The **Internet Protocol** *Journal*

March 2005 Volume 8, Number 1

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**

From The Editor

Internet Protocol Version 6 (IPv6) continues to be the focus of much work within the IETF as well as throughout the world in numerous deployment projects. The success of IPv6 depends not only on the protocol itself but also on its interaction with existing services such as the *Domain Name System* (DNS). In our first article, David Malone looks at some issues with DNS servers and IPv6. If you are interested in following the progress of IPv6 deployment, you might want to visit The IPv6 Forum's Website at: **`http://www.ipv6forum.org`**

A couple of years ago I signed up for GSM cellphone service and later added GPRS data service to my account. With my Bluetooth-enabled phone and laptop, I can access the Internet from almost anywhere in the world. The service is neither particularly fast nor inexpensive, but for occasional use it works very well, and has "saved the day" for me numerous times. However, GPRS is not the only wide-area wireless data network technology. Kostas Pentikousis gives an overview of the many alternatives.

The term "Internet Governance" is not well-defined, but it is being used more frequently when speaking about such organizations as the *Internet Corporation for Assigned Names and Numbers* (ICANN). The formation of the *World Summit on the Information Society* (WSIS) and its *Working Group on Internet Governance* (WGIG) has certainly brought the term into sharper focus. Although governance is certainly not a technical protocol issue, we still believe that it is important for our readers to follow both the debate about and the actual evolution of Internet Governance issues. However, we fully appreciate that this is an area where opinions differ—and that is why the article by Geoff Huston on this topic is labeled "Opinion."

We remind you to visit our Website, **`http://www.cisco.com/ipj`**, where you can find back issues of this journal, search the index files, or make changes to your subscription information. Your feedback is also very much appreciated, so drop us a line at **`ipj@cisco.com`**

—Ole J. Jacobsen, Editor and Publisher
`ole@cisco.com`

Misbehaving Name Servers and What They're Missing

by David Malone, Hamilton Institute, NUI Maynooth, Ireland

IPv6-capable hosts abound, and the number is growing. Evidence[1] shows that more than 2 million Windows XP machines are probing for 6to4[2] connectivity. When combined with deployments of Linux and BSD that have been shipping with IPv6 support enabled by default for some time, that is a sizable platform on which to build IPv6 applications. Most Web browsers (Internet Explorer, Mozilla, Opera) now support IPv6 if the underlying platform does, so that is a significant number of applications ready to start making IPv6 queries.

In fact, many of these applications are already looking for IPv6 addresses in the *Domain Name System* (DNS), even if IPv6 connectivity is not actually available. This usually does not result in a problem—the name server says there are no IPv6 records and the application falls back to IPv4. In a small number of cases, name servers running outdated or errant software are misbehaving when faced with a request for an IPv6 address.

The Problem

So, what problem are these name servers having with the request for IPv6 addresses? Well, the DNS stores different types of information, such as host names and addresses. Different types of data are stored using different record types. For example, IPv4 addresses are stored using a type "A" record and host names are stored using a type "PTR" record. Some new record types have been introduced for IPv6. The most important one is "AAAA," which is for storing IPv6 addresses. (Another type called "A6" was also introduced, but it is now consigned to experimental status because it proved too complicated in certain situations.)

When you issue a request to the DNS, you indicate the domain and type of record that you are interested in. If the server has records of that type for that domain, it replies, including those records. If the server has no records of that type, it should respond saying "there are no records of this type." If the domain does not exist, then the server should return a "no such domain" error.

However, the problems arise when the DNS server does something different, and some name servers behave badly when faced with a query for a type they do not explicitly know about. For the sake of simplicity, we will highlight three wrong reactions to an unknown query that have been observed. A more complete technical analysis of the problem can be found in[3].

The first reaction that people notice is that some name servers do not reply when faced with a query for an unknown type. In this case, the person who made the request waits a while before the request is reissued. Eventually the application falls back to IPv4. "Eventually" means anything from 10 seconds to 100 seconds, depending on the operating system and application—enough to irk the casual Web user.

The second reaction is more subtle. Here the name server returns a "no such domain" response. At first glance this may seem harmless enough—the query for an IPv4 address is issued quickly. However, DNS specifications say that the "no such domain" response may be cached. This means that the "A" query is never issued, and the system acts as if the domain does not exist.

The third reaction is that the server issues some other sort of incorrect response. Usually this is less serious than the two previous reactions, because other responses at worst result in a particular name server being considered "bad" and being avoided for future queries. This means that some better-behaved name server can answer the query.

The Extent of the Problem

Although sites with these problems are sometimes discussed on mailing lists, the extent of a problem is not always proportional to the coverage it receives. Historically, numerous online advertising companies have had load-balancing DNS servers that exhibit these symptoms. Because the content of an ad server is embedded in the Web pages of many organizations, this means a single errant DNS server can give the end user the impression that this problem is more widespread than it is.

To give some idea of the scale of the problem, Table 1 shows the results of querying the name servers for the names mentioned in a month's worth of Web proxy logs. The number of servers responding in each of the three ways mentioned (no reply, no such domain, or other error) is shown, along with a total. Also shown is the number of name servers that actually returned IPv6 addresses.

These results show that actually only a small number of name servers have this problem. Unfortunately, it also looks as if the number of name servers distributing IPv6 addresses is actually comparable. However, it does look like the proportion of problem name servers is decreasing over time.

Table 1: Responses to Name Queries

Nameservers that:	January 2004	April 2004	August 2004
Responded to type A	16838	20631	17934
Did not reply to type A	64 (0.38%)	49 (0.24%)	36 (0.20%)
Returned no such domain	11 (0.07%)	19 (0.09%)	11 (0.06%)
Returned other error	22 (0.13%)	39 (0.19%)	11 (0.06%)
Had any issue with AAAA	97 (0.58%)	107 (0.52%)	58 (0.32%)
Returned AAAA records	105 (0.62%)	123 (0.60%)	18 (0.66%)

Looking at Web logs to determine the size of the problem gives us a feeling for the number of name servers that need attention. Another interesting parameter to consider is the proportion of requests that might be subject to this problem. The answer would tell us how many queries might be mishandled if your name server cannot deal with new query types.

Looking at the queries for addresses at one authoritative name server shows that 65 percent of queries are for A records, 21 percent are for AAAA records, and 14 percent are for A6 records. Although this server is IPv6-capable and might attract more queries for AAAA records, even the root servers run by RIPE show that 10 percent of address queries are for IPv6 addresses.

The Solution

Some of the name servers that exhibit this problem are simply running old versions of DNS server software. If this is the case, then the fix is simple: *upgrade!*

A significant number of the remaining problem servers are running unusual name server software, and the only way to fix the problem is to have that software fixed. Where the name server software is maintained in house, there should be enough DNS expertise to resolve the issue when it is identified. Where DNS systems have been bought in, it can be difficult to get the relevant information to the developers who can make the necessary changes. Thus increasing awareness of the issue among DNS vendors and troubleshooters is important.

In some cases[5,6], discussions on Internet mailing lists has alerted those responsible for the server to the problem and the issue has been resolved. In other cases, feedback provided by users and customers has marked IPv6 conformance as an issue for future upgrades of a site's DNS infrastructure. Unfortunately, on some occasions, feedback has been ignored and the problem has persisted. This is maybe not so surprising because it is a subtle problem. The fact that it is IPv6-related means it is sometimes dismissed because the organization thinks "we have not begun IPv6 deployment yet, so it cannot affect us."

Where problems have persisted, people have resorted to various practical solutions (hacks?) to avoid the issue. Some people, who do not need IPv6 at this time, have just suppressed the AAAA queries. Others, when they discover a name server that times out, add it to a blacklist. This avoids any delays, but may make a site unavailable. Mozilla includes a more forgiving style of blacklisting, in the form of a "ipv4OnlyDomains" setting, that can be set to a list of domains known to have problems[7].

The long-term solution seems straightforward. As we have seen, the number of name servers exhibiting this problem is relatively small, though some do serve some often-queried domains. If we can ensure that no more servers with these problems get deployed, then as the existing servers are updated or retired the problem will be resolved.

To this end, it is worth testing new DNS deployments to make sure that they correctly respond to unusual query types[8]. This will smooth the path not just for IPv6, but also for other new technology such the *Domain Name System Security Extension* (DNSSEC)[9].

References

[1] "Observations of 6to4 Traffic on a 6to4 Router," Pekka Savola, preprint, October 2004.

[2] "Connecting IPv6 Routing Domains Over the IPv4 Internet," Carpenter, Moore, and Fink, *The Internet Protocol Journal,* Volume 3, No. 1, March 2000.

[3] "Common Misbehavior against DNS Queries for IPv6 Addresses," Y. Morishita and T. Jinmei, **draft-ietf-dnsop-misbehavior-against-aaaa-01.txt**, April 2004.

[4] K-Root information page, RIPE, **http://k.root-servers.org/**

[5] "**news.bbc.co.uk** NXDOMAIN problem fixed," itojun, Simon Lockhart, et al., 6bone mailing list, April 2002.

[6] "**ftp.perl.org** strangenes" thread, Mark Andrews, Ask Bjoern Hansen, et al., freebsd-stable mailing list, March 2004.

[7] "IPv6: Some IPv4 addresses won't resolve w/IPv6 OS," Mozilla bug 68796, **https://bugzilla.mozilla.org/show_bug.cgi?id=68796**

[8] "AAAA lookup checker," David Malone, **http://www.cnri.dit.ie/cgi-bin/check_aaaa.pl**

[9] "DNSSEC," Miek Gieben, *The Internet Protocol Journal,* Volume 7, No. 2, June 2004.

