The Internet Protocol Journal

December 2007 Volume 10, Number 4

A Quarterly Technical Publication for Internet and Intranet Professionals

In This Issue

From The Editor

Identity theft is a widely reported problem in today's world. Criminals can use numerous ways to obtain private information such as Social Security Numbers, credit card details, and other information that makes it possible for the perpetrator to successfully "pretend to be" someone else. A similar concept, albeit less personal, is so-called *IP Spoofing*, wherein fake IP datagrams can be generated and sent across the network in order to compromise remote systems in a variety of ways. Farha Ali gives an overview of IP Spoofing and explains ways in which the problem can be mitigated.

Our second article looks at numerous standards for information security management being developed by organizations such as the *International Organization for Standardization* (ISO), the *National Institute of Standards and Technology* (NIST), and others. The author of the article is William Stallings.

On November 2, 2007, Vint Cerf ended his term as chairman of the *Internet Corporation for Assigned Names and Numbers* (ICANN). At the same time he released a document entitled "Looking Toward the Future," which details ICANN's history, as well as outlining its challenges ahead. We've included the document in this issue and added some pointers for those readers who may not be familiar with the workings of ICANN.

In late October, the Internet technical community received the sad news that Dr. Junichiro Hagino, universally known as "Itojun" had passed away. Itojun played a very important role in the development of IPv6 and had many friends across the world. We asked one of them, Bob Hinden, to reflect on Itojun's life and compile some comments from those who knew him well.

We would like to remind you about our online adjunct to this journal. *The Internet Protocol Forum* (IPF) available at **http://www.ipjforum.org/** is a resource you can use to discuss articles and read additional material. Please take a moment to explore IPF.

—Ole J. Jacobsen, Editor and Publisher

ole@cisco.com

IP Spoofing

by Farha Ali, Lander University

The *Internet Protocol*, or IP, is the main protocol used to route information across the Internet. The role of IP is to provide best-effort services for the delivery of information to its destination. IP depends on upper-level TCP/IP suite layers to provide accountability and reliability. The heart of IP is the IP *datagram*, a packet sent over the Internet in a connectionless manner. An IP datagram carries enough information about the network to get forwarded to its destination; it consists of a *header* followed by bytes of *data*. The header contains information about the type of IP datagram, how long the datagram should stay on the network (or how many hops it should be forwarded to), special flags indicating any special purpose the datagram is supposed to serve, the destination and source addresses, and several other fields, as shown in Figure 1.

Figure 1: The IP Header

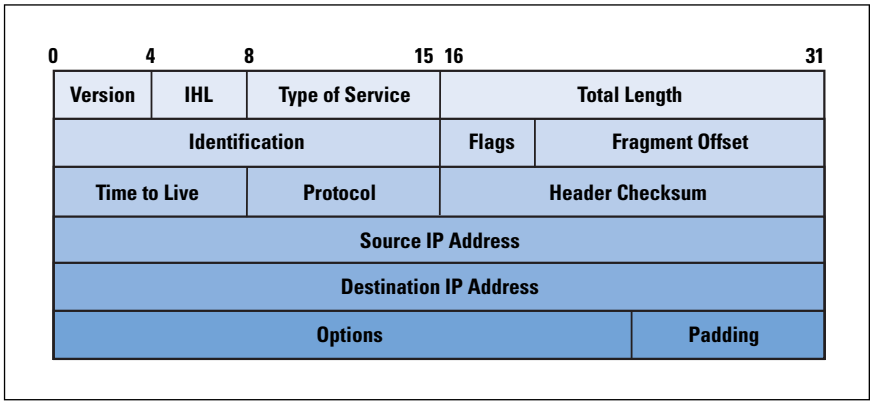

Layers above IP use the source address in an incoming packet to identify the sender. To communicate with the sender, the receiving station sends a reply by using the source address in the datagram. Because IP makes no effort to validate whether the source address in the packet generated by a node is actually the source address of the node, you can spoof the source address and the receiver will think the packet is coming from that spoofed address. Many programs for preparing spoofed IP datagrams are available for free on the Internet; for example, *hping* lets you prepare spoofed IP datagrams with just a one-line command, and you can send them to almost anybody in the world. You can spoof at various network layers; for example, you can use *Address Resolution Protocol* (ARP) spoofing to divert the traffic intended for one station to someone else. The *Simple Mail Transfer Protocol* (SMTP) is also a target for spoofing; because SMTP does not verify the sender's address, you can send any e-mail to anybody pretending to be someone else. This article focuses on the various types of attacks that involve IP spoofing on networks, and the techniques and approaches that experts in the field suggest to contend with this problem.

Spoofing IP datagrams is a well-known problem that has been addressed in various research papers. Most spoofing is done for illegitimate purposes—attackers usually want to hide their own identity and somehow damage the IP packet destination. This article discusses ways of spoofing IP datagrams, various attacks that involve spoofed IP packets, and techniques to detect spoofed packets and trace them back to their original source; spoofing concerns for IPv6 are briefly addressed.

Spoofing an IP Datagram

IP packets are used in applications that use the Internet as their communications medium. Usually they are generated automatically for the user, behind the scenes; the user just sees the information exchange in the application. These IP packets have the proper source and destination addresses for reliable exchange of data between two applications. The IP stack in the operating system takes care of the header for the IP datagram. However, you can override this function by inserting a custom header and informing the operating system that the packet does not need any headers. You can use raw sockets in UNIX-like systems to send spoofed IP datagrams, and you can use packet drivers such as *WinPcap* on Windows. Some socket programming knowledge is enough to write a program for generating crafted IP packets. You can insert any kind of header, so, for example, you can also create *Transmission Control Protocol* (TCP) headers. If you do not want to program or have no knowledge of programming, you can use tools such as *hping, sendip,* and others that are available for free on the Internet, with very detailed documentation to craft any kind of packet. Most of the time, you can send a spoofed address IP packet with just a one-line command.

Why Spoof the IP Source Address?

What is the advantage of sending a spoofed packet? It is that the sender has some kind of malicious intention and does not want to be identified. You can use the source address in the header of an IP datagram to trace the sender's location. Most systems keep logs of Internet activity, so if attackers want to hide their identity, they need to change the source address. The host receiving the spoofed packet responds to the spoofed address, so the attacker receives no reply back from the victim host. But if the spoofed address belongs to a host on the same subnet as the attacker, then the attacker can "sniff" the reply. You can use IP spoofing for several purposes; for some scenarios an attacker might want to inspect the response from the target victim (called "nonblind spoofing"), whereas in other cases the attacker might not care (blind spoofing). Following is a discussion about reasons to spoof an IP packet.

Scanning

An attacker generally wants to connect to a host to gather information about open ports, operating systems, or applications on the host. The replies from the victim host can help the attacker in gathering information about the system.

These replies might indicate open ports, the operating system, or several applications running on open ports. For example, a response for connection at port 80 indicates the host might be running a Web server. The hacker can then try to *telnet* to this port to see the banner and determine the Web server version and type, and then try to exploit any vulnerability associated with that Web server. In the scanning case, attackers want to examine the replies coming back from the host, so they need to see the returned packet. If the spoofed address is actually an address of a host on the attacker's subnet, then the attacker can use a sniffer to see the packets.

Sequence-Number Prediction

If you establish the connection between two hosts by using TCP, the packets exchanged between the two parties carry sequence numbers for data and acknowledgments. The protocol uses these numbers to determine out-of-order and lost packets, thus ensuring the reliable delivery to the application layer as promised by TCP. These numbers are generated pseudo-randomly in a manner known to both the parties. An attacker might send several spoofed packets to a victim to determine the algorithm generating the sequence numbers and then use that knowledge to intercept an existing session. Again it is important for the attacker to be able to see the replies.

