radical as the transformation of the electricity supply industry during the period of the construction of the national grid systems in the first half of the 20th century.

The necessary change implied here is from a high-value premium service provider dealing in individual transactions across the network to that of a high-volume undistinguished commodity utility operator. The architectural concepts of a minimal undistinguished network carriage role and the repositioning of service management into end-to-end applications is an intrinsic part of the architecture of the Internet itself. It is not a universally acclaimed step—and certainly not one that is particularly popular in today's carriage industry—but if we want to see long-term benefits from the use of the Internet in terms of positive economic outcomes and efficient exploitation of this technology in delivering goods and services, then it is a necessary step in the broader long-term public interest.

References

[0] Google's IPv6 statistics:
http://www.google.com/ipv6/statistics.html

[1] Connor Livingston, "A breakdown of Google's top advertisers," **http://www.techi.com/2012/03/a-breakdown-of-googles-top-advertisers/**

[2] Geoff Huston, "Anatomy: A Look inside Network Address Translators," *The Internet Protocol Journal*, Volume 7, No. 3, September 2004.

[3] Geoff Huston, "All Your Packets Belong to Us," July 2012, **http://www.potaroo.net/ispcol/2012-07/allyourpackets.html**

GEOFF HUSTON, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region. He has been closely involved with the development of the Internet for many years, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector. He is author of numerous Internet-related books, and was a member of the Internet Architecture Board from 1999 until 2005. He served on the Board of Trustees of the Internet Society from 1992 until 2001.
E-mail: **gih@apnic.net**

The Internet: Looking Forward

by Vint Cerf, Google

As I write, it is 2013 and 40 years have passed since the first drafts of the Internet design were written. The first published paper appeared in 1974[1] and the first implementations began in 1975. Much has happened since that time, but this essay is not focused on the past but, rather, on the future. Although the past is plainly prologue, our ability to see ahead is hampered by the unpredictable and the unknown unknowns that cloud and bedevil our vision. The exercise is nonetheless worth the effort, if only to imagine what might be possible.

Current trends reveal some directions. Mobile devices are accelerating access and applications. The economics of mobile devices have increased the footprint of affordable access to the Internet and the World Wide Web. Mobile infrastructure continues to expand on all inhabited continents. Speeds and functions are increasing as faster processors, more memory, and improved display technologies enhance the functions of these platforms. Cameras, microphones, speakers, sensors, multiple radios, touch-sensitive displays, and location and motion detection continue to evolve and open up new application possibilities. Standards and open source software facilitate widespread interoperability and adoption of applications. What is perhaps most significant is that these smart devices derive much of their power from access to and use of the extraordinary computing and memory capacity of the Internet. The Internet, cloud computing, and mobile devices have become hypergolic in their capacity to ignite new businesses and create new economic opportunities.

In the near term, the Internet is evolving. The *Domain Name System* (DNS) is expanding dramatically at the top level. Domain names can be written in non-Latin characters. The Internet address space is being expanded through the introduction of the IPv6 packet format, although the implementation rate among *Internet Service Providers* (ISPs) continues to be unsatisfactorily slow. This latter phenomenon may change as the so-called *Internet of Things*[2] emerges from its long incubation. Sensor networks, Internet-enabled appliances, and increasing application of artificial intelligence will transform the Internet landscape in ways that seem impossible to imagine. The introduction of IPv6 and the exhaustion of the older IPv4 address space have generated demand for application of the so-called *Network Address Translation* (NAT)[3] system. Geoff Huston has written and lectured extensively on this topic[4] and the potential futures involving their use. In some ways, these systems simultaneously interfere with the motivation to implement IPv6 and act as a bridge to allow both network address formats to be used concurrently.

Ironically, although most edge devices on the Internet today are probably IPv6-capable, as are the routers, firewalls, DNS servers, and other application servers, this advanced version of the Internet Protocol may not have been "turned on" by the ISP community. This situation is changing, but more slowly than many of us would like.

As the applications on the Internet continue to make demands on its capacity to transport data and to deliver low-latency services, conventional Internet technologies are challenged and new ideas are finding purchase in the infrastructure. The *OpenFlow*[5, 6] concept has emerged as a fresh look at packet switching in which control flow is segregated from data flow and routing is not confined to the use of address bits in packet headers for the formation and use of forwarding tables. Originally implemented with a central routing scheme to improve efficient use of network resources, the system has the flexibility to be made more distributed. It remains to be seen whether OpenFlow networks can be interconnected by using an extended form of the *Border Gateway Protocol* (BGP) so as to achieve end-to-end performance comparable to what has already been achieved in single networks.

Business models for Internet service play an important role here because end-to-end differential classes of service have not been realized, generally, for the current Internet implementations. Inter-ISP or edge-to-core commercial models also have not generally been perfected to achieve multiple classes of service. These aspirations remain for the Internet of the present day. Although it might be argued that increasing capacity in the core and at the edge of the Internet eliminates the need for differential service, it is fair to say that some applications definitely need lower delay, others need high capacity, and some need both (for example, for interactive video). Whether these requirements can be met simply through higher speeds or whether differential services must be realized at the edges and the core of the network is the source of substantial debate in the community. Vigorous experimentation and research continue to explore these topics.

Ubiquitous Computing

Mark Weiser[7] coined the term and concept of *Ubiquitous Computing*. He meant several things by this term, but among them was the notion that computers would eventually fade into the environment, becoming ever-present, performing useful functions, and operating for our convenience. Many devices would host computing capacity but would not be viewed as "computers" or even "computing platforms." Entertainment devices; cooking appliances; automobiles; medical, environmental, and security monitoring systems; our clothing; and our homes and offices would house many computing engines of various sizes and capacities. Many, if not all, would be interconnected in communication webs, responding to requirements and policies set by users or by their authorized representatives.

To this idyllic characterization, he implied there would be challenges: configurations of hundreds of thousands of appliances and platforms, privacy, safety, access control, information confidentiality, stability, resilience, and a host of other properties.

Even modest thought produces an awareness of the need for strong authentication to assure that only the appropriate devices and authorized parties are interacting, issuing instructions, taking data, etc. It is clear that multifactor authentication and some form of public key cryptography could play an important role in assuring limitations on the use and operation of these systems. Privacy of the information generated by these systems can be understood to be necessary to protect users from potential harm.

The scale of such systems can easily reach tens to hundreds of billions of devices. Managing complex interactions at such magnitudes will require powerful hierarchical and abstracting mechanisms. When it is also understood that our mobile society will lead to a constant background churn of combinations of devices forming subsets in homes, offices, automobiles, and on our persons, the challenge becomes all the more daunting. (By this I do not mean the use of mobile smartphones but rather a society that is geographically mobile and that moves some but not all its possessions from place to place, mixing them with new ones.) Self-organizing mechanisms, hierarchically structured systems, and systems that allow remote management and reporting will play a role in managing the rapidly proliferating network we call the Internet.

For further insight into this evolution, we should consider the position location capability of the *Global Positioning System* (GPS)[8]. Even small, low-powered devices (for example, mobile devices) have the ability to locate themselves if they have access to the proper satellite transmissions. Adding to this capability is geo-location using mobile cell towers and even known public Wi-Fi locations. In addition, we are starting to see appliances such as *Google Glass*[9] enter the environment. These appliances are portable, wearable computers that hear what we hear and see what we see and can respond to spoken commands and gestures. The Google self-driving cars[10] offer yet another glimpse into the future of computing, communication, and artificial intelligence in which computers become our partners in a common sensory environment—one that is not limited to the normal human senses. All of these systems have the potential to draw upon networked information and computing power that rivals anything available in history. The systems are potentially self-learning and thus capable of improvement over time. Moreover, because these devices may be able to communicate among themselves, they may be able to cooperate on a scale never before possible.