DAVID MALONE received B.A. (mod), M.Sc., and Ph.D. degrees from Trinity College Dublin. He has been involved with system administration since 1994 and has been slowly growing IPv6 networks since 1999, when he also became a FreeBSD committer. With Niall Murphy, he is the coatuthor of *IPv6 Network Administration,* ISBN 0-596-00934-8 published by O'Reilly and Associates, 2005. He is currently on secondment to the Hamilton Institute of NUI Maynooth. E-mail: **dwmalone@maths.tcd.ie**

Wireless Data Networks

by Kostas Pentikousis, VTT

Most IPJ readers are familiar with *Wireless Local-Area Networks* (WLANs; see, for example, IPJ Volume 5, No. 1). Some may even be familiar with recent developments in *Wireless Metropolitan-Area Networks* (WMANs), such as WiMAX. Although nonproprietary WMAN technologies are still in the standardization phase, the IEEE 802.11 family of protocols has reached maturity and rendered inexpensive (and often free) WLAN access increasingly popular. Both WLANs and WMANs provide high-speed connectivity (in the order of tens of Mbps), but user mobility is restricted. In fact, it is probably more appropriate to talk about "portability" rather than "mobility"[1] when referring to WLANs and WMANs.

Wireless wide-area networks (WWANs), on the other hand, allow full user mobility but at data rates typically in the order of tens of kbps. This will change to some extent when *third-generation* (3G) cellular networks are fully deployed. Still, 3G deployment is slower than originally anticipated, a development often attributed to the combination of high spectrum license costs, the recent economic downturn, and high equipment costs. As a result, both population and geographical coverage tend to be uneven. For example, in Finland, a forerunner in wireless communications, population coverage is well below the 35-percent level, and geographical coverage is even smaller

This article introduces several wireless network technologies, perhaps not so widely known, which deserve attention when considering how to provide mobile connectivity to field personnel, introduce *machine-to-machine* (M2M) communication, or deploy applications that require always-on connectivity. The approach taken in this article is a bit different from the one typically followed in the literature: We focus more on higher-level issues, the information that is essential for application developers, instead of modulation, channel coding, and other low-level details. Unlike WLANs and WMANs, none of the networks surveyed provide data rates in the order of tens of Mbps. Nevertheless, successful applications can be built even with stringent bandwidth limitations. For example, online gambling and several gaming applications can be served by really "thin" networks (and possibly "thick" clients).

Cellular Networks

The *Global System for Mobile Communications* (GSM) specifies a cellular, wide-area, circuit-switched, digital mobile phone network architecture[2]. Circuit-switched networks such as GSM and IS-95, commonly referred to as *Code Division Multiple Access* (CDMA) in the United States, can provide wireless data connectivity, cover a large area, and handle mobile host handovers efficiently[3]. Users can transfer data over, say, GSM, by establishing a "dialup" connection[4]. Mobile hosts can roam, even at high speeds, and remain connected throughout.

Communication is full-duplex at a radio data rate of 9.6 kbps or 14.4 kbps in GSM Phase 2+[5]. User throughput is always smaller than the nominal radio data rate.

While the user is connected using a wireless circuit-switched network, phone calls cannot be initiated or received whether data is being transferred or not. This is not much different from wire-line dialups over basic telephone service. The difference is that a dialup over a *Public Switched Telephone Network* (PSTN) takes up a resource, namely the wire-line local loop, which is dedicated to a single user, whereas a dialup over a cellular network such as GSM consumes a resource, the radio channel, which is shared among many users. Because of the *burstiness* that data traffic usually exhibits, circuit switching may lead to inefficient use of the network capacity. Establishing a GSM dialup connection usually takes several seconds, meaning that if the user has a small amount of data to send, a small e-mail message, for example, the overall experience is poor. Moreover, after the connection is established, the channel remains idle between traffic bursts and the allocated bandwidth is wasted. Packet switching is more efficient for bursty data transmission over a shared medium[6].

Another variable that favors packet-switching over circuit-switching, especially over slow wireless networks, is *billing*. Users of circuit-switched networks are usually charged based on the duration of a connection regardless of the amount of traffic transmitted or received. On the other hand, users of packet-switched networks can be charged based solely on the amount of data transferred—not how long they remain attached to the network. In short, introducing packet switching to wireless networks can lead to better use of network resources and attract more users as data transfers become more economical.

Two-way, packet-switched WWANs permit users to roam freely indoors and outdoors, even at relatively high speeds[7]. Most WWANs employ a cellular architecture to take advantage of frequency reuse and increase capacity while covering a larger area. Furthermore, because the coverage area of a single cell is generally large (cell diameters are typically in the order of dozens of kilometers), mobile hosts do not have to go through frequent and lengthy handovers. Hosts remain connected throughout after they attach to the network, permitting users to receive and transmit data on demand without having to dial up. The following sections survey some of the most widely deployed packet-switched wireless data networks.

Mobitex

Mobitex is the first digital data-only WWAN developed by Ericsson and Swedish Telecom. Not based on IP, Mobitex was introduced in Sweden in 1986 for emergency communications[8]. It uses a cellular architecture with cell diameters of up to 30 km. Each service area can operate 10–30 channels[9] and each base station is usually allocated 1 to 4 channels. Each channel is composed of a frequency pair: different frequencies are used for the uplink and the downlink.

Communication between the base station and a single mobile host is, nevertheless, effectively half-duplex. Although base stations can transmit and receive simultaneously, mobile nodes are unable to do so[10]. The Mobitex *Maximum Transmission Unit* (MTU) is 545 bytes, with up to 512 bytes of user data. Although the system has undergone several revisions, the raw transfer rate remains only 8 kbps. Effective user throughputs range from 4 kbps (for 125-byte packets) to 4.6 kbps (for 512-byte packets)[11], and round-trip times can be up to 10 seconds.

Mobitex deals with network lapses using a store-and-forward procedure: Packets destined for a mobile node outside the network coverage area are stored while awaiting delivery. When the mobile node reconnects, the stored packets are delivered. Mobitex uses a hierarchical routing architecture that prevents local traffic from being injected into the backbone network. In other words, packets destined for a node in the range of the same base station are switched locally[8]. Besides supporting unicast addressing, Mobitex allows hosts to send one packet to several recipients[10]. According to the *Mobitex Association* (**`www.mobitex.org`**), the technology features "true push functionality," whereby data can be pushed to both a single mobile node and a predefined group of nodes, a feature that can be very useful when trying to send an urgent message to field personnel. And, because the mobile host does not have to keep querying for pending data, network traffic can be kept to a minimum. All these features can also significantly boost battery life.

According to the Yankee Group, despite the limited data rates, a variety of applications have been developed based on Mobitex, including: burglar and fire alarm systems; paging, interactive messaging, e-mail, form-based applications, and access to databases; telemetry; credit card authorizations; field service; and fleet management. Virtually all of them require small and bursty transfers. Mobitex does not lend itself to large file transfers, e-mail with large attachments, or video transmission. In fact, file transfers of more than 20 KB used to be discouraged[8]. On the other hand, by using a slotted ALOHA[12] variation for channel access, Mobitex can provide message delivery delay guarantees and support hundreds of users within the same cell. Parsa[13] calculated that Mobitex can accommodate 2,000 users per channel, assuming two uplink and two downlink messages per hour. Other networks simply cannot provide tight delay bounds for such a large number of users. For example, the *Mobile Data Magazine* (No. 1, 2002) reported that a Korean operator launched real-time stock trading and horse gambling mobile applications with great commercial success, by guaranteeing delay bounds notwithstanding the low data rates.

DataTAC

DataTAC (also known as ARDIS in the United States) was developed by Motorola in the mid-1980s. DataTAC is also a non-IP based, wide-area, data-only message-oriented network. A single base station can cover an area exceeding 20 km in diameter[14]. Like Mobitex, communication between the base station and a single DataTAC mobile node is half-duplex, and mobile hosts have to compete to get access to transmit and receive data.

Unlike Mobitex, DataTAC was designed to provide optimal in-building coverage, and it uses a cellular architecture that does not take advantage of frequency reuse. Instead, a single frequency is used, increasing the probability that a packet transmission is successful (because the same transmission can be picked up by more than one base station), but at the expense of network capacity[8]. Bodsky notes that the U.S. DataTAC operator formerly recommended refraining from transferring files larger than 10 KB.

Although neither Mobitex nor DataTAC provides native IP support, middleware can take care of protocol translation and allow unmodified, off-the-shelf applications to communicate. The maximum Data-TAC message size is 2048 bytes[15], but the maximum over-the-air packet size depends on the link layer. For rural areas the maximum radio data rate is 4.8 kbps, and the maximum over-the-air packet size is 256 bytes. In metropolitan areas, the radio data rate is 19.2 kbps and the maximum packet size is 512 bytes[16]; end-user throughput does not exceed 10 kbps on average. Traditionally, DataTAC was used for dispatching and law enforcement applications. The *Worldwide Wireless Data Network Operators Group* (**`www.datatac.com`**) reports that DataTAC networks are also used for two-way messaging, wireless e-mail, telemetry, access to corporate databases, and package tracking by courier carriers.