Hijacking an Authorized Session

An attacker who can generate correct sequence numbers can send a reset message to one party in a session informing that party that the session has ended. After taking one of the parties offline, the attacker can use the IP address of that party to connect to the party still online and perform a malicious act on it. The attacker can thus use a trusted communication link to exploit any system vulnerability. Keep in mind that the party that is still online will send the replies back to the legitimate host, which can send a reset to it indicating the invalid session, but by that time the attacker might have already performed the intended actions. Such actions can range from sniffing a packet to presenting a shell from the online host to the attacker's machine.

Determining the State of a Firewall

A firewall is used to protect a network from Internet intruders. Packets entering a firewall are checked against an *Access Control List* (ACL). TCP packets sent by a source are acknowledged by acknowledgment packets. If a packet seems like an acknowledgement to a request or data from the local network, then a stateful firewall also checks whether a request for which this packet is carrying the acknowledgment was sent from the network. If there is no such request, the packet is dropped, but a stateless firewall lets packets enter the network if they seem to carry an acknowledgment for a packet. Most probably the intended receiver sends some kind of response back to the spoofed address. Again, for this process to work, the attacker should be able to see the traffic returning to the host that has the spoofed address—and the attacker generally knows how to use the returned packet to advantage.

Denial of Service

The connection setup phase in a TCP system consists of a *three-way handshake*. This handshake is done by using special bit combinations in the "flags" fields. If host A wants to establish a TCP connection with host B, it sends a packet with a SYN flag set. Host B replies with a packet that has SYN and ACK flags set in the TCP header. Host A sends back a packet with an ACK flag set, finishing the initial handshake. Then hosts A and B can communicate with each other, as shown in Figure 2.

Figure 2: A Normal TCP Connection Request from A to B

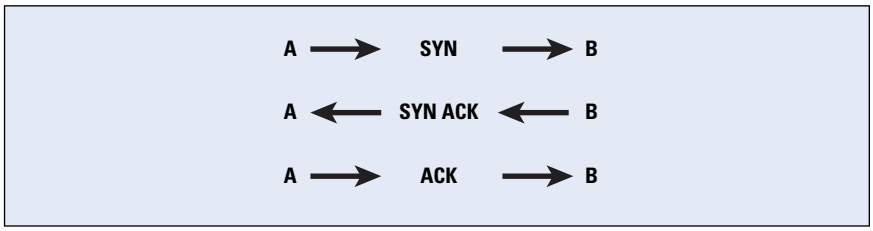

The three-way handshake must be completed in order to establish a connection. Connections that have been initiated but not finished are called *half-open connections*. A finite-size data structure is used to store the state of the half-open connections. An attacking host can send an initial SYN packet with a spoofed IP address, and then the victim sends the SYN-ACK packet and waits for a final ACK to complete the handshake. If the spoofed address does not belong to a host, then this connection stays in the half-open state indefinitely, thus occupying the data structure. If there are enough half-open connections to fill the state data structure, then the host cannot accept further requests, thus denying service to the legitimate connections (Figure 3).

Figure 3: Half-Open TCP Connection

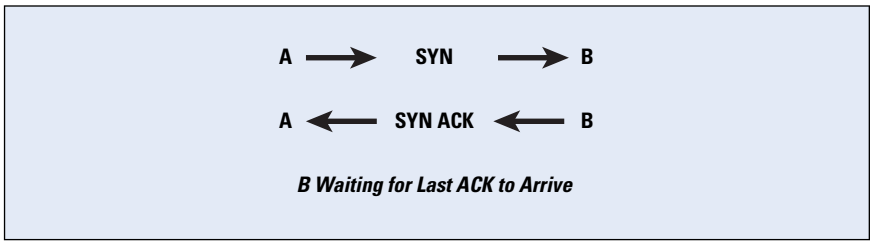

Setting a time limit for half-open connections and then erasing them after the timeout can help with this problem, but the attacker may keep continuously sending the packets. The attacked host will not have space to accept new incoming legitimate connections, but the connection that was established before the attack will have no effect. In this type of attack, the attacker has no interest in examining the responses from the victim. When the spoofed address does belong to a connected host, that host sends a reset to indicate the end of the handshake.

Flooding

In this type of attack an attacker sends a packet with the source address of the victim to multiple hosts. Responses from other machines flood the victim. For example, if an attacker uses the IP address of source A and sends a broadcast message to all the hosts in the network, then all of them will send a reply back to A, hence flooding it. The well-known *Smurf* and *fraggle* attacks used this technique.

Countermeasures for IP Spoofing

IP spoofing countermeasures include detecting spoofed IP packets and then tracing them back to the originating source. Detection of spoofed IP packets requires support of routers, host-based methods, and administrative controls, whereas tracing of IP packets involves special traceback equipment or traceback features in routers. The following section discusses both IP spoofing detection and IP spoofing traceback techniques.

Spoofed Packet Detection

Detection of a spoofed packet can start as early as at Layer 2. Switches with the *IP Source Guard* feature[8] match the MAC address of the host with a *Dynamic Host Configuration Protocol* (DHCP)-assigned dynamic or administratively assigned static IP address. Packets that do not have the correct IP source address for that particular MAC address are dropped, thereby limiting the ability of hosts connected to such a switch to send a packet with their neighbor's address. The IP Source Guard feature works very well for interfaces with a single IP address, but one interface can be assigned multiple IP addresses, and that may cause problems. The same problems can occur with *Network Address Translation* (NAT), where hosts might get different IP addresses several times. Routers work at Layer 3 in networks, and they know which interface a network is connected to and what network addresses can be expected to come from that network. If the outgoing packet from an interface does not have the network address of that interface, then the packet is spoofed and the router can stop that packet at that point; however, if the attacker is spoofing an IP address of a host on the same network (most likely in the attacks where they will be sniffing the replies), then this technique is not really helpful. The same logic can be used for an incoming packet; if a packet destined for an interface has a source address of the same network as the interface, then it is a spoofed packet. Routers can detect spoofed packets only when the packets pass through them, and if the target and attacker are both on the same subnet then this technique does not work.

Hosts receiving a suspicious packet can also use certain techniques to determine whether or not the IP address is spoofed. The first (and easiest) one is to send a request to the address of the packet and wait for the response; most of the time the spoofed addressees do not belong to active hosts and hence no response is sent.

Another method is to check the *Time to Live* (TTL) value of the packet, and then send a request to the spoofed host. If the reply comes, you can compare the TTL of both packets. Most probably the TTL values will not match. But of course it is also possible that these TTL values are the same but the packet is coming from a different source, and conversely. Packets generated by different operating systems differ slightly in values of certain fields; for example, in *Internet Control Message Protocol* (ICMP) *ping* packets, you can examine the data payload to determine the operating system. Windows fills the packet with letters of the alphabet, whereas Linux puts numbers in the data portion. If the suspicious packet does not have the same characteristics as the legitimate packet, that is evidence it was not sent from the IP address that is in its source address field. You can also use IP identification numbers to determine whether a packet is actually coming from the said source. For legitimate packets the IP ID is close in value, but this method is not reliable because the attacker can ping the said source and determine the IP ID that it is using, and then craft packets that will seem legitimate. In all these techniques we are trying to determine only whether or not a packet is spoofed, and taking all these steps for all packets would be prohibitive from an overhead standpoint. Thus you should either randomly check packets or determine some suspicious activity that would trigger further investigation for spoofed-packet detection. The next section addresses measures you can take to trace a spoofed packet back to its real source.

Tracing Spoofed IP Packets

IP traceback technology plays an important role in discovering the source of spoofed packets. Hop-by-hop traceback and logging of suspicious packets in routers are the two main methods for tracing the spoofed IP packets back to their source.

When a node detects that it is a victim of flood attack, it can inform the *Internet Service Provider* (ISP). In flood attacks the ISP can determine the router that is sending this stream to the victim, and then it can determine the next router, and so on. It reaches either to the source of the flood attack or the end of its administrative domain; for this case it can ask the ISP for the next domain to do the same thing. This technique is useful only if the flood is ongoing.