Even now we can see the outlines of a potential future in which virtually all knowledge can be found for the asking; in which the applications of the Internet continue to evolve; in which devices and appliances of all kinds respond and adapt to our needs, communicate with each other, learn from each other, and become part of an integrated and global environment.

Indeed, our day-to-day environment is very likely to be filled with information and data gathered from many sources and subject to deep analysis benefitting individuals, businesses, families, and governments at all levels. Public health and safety are sure to be influenced and affected by these trends.

Education

It is often noted that a teacher from the mid-19th century would not feel out of place in the classroom of the 21st, except, perhaps, for subject matter. There is every indication that this situation may be about to change. In 2012, two of my colleagues from Google, Peter Norvig and Sebastian Thrun, decided to use the Internet to teach an online class in artificial intelligence under the auspices of Stanford University. They expected about 500 students, but 160,000 people signed up for the course! There ensued a scramble to write or revise software to cope with the unexpectedly large scale of the online class. This phenomenon has been a long time in coming. Today we call such classes "MOOCs" (*Massive, Open, OnLine Classes*). Of the 160,000 who signed up, something like 23,000 actually completed the class. How many professors of computer science can say they have successfully taught 23,000 students?

The economics of this form of classroom are also very intriguing. Imagine a class of 100,000 students, each paying $10 per class. Even one class would produce $1,000,000 in revenue. I cannot think of any university that regularly has million dollar classes! There are costs, but they are borne in part by students (Internet access, equipment with which to reach the Internet, etc., for example) and in part by the university (Internet access, multicast or similar capability, and salaries of professors and teaching assistants). In some cases, the professors prepare online lectures that students can watch as many times as they want to—whenever they want to because the lectures can be streamed. The professors then hold classroom hours that are devoted to solving problems, in an inversion of the more typical classroom usage. Obviously this idea could expand to include nonlocal teaching assistants. Indeed, earlier experiments with video-taped lectures and remote teaching assistants were carried out with some success at Stanford University when I served on the faculty in the early 1970s.

What is potentially different about MOOCs is *scale*. Interaction and examinations are feasible in this online environment, although the form of exams is somewhat limited by the capabilities of the online platform used. Start-ups are experimenting with and pursuing these ideas (refer to **www.udacity.com** and **www.coursera.org**).

People who are currently employed also can take these courses to improve their skills, learn new ones, and position themselves for new careers or career paths. From young students to retired workers, such courses offer opportunities for personal expansion, and they provide a much larger customer base than is usually associated with a 2- or 4-year university or college program. These classes can be seen as re-invention of the university, the short course, the certificate program, and other forms of educational practice. It is my sense that this state of affairs has the potential to change the face of education at all levels and provide new options for those who want or need to learn new things.

The Information Universe

It is becoming common to speak of "big data" and "cloud computing" as indicators of a paradigm shift in our view of information. This view is not unwarranted. We have the ability to absorb, process, and analyze quantities of data beyond anything remotely possible in the past. The functional possibilities are almost impossible to fully fathom. For example, our ability to translate text and spoken language is unprecedented. With combinations of statistical methods, hierarchical hidden Markov models, formal grammars, and Bayesian techniques, the fidelity of translation between some language pairs approaches native language speaker quality. It is readily predictable that during the next decade, real-time, spoken language translation will be a reality.

One of my favorite scenarios: A blind German speaker and a deaf *American Sign Language* (ASL) signer meet, each wearing Google Glass. The deaf signer's microphone picks up the German speaker's words, translates them into English, and displays them as captions for the deaf participant. The blind man's Glass video camera sees the deaf signer's signs, translates the signs from ASL to English and then to German, and then speaks them through the bone conduction speaker of the Google Glass. We can do all of this now except for the correct interpretation of ASL. This challenge is not a trivial one, but it might be possible in the next 10 to 15 years.

The World Wide Web continues to grow in size and diversity. In addition, large databases of information are being accumulated, especially from scientific disciplines such as physics, astronomy, and biology. Telescopes (ground and space-based), particle colliders such as the Large Hadron Collider[11], and DNA sequencers are producing petabytes and more—in some cases on a daily basis!

We seem to be entering a time when much of the information produced by human endeavor will be accessible to everyone on the planet. Google's motto: "To organize the world's information and make it universally accessible and useful," might be nearly fulfilled in the decades ahead. Some tough problems lie ahead, however. One I call "bit rot."

By using this term, I do not mean the degradation of digital recordings on various media, although this is a very real problem. The more typical problem is that the readers of the media fall into disuse and disrepair. One has only to think about 8-inch Wang disks for the early Wang word processor, or 3.5-inch floppy disks or their 5 ¼-inch predecessors. Now we have CDs, DVDs, and Blu-Ray disks, but some computer makers—Apple for example—have ceased to build in readers for these media.

Another, more tricky problem is that much of the digital information produced requires software to correctly interpret the digital bits. If the software is not available to interpret the bits, the bits might as well be rotten or unreadable. Software applications run over operating systems that, themselves, run on computer hardware. If the applications do not work on new versions of the operating systems, or the applications are upgraded but are not backward-compatible with earlier file and storage formats, or the maker of the application software goes out of business and the source code is lost, then the ability to interpret the files created by this software may be lost. Even when open source software is used, it is not clear it will be maintained in operating condition for thousands of years. We already see backward-compatibility failures in proprietary software emerging after only years or decades.

Getting access to source code for preservation may involve revising notions of copyright or patent to allow archivists to save and make usable older application software. We can imagine that "cloud computing" might allow us to emulate hardware, run older operating systems, and thus support older applications, but there is also the problem of basic input/output and the ability to emulate earlier media, even if the physical media or their readers are no longer available. This challenge is a huge but important one.

Archiving of important physical data has to be accompanied by archiving of metadata describing the conditions of collections, calibration of instruments, formatting of the data, and other hints at how to interpret it. All of this work is extra, but necessary to make information longevity a reality.

The Dark Side

To the generally optimistic and positive picture of Internet service must be added a realistic view of its darker side. The online environment and the devices we use to exercise it are filled with software. It is an unfortunate fact that programmers have not succeeded in discovering how to write software of any complexity that is free of mistakes and vulnerabilities.

Despite the truly remarkable and positive benefits already delivered to us through the Internet, we must cope with the fact that the Internet is not always a safe place.

The software upon which we rely in our access devices, in the application servers, and in the devices that realize the Internet itself (routers, firewalls, gateways, switches, etc.) is a major vulnerability, given the apparently inescapable presence of bugs.

Not everyone with access to the Internet has other users' best interests at heart. Some see the increasing dependence of our societies on the Internet as an opportunity for exploitation and harm. Some are motivated by a desire to benefit themselves at the expense of others, some by a desire to hurt others, some by nationalistic sentiments, some by international politics. That Shakespeare's plays are still popular after 500 years suggests that human frailties have not changed in the past half millennium! The weaknesses and vulnerabilities of the Internet software environment are exploited regularly. What might the future hold in terms of making the Internet a safer and more secure place in which to operate?