CDPD

Cellular Digital Packet Data (CDPD) was designed by IBM and McCaw Cellular Communications in the early 1990s to take advantage of channels that do not carry voice traffic in the *Advanced Mobile Phone Service* (AMPS), the first-generation analog cellular network[17]. Data channels are allocated dynamically, sharing the network capacity with AMPS voice traffic, which is quite different from Mobitex and DataTAC. This, for example, might mean that data can be transmitted and received only when phone calls do not consume all available capacity. One could argue that CDPD considers data traffic less important than voice. However, the standard allows network operators to specifically assign channels to data traffic only. In theory, deployment can be more economical than it is for other WWANs because CDPD takes advantage of existing AMPS infrastructure and does not require licensing new spectrum. Original projections anticipated that as CDPD gained popularity—and AMPS became obsolete—more CDPD dedicated channels would be allocated. With time, CDPD would have taken over the existing AMPS bandwidth, effectively becoming a data-only WWAN.

CDPD is based on a *Carrier Sense Multiple Access* (CSMA) variant called *Digital Sense Multiple Access*[14] and transparently provides IP services, constituting a great advantage. CDPD allows for an MTU of 2048 bytes. However, one has to account for the TCP/*User Datagram Protocol* (UDP) and IP headers that are used to encapsulate the application payload before sending it over the CDPD network and also for the fact that CDPD user data is transmitted in much smaller blocks. Although the CDPD raw data rate is 19.2 kbps, the effective throughput is in the order of 10 kbps and response times have been reported to be in the order of 4 seconds[18].

GPRS

The *General Packet Radio Service* (GPRS) is overlaid on a GSM network in a fashion similar to the way CDPD is embedded in AMPS: Voice and data traffic share the same bandwidth and network infrastructure[14]. In other words, GPRS is an add-on to GSM networks, and it requires certain hardware and software upgrades and introduces packet switching to a circuit-switched architecture. GSM voice traffic is oblivious to the presence of GPRS data traffic. Similar to CDPD, GPRS is designed to appear as a regular IP subnetwork both to hosts attached over the air interface and to hosts outside the GPRS network.

The GPRS standard was finalized by the *European Telecommunications Standards Institute* (ETSI) in late 1997 as part of GSM Phase 2+[5]. It is regarded as a transitional technology toward 3G networks[19], and is commonly referred to as 2.5G. One of its main advantages is that the same device can be used to transmit and receive data, and initiate and accept phone calls. GPRS defines three classes with respect to simultaneous usage of voice and data. Class A mobile hosts can transmit and receive voice and data at the same time. Class B hosts can transmit and receive either voice or data but not both simultaneously. Finally, class C hosts have the user manually select if the host should be attached to the GSM (voice) or GPRS (data) network. When compared to Mobitex, DataTAC, and CDPD, GRPS class A devices can have simultaneous access to a packet-switched and circuit-switched network. Of course, GSM-only devices do not have this capability either, as mentioned earlier.

GSM uses a combination of *Frequency Division Multiple Access* (FDMA) and *Time Division Multiple Access* (TDMA) for channel allocation, as explained in detail in[5]. In short, each frequency channel carries eight TDMA channels. Each of these channels is essentially a time slot in a TDMA frame. Thus, any GSM frequency channel can carry up to eight circuit-switched connections with each slot reserved for a single connection (read *voice call*). In GPRS, each slot is treated as a shared resource and any mobile host can use it to transmit or receive data. In addition, a mobile host can be allocated more than one of the eight available slots in the same TDMA frame. In other words, GPRS can multiplex different traffic sources in one channel and allocate several channels to the same traffic source.

GPRS defines four different channel coding schemes[20], namely CS1, CS2, CS3, and CS4, with radio data rates 8.8 kbps, 13.3 kbps, 15.6 kbps, and 21.4 kbps, respectively. CS1 is the most “conservative” (includes more error correction bits) and is used for signaling packets and when poor channel conditions prevail. CS4 is the most “optimistic” (includes minimal error correction bits), and, assuming excellent channel conditions, allows operators to advertise a maximum radio data rate of 171.2 kbps per 200-kHz frequency channel (or TDMA frame).

In practice, CS4 is rarely used because it can lead to frequent retransmissions of lost packets and overall network underperformance. CS3 is commonly used, providing 124.8 kbps per frequency channel. Because a mobile host can be allocated multiple slots, user throughputs can range between 40 and 60 kbps. Mobile hosts typically use an MTU of 1500 bytes.

Communication between the base station and any given mobile host is full-duplex but can be *asymmetric;* that is, the downlink and uplink capacities need not be the same. The *GSM Association* has defined 12 multislot classes for GPRS. Each class is associated with a maximum number of uplink and downlink slots that can be allocated to a single mobile host. The slot allocation is usually written as $M + N$, where M is the maximum number of downlink slots and N is the maximum number of uplink slots. For example, class 1 is "1 + 1" (one downlink slot plus one uplink slot); class 2 is "2 + 1"; . . . ; and class 12 is "4 + 4" (four downlink and four uplink slots). In addition, each multislot class has an active slot constraint: A mobile host cannot use more than K active slots simultaneously. Given the number of slots and the channel coding scheme, one can calculate the peak rate. For example, for a class 12 device the sum of the physical downlink and uplink rates cannot exceed 124.8 kbps, if CS3 is used. However, the active slot constraint limits this rate even further. In the case of a class 12 mobile node, $K = 5$, that is, only "4 + 1", "3 + 2", "2 + 3", or "1 + 4" slots can be used simultaneously. See **www.gsmworld.com**

EDGE and Beyond

Enhanced Data for GSM Evolution (EDGE), also known as Enhanced GPRS, builds on the changes introduced by GPRS to GSM. EDGE essentially increases the radio data rates by using a more efficient modulation scheme[21], namely *8-Phase Shift Keying* (8-PSK) instead of the *Gaussian Minimum Shift Keying* (GMSK) used by both GSM and GPRS. EDGE defines nine modulation coding schemes named MCS1 to MCS9. MCS1 to MCS4 use GMSK with radio data rates similar to the four GPRS coding schemes. The real throughput improvements come from MCS6 (29.6 kbps per slot) through MCS9 (59.2 kbps per slot). The data rate usually associated with EDGE is a (shared) 384 kbps. This corresponds to using MCS7 for all 8 TDMA slots. Higher data rates are theoretically possible (up to 473 kbps using MCS9) but are not commonly deployed.

EDGE improves not only on the high end of data rates but also on the low end[22]. First, the greater diversity of coding schemes permits an EDGE network to choose the most appropriate one depending on channel conditions. Changing coding schemes is dynamic. Second, EDGE supports *packet resegmentation:* Packets that failed to be transmitted successfully can be resegmented and retransmitted using a more "conservative" coding scheme.

Table 1 summarizes the main high-level features for the WWANs surveyed.

Table 1: WWAN Characteristics

	Transmit/ Receive	Radio Data Rate	User Throughput	MTU
Mobitex	*Half duplex*	*8.0 kbps*	*<4.6 kbps*	*512 B*
DataTAC	*Half duplex*	*19.2 kbps*	*<10 kbps*	*2048 B**
CDPD	*Full duplex*	*19.2 kbps*	*<10 kbps*	*2048 B*
GPRS	*Full duplex*	*<171 kbps*	*40–60 kbps*	*1500 B*
EDGE	*Full duplex*	*<473 kbps*	*50–60 kbps*	*1500 B*

* Typically 512 B

Discussion and Trends

Among the WWANs presented, Mobitex and GPRS can be singled out as the most widely deployed; they also have enjoyed significant gains in the number of users and traffic volume in recent years. The popularity of enterprise wireless e-mail (due in part to the success of the Research in Motion BlackBerry devices) allowed Mobitex and DataTAC operators to revive their business models briefly. Worldwide, however, GSM dwarfs all other technologies: There are more than 1 billion GSM subscribers compared to the 1 million Mobitex users. DataTAC enjoys an even smaller user base. Even if a small percentage of GSM subscribers use GPRS and EDGE, the potential market for wireless applications is tremendous. On the other hand, subscribers who do not take advantage of GPRS or EDGE do use the inexpensive, (two-way) *Short Message Service* (SMS), which is built in GSM. Two-way messaging was available for many years but was certainly popularized by less-affluent and younger GSM users in the late 1990s. SMS is now commonplace, and in many countries it is more popular than e-mail. Dedicated data-only networks such as Mobitex have to look elsewhere for their niche.

For some, Mobitex, let alone DataTAC and CDPD, is virtually moribund. In the United States, for example, Cingular sold its Mobitex network and is investing heavily on GRPS and EDGE. DataTAC and CDPD are phased out by service providers in the United States in favor of newer technologies. Low-speed packet radio is considered lackluster and is not popular with younger crowds. After all, narrowband WWANs had their chance and failed to attract large numbers of subscribers. Recent pricing trends, too, reveal a heavy operator push in favor of GRPS and EDGE. In Finland, for example, 100 MB over GPRS costs less than 18 euros (approximately $24). Compare that to the $30–50 that 1 MB of traffic costs over Mobitex. Service and product popularity create economies of scale that cannot be ignored.

Nonetheless, open standards, an explicit focus on business applications with *Quality-of-Service* (QoS) guarantees in service response times, and narrowband M2M communication may well keep Mobitex going for years to come. Besides, bundling Mobitex with a wireless network that features fast and inexpensive connectivity, for example, WLAN or Bluetooth, might be promising: Large downloads and software updates can be done over the high-speed wireless network and critical messages can always reach the user through the WWAN.