As mentioned earlier, a router has an idea of the IP addresses that should be arriving at its interfaces. If it sees any packet that does not seem to belong to the address range for its interface, it can log the packet as suspicious. Appropriately timed broadcasts among different domains to detect spoofed packets can help administrators of different networks trace spoofed IP packets back to their source.

IP Spoofing and IPv6

IP spoofing detection, or in other words validating the source address of an IPv6 packet, is a little more complicated than the process for IPv4. A host using IPv6 may potentially have multiple addresses. Again the problem inside the Local Area Network is to associate the IPv6 address with the Layer 2 or MAC address. Among peers on the same network, you can use *Neighbor Discovery* or *Secure Neighbor Discovery* (SEND) advertisements to verify the source address in a packet. You can verify source addresses of packets arriving from nodes outside the network by using the *Authentication Header* (AH) in IPv6 datagrams. You can use agreed-upon parameters between source and destination to calculate authentication information on header fields that does not change during transit. Although this process will not prevent someone from signing a spoofed address, it does provide a means to authenticate the identity of the source.

IPv6 and IPv4 network interconnections will likely face spoofing problems. IPv6 packets are usually encapsulated in IPv4 packets to travel across the non-IPv6 supporting networks. The IPv6 interim mechanism "6to4"[10, 11] uses automatic IPv6-to-IPv4 tunneling to interconnect networks using different IP versions. This mechanism uses 6to4 routers and 6to4 *Relay Routers* that accept and decapsulate IPv4 traffic from anywhere. There are no constraints on such embedded packets. Relay routers act as bridges between IPv6 and 6to4 networks and can be tricked into sending spoofed traffic anywhere. Also, anyone can send tunneled spoofed traffic to a 6to4 router, and the router will believe that it is coming from a legitimate relay. There is no simple way to prevent such attacks, and longer-term solutions are needed in both IPv6 and IPv4 networks.

Conclusion

IP spoofing is a difficult problem to tackle, because it is related to the IP packet structure. IP packets can be exploited in several ways. Because attackers can hide their identity with IP spoofing, they can make several network attacks. Although there is no easy solution for the IP spoofing problem, you *can* apply some simple proactive and reactive methods at the nodes, and use the routers in the network to help detect a spoofed packet and trace it back to its originating source.

References

[1] Alaaeldin A. Aly, "Tracking and Tracing Spoofed IP Packets to Their Sources," Proceedings of 6th annual conference, UAEU April 2005.

[2] S.J. Templeton and K.E. Levitt, "Detecting Spoofed Packets," DARPA Information Survivability Conference and Exposition, 2003.

[3] "IP Spoofing an Introduction," **http://www.securityfocus.com/infocus/1674**

[4] **http://www.phrack.org/issues.html?issue=48&id=14#article**

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

[6] **http://www.insecure.org/nmap**

[7] **http://www.ietf.org/internet-drafts/draft-baker-sava-operational-00.txt**

[8] **http://tools.ietf.org/html/draft-baker-sava-cisco-ip-source-guard-00**

[9] **http://tools.ietf.org/id/draft-baker-sava-implementation-00.txt**

[10] **http://tools.ietf.org/html/draft-ietf-v6ops-6to4-security-04**

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

[12] Wesley Eddy, "Defenses Against TCP SYN Flooding Attacks," *The Internet Protocol Journal,* Volume 9, No. 4, December 2006.

FARHA ALI holds a BE in Computer Engineering from NED University, Pakistan, and an MS in Computer Engineering from Clemson University, South Carolina, with a focus area in Computer Communications. She is a member of American Mensa and ACM. Her research papers (co-authored with her advisor) as a PhD student at Clemson University's Computer Science Department were published in IEEE's Conferences on Web Intelligence and Web Services. She is a Sun Certified Java Programmer and a Certified Ethical Hacker. Her main interests are Distributed Computing, Network Security, and Semantic Web. Currently she is working as a faculty member at Lander University's Department of Mathematics and Computing. She teaches mainly Networking and Programming courses.
E-mail: **fali@lander.edu**

Standards for Information Security Managment

by William Stallings

To effectively assess the security needs of an organization and to evaluate and choose various security products and policies, the manager responsible for security needs some systematic way of defining the requirements for security and characterizing the approaches to satisfy those requirements. This process is difficult enough in a centralized data processing environment; with the use of local- and wide-area networks (LANs and WANs, respectively), the problems are compounded.

The challenges for management in providing information security are formidable. Even for relatively small organizations, information system assets are substantial, including databases and files related to personnel, company operation, financial matters, and so on. Typically, the information system environment is complex, including a variety of storage systems, servers, workstations, local networks, and Internet and other remote network connections. Managers face a range of threats always growing in sophistication and scope. And the range of consequences for security failures, both to the company and to individual managers, is substantial, including financial loss, civil liability, and even criminal liability.

Standards for providing information system security become essential in such circumstances. Standards can define the scope of security functions and features needed, policies for managing information and human assets, criteria for evaluating the effectiveness of security measures, techniques for ongoing assessment of security and for the ongoing monitoring of security breaches, and procedures for dealing with security failures.

Figure 1, based on [1], suggests the elements that, in an integrated fashion, constitute an effective approach to information security management. The focus of this approach is on two distinct aspects of providing information security: process and products. *Process security* looks at information security from the point of view of management policies, procedures, and controls. *Product security* focuses on technical aspects and is concerned with the use of certified products in the IT environment when possible. In Figure 1, the term *technical standards* refers to specifications that refer to aspects such as IT network security, digital signatures, access control, nonrepudiation, key management, and hash functions. Operational, management, and technical *procedures* encompass policies and practices that are defined and enforced by management. Examples include personnel screening policies, guidelines for classifying information, and procedures for assigning user IDs. *Management system audits, certification, and accreditation* deals with management policies and procedures for auditing and certifying information security products.

Codes of practice refer to specific policy standards that define the roles and responsibilities of various employees in maintaining information security. *Assurance* deals with product and system testing and evaluation. *Cultural, ethical, social, and legal issuers* refer to human factors aspects related to information security.

Figure 1: Information Security Management Elements

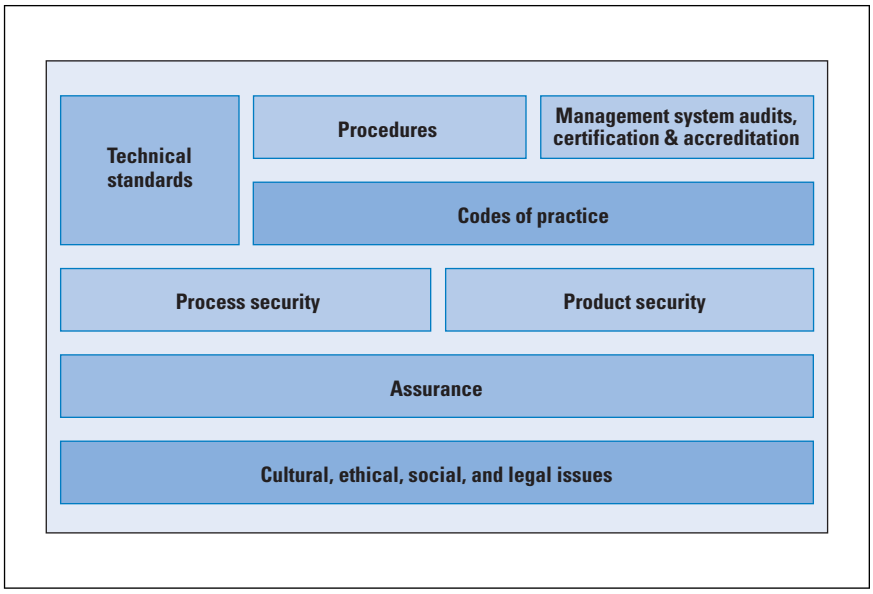

Many standards and guideline documents have been developed in recent years to aid management in the area of information security. The two most important are *ISO 17799*, which deals primarily with process security, and the *Common Criteria*, which deals primarily with product security. This article surveys these two standards, and examines some other important standards and guidelines as well.