It is clear that simple usernames and passwords are inadequate to the task of protecting against unauthorized access and that multifactor and perhaps also biometric means are going to be needed to accomplish the desired effect. We may anticipate that such features might become a part of reaching adulthood or perhaps a rite of passage at an earlier age. Purely software attempts to cope with confidentiality, privacy, access control, and the like will give way to hardware-reinforced security. Digitally signed *Basic Input/Output System* (BIOS), for example, is already a feature of some new chipsets. Some form of trusted computing platform will be needed as the future unfolds and as online and offline hazards proliferate.

Governments are formed that are, in principle, kinds of social contracts. Citizens give up some freedoms in exchange for safety from harm. Not all regimes have their citizens' best interests at heart, of course. There are authoritarian regimes whose primary interest is staying in power. Setting these examples aside, however, it is becoming clear that the hazards of using computers and being online have come to the attention of democratic as well as authoritarian regimes. There is tension between law enforcement (and even determination of what the law should be) and the desire of citizens for privacy and freedom of action. Balancing these tensions is a nontrivial exercise. The private sector is pressed into becoming an enforcer of the law when this role is not necessarily an appropriate one. The private sector is also coerced into breaching privacy in the name of the law.

"Internet Governance" is a broad term that is frequently interpreted in various ways depending on the interest of the party desiring to define it for particular purposes. In a general sense, Internet Governance has to do with the policies, procedures, and conventions adopted domestically and internationally for the use of the Internet. It has not only to do with the technical ways in which the Internet is operated, implemented, and evolved but also with the ways in which it is used or abused.

In some cases it has to do with the content of the Internet and the applications to which the Internet is put. It is evident that abuse is undertaken through the Internet. Fraud, stalking, misinformation, incitement, theft, operational interference, and a host of other abuses have been identified. Efforts to defend against them are often stymied by lack of jurisdiction, particularly in cases where international borders are involved. Ultimately, we will have to reach some conclusions domestically and internationally as to which behaviors will be tolerated and which will not, and what the consequences of abusive behavior will be. We will continue to debate these problems well into the future.

Our societies have evolved various mechanisms for protecting citizens. One of these mechanisms is the Fire Department. Sometimes volunteer, this institution is intended to put out building or forest fires to minimize risks to the population. We do not have a similar institution for dealing with various forms of "cyberfires" in which our machines are under attack or are otherwise malfunctioning, risking others by propagation of viruses, worms, and Trojan horses or participation in botnet denial-of-service or other forms of attacks. Although some of these matters may deserve national-level responses, many are really local problems that would benefit from a "Cyber Fire Department" that individuals and businesses could call upon for assistance. When the cyber fire is put out, the question of cause and origin could be investigated as is done with real fires. If deliberately set, the problem would become one of law enforcement.

Intellectual property is a concept that has evolved over time but is often protected by copyright or patent practices that may be internationally adopted and accepted. These notions, especially copyright, had origins in the physical reproduction of content in the form of books, films, photographs, CDs, and other physical things containing content. As the digital and online environment penetrates more deeply into all societies, these concepts become more and more difficult to enforce. Reproduction and distribution of digital content gets easier and less expensive every day. It may be that new models of compensation and access control will be needed in decades ahead.

Conclusion

If there can be any conclusion to these ramblings, it must be that the world that lies ahead will be immersed in information that admits of extremely deep analysis and management. Artificial intelligence methods will permeate the environment, aiding us with smart digital assistants that empower our thought and our ability to absorb, understand, and gain insight from massive amounts of information.

It will be a world that is also at risk for lack of security, safety, and privacy—a world in which demands will be made of us to think more deeply about what we see, hear, and learn. While we have new tools with which to think, it will be demanded of us that we use them to distinguish sound information from unsound, propaganda from truth, and wisdom from folly.

References

[1] Vinton G. Cerf and Robert E. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Transactions on Communications,* Vol. Com-22, No. 5, May 1974.

[2] David Lake, Ammar Rayes, and Monique Morrow, "The Internet of Things," *The Internet Protocol Journal,* Volume 15, No. 3, September 2012.

[3] Geoff Huston, "Anatomy: A Look inside Network Address Translators," *The Internet Protocol Journal,* Volume 7, No. 3, September 2004.

[4] Geoff Huston and Mark Kosters, "The Role of Carrier Grade NATs in the Near-Term Internet," TIP 2013 Conference, **http://events.internet2.edu/2013/tip/agenda.cfm?go=session&id=10002780**

[5] **http://www.openflow.org/**

[6] William Stallings, "Software-Defined Networks and OpenFlow," *The Internet Protocol Journal,* Volume 16, No. 1, March 2013.

[7] **http://en.wikipedia.org/wiki/Mark_Weiser**

[8] **http://en.wikipedia.org/wiki/Global_Positioning_System**

[9] **http://www.google.com/glass/start/**

[10] **http://en.wikipedia.org/wiki/Google_driverless_car**

[11] **home.web.cern.ch**

[12] Cerf, V., "Looking Toward the Future," *The Internet Protocol Journal,* Volume 10, No. 4, December 2007.

[13] Vint Cerf, "A Decade of Internet Evolution," *The Internet Protocol Journal,* Volume 11, No. 2, June 2008.

[14] Geoff Huston, "A Decade in the Life of the Internet," *The Internet Protocol Journal,* Volume 11, No. 2, June 2008.

VINTON G. CERF is vice president and chief Internet evangelist for Google. Cerf has held positions at MCI, the Corporation for National Research Initiatives, Stanford University, UCLA, and IBM. He served as chairman of the board of the Internet Corporation for Assigned Names and Numbers (ICANN) and was founding president of the Internet Society. Cerf was appointed to the U.S. National Science Board in 2013. Widely known as one of the "Fathers of the Internet," he received the *U.S. National Medal of Technology* in 1997, the *Marconi Fellowship* in 1998, and the *ACM Alan M. Turing Award* in 2004. In November 2005, he was awarded the *Presidential Medal of Freedom,* in April 2008 the *Japan Prize,* and in March 2013 the *Queen Elizabeth II Prize for Engineering*. He is a Fellow of the IEEE, ACM, and AAAS, the American Academy of Arts and Sciences, the American Philosophical Society, the Computer History Museum, and the National Academy of Engineering. Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA, and he holds 21 honorary degrees from universities around the world.
E-mail: **vint@google.com**

Optimizing Link-State Protocols for Data Center Networks

by Alvaro Retana, Cisco Systems, and Russ White, Verisign

With the advent of cloud computing[6, 7], the pendulum has swung from focusing on wide-area or global network design toward a focus on *Data Center* network design. Many of the lessons we have learned in the global design space will be relearned in the data center space before the pendulum returns and wide-area design comes back to the fore.

This article examines three extensions to the *Open Shortest Path First* (OSPF) protocol that did not originate in the data center field but have direct applicability to efficient and scalable network operation in highly meshed environments. Specifically, the application extensions to OSPF to reduce flooding in *Mobile Ad Hoc Networks* (MANET)[1], demand circuits designed to support on-demand links in wide-area networks[2], and OSPF stub router advertisements designed to support large-scale *hub and spoke* networks[3] are considered in a typical data center network design to show how these sorts of protocol improvements could affect the scaling of data center environments.

Each of the improvements examined has the advantage of being available in shipping code from at least one major vendor. All of them have been deployed and tested in real-world networks, and have proven effective for solving the problems they were originally designed to address. Note, as well, that OSPF is used throughout this article, but each of these improvements is also applicable to *Intermediate System-to-Intermediate System* (IS-IS), or any other link-state protocol.

Defining the Problem

Figure 1 illustrates a small Clos[0] fabric, what might be a piece of a much larger network design. Although full-mesh fabrics have fallen out of favor with data center designers, Clos and other styles of fabrics are in widespread use. A Clos fabric configured with edge-to-edge Layer 3 routing has three easily identifiable problems.