Bundling several functions in a single handheld device is, after all, a major trend in the industry. Vendors scramble to integrate *Personal Information Managers* (PIMs), voice and data communications, as well as entertainment features (digital camera, games, or digital music players) in a single product. This is quite different from earlier mobile devices, which tended to be either single-purpose or tied to a particular set of applications. Even the BlackBerry devices still work, to some extent, in a closed architecture. Enterprise e-mail systems need to be supported by and integrated with BlackBerry servers in order to be accessible over the WWAN. Yet, one of the main objectives in 2.5G and 3G is to allow mobile users to use standard Internet protocols on a mobile radio network at significantly higher bit rates than other systems. In particular, GPRS was designed with certain office applications in mind and can support consumer and enterprise mobile communications alike, without being tied to any given platform or application servers. I expect that functionality bundling and 2.5G and 3G WWANs will allow for more open systems and will expedite the transformation of WWAN operators from integrated application providers to wireless ISPs.

For Further Reading

[1] Charles Perkins, *Mobile IP Design Principles and Practices,* ISBN 0201634694, Addison-Wesley, 1998.

[2] Joachim Tisal, *GSM Cellular Radio Telephony,* ISBN 0471968269, John Wiley & Sons, 1998.

[3] Tero Ojanpera and Ramjee Prasad, *WCDMA: Towards IP Mobility and Mobile Internet,* ISBN B000066OB4, Artech House, 2001.

[4] R. Ludwig, B. Rathonyi, A. Konrad, K. Oden, and A. Joseph, "Multi-layer tracing of TCP over a Reliable Wireless Link," presented at ACM SIGMETRICS 1999.

[5] C. Bettstetter, H.-J. Vogel, and J. Eberspacher, "GSM phase 2+ General Packet Radio Service GPRS: Architecture, Protocols, and Air Interface," *IEEE Communications Surveys & Tutorials,* Vol. 2(3), pp. 2–14, 1999.

[6] Larry L. Peterson and Bruce S. Davie, *Computer Networks: A Systems Approach,* 3rd ed., ISBN 155860832X, Morgan-Kaufmann, 2003.

[7] Rudi Bekkers and Jan Smits, *Mobile Telecommunications: Standards, Regulation, and Applications,* ISBN 0890068062, Artech House, 1999.

[8] Ira Brodsky, *Wireless: The Revolution in Personal Telecommunications*, ISBN 089006717, Artech House, 1995.

[9] Nathan J. Muller, *Wireless Data Networking*, ISBN 0890067538, Artech House, 1995.

[10] A. K. Salkintzis and C. Chamzas, "Mobile Packet Data Technology: An insight into Mobitex Architecture," *IEEE Personal Communications*, Vol. 4(1), pp. 10–18, 1997.

[11] M. S. Taylor, M. Banan, W. Waung, and M. Taylor, *Internetwork Mobility: The CDPD Approach*, ISBN 0132096935, Prentice-Hall, 1996.

[12] Andrew S. Tanenbaum, *Computer Networks, 4th ed.*, ISBN 0130661023, Pearson Education, 2003.

[13] K. Parsa, "The Mobitex packet-switched Radio Data System," presented at IEEE PIMRC '92, 1992.

[14] Sami Tabbane, *Handbook of Mobile Radio Networks*, ISBN 1580530095, Artech House, 2000.

[15] J. Rodriguez, W. Schollenberger, M. Anzib, and B. Widyarso, *Mobile Computing: The eNetwork Wireless Solution*, ISBN 0738412856, IBM Redbooks, 1999.

[16] Research in Motion, "Developer's guide for BlackBerry and RIM Wireless Handhelds—Radio API (DataTAC) Version 2.1," 2001.

[17] John Agosta and Travis Russel, *CDPD: Cellular Digital Packet Data Standards and Technology*, ISBN 0070006008, McGraw-Hill, 1996.

[18] P. Sinha, N. Venkitaraman, T. Nandagopal, R. Sivakumar, and V. Bharghavan, "A Wireless Transmission Control Protocol for CDPD," presented at IEEE WCNC '99, 1999.

[19] A. K. Salkintzis, "A survey of mobile data networks," *IEEE Communications Surveys & Tutorials*, Vol. 2(3), pp. 2–18, 1999.

[20] L. F. Chang, "Wireless Internet—Networking Aspect," in *Wireless Communication Technologies*, New Multimedia Systems, N. Morinaga, R. Kuhno, and S. Sampei, Eds., Kluwer Academic Publishers, 2000, pp. 215–244.

[21] Behrouz A. Forouzan, *Data Communications and Networking, 2nd ed.* Update, ISBN 0072822945, McGraw-Hill, 2002.

[22] Alexander J. Huber and Josef F. Huber, *UMTS and Mobile Computing*, ISBN B000089CJ3, Artech House, 2002.

KOSTAS PENTIKOUSIS, PhD, studied computer science at Aristotle University of Thessaloniki and Stony Brook University. He is an ERCIM Fellow at VTT, The Technical Research Center of Finland, and currently resides in Oulu, Finland. For more about his research and publications visit: **www.cs.stonybrook.edu/~kostas**. The best way to reach him is via skype. E-mail: **kostas@cs.sunysb.edu**

Opinion: ICANN, the ITU, WSIS, and Internet Governance

by Geoff Huston, APNIC

This is an opinion piece, intended primarily to provoke thought and comment. The author does not claim to personally hold any of the opinions expressed in this article.

It may have taken some three decades to get here, but there is no doubt that the Internet is now a major public communications utility. That is hardly the most important piece of news you are likely to read today, but the implication of this public role is that there are legitimate issues of public policy to consider when looking at the broad topic of coordination of various aspects of Internet infrastructure. In other words, "Internet Governance" is a matter of significant concern to many.

This opinion piece looks at the various range of views about the *Internet Corporation for Assigned Names and Numbers* (ICANN)[1] and its rationale and role over its brief history. Of course, no look at Internet Governance would be complete without also looking at the role of the *International Telecommunications Union* (ITU), as well as the broader background to this topic. It is a large topic and it has already been the catalyst for numerous articles.

Data Networking and Public Networks

Whether it was because of its antecedents in the research community, or simply because it was not originally envisaged that the Internet would become a global communications platform in its own right, or for whatever reasons, the administration of the Internet infrastructure was not originally crafted with conventional public network coordination in mind. The retrofitting of a model that incorporates considerations of a public utility role is proving to be a rather complicated process.

For example, the original hierarchical name space for the Internet used a set of generic top-level root zone names of "**`edu`**," "**`net`**," "**`com`**," "**`gov`**," and "**`mil`**." Adding country codes to the root of the name space was a later modification. Even then the original country code delegations were undertaken to individuals or entities who appeared to have some form of link to the national Internet community, rather than specifically seeking out an appropriate office of the national administration of communications services as the point of delegation. Similarly, IP addresses were structured without any form of national prefix, nor were IP addresses distributed along any national lines. In these respects the Internet was really no different from any other computing networking protocols of the 1980s, such as *DECnet*, the *Xerox Network System* (XNS), *AppleTalk*, or IBM's *Systems Network Architecture* (SNA), where names and addresses were defined in a limited context of the scope of the network, rather than within some broader public name framework.

There were two notable exceptions to this characterization of computer network protocols, and both were designed with a public communications utility as their primary objective, namely *X.25* and the *Open Systems Interconnection* (OSI) model. They can be regarded as offerings from the data services sector of the established telephone industry. X.25, the earlier of these two protocols, had a very obvious relationship to telephony, complete with the notion of a "call" as the means of establishing a data connection and as the unit of a transaction. The addressing scheme used a structured space that drew heavily on the telephone number structure. Like telephony, there was no associated name scheme and endpoints were identified by their numeric X.25 protocol address. OSI represented a later effort to design a packet-switched network architecture that was intended to reflect an increasing level of experience with this technology, but nevertheless continued to draw heavily on telephony design. Much was written about OSI at the time, and it would be a diversion to explore it in depth here. However, the salient observation here is that despite the extensive effort invested into its promotion, OSI was a market failure, and whatever its technical merits it was simply not accepted by the communications industry.

OSI was heavily supported by the ITU, and by virtue of this very active sponsorship of this technology, the implication of the aftermath of OSI was that the ITU was seen as being simply out of touch with data networking. It was often portrayed that the ITU was coming from a mindset that was incapable of engaging with either the data communications industry or the broader consumer market for data services. From the perspective of data networking, the failure of OSI was seen as a failure of the ITU itself.

The ITU and the Internet

The ITU is certainly one of the more venerable institutions in the communications sector. It can trace its origins to May 1865, when the first *International Telegraph Convention* was signed by 20 founding national members, and the *International Telegraph Union* was established to facilitate subsequent amendments to this initial agreement. Two decades later, in 1885, the ITU drafted international legislation governing telephony. With the invention in 1896 of wireless telegraphy, similar coordinating measures were adopted by the *International Radiotelegraph Convention*. In 1932 the Union combined the International Telegraph Convention of 1865 and the International Radiotelegraph Convention of 1906 to form the *International Telecommunication Convention*. The name of the body was changed to *International Telecommunication Union* to properly reflect the full scope of the Union's responsibilities, which by this time covered all forms of wireline and wireless communication.

In 1947 the ITU, under an agreement with the newly created United Nations, became an agency of the United Nations, with responsibilities in international telephony, telegraphy, and radio communications. Over the next four decades the ITU oversaw a system of international interconnection of telephony and data systems that became an industry in and of itself.