ISO 17799

An increasingly popular standard for writing and implementing security policies is *ISO 17799* "Code of Practice for Information Security Management." (ISO 17799 will eventually be reissued as ISO 27002 in the new ISO 27000 family of security standards). ISO 17799 is a comprehensive set of controls comprising best practices in information security. It is essentially an internationally recognized generic information security standard. Table 1 summarizes the area covered by this standard and indicates the objectives for each area.

Table 1: ISO 17799 Areas and Objectives

Security Policy Provide management direction and support for information security in accordance with business requirements and relevant laws and regulations. **Organization of Information Security** Manage information security within the organization. Maintain the security of the organization's information and information processing facilities that are accessed, processed, communicated to, or managed by external parties. **Asset Management** Achieve and maintain appropriate protection of organizational assets. Ensure that information receives an appropriate level of protection. **Human Resources Security** Ensure that employees, contractors, and third-party users (1) understand their responsibilities and are suitable for the roles they are considered for; (2) are aware of information security threats and concerns; (3) exit an organization or change employment in an orderly manner. **Physical and Environmental Security** Prevent unauthorized physical access, damage, and interference to the organization's premises and information. Prevent loss, damage, theft, or compromise of assets and interruption to the organization's activities. **Communications and Operations Management** Develop controls for operational procedures, third-party service delivery management, system planning, malware protection, backup, network security management, media handling, information exchange, e-commerce services, and monitoring.	**Access Control** Develop controls for business requirements for user access, user responsibilities, network access control, OS access control, application access control, and information access control. **Information Systems Acquisition, Development, and Maintenance** Develop controls for correct processing in applications, cryptographic functions, system file security, support process security, and vulnerability management. **Information Security Incident Management** Ensure information security events and weaknesses associated with information systems are communicated in a manner that allows timely corrective action to be taken. Ensure a consistent and effective approach is applied to the management of information security incidents. **Business Continuity Management** Counteract interruptions to business activities to protect critical business processes from the effects of major failures of information systems or disasters and to ensure their timely resumption. **Compliance** Avoid breaches of any law, statutory, regulatory, or contractual obligations, and of any security requirements. Ensure compliance of systems with organizational security policies and standards. Maximize the effectiveness of and minimize interference to and from the information systems audit process.

With the increasing interest in security, ISO 17799 certification, provided by various accredited bodies, has been established as a goal for many corporations, government agencies, and other organizations around the world. ISO 17799 offers a convenient framework to help security policy writers structure their policies in accordance with an international standard.

Much of the content of ISO 17799 deals with security controls, which are defined as practices, procedures, or mechanisms that may protect against a threat, reduce a vulnerability, limit the effect of an unwanted incident, detect unwanted incidents, and facilitate recovery. Some controls deal with security management, focusing on management actions to institute and maintain security policies. Other controls are operational; they address the correct implementation and use of security policies and standards, ensuring consistency in security operations and correcting identified operational deficiencies. These controls relate to mechanisms and procedures that are primarily implemented by people rather than systems.

Finally, there are technical controls; they involve the correct use of hardware and software security capabilities in systems. These controls range from simple to complex measures that work together to secure critical and sensitive data, information, and IT systems functions. This concept of controls cuts across all the areas listed in Table 1.

To give some idea of the scope of ISO 17799, we examine several of the security areas discussed in that document. *Auditing* is a key security management function that is addressed in multiple areas within the document. First, ISO 17799 lists key data items that should, when relevant, be included in an audit log:

- User IDs
- Dates, times, and details of key events, for example, log-on and log-off
- Terminal identity or location if possible
- Records of successful and rejected system access attempts
- Records of successful and rejected data and other resource access attempts
- Changes to system configuration
- Use of privileges
- Use of system utilities and applications
- Files accessed and the kind of access
- Network addresses and protocols
- Alarms raised by the access control system
- Activation and deactivation of protection systems, such as antivirus systems and intrusion detection systems

It provides a useful set of guidelines for implementation of an auditing capability:

1. Audit requirements should be agreed upon by appropriate management.
2. The scope of the checks should be agreed upon and controlled.
3. The checks should be limited to read-only access to software and data.
4. Access other than read-only should be allowed only for isolated copies of system files, which should be erased when the audit is completed or given appropriate protection if there is an obligation to keep such files under audit documentation requirements.
5. Resources for performing the checks should be explicitly identified and made available.
6. Requirements for special or additional processing should be identified and agreed upon.

7. All access should be monitored and logged to produce a reference trail; the use of timestamped reference trails should be considered for critical data or systems.
8. All procedures, requirements, and responsibilities should be documented.
9. The person(s) carrying out the audit should be independent of the activities audited.

Under the area of communications and operations management, ISO 17799 includes *network security management*. One aspect of this management is concerned with network controls for networks owned and operated by the organization. The document provides implementation guidance for these in-house networks. An example of a control follows: Restoration procedures should be regularly checked and tested to ensure that they are effective and that they can be completed within the time allotted in the operational procedures for recovery. Similarly, the document provides guidance for security controls for network services provided by outside vendors. An example of guidance in this area follows: The ability of the network service provider to manage agreed-upon services in a secure way should be determined and regularly monitored, and the right to audit should be agreed upon.

As can be seen, some ISO 17700 specifications are detailed and specific, whereas others are quite general.

Common Criteria

The *Common Criteria for Information Technology Security Evaluation* (CC) is a joint international effort by numerous national standards organizations and government agencies[3,4,5]. U.S. participation is by the *National Institute of Standards and Technology* (NIST) and the *National Security Agency* (NSA). CC defines a set of IT requirements of known validity that can be used in establishing security requirements for prospective products and systems. The CC also defines the *Protection Profile* (PP) construct that allows prospective consumers or developers to create standardized sets of security requirements that will meet their needs.

The aim of the CC specification is to provide greater confidence in the security of IT products as a result of formal actions taken during the process of developing, evaluating, and operating these products. In the development stage, the CC defines sets of IT requirements of known validity that can be used to establish the security requirements of prospective products and systems. Then the CC details how a specific product can be evaluated against these known requirements, to provide confirmation that it does indeed meet them, with an appropriate level of confidence. Lastly, when in operation the evolving IT environment may reveal new vulnerabilities or concerns. The CC details a process for responding to such changes, and possibly reevaluating the product.

Following successful evaluation, a particular product may be listed as CC certified or validated by the appropriate national agency, such as NIST or NSA in the United States. That agency publishes lists of evaluated products, which are used by government and industry purchasers who need to use such products.

The CC defines a common set of potential *security requirements* for use in evaluation. The term *Target of Evaluation* (TOE) refers to that part of the product or system that is subject to evaluation. The requirements fall into two categories:

- *Functional requirements:* Define desired security behavior. CC documents establish a set of security functional components that provide a standard way of expressing the security functional requirements for a TOE.
- *Assurance requirements:* The basis for gaining confidence that the claimed security measures are effective and implemented correctly. CC documents establish a set of assurance components that provide a standard way of expressing the assurance requirements for a TOE.

Both functional requirements and assurance requirements are organized into classes: A *class* is a collection of requirements that share a common focus or intent. Each of these classes contains numerous families. The requirements within each *family* share security objectives but differ in emphasis or rigor. For example, the audit class contains six families dealing with various aspects of auditing (for example, audit data generation, audit analysis, and audit event storage). Each family, in turn, contains one or more components. A *component* describes a specific set of security requirements and is the smallest selectable set of security requirements for inclusion in the structures defined in the CC.

For example, the cryptographic support class of functional requirements includes two families: cryptographic key management and cryptographic operation. The cryptographic key management family has four components, which are used to specify key generation algorithm and key size; key distribution method; key access method; and key destruction method. For each component, a standard may be referenced to define the requirement. Under the cryptographic operation family, there is a single component, which specifies an algorithm and key size based on a an assigned standard.