The flooding rate is the first problem a link-state protocol used in this configuration must deal with. Router B (and the other routers in spine 2), for instance, will receive four type 1 *Link State Advertisements* (LSAs) from the four routers in spine 1. Each of the routers in spine 2 will reflood each of these type 1 LSAs into spine 3, so the other routers in spines 3, 4, and 5 will each receive four copies of each type 1 LSA originated by routers in spine 1, a total of 16 type 1 LSAs in all.

Figure 1: A Clos Fabric with Layer 3 to the Top of Rack

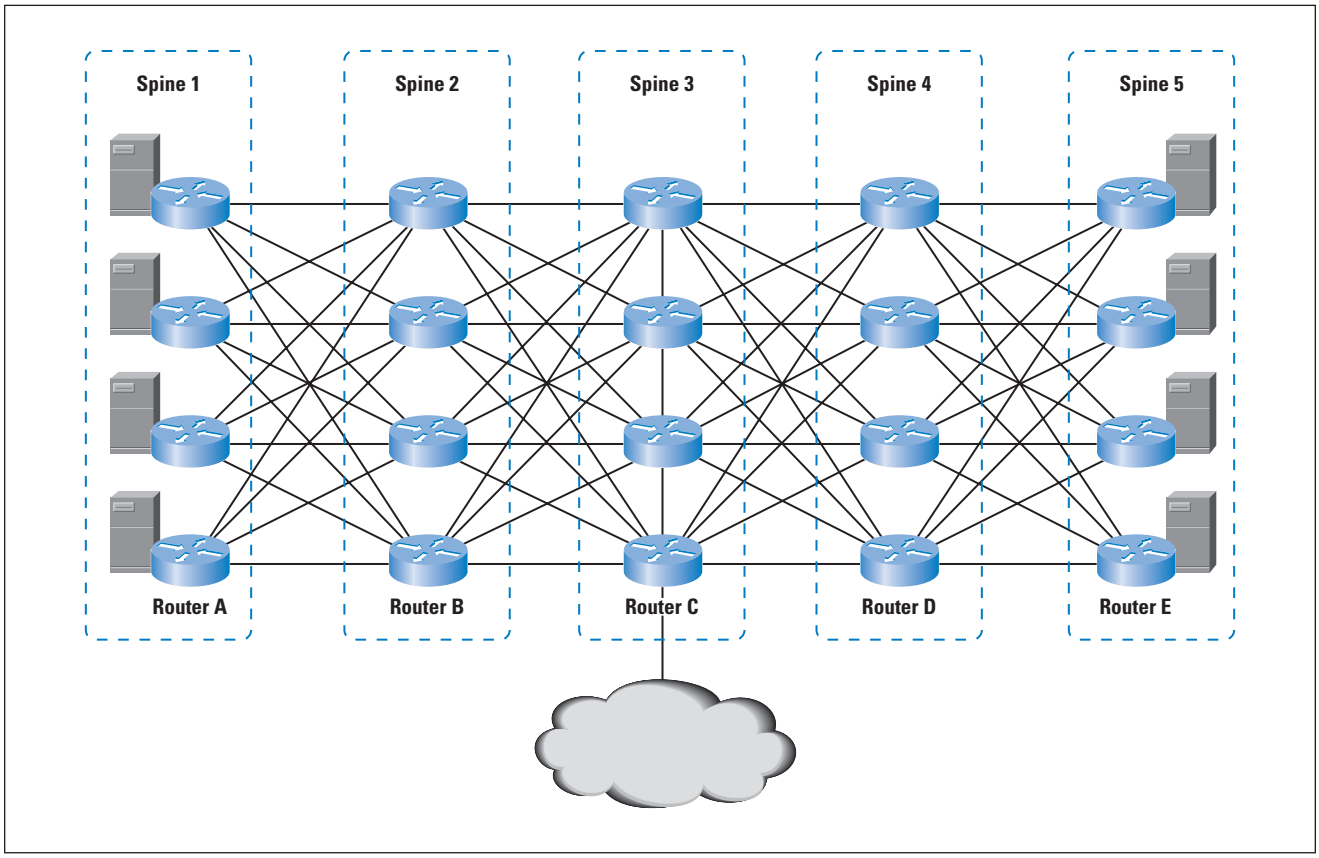

To make matters worse, OSPF is designed to time out every LSA originated in the network once every 20 to 30 minutes. This feature was originally put in OSPF to provide for recovery from bit and other transmission errors in older transport mechanisms with little or no error correction. So a router in spine 5 will receive 16 copies of each type 1 LSA generated by routers in spine 1 every 20 minutes. A single link failure and recovery can also cause massive reflooding. The process of bringing the OSPF adjacency back into full operation requires a complete exchange of local link-state databases. If the link between router A and router B fails and then is recovered, the entire database must be transferred between the two routers, even though router B clearly has a complete copy of the database from other sources.

Finally, the design of this network produces some challenges for the *Shortest Path First* (SPF) algorithm, which link-state protocols use to determine the best path to each reachable destination in the network. Every router in spine 1 appears to be a transit path to every other destination in the network. This outcome might not be the intent of the network designer, but SPF calculations deal with available paths, not intent.

This set of problems has typically swayed network designers away from using link-state protocols in such large-scale environments. Some large cloud service providers use the *Border Gateway Protocol* (BGP) (see [4]), with each spine being a separate Autonomous System, so they can provide scalable Layer 3 connectivity edge-to-edge in large Clos network topologies. Others have opted for simple controls, such as removing all control-plane protocols and relying on reverse-path-forwarding filters to prevent loops.

The modifications to OSPF discussed in this article, however, make it possible for a link-state protocol to not only scale in this type of environment, but also to be a better choice.

Reducing Flooding Through MANET Extensions

MANET networks are designed to be "throw and forget;" a collection of devices is deployed into a quickly fluid situation on the ground, where they connect over short- and long-haul wireless links, and "just work." One of the primary scaling (and operational) factors in these environments is an absolute reduction of link usage wherever possible, including for the control plane.

The "Extensions to OSPF to Support Mobile Ad Hoc Networking,"[1] were developed to reduce flooding in single-area OSPF networks to the minimal necessary, while providing fast recovery and guaranteed delivery of control-plane information. The idea revolves around the concept of an overlapping relay, which reduces flooding by accounting for the network topology, specifically groups of overlapping nodes.

Let's examine the process from the perspective of router A shown in Figure 2.