The ITU assumed a role of facilitating what was asserted to be a balanced international environment where the costs of running the international system were fairly apportioned between national service providers. In practice these lofty goals were not achieved very efficiently, and international facilities were priced at levels that were considerably higher than the associated costs of actual service provision. When attempts were made to redress the imbalances between large and small national carriers, the outcomes included collective action on the part of the national carriers that operated in ways not dissimilar to a cartel.

In 1992 the ITU was restructured into three sectors, corresponding to its three main areas of activity, namely the standardization of telecommunications technologies in the ITU-T, the coordination of radiocommunications in the ITU-R, and telecommunication development in the ITU-D. In 1994 the ITU established the *World Telecommunication Policy Forum* (WTPF), a group that encouraged the exchange of ideas and information about emerging policy issues arising from the changing telecommunication environment. The first WTPF was held in 1996 on the theme of global mobile personal communications by satellite, and the second in 1998, on trade in telecommunication services.

The ITU was heavily criticized over the ponderous amount of time taken to generate telecommunications standards, the nature of the process used in developing these standards in a closed set of forums, the marginal relevance of these standards, and the final indignity, that the ITU charged for paper and electronic copies of these standards. As some critics pointed out, perhaps harshly, this was not just a case of paperware about vapourware, it was a case of very expensive paperware about vapourware!

More recently, the ITU has focused on attempting to strengthen the participation of the private sector in the work of the Union, as well as streamlining the ITU's processes to reduce the level of delay and amount of process overhead in standardization of technology and operational practices. The ITU has sponsored the establishment of the *World Summit on the Information Society* (WSIS)[2], and has been attempting to position itself more centrally in the process of further evolution of the Internet as part of its overall charter.

The Internet has posed a severe challenge to the ITU. Not only was the ITU often perceived as being out of touch with the data communications sector, more critically it had been perceived as being incapable of making the necessary reforms to its mode of operation and policy setting to bring it back into relevance for the rapidly changing communications industry. The inference was being drawn that the ITU was apparently in a state of denial over progressive deregulation of national communications sectors. In many cases the national position had already moved to a position of lightweight regulation, relying on strong competitive pressures in the private sector to enforce regimes of efficiency and effectiveness in the supply of communications services to consumers. The ITU, as an intergovernmental organization, was being seen in some quarters as an anachronistic recalcitrant relic of an earlier era of communications service provision.

It was also evident that this critical view of the ITU was most strongly held within the United States, and in particular those parts of the U.S. administration and industry that were involved with the growth of the Internet. It was perhaps no coincidence that in these growth industries of personal computer technologies and the related Internet industry it was U.S. enterprises that were the "poster children" of this new model of industry-led deregulated communications services. Their consequent rapid expansion into a massive global undertaking of the global Internet was perhaps the most eloquent form of statement about the effectiveness of deregulation, and the degree to which the previous regulatory model had simply not managed to encompass the burgeoning demand for data services in a timely fashion.

From this perspective it should be no surprise to observe that when the transition of the *Internet Assigned Numbers Authority* (IANA) function from a fully federally funded research activity to some form of new foundational base was being considered by the U.S. administration, it appears that the ITU was never seriously contemplated as a viable home for this function. If the Internet was a child of deregulation and industry initiative taking on the outcomes of research activity, then the appropriate progression of the IANA function was also from a research context into an enterprise context. IANA should be responsive to industry needs, and to best achieve this the IANA function itself should be undertaken as a task housed within the deregulated private enterprise sector, rather than establishing yet another public bureaucracy, or using existing bureaucracies for the role. ICANN was the embodiment of this aspiration on the part of the U.S. administration, and to pass the effective levers of control of the Internet to the ITU was seen as denying the Internet any form of a productive, innovative, and successful future.

The Formation of ICANN

Whatever the original motivation in creating ICANN to administer the IANA responsibilities, it is now apparent that ICANN was deliberately structured to provide the industry with an alternative structure of coordination and regulation within national and international communications sectors to that of the ITU. The critical difference is that ICANN had not placed governments at the forefront of visible activity, but instead placed industry needs and the operation of a competitive deregulated international communications sector as being the major thrust of coordination activities.

As with any novel model of public policy determination, ICANN's acceptance ranged from cautious approval to advanced skepticism. Even within the U.S. administration ICANN has yet to be "unleashed," and it currently operates under the terms of a Cooperative Agreement with the *National Telecommunications and Information Administration* of the U.S. Department of Commerce under a sole source cooperative agreement. In this light ICANN appears to be a cautious step in a bold direction.

ICANN undertakes activities of management of Internet Protocol infrastructure in the areas of the content of the root of the *Domain Name System* (DNS) and the identification of parties to whom are delegated administrative and operational control of the top-level domains and the associated specification of terms and conditions of this delegation. ICANN, through IANA, also manages the pool of unallocated IP addresses (IPv4 and IPv6 addresses and Autonomous System numbers), and also manages the protocol parameter registries as defined by IETF Standards Actions.

ICANN MkI

The initial structure of ICANN had three "supporting organizations," focusing on:

- Coordination of the DNS with the *Names Supporting Organization* (NSO)
- Coordination of address policies with the *Address Supporting Organization* (ASO)
- Operation of Internet Protocol parameter registries with the assistance of the *Protocol Supporting Organization* (PSO)

The intended role of these supporting organizations was to provide a venue where interested parties could develop and consider policy proposals, leaving the task of ultimate identification of broad support for particular policy initiatives to the ICANN Board.

As has been evident to any observer of the ICANN process, things did not proceed within the parameters of that plan. The NSO met problems due to the diversity of interests that were encompassed with the DNS domain, including emerging national and regional interests in the country code top-level domains, the operators of the generic top-level domains, the trademark and intellectual property collection of interests, the emerging industry of registrars, and a continual interest of individuals who maintained that they had legitimacy of inclusion by virtue of their representation of interests of end users and consumers, or, to use an emerging ICANN lexicon, the "at large" constituency.

The ASO was formed within the parameters of a different model. The *Regional Internet Registries* (RIRs) had already developed a considerable history of working within their communities, and being widely accepted by these communities as an appropriate means of coordination of activity in the role of number resource administration and distribution. The ASO was formed with membership of the associated council based on processes determined by each RIR. Even then it was unclear as to the relationship between the RIRs' already well-established open policy development process and the ASO and ICANN. The RIRs were unwilling to pass all regionally developed policies to ICANN for a second round of consideration and potential alteration. They insisted that only those policies that were considered to be "global," in that they were common to all the RIRs, would be passed into this ICANN sphere.

The PSO was placed under strong pressure to include the ITU-T and the *European Telecommunications Standards Institute* (ETSI), and the *World Wide Web Consortium* (W3C) was also enlisted, in addition to the IETF. If the objective of the PSO was oversight and policy formulation concerning the role of protocol parameter registration of IETF protocols, then this enlarged membership of the PSO was unwarranted. Even within the terms of consideration of the PSO as a source of standards-based technical advice to the ICANN Board, the presence of these additional organizations was somewhat puzzling in terms of the match of resultant structure of the PSO to its intended role. The PSO, however, had a role in seating individuals onto the board of ICANN, and it was likely that this aspect of the PSO had been part of the reason for the interest in broader institutional membership. Uncertainty about the extent of the role of ICANN saw many groups attempting to gain access to board seats.

Missing from this mosaic of diverse interests was the inclusion of various national public communications sector entities who also felt that they had clear legitimacy to undertake an active role within the ICANN policy development process, and, in response, the *Government Advisory Committee* (GAC) was formed.

ICANN Evolution and Reform

If a camel is a horse designed by a committee, then it is unclear whether ICANN was a three-humped camel or a three- and three-quarter-humped camel as a result of all this, but camel it undoubtedly was.

The PSO was dysfunctional and missing any tangible agenda of activity. A fracture was apparent in the relationship between ICANN and the IETF. Attempts to create an agreement between ICANN and the IETF over the IANA function were not recognized by the U.S. administration, who continued to insist that, formally, the IANA function for the IETF was undertaken at the behest of the U.S. Department of Commerce rather than the IETF. This view was not shared by the IETF.

The ASO was criticized by ICANN itself of being insufficiently "representative" of the addressing community, and the ICANN Board established its own temporary advisory committee on addresses, and in so doing alienated the RIR community from the entire ICANN framework.

The NSO was hopelessly wedged into factional-based politics.

The GAC decided at the outset that it would operate behind closed doors, in contrast to ICANN's continuing efforts to operate in an open and transparent manner.

The "At Large" election process undertaken by ICANN appeared to be of dubious validity because of problems in establishing a reliable constituency of individuals who had an interest in ICANN, and a direct election process was attempted only once.

Not surprisingly, ICANN fell into some disarray under these pressures, and by early 2002 the CEO of ICANN at the time, Stuart Lynn[3], was warning all who cared to listen that ICANN was paralyzed, dysfunctional, and in danger of an imminent demise. Whether this was a message directed to the ICANN Board or to a fractious set of communities that had some intersection with ICANN, or to the U.S. administration who had been influential in determining the original ICANN structure was not entirely clear to any observer of the process.