Sets of functional and assurance components may be grouped together into reusable packages, which are known to be useful in meeting identified objectives. An example of such a package would be functional components required for *Discretionary Access Controls*.

The CC also defines two kinds of documents that can be generated using the CC-defined requirements.

- *Protection profiles (PPs):* Define an implementation-independent set of security requirements and objectives for a category of products or systems that meet similar consumer needs for IT security. A PP is intended to be reusable and to define requirements that are known to be useful and effective in meeting the identified objectives. The PP concept has been developed to support the definition of functional standards and as an aid to formulating procurement specifications. The PP reflects user security requirements.
- *Security targets (STs):* Contain the IT security objectives and requirements of a specific identified TOE and define the functional and assurance measures offered by that TOE to meet stated requirements. The ST may claim conformance to one or more PPs and forms the basis for an evaluation. The ST is supplied by a vendor or developer.

Figure 2 illustrates the relationship between requirements on the one hand and profiles and targets on the other. For a PP, a user can select many components to define the requirements for the desired product. The user may also refer to predefined packages that assemble numerous requirements commonly grouped together within a product requirements document. Similarly, a vendor or designer can select numerous components and packages to define an ST.

Figure 2: Organization and Construction of Common Criteria Requirements

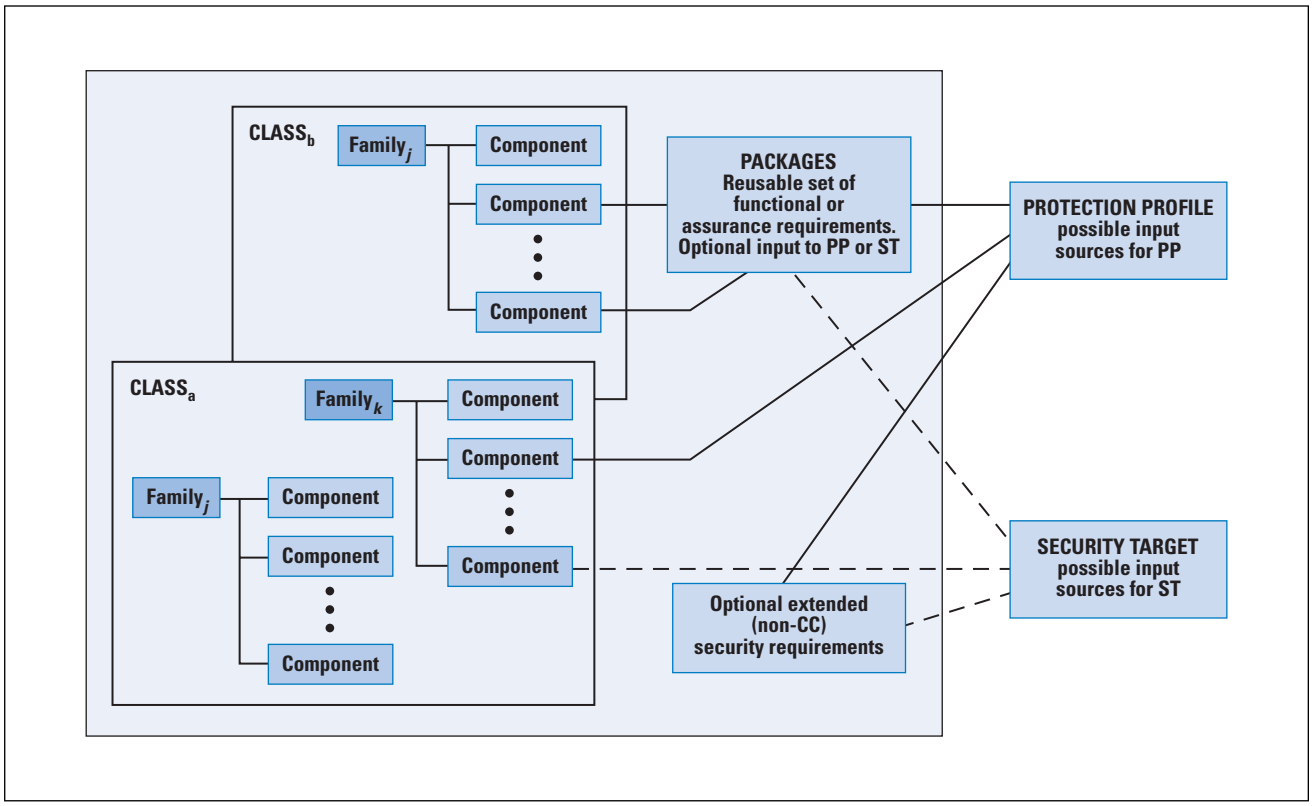

As an example for the use of the CC, consider the smart card. The protection profile for a smart card, developed by the *Smart Card Security User Group,* provides a simple example of a PP. This PP describes the IT security requirements for a smart card to be used in connection with sensitive applications, such as banking industry financial payment systems. The assurance level for this PP is *Evaluation Assurance Level* (EAL) 4, which is described subsequently. The PP lists *threats* that must be addressed by a product that claims to comply with this PP. The threats include the following:

- *Physical probing:* May entail reading data from the TOE through techniques commonly employed in *Integrated Circuit* (IC) failure analysis and IC reverse engineering efforts.
- *Invalid input:* Invalid input may take the form of operations that are not formatted correctly, requests for information beyond register limits, or attempts to find and execute undocumented commands. The result of such an attack may be a compromise in the security functions, generation of exploitable errors in operation, or release of protected data.
- *Linkage of multiple operations:* An attacker may observe multiple uses of resources or services and, by linking these observations, deduce information that may reveal security function data.

Following a list of threats, the PP turns to a description of *security objectives*, which reflect the stated intent to counter identified threats or comply with any organizational security policies identified. Nineteen objectives are listed, including the following:

- *Audit:* The system must provide the means of recording selected security-relevant events, so as to assist an administrator in the detection of potential attacks or misconfiguration of the system security features that would leave it susceptible to attack.
- *Fault insertion:* The system must be resistant to repeated probing through insertion of erroneous data.
- *Information leakage:* The system must provide the means of controlling and limiting the leakage of information in the system so that no useful information is revealed over the power, ground, clock, reset, or I/O lines.

Security requirements are provided to thwart specific threats and to support specific policies under specific assumptions. The PP lists specific requirements in three general areas: TOE security functional requirements, TOE security assurance requirements, and security requirements for the IT environment. In the area of *security functional requirements,* the PP defines 42 requirements from the available classes of security functional requirements.

For example, for security auditing, the PP stipulates what the system must audit; what information must be logged; what the rules are for monitoring, operating, and protecting the logs; and so on. Functional requirements are also listed from the other functional requirements classes, with specific details for the smart card operation.

The PP defines 24 *security assurance requirements* from the available classes of security assurance requirements. These requirements were chosen to demonstrate:

- The quality of the product design and configuration
- That adequate protection is provided during the design and implementation of the product
- That vendor testing of the product meets specific parameters
- That security functions are not compromised during product delivery
- That user guidance, including product manuals pertaining to installation, maintenance, and use, are of a specified quality and appropriateness

The PP also lists *Security Requirements of the IT Environment*. They cover the following topics:

- Cryptographic key distribution
- Cryptographic key destruction
- Security roles

The final section of the PP (excluding appendices) is a lengthy rationale for all the selections and definitions in the PP. The PP is an industrywide effort designed to be realistic in its ability to be met by a variety of products with a variety of internal mechanisms and implementation approaches.

The concept of *Evaluation Assurance* is a difficult one to define. Further, the degree of assurance required varies from one context and one function to another. To structure the need for assurance, the CC defines a scale for rating assurance consisting of seven EALs ranging from the least rigor and scope for assurance evidence (EAL 1) to the most (EAL 7). The levels are as follows:

- *EAL 1: Functionally tested:* For environments where security threats are not considered serious. It involves independent product testing with no input from the product developers. The intent is to provide a level of confidence in correct operation.
- *EAL 2: Structurally tested:* Includes a review of a high-level design provided by the product developer. Also, the developer must conduct a vulnerability analysis for well-known flaws. The intent is to provide a low to moderate level of independently assured security.