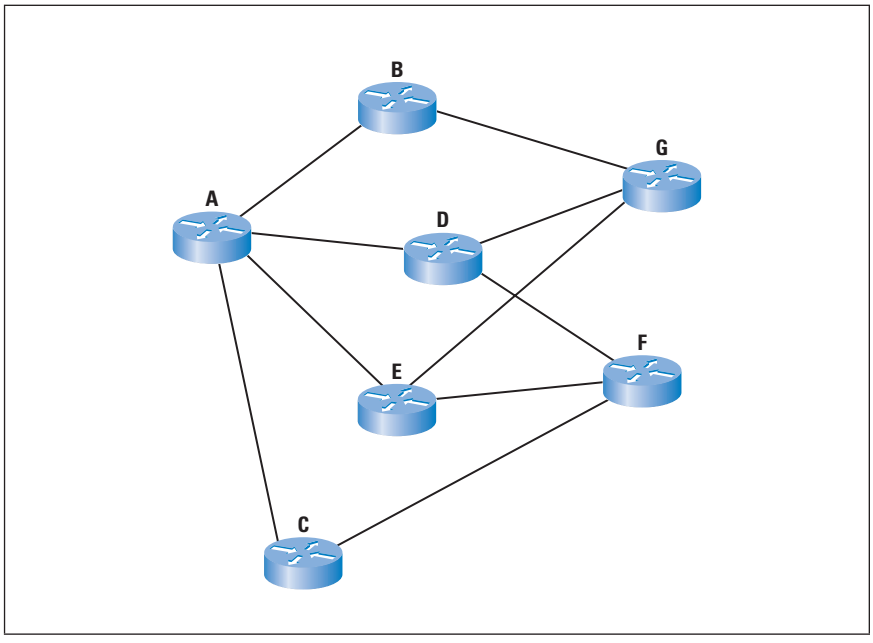

Figure 2: Ad Hoc Extensions to OSPF

Router A begins the process by not only discovering that it is connected to routers B, C, D, and E, but also that its *two-hop neighborhood* contains routers F and G. By examining the list of two-hop neighbors, and the directly connected neighbors that can reach each of those two-hop neighbors, router A can determine that if router D refloods any LSAs router A floods, every router in the network will receive the changes. Given this information, router A notifies routers B, C, and E to delay the reflooding of any LSAs received from router A itself.

When router A floods an LSA, router D will reflood the LSA to routers F and G, which will then acknowledge receiving the LSA to routers B, C, D, and E. On receiving this acknowledgement, routers B, C, and E will remove the changed LSA from their reflood lists.

Routers F and G, then, will receive only one copy of the changed LSA, rather than four.

Applying this process to the Clos design in Figure 1 and using this extension would dramatically reduce the number of LSAs flooded through the network in the case of a topology change. If router A, for instance, flooded a new type 1 LSA, the routers in spine 2 would each receive one copy. The routers in spines 3, 4, and 5 would also receive only one copy each, rather than 4 or 16.

Reducing Flooding Through Demand Circuits

Network engineers have long had to consider links that are connected only when traffic is flowing in their network and protocol designs. Dial-up links, for instance, or dynamically configured *IP Security* (IPsec) tunnels, have always been a part of the networking landscape. Part of the problem with such links is that the network needs to draw traffic to destinations reachable through the link even though the link is not currently operational.

With protocols that rely on neighbor adjacencies to maintain database freshness, such as OSPF, links that can be disconnected in the control plane and yet still remain valid in the data plane pose a unique set of difficulties. The link must appear to be available in the network topology even when it is, in fact, not available.

To overcome this challenge, the OSPF working group in the IETF extended the protocol to support demand links. Rather than attacking the problem at the adjacency level, OSPF attacks the problem at the database level. Any LSA learned over a link configured as a demand link is marked with the *Do Not Age* (DNA) bit; such LSAs are exempt from the normal aging process, causing LSAs to be removed from the link-state database periodically.

How does this situation relate to scaling OSPF in data center network design?

Every 20 minutes or so, an OSPF implementation will time out all the locally generated LSAs, replacing them with newly generated (and identical) LSAs. These newly generated LSAs will be flooded throughout the network, replacing the timed-out copy of the LSA throughout the network. In a data center network, these refloods are simply redundant; there is no reason to refresh the entire link-state database periodically.

To reduce flooding, then, data center network designers can configure all the links in the data center as demand circuits. Although these links are, in reality, always available, configuring them as demand circuits causes the DNA bit to be set on all the LSAs generated in the network. This process, in turn, disables periodic reflooding of this information, reducing control-plane overhead.

Reducing Control-Plane Overhead by Incremental Database Synchronization

When a link fails and then recovers, the OSPF protocol specifies a lengthy procedure through which the two newly adjacent OSPF processes must pass to ensure their databases are exactly synchronized. In the case of data center networks, however, there is little likelihood that a single link failure (or even multiple link failures) will cause two adjacent OSPF processes to have desynchronized databases.

For instance, in Figure 1, if the link between routers A and B fails, routers A and B will still receive any and all link-state database updates from some other neighbor they are still fully adjacent with. When the link between routers A and B is restored, there is little reason for routers A and B to exchange their entire databases again.

This situation is addressed through another extension suggested through the MANET extensions to OSPF called *Unsynchronized Adjacencies*. Rather than sending an entire copy of the database on restart and waiting until this exchange is complete to begin forwarding traffic on link recovery, this extension states that OSPF processes do not need to synchronize their databases if they are already synchronized with other nodes in the network. If needed, the adjacency can be synchronized out of band at a later time.

The application of the MANET OSPF extensions[1] to a data center network means links can be pressed into service very quickly on recovery, and it provides a reduction in the amount of control-plane traffic required for OSPF to recover.

Reducing Processing Overhead Through Stub Routers

The SPF calculation that link-state protocols use to determine the best path to any given destination in the network treats all nodes and all edges on the graph as equal. Returning to Figure 2, router B will calculate a path through router A to routers D, E, and C, even if router A is not designed to be a transit node in the network. This failure to differentiate between transit and nontransit nodes in the network graph increases the number of paths SPF must explore when calculating the shortest-path tree to all reachable destinations.

Although modern implementations of SPF do not suffer from problems with calculation overhead or processor usage, in large-scale environments, such as a data center network with tens of thousands of nodes in the shortest-path tree and virtualization requirements that cause a single node to run SPF hundreds or thousands of times, small savings in processing power can add up.

The "OSPF Stub Router Advertisement"[3] mechanism allows network administrators to mark an OSPF router as nontransit in the shortest-path tree. This mechanism would, for instance, prevent router A in Figure 1 from being considered a transit path between router B and some other router in spine 2. You would normally want to consider this option only for any actual edge routers in the network, such as the top-of-rack routers shown here. Preventing these routers from being used for transit can reduce the amount of redundancy available in the network, and, if used anyplace other than a true edge, prevent the network from fully forming a shortest-path tree.

Advantages and Disadvantages of Link-State Protocols in the Data Center

Beyond the obvious concerns of convergence speed and simplicity, there is one other advantage to using a link-state protocol in data center designs: equal-cost load sharing. OSPF and IS-IS both load share across all available equal-cost links automatically (subject to the limitations of the forwarding table in any given implementation). No complex extensions (such as [5]), are required to enable load sharing across multiple paths.

One potential downside to using a link-state protocol in a data center environment must be mentioned, however—although BGP allows route filtering at any point in the network (because it is a path vector–based protocol)—link-state protocols can filter or aggregate reachability information only at flooding domain boundaries. This limitation makes it more difficult to manage traffic flows through a data center network using OSPF or IS-IS to advertise routing information. This problem has possible solutions, but this area is one of future, rather than current, work.

Conclusion

Many improvements have been made to link-state protocols over the years to improve their performance in specific situations, such as MANETs, and when interacting with dynamically created links or circuits. Many of these improvements are already deployed and tested in real network environments, so using them in a data center environment is a matter of application rather than new work. All of these improvements are applicable to link-state control planes used for Layer 2 forwarding, as well as Layer 3 forwarding, and they are applicable to OSPF and IS-IS.

These improvements, when properly applied, can make link-state protocols a viable choice for use in large-scale, strongly meshed data center networks.

References

[0] **http://en.wikipedia.org/wiki/Clos_network**

[1] Roy, A., "Extensions to OSPF to Support Mobile Ad Hoc Networking," RFC 5820, March 2010.

[2] Abhay Roy and Sira Panduranga Rao, "Detecting Inactive Neighbors over OSPF Demand Circuits (DC)," RFC 3883, October 2004.