However, given that ICANN had been set up as an example of a new form of international coordination of communication infrastructure support activities that was based on private-sector activity rather than governmental fiat, this message of imminent failure was widely interpreted both as a potential failure of ICANN and a sign of failure of this new model of coordination of international activity. ICANN was seen as a point of vulnerability with respect to the U.S. administration's diplomatic efforts to reform this international activity sector. The ITU-T's activities in this same area was reinvigorated, with considerable support from national sectors who saw their national interests being potentially advantaged in a ITU-led international environment.

ICANN MkII

Although still firmly positioned as a private-sector activity, and although still making no concessions in the direction of the ITU, ICANN has managed to reorganize its structure through a protracted evolution and reform process.

With respect to the ASO, The Regional Internet Registries formed its own coordination entity, the *Number Resource Organization* (NRO)[4], and has proposed this entity to ICANN as the means of interfacing between the addressing community and ICANN's policy-development activities.

The PSO was abolished, to be replaced by a *Technical Liaison Group* that, apart from its function of seating an individual on the ICANN Board, is a group without an obvious role or agenda.

The NSO was forced to recognize the fundamental difference between the generic top-level domains, which fall under a more direct relationship with ICANN and its processes, and the country code domains (ccTLDs), which have from the outset been quite wary of ICANN. From the ICANN reform process emerged the *Country Code Name Supporting Organization* (CCNSO) and the *Generic Names Supporting Organization* (GNSO), as a recognition that these two groupings are so dissimilar that they have almost nothing in common.

In addition, an *At Large Advisory Committee* was formed.

The reform process has had some more tangible outcomes, in that formal open meetings of the ICANN Board of Directors have managed to be progressively refined from efforts at direct dialogue and open debate into highly structured events with many formalisms and appropriate quantities of ceremony.

ICANN Today

Despite the effort to encompass coordination activities in the areas of names, addresses, and protocol parameters, ICANN has been largely captured by the names industry, and ICANN's agenda, activity focus, and outcomes are concentrated mostly in the name domain.

In this activity domain, the track record of ICANN is very mixed. To its credit, it has managed to dismantle the most objectionable parts of the monopoly hold over the *generic Top-Level Domains* (gTLDs), create an operational model that makes a clear distinction between registry operators and registrars, impose price and business controls on the registry operation as a means of controlling the natural tendency for the registry operation to reflect its unique position in the form of monopoly rentals, and assist in the creation of a global network of competitive enterprises, with the expectation that competition will instill operational and price efficiency in the registrar business.

In addition, ICANN has been successful in not only introducing new gTLDs to complete with the established brands of **`.com`**, **`.net`**, and **`.org`**, but also in moving **`.org`** and **`.net`** to new registry operations (**`.net`** is under way at the time of writing of this article). Despite these positive achievements, it is not clear that this new regime has been entirely successful.

True competition in the name space is still some way off, and the recently introduced gTLD brands have failed to gain any leverage within the market. The name market itself remains one where the role of name speculators continues to play a significant role in terms of proportion of registered names. The overarching dominance of **`.com`** as a brand has continued, and the advantaged position of the U.S.-based registrar of this zone continues.

The obscure nature of the relationships between the IETF, ICANN, and the U.S. administration over the protocol parameter registries remains unresolved. The IETF is clearly not in control of its own protocol parameters, and has abrogated this role to ICANN. Standards making entirely divorced from any effective engagement with deployment tends to result in a standards body of dubious long-term validity, and despite its impressive track record in the past, the IETF is clearly already well-distanced from current technology directions in the industry—and the gap continues to widen.

The DNS *Root Server Operators* continue to operate as an independent group. The recent moves to dramatically increase the number of DNS root servers and improve the overall robustness of DNS resolution through anycasting root servers and distributing anycast instances across the globe has been a well-received initiative. The fact this has occurred without any form of ICANN involvement is an interesting commentary on the ability of ICANN to engage with the operational parts of the infrastructure of the Internet. Comparable activities to improve the DNS in terms of resolution services within the ICANN sphere have become protracted exercises that impose a very heavy burden on the patience of the players.

The moves to introduce IPv6 AAAA records into the DNS root have been anticipated for many years, and the response to the recent ICANN announcement is, in general, of the tenor "why didn't this happen some years ago?" The continuing frustration to get the DNS root to include *Secure DNS* (DNSSEC)[5] important information continues to illustrate a perspective that the ICANN process appears to be unresponsive to technical needs and end-user imperatives.

The situation today is that ICANN appears to enjoy a mixed level of success. It has managed to establish itself as a means of administering the infrastructure elements of the Internet Protocol in a manner that is reflective of the deregulated nature of the Internet industry. It has managed to reform parts of the landscape and generate an industry structure that uses open competition as the major control mechanism. ICANN has managed to bring much of the discussion about the administration of Internet infrastructure out into the open. All these are major milestones, and it is to the credit of many dedicated individuals that ICANN has managed these impressive outcomes. However, it has been able to achieve all this with the continued sponsorship of the U.S. administration, and the question of whether it can firmly establish itself in its own right in the coming years remains today perhaps a matter of hope rather than absolute certainty.

There are still the lingering concerns that if ICANN, as a private-sector entity, were to once more explore positioning itself on the brink of imminent demise, the collective task of picking up the pieces and continuing to support the operation of the Internet is one that appears to have a very uncomfortable level of uncertainty. In addition, the perception of ICANN as an entity whose single purpose is to maintain an entrenched advantaged position of the United States and of U.S.-based enterprises in the global Internet has been widely promulgated. It is often portrayed that ICANN offers no viable mechanisms for other national or regional interests at a governmental level to alter this somewhat disturbing picture of international imbalance. Although other aspects of international activity fall under various political or trading frameworks, and national and regional interests and positions can be collectively considered and negotiated, critics of ICANN point out that the message ICANN sends to the rest of the world is that the United States is withholding the Internet from conventional international governance processes. Skeptical commentators interpret the U.S. administration's use of ICANN as at best a delaying technique to gain time to further strengthen the position of U.S.-based enterprises across a lucrative global Internet market, aided and abetted by a compliant industry body that masquerades as an international standards organization.

Such a critical perspective also points to ICANN's tenuous lines of authority, its lack of performance in many aspects of the domain name enterprise, its seeming obsession with the registrar sector to the apparent exclusion of any other activity, its burgeoning costs, and its lack of acceptance, particularly as it relates to the acceptance of ICANN by the various country code DNS administrators, to name but a few factors.

Accompanying this strident criticism is the line of argument that the Internet does not actually represent a viable challenge to existing mechanisms for coordination of international activity. At both a national and international level, the Internet should not require novel and untested regulatory mechanisms as a means of expressing public interest and public policies. The line of argument from this perspective is that there is neither the demonstrated need, nor any appropriate level of international support at a governmental level to sustain the argument that a private-sector, nonprofit corporation is the best, or even the only viable model of coordination of Internet activity. If "Internet Governance" is the question, then, the line of argument goes, the model upon which ICANN is based is definitely not the best answer we can devise. This very critical line of reasoning has become particularly prominent in the WSIS process, and lies behind much of the continual fascination of the topic of "Internet Governance" in WSIS meetings.

WSIS and Internet Governance

The WSIS has been a long time coming, and it represents a move on the part of the ITU to formulate a revised role for the ITU to engage with a world richly populated by all manner of information services layered upon a highly diverse and capable communications environment. This summit was planned in two phases. The first summit was held in Geneva December 10–12, 2003, where the foundations were laid by reaching agreement on a *Declaration of Principles* and a *Plan of Action*. The second phase will be held in Tunis, November 16–18, 2005, to implement the agenda leading up to achievable targets by 2015, and to agree on unfinished business, most importantly on the question of Internet governance and of financing mechanisms.

Irrespective of any particular political perspective here, the universal observation is that the Internet has heralded a revolutionary change to the global communications enterprise. Markets for communications services are changing, the technology base is changing, the economic models of communication are changing, and the models of interaction at the provider level are changing. The challenge from the public-policy perspective at a world level is to create a framework that ensures that the benefits of this change, in both social and economic terms, are accessible to all, rather than to a subset of the world's population. It is within this broad framework that WSIS has been positioned.

These are lofty and ambitious goals, and the task before WSIS is certainly as challenging as any in this environment. The hope is that the myriad of participants in this process includes sufficient resources to engage in the agenda in a meaningful way.

However, the underlying issue is that of the progressive change in the role of communications infrastructure from a predominately public-sector activity to a very diverse spectrum of public- and private-sector activity. We appear to have become increasingly reliant on private-sector investment and private enterprise to support the public communications enterprise. But is this necessarily the appropriate model for the entire world, or even any part of the world?

As many recently privatized industries could attest, private-sector activity has entirely different investment motivations and entirely different service objectives. If the nature of the activity is one that requires long-term investment in infrastructure with low returns, then private-sector activity tends to use the existing infrastructure base without necessarily making adequate longer-term replenishment investments. Private activity also tends to concentrate service delivery to the most lucrative sectors of the market, and, if possible, will deliberately avoid establishing services in areas that are less financially attractive. The task of structural cross-subsidization that makes ubiquitous equity of access possible is not seen as a private enterprise outcome, and aspects of communications such as universal service obligations and equity of access are seen as public regulatory functions rather than natural market outcomes of a deregulated industry.

The Internet today is anything but a level and balanced environment. There are concentrations of investment capability, concentrations of technical knowledge and logistical capability, concentrations of intellectual wealth, and concentrations of power and influence. How to create from this current diverse environment some form of structural cross-subsidization that extends the basic means of access to all is the appropriately lofty goal of the WSIS endeavor. There is also the more focused investigation of "Internet Governance" and the agenda of establishing to what extent the perception of the advantaged position of a small number of national entities in all this can be balanced by measures that allow other national economies to invest in this space on terms and conditions that do not involve a continuing flow of money and a ceding of power to these existing advantaged national interests.