- *EAL 3: Methodically tested and checked:* Requires a focus on the security features, including requirements that the design separate security-related components from those that are not; that the design specifies how security is enforced; and that testing be based both on the interface and the high-level design, rather than a black box testing based only on the interface. It is applicable where the requirement is for a moderate level of independently assured security, with a thorough investigation of the TOE and its development without incurring substantial reengineering costs.
- *EAL 4: Methodically designed, tested, and reviewed:* Requires both a low-level and a high-level design specification; requires that the interface specification be complete; requires an abstract model that explicitly defines security for the product; and requires an independent vulnerability analysis. It is applicable in those circumstances where developers or users require a moderate to high level of independently assured security in conventional commodity TOEs, and there is willingness to incur some additional security-specific engineering costs.
- *EAL 5: Semiformally designed and tested:* Provides an analysis that includes all of the implementation. Assurance is supplemented by a formal model and a semiformal presentation of the functional specification and high-level design and a semiformal demonstration of correspondence. The search for vulnerabilities must ensure resistance to penetration attackers with a moderate attack potential. Covert channel analysis and modular design are also required.
- *EAL 6: Semiformally verified design and tested:* Permits a developer to gain high assurance from application of specialized security engineering techniques in a rigorous development environment, and to produce a premium TOE for protecting high-value assets against significant risks. The independent search for vulnerabilities must ensure resistance to penetration attackers with a high attack potential.
- *EAL 7: Formally verified design and tested:* The formal model is supplemented by a formal presentation of the functional specification and high-level design, showing correspondence. Evidence of developer "white box" testing of internals and complete independent confirmation of developer test results are required. Complexity of the design must be minimized.

The first four levels reflect various levels of commercial design practice. Only at the highest of these levels (EAL 4) is there a requirement for any source code analysis, and this analysis is required only for a portion of the code. The top three levels provide specific guidance for products developed using security specialists and security-specific design and engineering approaches.

National Institute of Standards and Technology

NIST has produced a large number of *Federal Information Processing Standards Publications* (FIPS PUBs) and special publications (SPs) that are enormously useful to security managers, designers, and implementers. Following are a few of the most significant and general. *FIPS PUB 200* "Minimum Security Requirements for Federal Information and Information Systems," is a standard that specifies minimum security requirements in 17 security-related areas with regard to protecting the confidentiality, integrity, and availability of federal information systems and the information processed, stored, and transmitted by those systems[6].

NIST *SP 800-100* "Information Security Handbook: A Guide for Managers," provides a broad overview of information security program elements to assist managers in understanding how to establish and implement an information security program[7]. Its topical coverage overlaps considerably with ISO 17799.

Several other NIST publications are of general interest. *SP 800-55* "Security Metrics Guide for Information Technology Systems," provides guidance on how an organization, through the use of metrics, identifies the adequacy of in-place security controls, policies, and procedures[8]. *SP 800-27* "Engineering Principles for Information Technology Security (A Baseline for Achieving Security)," presents a list of system-level security principles to be considered in the design, development, and operation of an information system[9]. *SP 800-53* "Recommended Security Controls for Federal Information Systems," lists management, operational, and technical safeguards or countermeasures prescribed for an information system to protect the confidentiality, integrity, and availability of the system and its information[10].

Other Standards and Guidelines

Another important set of standards is the *Control Objectives for Information and Related Technology* (COBIT)[11], a business-oriented set of standards for guiding management in the sound use of information technology. It has been developed as a general standard for information technology security and control practices and includes a general framework for management, users, IS audit, and security practitioners. COBIT also has a process focus and a governance flavor; that is, management's need to control and measure IT is a focus point. COBIT was developed under the auspices of a professional organization, the *Information Systems Audit and Control Association* (ISACA). The documents are quite detailed and provide a practical basis for not only defining security requirements but also implementing them and verifying compliance.

Another excellent source of information is "The Standard of Good Practice for Information Security" from the *Information Security Forum*. The standard is designed as an aid to organizations in understanding and applying best practices for information security. Because it addresses security from a business perspective, The Standard appropriately recognizes the intersection between organizational factors and security factors.

In addition to these standards, numerous informal guidelines are widely consulted by organizations in developing their own security policy. The *CERT Coordination Center* (**www.cert.org**) has an Evaluations and Practices section of its Website with a variety of documents and training aids related to information security for organizations. The *Chief Information Officers Council* (**cio.gov**) has published a collection of Best Practices and other documents related to organizational security.

References

[1] Eloff, J., and Eloff, M., "Information Security Management," *Proceedings of SAICSIT 2003,* South African Institute of Computer Scientists and Information Technologists, 2003.

[2] International Organization for Standardization, "ISO/IEC 27001 – Information technology – Security Techniques – Information security management systems – Requirements," June 2005.

[3] Common Criteria Project Sponsoring Organisations, "Common Criteria for Information Technology Security Evaluation, Part 1: Introduction and General Model," CCIMB-2004-01-001, January 2004.

[4] Common Criteria Project Sponsoring Organisations, "Common Criteria for Information Technology Security Evaluation, Part 2: Security Functional Requirements," CCIMB-2004-01-002, January 2004.

[5] Common Criteria Project Sponsoring Organisations, "Common Criteria for Information Technology Security Evaluation, Part 3: Security Assurance Components," CCIMB-2006-09-003, September 2006.

[6] National Institute of Standards and Technology, "Minimum Security Requirements for Federal Information and Information Systems," FIPS PUB 200, March 2006.

[7] National Institute of Standards and Technology, "Information Security Handbook: A Guide for Managers," NIST Special Publication 800-100, October 2006.

[8] "Security Metrics Guide for Information Technology Systems," NIST Special Publication 800-55, July 2003.

[9] National Institute of Standards and Technology, "Engineering Principles for Information Technology Security (A Baseline for Achieving Security)," NIST Special Publication 800-27, June 2004.

[10] National Institute of Standards and Technology, "Recommended Security Controls for Federal Information Systems," NIST Special Publication 800-53, February 2005.

[11] IT Governance Institute, "COBIT 4.0.," USA, 2005.

[12] Information Security Forum, "The Standard of Good Practice for Information Security," 2005.

WILLIAM STALLINGS is a consultant, lecturer, and author of more than a dozen books on data communications and computer networking. His latest book, with Lawrie Brown, is *Computer Security: Principles and Practice* (Prentice Hall, 2007). He maintains a computer science resource site for computer science students and professionals at **`WilliamStallings.com/StudentSupport.html`** and is on the editorial board of *Cryptologia*. He has a Ph.D. in computer science from M.I.T. He can be reached at **`ws@shore.net`**.

Looking Toward the Future

by Vint Cerf, Google

The *Internet Corporation for Assigned Names and Numbers* (ICANN) was formed 9 years ago, following a period of considerable debate about the institutionalization of the basic functions performed by the *Internet Assigned Numbers Authority* (IANA)[1]. Nearly simultaneous with the inauguration of ICANN in September 1998 came the unexpected and untimely death of the man, Jonathan B. Postel[2], who had responsibility for these functions for more than a quarter century. The organization began with very limited sources of funds, a small and overworked staff, and contentious debate about its organizational structure, policy apparatus, and operational procedures. The organization underwent substantial change through its *Evolution and Reform Process* (ERP)[3]. Among the more difficult constituencies to accommodate in the organization's policy-making process was the general public. An *At-Large Advisory Committee* (ALAC)[4] emerged from the ERP and has recently formed *Regional At-Large Organizations* (RALOs) in all of ICANN's five regions.