[3] Alvaro Retana, Danny McPherson, Russ White, Alex D. Zinin, and Liem Nguyen, "OSPF Stub Router Advertisement," RFC 3137, June 2001.

[4] Petr Lapukhov and Ariff Premji, "Using BGP for routing in large-scale data centers," Internet Draft, work in progress, April 2013, **draft-lapukhov-bgp-routing-large-dc-04**

[5] Daniel Walton, John Scudder, Enke Chen, and Alvaro Retana, "Advertisement of Multiple Paths in BGP," Internet Draft, work in progress, December 2012, **draft-ietf-idr-add-paths-08**

[6] T. Sridhar, "Cloud Computing—A Primer, Part 1: Models and Technologies," *The Internet Protocol Journal,* Volume 12, No. 3, September 2009.

[7] T. Sridhar, "Cloud Computing—A Primer, Part 2: Infrastructure and Implementation Topics," *The Internet Protocol Journal,* Volume 12, No. 4, December 2009.

RUSS WHITE is a Principle Research Engineer at Verisign, where he works on the intersection of naming and routing. In the more than 20 years since he first began working in computer networking, he has co-authored 8 technical books and more than 30 patents; he has participated in the writing, editing, and guiding of numerous Internet Standards, and he has written a fiction novel. He is currently working on *The Art of Network Architecture,* to be published by Cisco Press in 2013. Russ splits his time between the Raleigh, N.C., area and Oak Island, N.C.; he teaches in a local homeschool coop and attends Shepherds Theological Seminary. He is a regular blogger and guest on the *Packet Pushers* podcast.
E-mail: **riwhite@verisign.com**

ALVARO RETANA is a Distinguished Engineer in Cisco Technical Services, where he works on strategic customer enablement. Alvaro is widely recognized for his expertise in routing protocols and network design and architecture; he has CCIE® and CCDE® certifications, and he is one of a handful of people who have achieved the CCAR® certification. Alvaro is an active participant in the IETF, where he co-chairs the Routing Area Working Group (rtgwg), is a member of the Routing Area Directorate, and has authored several RFCs on routing technology. Alvaro has published 4 technical books and has been awarded more than 35 patents by the U.S. Patent and Trademark Office. His current interests include software-defined networking, energy efficiency, infrastructure security, routing protocols, and other related topics. E-mail: **aretana@cisco.com**

Letter to the Editor

Dear Editor,

I enjoyed reading the article on "Address Authentication" in the March 2013 edition of *The Internet Protocol Journal* (Volume 16, No. 1), but I couldn't help thinking to myself how the widespread adoption of the use of *IPv6 Privacy Addresses* (RFC 4941) would affect some of the assertions in the article about the relative merits of using IPv6 addresses for authentication. With both Microsoft and Apple operating systems now implementing IPv6 Privacy Addresses, it is now effectively impossible for any user authentication service to assume that a presented IPv6 address is going to remain constant over time. It is probably safer to assume that such IPv6 addresses are in fact not constant at all, and not to use them in any context of authentication. Given that the widespread use of NATs in IPv4 leads one to the same basic conclusion about using IPv4 addresses for authentication, isn't the best advice these days to avoid "Address Authentication" as it is applied to Internet end users?

Regards,

—*Geoff Huston*
gih@apnic.net

The author responds:

I agree with Geoff's comments. My article explores the idea that IPv6 may be more "trustworthy," but it concludes by recommending against using any IP address as a form of authentication.

IPv4 addresses will be far less "trustworthy" with the introduction of *Carrier-Grade NATs* or *Large-Scale NATs*. We will not be able to trust IPv6 addresses if the interface identifier changes frequently. My expectation is that most enterprises would prefer *Dynamic Host Configuration Protocol for IPv6* (DHCPv6) with randomized interface identifiers, but most broadband Internet access subscribers will use a *Customer Premises Equipment* (CPE) that uses *Stateless Address Autoconfiguration* (SLAAC) and Stateless DHCPv6. IPv6 offers the ability to perform traceback to the /64 subnet level. That feature is only slightly better than IPv4 traceback.

—*Scott Hogg*
scott@hoggnet.com

Book Review

Network Geeks

Network Geeks: How They Built the Internet, by Brian E. Carpenter, Copernicus Books, ISBN 978-1-4471-5024-4, 2013.

The movie opens on a familiar scene, toward the end of a congenial dinner party at the plush home of an august personage. Conversation has been casual and wide-ranging. The group retires to the library for brandy, cigars, and more conversation. Because you are new to your profession and the august personage was involved in its early years, you ask him what it was like. As he begins his recitation, the scene fades to an earlier time... "My great-grandfather, John Winnard, was born in Wigan..."

Such is the style of Brian Carpenter's book, *Network Geeks: How They Built the Internet*. Although indeed many other people are cited, the book really is Brian's personal memoir, complete with his own photographs. It explores his background and work, providing a fascinating travelogue of one person's arc through recent history. Given the breadth and scale of the 50-year process of invention and development of the global Internet, we need perhaps a thousand more such reminiscences to provide sufficiently rich detail about the many actors and acts that contributed to its success.

Brian's experiences within that global history are certainly worthy of note. His writing paints pictures of places and topics such as the forces and attractions that drew him to computer networking; in those days, it was an outlier technical topic and people often happened into it, rather than setting out with a plan. Indeed, Brian's doctoral work was in computer speech understanding—not networking. However, he has played a key role in many significant Internet activities. His frequent employer, the Swiss CERN[0], was a focal point for much of the early European networking activity—as well as being the birthplace of the World Wide Web—and Brian's various leadership roles in the *Internet Engineering Task Force* (IETF) came at pivotal times. Other popular references to Internet history tend to emphasize its American basis, making Brian's primarily European perspective refreshing and helpful.

The book is short, just 150 pages. Although Brian makes some terse references early in the book, he does not get fully into gear talking about the Internet until a third of the way through it. He started in physics, coming fully to computer science only in graduate school. Over the course of the memoir, we hear quite a bit about his physics work at CERN and elsewhere, as well as his activities with the early European deployment of Internet services, his eventual work with Internet standards, and the like.

The IETF

Brian's reference to his great-grandfather does appear, but not until page 10 in a chapter that extensively details his family history and his own upbringing—how many other books on Internet history are likely to include an inset distinguishing the English Baptist church from the American Southern Baptist? Rather, the book begins with a description of a prototypical IETF plenary session at the thrice-annual standards meeting, and he paints the picture well enough to have prompted a guessing game about the person he was describing. IETF meetings, including the plenaries, have a great deal of audience participation, because these meetings are working meetings, not conferences. I particularly enjoyed Brian's turn of phrase when describing one participant, "...who had given several articulate but incomprehensible arguments at the microphone." Later in the book he also equitably describes a colleague as "a wise leader, decisive or even pig-headed, but willing to listen..."

After its opening sequences, the book follows Brian's life chronology, including extended periods in England, Switzerland, the United States, and New Zealand, most recently landing at the last. His employment has variously been university, research, and corporate, including roles as researcher, manager, chair, and teacher.

This book is a memoir, so Brian casually and regularly moves between discussion of personal and professional developments. From one paragraph to the next, he might describe structural aspects of an Internet organization, insulation of housing in New Zealand, the next effort at particle physics, optimizing travel when flying out of southeast England, the nature of a computer networking technology, or the personal style of a co-worker.