As the WSIS documentation points out, "... building the foundations for an Information Society is a complex task. The digital revolution is already impacting the world in deeply intrinsic ways, perhaps more profoundly than even the industrial revolution itself. Yet, while the digital revolution has extended the frontiers of the global village, the vast majority of the world remains unhooked from this unfolding phenomenon."

The Secretary General of the UN chartered a smaller group to examine Internet Governance, in particular, the *Working Group on Internet Governance,* or WGIG. Its nine-month brief is to glean these issues of public policy in an environment that has very significant private-sector interest. Indeed from an international perspective, where regulatory powers, even of a reserve nature, are in a very real sense ephemeral, the work in WGIG to date with its discussion papers has done little. The discussion papers have illustrated the broad nature of the topics raised in the context of Internet Governance, but their poor depth, visibly poor levels of research, and lack of any real analysis of the selected topics only highlights the complexity of the underlying interplay of public- and private-sector interests within a domain that is also bounded by technical considerations.

At the same time the poor quality of these reports highlights the inability of WGIG to engage directly into the heart of this exercise, given their obvious constraints of time and resources. It is not surprising to observe that, following its February meeting WGIG has decided to abandon this set of discussion papers. If a fresh start is being contemplated for WGIG, then perhaps it is time to note that only half of the group's allocated time remains, and the topic is getting no easier with the passing of the days.

For those interests who wanted the ITU to become engaged in the Internet, hope has now been passed to the WSIS process and the related WGIG study into Internet Governance issues. This is seen as being a means of opening up the control of the Internet into a more conventional international process that dismantles what they see as the current position of global taxation that U.S. national interests have imposed on the rest of the world's population in the adoption of Internet-based services. For those who think the ITU remains an unreformed vehicle for the imposition of anachronistic, inappropriate regulatory measures that stultify any form of innovation and progress in telecommunications, the WSIS process is yet another venue to parade the stark contrast between the rather impressive track record of a deregulated market-driven approach to coordination of telecommunications services, as seen with the Internet, and the ineffectual outcomes from the international public regulatory sector.

Looking Forward

One view of this process is that this is a negotiation of national roles of influence and power over the coming century or more, and that this process requires some considerable care and attention at an international level.

This topic is one that places a model of deregulated private sector-led activity, with its market-based disciplines, into direct contrast with a more traditional model of the balancing of various national interests through common regulatory measures undertaken within each national regime as a regulated public-sector process. The proponents of a deregulated approach argue that the Internet is a child of the progressive position of deregulation of communications markets in many national environments, and it is the dynamic and creative impetus of highly competitive markets that has led to the rapid spread of the Internet and the consequent improvements in the efficiency and effectiveness of national and international communications systems. None of these outcomes would have been achievable, they argue, in a regulated regime where innovation and competition for the consumer were completely stifled by the deadening weight of regressive regulation.

Like many bold innovative experiments in international coordination and the establishment of new world orders, ICANN stands a strong risk of falling foul of an inherent conservatism in international politics, where the careful balancing of national interests is seen as being far more critical an objective than any actual outcomes that may be achieved from the process.

From this perspective, ICANN is critically reliant on its acceptance by all players of its legitimacy to operate in this space, and also critically reliant on acceptance of the proposition that these issues are best addressed in open forums of debate. This task is difficult, and the limited set of outcomes that ICANN can point to as being products of this process do not install a high degree of confidence that this process is stable, scalable, well-founded, and sustaining. Currently the proposition is not that ICANN represents the most appropriate enduring framework here, but that the track record of the alternative has failed in the past and nothing has changed to prevent the historical alternative framework making similar flawed decisions in the future.

The opposite end of the spectrum of views argues that nothing has really changed with the introduction of the Internet, and the international regime remains one where various national interests need to be resolved in a coordinated and equitable fashion. Without some form of common regulatory constraint, there are inevitable market distortions where the expression of vigorous national aspirations results in an advantaged position in the international domain. Public communications is a public-sector activity, they argue, and, ultimately, the only points of control rest within national regulatory regimes, and internationally it is a case where national interests must be balanced through a process that recognizes political realities of coordination and compromise. From this perspective it is asserted that the ITU is the intergovernmental venue for this activity as it relates to the communications sector, and it is to the ITU that national interests must look to redress distortions where one national entity or one region holds a contrived privileged position with respect to international communications.

In looking at these two extremes of perspective, an obvious question is what then is the role of international public policy setting? In this form of market-mediated service supply functions, are international issues being progressively transformed into aspects of international trade? Does such an environment provide adequate protection for developing economies? Are common social priorities being adequately considered in such a framework?

This leads to a more basic question of whether the existing international institutions, such as the ITU, are appropriately positioned to meet these public policy challenges, or should we be considering changes here in order to bring the international institutional framework into better alignment with the emerging information society?

These are certainly difficult positions to attempt to reconcile, and perhaps it is being impatient to expect clear outcomes in the near future, and certainly very difficult to expect that in a few short months WGIG and WSIS will be able to deliver a balanced, considered, and generally acceptable outcome in this space. It is also a natural concern in looking at these rather aggressive schedules for WSIS that short-term political expediency will obstruct genuine attempts to truly understand the fundamental nature of the changes that are happening with the differing model of communications that are heralded by the Internet model.

References

[1] Internet Corporation for Assigned Names and Numbers (ICANN): **`http://www.icann.org`**

[2] The World Summit on the Information Society (WSIS): **`http://www.itu.int/wsis/`**

[3] M. Stuart Lynn, "A Unique, Authoritative Root for the DNS," *The Internet Protocol Journal,* Volume 4, No. 3, September 2001.

[4] Number Resource Organization (NRO): **`http://www.nro.net/`**

[5] Miek Gieben, "DNSSEC: The Protocol, Deployment, and a Bit of Development," *The Internet Protocol Journal,* Volume 7, No. 2, June 2004.

GEOFF HUSTON holds a B.Sc. and a M.Sc. from the Australian National University. He has been closely involved with the development of the Internet for the past decade, particularly within Australia, where he was responsible for the initial build of the Internet within the Australian academic and research sector, and has served his time with Telstra, where he was the Chief Scientist in the company's Internet area. Geoff is currently the Internet Research Scientist at the Asia Pacific Network Information Centre (APNIC). He is also the Executive Director of the Internet Architecture Board, and is a member of the Board of the Public Interest Registry. He is author of *The ISP Survival Guide,* ISBN 0-471-31499-4, *Internet Performance Survival Guide: QoS Strategies for Multiservice Networks,* ISBN 0471-378089, and co-author of *Quality of Service: Delivering QoS on the Internet and in Corporate Networks,* ISBN 0-471-24358-2, a collaboration with Paul Ferguson. All three books are published by John Wiley & Sons. E-mail: **`gih@apnic.net`**

Book Review

Unix Network Programming

Unix Network Programming, 3rd Edition, by W. Richard Stevens, Bill Fenner, Andrew M. Rudoff, ISBN 0131411551, Addison-Wesley Professional, 2003.

It would be difficult to put value on a book that has been a classic text and a reference in academia and in the real world in the context of network programming for over a decade. Richard Stevens published the ever-popular *Unix Network Programming* [UNP] back in 1990, and the second edition followed in 1998. With a dedication to the memory of R. Stevens, the UNP book found itself two new authors, Bill Fenner and Andrew M. Rudoff, who would write the third edition of this book. The third edition has many updates, a new look and feel and many of new chapters that cover the topics more applicable these days. In my opinion, it is still the most valuable and profound text in the context of network programming.

Changes and Updates

For those of us who have the first two editions of this book, the third edition has the following changes:

- IPv6 updates. In the second version of the book, IPv6 was merely a draft, and the sections covering IPv6 have been updated to reflect these changes.
- POSIX updates. The functions/APIs and examples have been updated to reflect the changes to the latest version of the POSIX specification (1003.1-2001).
- SCTP coverage. Three new chapters that cover this new reliable, message-based transport protocol have been added.
- Key Management Sockets coverage. Network security and its applicability and use with IPsec are covered.
- The Operating Systems and machines that are used for the examples have been updated.
- Some topics such as Transaction TCP and X/Open Transport Interface have been dropped.

Many topics and sections have been updated with the authors' comments. These comments even though simple for someone new to the profession, are extremely useful because they are like hints and tips from one developer to the next to help you in your next programming assignment.

Unix Focus

If this is the only edition of the book that you will read, you are in for a treat. Topics in Network Programming are covered in detail, using concrete programming examples that all of us can relate to—all Unix, but what else is there?!

All kidding aside, the topics are covered well enough that they are useful information under any operating system. The concepts don't change; sockets are sockets under any operating system. The function call is different, but one needs to go through the same steps under any environment.

Being the most popular networking protocol, TCP/IP is covered in Part I of the book. You need to have prior understanding of the TCP/IP protocol and the OSI model, however. If this is the first time you are looking at the programming aspects of networking protocols, Part I of this book covers the basics. It begins with a couple of simple examples such as such as daytime client and a daytime server and it builds on that. TCP, UDP, and SCTP (*Stream Control Transmission Protocol*) are covered in brief in Part I, and basic concepts such as the three-way handshake of TCP and the four-way handshake of SCTP are depicted.