Today, ICANN is larger, more capable, more international, and better positioned to fulfill its mandate. It stands for one global, interoperable Internet, and the model of stakeholder representation has worked. But the Internet and its vast user population have grown during the same time by a factor of more than 20 in all dimensions. The 50 million users of 1997 have become nearly 1.2 billion users today. The 22 million hosts on the network have increased to nearly 500 million today. The bandwidth of the core data circuits in the Internet have grown from 622 million bits per second to between 10 and 40 billion bits per second. This dramatic growth in physical size has been accompanied by an equally dramatic growth in the number and diversity of applications running on the Internet. All forms of media now appear on and are carried by Internet packets. Consumers of information are producing more and more of it themselves with e-mail, blogs, instant messaging, social and game-playing Websites, video uploads, and podcasts. The Internet continues to evolve and although ICANN has achieved more than most people realize, it must continue to evolve along with it.

Operational Priorities

ICANN's primary responsibility is to contribute to the security and stability of the Internet system of unique identifiers. In the most direct way, it carries out this mandate through its operation of the IANA. There is no doubt that the conduct of this function in an exemplary fashion is essential not only to ICANN's mission but also to inspiring confidence in ICANN as an organization.

But ICANN's role in the Internet goes beyond these specific IANA functions. ICANN is an experiment in the balancing of multiple stakeholder interests in policy about the implementation, operation, and use of the *Domain Name System* (DNS) and the address spaces of the Internet. Its policy choices can directly affect the business models of operating entities involved in the management of domain names and Internet addresses. The privacy and Internet-related rights of registrants and, more generally, Internet users may also be directly affected. Some policy choices raise public policy concerns in the view of governments and methods and will be needed to factor such concerns into the making of ICANN policy.

Effective, fair, and timely policy development should be a priority for ICANN. That this policy development needs to be achieved in a global setting is simply another challenge to be met. ICANN leadership and staff must seek to maintain and improve the ability of all of ICANN's many constituencies to achieve consensus or at least to prepare the ICANN Board to make choices when consensus may not be forthcoming. Because policies often have technical, economic, social, and governance implications, it is vital that ICANN's practices draw on expertise in all these domains.

Clarity in the roles and responsibilities of the many participants in the Internet arena, especially those with specific interest in ICANN policies and practices, will be helpful and should be documented. In some cases, the documentation might take the form of relatively formal relationships such as the contracts between ICANN and domain-name registries and registrars. In other cases, they may need only to characterize in plain terms the roles that each party plays.

In some areas, such as root-zone operation, excellence can be measured in such terms as responsiveness, scalability, resilience to disruption, and ability to adapt to changing needs such as *Domain Name System Security* (DNSSEC)[5], *Internationalized Domain Names* (IDNs)[6], and the addition of IPv6 records to the root zone. Many parties currently play a role in the maintenance of the root-zone file, and clear documentation of responsibility and lines of authority will be beneficial. As the technology of the Internet continues to evolve, the roles of various parties may need to change to meet the objectives of stability and security of the Internet system of unique identifiers. Managing the evolution of these roles represents another priority for policy development and implementation.

Because of the potential effect of decisions made through the ICANN policy process, it is important to implement checks and balances that make all aspects of ICANN's operation accountable and transparent. Work is still necessary in this area so that independent review of legitimate concerns arising out of policy making is possible when deemed necessary.

At the same time, it is vital that the mechanisms chosen do not have the effect of locking up the policy-making process and preventing any decisions from being made. We need to seek a balance between a potentially unfair tyranny of the majority and an equally unacceptable tyranny of the minority.

The general success of the *Uniform Dispute Resolution Process* (UDRP)[7] suggests that ICANN should seek mechanisms for resolving disputes arising in connection with implementing ICANN policies that scale, permit choice without abusive "forum shopping," and make efficient use of ICANN resources.

Outreach, transparency, and broadly participatory processes on an international basis are not inexpensive. It is vital for ICANN to continue to refine its models for sustainable operation, accounting for the economics of the various actors in the Internet arena that rely on ICANN's operation, and fairly apportioning costs of ICANN operation to appropriate sources of support. Not all of the beneficiaries of ICANN's work derive the same level of revenue from the Internet (and some, none at all). ICANN must account for this discrepancy when devising mechanisms for supporting its operation, and it should work to make transparent the need to provide services to parties who may not be able to contribute commensurate with cost. Adequate and stable funding for ICANN is necessary if ICANN is to fulfill its charter. Over the past several years, ICANN has significantly increased its ability to staff vital functions, contributing to the effectiveness of the organization. It should be a priority to assure adequate reserves to weather unanticipated expenses or periods of decreased income.

Organizational Perspectives

ICANN is a multistakeholder institution operating in the private sector but including the involvement of governments. Throughout its history, ICANN has sought to draw on international resources and to collaborate, coordinate, and cooperate with institutions whose expertise and responsibilities can assist ICANN in the achievement of its goals. ICANN should seek to establish productive relationships with these institutions, cementing its own place in the Internet universe while confining its role to its principal responsibilities.

As part of its normal operation, ICANN engages in self-examination and external review of the effectiveness of its organizational structure and processes. Improvements in all aspects of ICANN operation and structure will increase confidence in the organization and its ability to sustain long-term operation.

Finding and engaging competent participants and leaders in each of ICANN's constituent parts must be a priority. ICANN should seek to improve its ability to identify from around the world and attract highly qualified staff, executive leadership, board, and supporting organization participants. It is possible and even likely that improvements in the processes by which this process is done today will have significant payoff in the future.

Although ICANN does not bear a specific responsibility for achieving the *Millennium Development Goals* (MDGs) developed during the conduct of the *World Summit on the Information Society* (WSIS)[8], it has an opportunity to contribute to them both directly and indirectly. Its operation of its IANA functions and support for actors in the domain-name, Internet-address, and standards-development areas provides ICANN with a specific opportunity. Participation in forums dedicated to developing policies for Internet expansion and use offer indirect ways for ICANN to draw upon and provide expertise in these areas.

It has been demonstrated that the presence of ICANN staff in various regions and time zones around the world and familiarity with local languages and customs has been beneficial to parties reliant on ICANN for its services. ICANN should continue to seek ways to improve its effectiveness in this area. The introduction of the Fellowship program that supports the participation of qualified candidates in ICANN-related activities is a vital step in facilitating ICANN's outreach to the developing world. We should pursue expansion of this program through partnerships with other like-minded organizations in the interest of the globalization of ICANN.

It is possible that the present formulation of ICANN as a not-for-profit, charitable research and education entity under California law could be beneficially adapted to a more international framework. As part of its long-term strategic development, ICANN should evaluate a variety of alternatives on the possibility that a change could increase the effectiveness of its operation.

The successful creation of five Regional At-Large Organizations, one in each of ICANN's five regions, needs to be followed by a serious effort to engage these entities in the formulation of ICANN policies and in dialog with the general user community. The various constituency reviews that form part of ICANN's normal processes should address the role of these entities in the conduct of ICANN business. To the extent that civil society is not fully represented through the *Governmental Advisory Committee* (GAC)[9] and the ALAC/RALO system, an organizational home may be needed to accommodate the interests of that constituency.

The five *Regional Internet Registries* (RIRs)[10] represent a key element in the Internet and ICANN pantheon. The RIRs have responsibility for allocating IP address space to Internet service providers and sometimes individual end-user organizations. They are the means by which bottom-up global policy is developed and recommended, through the *Number Resource Organization* (NRO)[11], to ICANN. It will require substantial coordination and cooperation between the RIRs and ICANN to work through the coming years of depletion of available new IPv4 address space and the rising implementation of the new IPv6 address space.

There is little doubt that economic incentives will emerge that will distort fair and neutral IPv4 address-space allocations as the available space is depleted. Minimizing the effect of this transition will be the joint responsibility of ICANN and the RIRs.

Similarly, ICANN's cooperative relationship with the *Root Server Operators*[12] will also demand coordination and capacity building as IPv4 and IPv6 addresses are associated with old and new domain names and as the IPv6 infrastructure grows. A vital objective is to assure that the IPv6 Internet and the IPv4 Internet are, to the extent possible, completely and totally coterminous. Every termination needs to be reachable through both address spaces. In the absence of this uniformity, some IPv6 addresses may be unreachable from others, defeating the goal of a single, interoperable, and fully reachable network.