In particular, this work is not a tutorial on Internet technology or on its invention. Although Brian does discuss many aspects of the technologies, the pedagogy suits an after-dinner evening's reminiscences, not a classroom lecture. Some concepts are explained in great detail, while others are merely cited. For example, his early discussion of computer networking references the fact that it enables mesh topologies, in contrast to then-common star configurations, but he doesn't give much sense of what "mesh" means in technical terms. Also, the core technology of networking is *packet-switching* and although his discussion on the page after the mesh reference cites queuing theory, he never introduces the motivating design construct of "store and forward."

His discussion of addressing suggests his hardware background, and misses the essence that a name at one level of architecture is often an address at the next level up. So although `www.example.com` is the "name" of a host system attached to the Internet, it has the role of "address" in a URL, because it specifies where to go to resolve the remainder of the URL.

That said, quibbling with such an issue in a tutorial might be reasonable, but it is entirely inappropriate for a memoir. These are Brian's recollections. If they prompt the reader to explore things later, so much the better; but arguing his view will not do. Perhaps reflexively, it is convenient that the Internet makes such exploration quite easy...

NATs

Except that I remain sorry to see that Brian still has such a strikingly purist view about *Network Address Translation* (NAT)[1, 2] devices, which map between internal (private) IP addresses and public ones. The purist view is that they are an abomination that breaks the elegance of the "end-to-end" design principle of the Internet. The principle is powerful, because it tends to greatly simplify the communications infrastructure and greatly enable innovation at the endpoints. The problem is that the real world imposes organizational and operational models that are more complex than easily supported by the basic end-to-end construct, at the least needing to include enterprise-level policies. NATs do cause problems, by replacing one IP address for another, and some mechanisms do cease to work because of these replacements. However, the operational world views NATs as being useful against multiple problems. One is address space constraints, which is the formal justification for creating the mechanism: an enterprise uses far fewer public IP addresses—a reality that is now essential as IPv4 addresses have grown scarce. Another justification is the misguided view that they improve enterprise security, and the other is the legitimate view that they simplify enterprise network administration. After more than 20 years of extensive deployment, these devices might be expected to have become tolerable to a pragmatist, possibly even forcing consideration of a more elaborate architectural model for the Internet. Yet Brian suffers no such weakness; NATs are evil.

One of the technical points that intrigued me was Brian's repeated discussion of the *Remote Procedure Call* (RPC). This mechanism makes network interaction for an application look like little more than a subroutine invocation. It was hoped that it would greatly simplify network-oriented programming and make it accessible to any software developer, rather than requiring the developer to have a deep understanding of networking interfaces and dynamics. Brian cites the mechanism as having been "invented by the ARPANET community in the mid-1970s..." and used at CERN in a programming language shortly after that. But my own recollection is of hearing a Xerox *Palo Alto Research Center* (PARC) manager in 1980 proudly announce that one of his summer interns had just developed the idea. Indeed, Wikipedia credits the late Bruce Jay Nelson, a Carnegie Mellon University graduate student who was working at PARC.[3]

And that is the essence of a memoir. It is the remembrances of the speaker, not the formal work of a historian or journalist. It is not the diligent unfolding of a researched history, such as in *Where Wizards Stay up Late*[4], nor the tourist approach of *Exploring the Internet: A Technical Travelogue*[5] that seeks to name every possible person active at the time—although Brian does sometimes invoke that latter template. Instead it shares one person's sense of what happened—what he remembers doing and seeing.

Railing against architectural biases or historical nuances is essential when evaluating formal professional writing, and we do need such judicious efforts to capture the history of the Internet. But had Brian sought to produce such a tome, it would not have been as rich or as personal.

References

[0] European Organization for Nuclear Research, Geneva, **http://home.web.cern.ch/**

[1] Network Address Translation, **https://en.wikipedia.org/wiki/Network_address_translation**

[2] Geoff Huston, "Anatomy: A Look inside Network Address Translators," *The Internet Protocol Journal*, Volume 7, No. 3, September 2004.

[3] Remote Procedure Call, **http://en.wikipedia.org/wiki/Remote_procedure_call**

[4] Katie Hafner and Matthew Lyon, *Where Wizards Stay Up Late: The Origins of the Internet,* Simon & Schuster, ISBN 0-684-81201-0, 1996.

[5] Carl Malamud, *Exploring the Internet: A Technical Travelogue,* Prentice-Hall, Inc., ISBN 0-13-296898-3 1992/1997, **http://museum.media.org/eti/**

—Dave Crocker,
dcrocker@bbiw.net

Number of IPv6-Connected Internet Users Doubles

The *Internet Society* (ISOC) recently reported that the number of IPv6-connected users has doubled since *World IPv6 Launch* began on June 6, 2012, when thousands of *Internet Service Providers* (ISPs), home networking equipment manufacturers, and Web companies around the world came together to permanently enable the next generation of *Internet Protocol Version 6* (IPv6) for their products and services. This marks the third straight year IPv6 use on the global Internet has doubled. If current trends continue, more than half of Internet users around the world will be IPv6-connected in less than 6 years.

"The year since World IPv6 Launch began has cemented what we know will be an increasing reality on the Internet: IPv6 is ready for business," said Leslie Daigle, the Internet Society's Chief Internet Technology Officer. "Forward-looking network operators are successfully using IPv6 to reduce their dependency on expensive, complex network address translation systems (*Carrier Grade Network Address Translators*) to deal with a shortage of IPv4 addresses. Leaders of organizations that aspire to reach all Internet users must accelerate their IPv6 deployment plans now, or lose an important competitive edge."

As IPv6 adoption continues to grow, members of the worldwide Internet community are contributing to its deployment. Statistics reported by World IPv6 Launch participants underscore the increasing deployment of IPv6 worldwide:

- Google reports the number of visitors to its sites using IPv6 has more than doubled in the past year.
- The number of networks that have deployed IPv6 continues to grow, with more than 100 worldwide reporting significant IPv6 traffic.
- Australian ISP Internode reports that 10 percent of its customers now use IPv6 to access the Internet.
- Akamai reports that it is currently delivering approximately 10 billion requests per day over IPv6, which represents a 250 percent growth rate since June of last year.
- KDDI measurement shows that the number of IPv6 users of KDDI has doubled and that IPv6 traffic has increased approximately three times from last year.

World IPv6 Launch participants have worked together to help drive adoption, leading to the creation of *World IPv6 Day* in 2011, in which hundreds of websites joined together for a successful global 24-hour test flight of IPv6.

This was followed by World IPv6 Launch in 2012, in which more than a thousand participants permanently enabled IPv6 for their products and services, including four of the most visited websites: Google, Facebook, YouTube, and Yahoo!.

As a platform for innovation and economic development, the Internet plays a critical role in the daily lives of billions. This momentum has not slowed—IPv6 adoption continues to skyrocket, fast establishing itself as the "new normal" and a must-have for any business with an eye towards the future.

For more information about companies that have deployed IPv6, as well as links to useful information for users and how other companies can participate in the continued deployment of IPv6, please visit: **http://www.worldipv6launch.org**

IPv4 has approximately four billion IP addresses (the sequence of numbers assigned to each Internet-connected device). The explosion in the number of people, devices, and web services on the Internet means that IPv4 is running out of space. IPv6, the next-generation Internet protocol which provides more than 340 trillion, trillion, trillion addresses, will connect the billions of people not connected today and will help ensure the Internet can continue its current growth rate indefinitely.