Part II of the book covers *sockets* and socket programming. Topics such as the socket Address Structure in IPv4 and IPv6 for TCP, UDP and SCTP are covered and examples (the same daytime client/server) are given to convey the point. It is important to mention here that all the topics and concepts are depicted for the three transport protocols: TCP, UDP and SCTP. Every socket API under the Unix programming environment is covered and examples are given for each function call to show the reader how the function can be utilized. Much attention is dedicated to Socket Options and how they are used or can be used for best results. Hints are given throughout the chapter about the pitfalls and best practices of each option.

After the basics are been covered, various I/O models are depicted in detail and examples are shown to convey the advantages and disadvantages of each I/O model. The five I/O models used through the book (and available under the Unix environment) follow:

- Blocking I/O
- Non-blocking I/O
- I/O Multiplexing (using select and poll)
- Signal driven I/O
- Asynchronous I/O

The *Stream Control Transmission Protocol* (SCTP), a new IETF standard is also covered in detail—from the basics to the advanced. The two interface models of SCTP (one-to-one and one-to-many) are covered in detail, and their differences with TCP are also explained in full. The client/server example used throughout the book is ported to use the new SCTP protocol. The authors then explain in detail the problems that SCTP solves over TCP and where and how it would be useful to use SCTP.

Advanced topics such as IPv4 and IPv6 portability, Unix Domain Protocols, Multicasting and advanced Socket programming for UDP, TCP and SCTP cover the rest of the chapters in this book.

Various options for interoperability between IPv4 and IPv6 are discussed in the last section of the book. Advanced I/O functions bring us a new perspective of how complicated Network Programming can become. Benefits and examples of nonblocking I/O are covered in detail—the authors give examples to show us how, with very few modifications, the performance of a socket application can improve dramatically. Various methods on how to control socket operations are discussed including the use of an alarm along with SIGALRM, the use of select and various timeout options that are available in the API.

The chapters that discuss Multicasting and adding reliability to UDP are my favorite chapters in this book. The Time Server used throughout the book is re-coded to become a multicast application. Some issues that arise when designing multicast applications such as multicast on a WAN are also discussed.

As Good as Ever

The third edition of *Unix Network Programming* is as good as ever. The updates truly reflect solutions to today's challenges in network programming. Bill Fenner and Andrew Rudoff did an amazing job continuing the work of a true legend in the field of Computer Science.

—*Art Sedighi*
asedighi@tibco.com

Read Any Good Books Lately?

Then why not share your thoughts with the readers of IPJ? We accept reviews of new titles, as well as some of the "networking classics." In some cases, we may be able to get a publisher to send you a book for review if you don't have access to it. Contact us at **ipj@cisco.com** for more information.

Fragments

Internet Pioneers Cerf and Kahn to Receive ACM Turing Award

The *Association for Computing Machinery* (ACM), has named Vinton G. Cerf and Robert E. Kahn the winners of the 2004 *A.M. Turing Award,* considered the "Nobel Prize of Computing," for pioneering work on the design and implementation of the Internet's basic communications protocols. The Turing Award, first awarded in 1966, carries a $100,000 prize, with financial support provided by Intel Corporation. Cerf and Kahn developed TCP/IP, a format and procedure for transmitting data that enables computers in diverse environments to communicate with each other. This computer networking protocol, widely used in information technology for a variety of applications, allows networks to be joined into a network of networks now known as the Internet.

ACM President David Patterson said the collaboration of Cerf and Kahn in defining the Internet architecture and its associated protocols represents a cornerstone of the information technology field. "Their work has enabled the many rapid and accessible applications on the Internet that we rely on today, including e-mail, the World Wide Web, Instant Messaging, Peer-to-Peer transfers, and a wide range of collaboration and conferencing tools. These developments have helped make IT a critical component across the industrial world," he said.

"The Turing Award is widely acknowledged as our industry's highest recognition of the scientists and engineers whose innovations have fueled the digital revolution," said Intel's David Tennenhouse, Vice President in the Corporate Technology Group and Director of Research. "This award also serves to encourage the next generation of technology pioneers to deliver the ideas and inventions that will continue to drive our industry forward. As part of its long-standing support for innovation and incubation, Intel is proud to sponsor this year's Turing Award. As a fellow DARPA alumnus, I am especially pleased to congratulate this year's winners, who are outstanding role models, mentors and research collaborators to myself and many others within the network research community."

In 1973, Cerf joined Kahn in a *Defense Advanced Research Projects Agency* (ARPA, now called DARPA) project to link three independent networks into an integrated "network of networks." They sought to develop an open-architecture network model for heterogeneous networks to communicate with each other independent of individual hardware and software configuration, with sufficient flexibility and end-to-end reliability to overcome transmission failures and disparity among the participating networks. Their collaboration led to the realization that a "gateway" (now known as a *router*) was needed between each network to accommodate different interfaces and route packets of data. This meant designating host computers on a global Internet, for which they introduced the notion of an *Internet Protocol* (IP) address.

As a graduate student at the University of California at Los Angeles, Cerf had contributed to a host-to-host protocol for ARPA's fledgling packet-switching network known as ARPANET. Kahn, prior to his arrival at ARPA, led the architectural development of the ARPANET packet switches while at Bolt Beranek and Newman (BBN), and had showcased the ARPANET in 1972, at the first International Conference on Computer Communications. ARPANET had already connected some 40 different computers and demonstrated the world's first networked e-mail application.

In May 1974, they published a paper describing a new method of communication called *Transmission Control Protocol* (TCP) to route messages or packets of data. Like an envelope containing a letter, TCP broke serial streams of information into pieces, enclosed these pieces in envelopes called "datagrams" marked with standardized "to and from" addresses, and passed them through the underlying network to deliver them to host computers. Only the host computers would "open" the envelope and read the contents.

This networking arrangement allowed for a three-way "handshake" that introduced distant and different computers to each other and confirmed their readiness to communicate in a virtual space. In 1978, Cerf and several colleagues split the original protocol into two parts, with TCP responsible for controlling and tracking the flow of data packets ("letters"), and IP responsible for addressing and forwarding individual packets ("envelopes"). The new protocol, TCP/IP, has since become the standard for all Internet communications.

Vinton Cerf and Robert Kahn share a number of awards, including the 1991 ACM Software System Award, the 2001 Charles Stark Draper Prize from the National Academy of Engineering, the 2002 Prince of Asturias Award, and the 1997 National Medal of Technology from President Bill Clinton. They are both the recipients of numerous honorary degrees. ACM will present the Turing Award at the annual ACM Awards Banquet on June 11, 2005, in San Francisco, CA.

The A.M. Turing Award was named for Alan M. Turing, the British mathematician who articulated the mathematical foundation and limits of computing, and who was a key contributor to the Allied cryptanalysis of the German Enigma cipher during World War II. Since its inception, the Turing Award has honored the computer scientists and engineers who created the systems and underlying theoretical foundations that have propelled the information technology industry.

For additional information see:
`http://www.acm.org/awards/taward.html`

New Administrative Structure for the IETF

The *Internet Engineering Task Force* (IETF) is well advanced in the process of making a significant change to the administrative structure that supports the world's leading Internet standards development group. The creation of an *IETF Administrative Support Activity* (IASA) is an important move designed to help the IETF maintain and expand the unique open processes that have enabled the development of Internet standards since 1986.

The new structure will allow the IETF to take full responsibility for managing the resources required to accomplish its work—giving the IETF a solid foundation on which future operations will be based.

This is the first time that all the IETF's administrative and support functions will be managed directly by the IETF as one fully integrated entity. Until now, administration of the IETF has been carried out exclusively by helper organizations and volunteers. The new IASA will be formally structured as an activity within the *Internet Society* (ISOC)—the organizational home of the IETF—and an *IASA Administrative Director* (IAD) will be appointed to provide central management of IETF administration.

The decision to move forward with the new structure was taken after extensive consultations with the Internet community. A number of key prerequisites for efficient administrative operations were identified, including the need for the IETF to have budgetary autonomy. The IETF is currently supported by funding from multiple sources, including meeting fees, donations from interested corporate and non-corporate entities, and donations in kind of equipment or manpower. The IASA will allow the IETF to be able to consider all sources of income, and all expenses involved in running the IETF, as pieces of one budget.

The IASA will also be responsible for defining clear contractual relationships with other organizations that will continue to provide basic services, including meeting organization, secretarial services, IT services, etc. The new structure also gives the IETF flexibility in how it chooses to fund and develop any additional services that may be required.

The IETF is a large open international community of network designers, operators, vendors, and researchers concerned with the evolution of the Internet architecture and the smooth operation of the Internet. It is open to any interested individual. See: **`http://www.ietf.org`**

ISOC is a non-governmental international organization for global cooperation and coordination for the Internet and its internetworking technologies and applications. Members comprise commercial companies, governmental agencies, foundations, and individuals. ISOC has 82 Chapters in over 60 countries around the world. For more information see: **`http://www.isoc.org`**

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, trouble-shooting, 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 will contain 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 is published quarterly by the Chief Technology Office, Cisco Systems, Inc.
www.cisco.com
Tel: +1 408 526-4000
E-mail: ipj@cisco.com

The Internet Protocol Journal, Cisco Systems
170 West Tasman Drive, M/S SJ-7/3
San Jose, CA 95134-1706
USA

ADDRESS SERVICE REQUESTED

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