The Internet Society is the trusted independent source for Internet information and thought leadership from around the world. With its principled vision and substantial technological foundation, the Internet Society promotes open dialogue on Internet policy, technology, and future development among users, companies, governments, and other organizations. Working with its members and Chapters around the world, the Internet Society enables the continued evolution and growth of the Internet for everyone. For more information, visit: **http://www.internetsociety.org**

RIPE NCC Report on ITU WTPF-13

The RIPE NCC has published a report on the recent *ITU World Telecommunications/ICT Policy Forum* (WTPF-13). The report is available from the following URL:

https://www.ripe.net/internet-coordination/news/ripe-ncc-report-on-the-itu-wtpf-13

Any comments or questions are welcome on the RIPE Cooperation Working Group mailing list:

https://www.ripe.net/ripe/mail/wg-lists/cooperation

Google.org Awards Grant to ISOC to Advance IXPs in Emerging Markets

The *Internet Society* (ISOC) recently announced that it has been awarded a grant by Google.org to extend its *Internet Exchange Point* (IXP) activities in emerging markets. The grant will build on the Internet Society's previous efforts and will establish a methodology to assess IXPs, provide training for people to operate the IXPs, and build a more robust local Internet infrastructure in emerging markets.

IXPs play an important role in Internet infrastructure that allows *Internet Service Providers* (ISPs) and other network operators to exchange traffic locally and more cost effectively, which can help lower end-user costs, speed-up transmissions, increase Internet performance, and decrease international Internet connectivity costs. The Internet Society and Internet technical experts have been working for several years to bring IXPs to emerging markets. These efforts have resulted in locally trained experts and facilitated the development of local and regional technical infrastructures. An additional benefit of IXP development is the expansion of community governance models as well as building local Internet expertise.

Google.org, a team within Google focused on social impact, develops and supports technology solutions that can address global challenges, such as expanding Internet access to more of the world's seven billion people.

"The Internet Society has proved to be one of the most effective institutions in the Internet community," said Vint Cerf, vice president and Chief Internet Evangelist at Google. "I am confident that they will apply their grant wisely to extend their work to increase Internet access for everyone, including those in emerging markets."

Lynn St. Amour, President and CEO of the Internet Society, stated, "We are very excited to receive this grant from Google.org. With support to extend our IXP development and improvement projects, we can more quickly bring core Internet infrastructure to underserved countries and assist in building key human and governance capabilities. We will also be able to extend the Internet Society's mission to ensure the open development, evolution, and use of the Internet for the benefit of people everywhere. We look forward to working with Google.org, and we are committed to collaborating with Internet community partners around the world on this important project."

IPJ Subscription FAQ

What is my "Subscription ID" for The Internet Protocol Journal (IPJ) and where do I find it?

Your Subscription ID is a unique combination of letters and numbers used to locate your subscription in our database. It is printed on the back of your IPJ issue or on the envelope. You will also find information about your subscription expiration date near your Subscription ID. Here is an example:

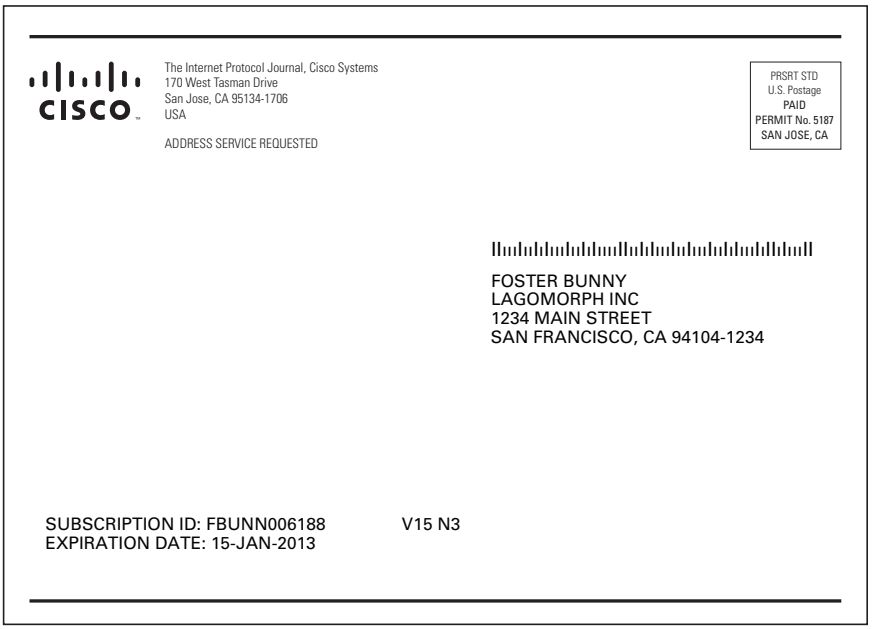

How do I renew or update my subscription?

From the IPJ homepage (`www.cisco.com/ipj`) click "Subscriber Service" and then enter your Subscription ID and your e-mail address in the boxes. After you click "Login" the system will send you an e-mail message with a unique URL that allows access to your subscription record. You can then update your postal and e-mail details, change delivery options, and of course *renew* your subscription.

What will you use my e-mail address and postal address for?

This information is used *only* to communicate with you regarding your subscription. You will receive renewal reminders as well as other information about your subscription. We will never use your address for any form of marketing or unsolicited e-mail.

I didn't receive the special URL that allows me to renew or update my Subscription. Why?

This is likely due to some form of spam filtering. Just send an e-mail message to `ipj@cisco.com` with your Subscription ID and any necessary changes and we will make the changes for you.

Do I need my Subscription ID to read IPJ online? What is my username and password?

Your Subscription ID is used *only* for access to your subscription record. No username or password is required to read IPJ. All back issues are available for online browsing or for download at `www.cisco.com/ipj`

I can't find my Subscription ID and I have since changed e-mail address anyway; what do I do now?

Just send a message to `ipj@cisco.com` and we will take care of it for you.

Call for Papers

The Internet Protocol Journal (IPJ) is published quarterly by Cisco Systems. The journal is not intended to promote any specific products or services, but rather is intended to serve as an informational and educational resource for engineering professionals involved in the design, development, and operation of public and private internets and intranets. The journal carries tutorial articles ("What is...?"), as well as implementation/operation articles ("How to..."). It provides readers with technology and standardization updates for all levels of the protocol stack and serves as a forum for discussion of all aspects of internetworking.

Topics include, but are not limited to:

- Access and infrastructure technologies such as: ISDN, Gigabit Ethernet, SONET, ATM, xDSL, cable, fiber optics, satellite, wireless, and dial systems
- Transport and interconnection functions such as: switching, routing, tunneling, protocol transition, multicast, and performance
- Network management, administration, and security issues, including: authentication, privacy, encryption, monitoring, firewalls, troubleshooting, and mapping
- Value-added systems and services such as: Virtual Private Networks, resource location, caching, client/server systems, distributed systems, network computing, and Quality of Service
- Application and end-user issues such as: e-mail, Web authoring, server technologies and systems, electronic commerce, and application management
- Legal, policy, and regulatory topics such as: copyright, content control, content liability, settlement charges, "modem tax," and trademark disputes in the context of internetworking

In addition to feature-length articles, IPJ contains standardization updates, overviews of leading and bleeding-edge technologies, book reviews, announcements, opinion columns, and letters to the Editor.

Cisco will pay a stipend of US$1000 for published, feature-length articles. Author guidelines are available from Ole Jacobsen, the Editor and Publisher of IPJ, reachable via e-mail at `ole@cisco.com`

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

PRSRT STD
U.S. Postage
PAID
PERMIT No. 5187
SAN JOSE, CA

The Internet Protocol Journal is published quarterly by the Chief Technology Office, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com