Meeting the Challenges

As ICANN approaches the close of its first decade, the operational Internet will be turning 25. In the course of its evolution, it has become a global digital canvas on which a seemingly endless array of applications has been painted. Despite the broad swath of its current applications, it is almost certain that many, many more will be invented. All of them will rely, for the foreseeable future, on the basic architecture of the system, including the global Internet address space and the DNS. But the structure will become more complex. Two parallel address spaces, IPv4 and IPv6, will be in use. ICANN needs to promote the adoption of IPv6 so as to limit the side effects of the exhaustion of the unique address space provided by IPv4.

A vast and new range of non-Latin, internationalized domain names may be registered, certainly at the second or lower levels in the domain-name hierarchy, and many will be proposed for the top level. Their diversity will create new challenges for the protection of users from confusing and potentially abusive registrations. New dispute resolution principles may be needed to deal with domain-name registrations and delegations of new top-level domains. The exposure of ASCII *punycode* strings in browsers or other applications may produce additional stresses in the intellectual property arena (for example **`xn--cocacola`**).

Digital signatures will play an increasingly important role in validating the assignment of domain names and Internet addresses, and new protocols are certain to be invented and their parameters recorded by the IANA. Infrastructure for the management of digital certificates or other authentication mechanisms will be needed to realize the value of the DNSSEC concept.

More generally, the multilayer architecture of the Internet shows vulnerabilities of various kinds that demand redress. Attacks against the DNS root servers, name resolvers, and general name servers at all levels must be mitigated.

Some of the components of the DNS are actually used to exacerbate the effects of *Denial-of-Service* (DoS) attacks. Although ICANN does not have responsibility for developing the Domain Name technology, it can use its visibility and area of responsibility to highlight the need for increased security measures for the protection of the technical infrastructure of the Internet and to facilitate its implementation where ICANN has a direct involvement in its operation.

An increasing number of mobile devices will become Internet-enabled, as will appliances of all kinds. Access speeds will increase, enabling many new applications and enhancing older ones. All of this activity will contribute to increasing reliance on the Internet for a wide range of functions by an increasingly larger user population. Electronic commerce will continue to expand, placing high priority on the stable, secure, and reliable operation of all aspects of the Internet, including those within ICANN's purview.

Although some of these aspects of the evolution of the Internet will be of direct concern to ICANN, the ICANN organization and processes will need to pay attention to additional matters as well. The business processes that sustain the management of the Internet address space and domain names will almost certainly need to adapt to account for new applications. Some of these applications will monetize various aspects of the Internet in unexpected and innovative ways that will challenge existing policy and procedures. It will be extremely important for ICANN to evolve and strengthen its implementation of multistakeholder policy development. The interests of a wide range of entities must be balanced in the process.

Although adherence to a set of technical standards has allowed millions of component networks and systems to interwork on the Internet, it is also the case that many varying business models have sustained their operation. The richness and diversity of these models is one of the reasons that the Internet has proved to be so resilient in many dimensions. ICANN's policy-development processes need to account for an informed understanding of the economics of these varying business models and the ways in which ICANN policy may affect them.

On the Domain Name side, the development of market-savvy rules of operation for operators will be essential. ICANN needs to assure compliance with policies developed through the ICANN consensus process to establish confidence in the policy processes and their execution. Clear rules for the creation of new *Top Level Domains* (TLDs) of all kinds must be adopted and enforced.

The roles of registrars, registries, wholesale registry operators, root-server operators, regional Internet address registries, governments, and standards and technical research and development bodies, among others, need to be characterized so as to set expectations and permit the establishment of practical working relationships. The documentation of best practices will be beneficial, especially where the introduction of the Internet is new.

In matters of public policy—including but not limited to public safety, security, privacy, law enforcement, conduct of electronic commerce, protection of digital property, and freedom of speech—broad and international agreements may be needed if the Internet is to serve as a useful, global infrastructure. Many of these matters lie outside the formal purview of ICANN, but some ICANN policies and resulting operational practices will contribute to the global framework for life online. ICANN must seek to contribute to public confidence in the Internet and the processes that govern its operation. It cannot accomplish this objective alone. The coordinated and cooperative efforts of many distinct entities will be essential to achieving this goal. At the same time, ICANN must protect its processes from capture or abuse by interests that are inimical to the openness and accessibility of the Internet for everyone.

A Collective Goal

As of this writing, only about 1.2 billion people around the world use the Internet. Over the course of the next decade, that number could conceivably quintuple to 6 billion, and users will be depending on ICANN, among many others, to do its part to make the Internet a productive infrastructure that invites and facilitates innovation and serves as a platform for egalitarian access to information. It should be a platform that amplifies voices that might otherwise never be heard and creates equal opportunities for increasing the wealth of nations and their citizens.

ICANN's foundation has been well and truly fashioned. It is the work of many heads and hands. It represents a long and sometimes difficult journey that has called for personal sacrifices from many colleagues and bravery from others. It has demanded long-term commitments, long hours, days, months, and years. It has called upon many to transform passion and zeal into constructive and lasting compromises. ICANN has earned its place in the Internet universe. To those who now guide its path into the future comes the challenge to fashion an enduring institution on this solid foundation. I am confident that this goal is not only attainable but now also necessary. The opportunity is there: make it so.

For Further Reading

[1] **http://www.iana.org**

[2] Vint Cerf, "I Remember IANA," *The Internet Protocol Journal*, Volume 1, No. 3, December 1998. Also published as RFC 2468, October 1998.

[3] **http://www.icann.org/committees/evol-reform/**

[4] **http://alac.icann.org/**

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

[6] **http://icann.org/topics/idn/**

[7] **http://icann.org/udrp/**

[8] **http://www.itu.int/wsis/index.html**

[9] **http://gac.icann.org/web/index.shtml**

[10] Daniel Karrenberg, Gerard Ross, Paul Wilson, and Leslie Nobile, "Development of the Regional Internet Registry System," The Internet Protocol Journal, Volume 4, No. 4, December 2001.

[11] **http://nro.org/**

[12] **http://www.root-servers.org/**

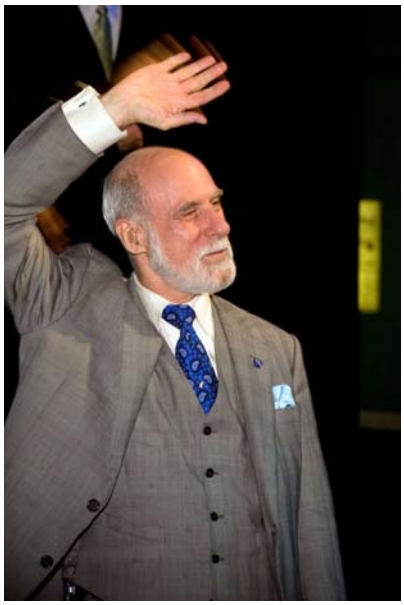

Photographer: Vanessa Stump

VINTON G. CERF is vice president and chief Internet evangelist for Google. In this role, he is responsible for identifying new enabling technologies to support the development of advanced Internet-based products and services from Google. He is also an active public face for Google in the Internet world. Cerf is the former senior vice president of Technology Strategy for MCI. In this role, he helped guide corporate strategy development from a technical perspective. Previously, he served as MCI's senior vice president of Architecture and Technology, leading a team of architects and engineers to design advanced networking frameworks, including Internet-based solutions for delivering a combination of data, information, voice, and video services for business and consumer use.

Widely known as one of the "Fathers of the Internet," Cerf is the co-designer of the TCP/IP protocols and the architecture of the Internet. In December 1997, President Clinton presented the U.S. National Medal of Technology to Cerf and his colleague, Robert E. Kahn, for founding and developing the Internet. Kahn and Cerf were named the recipients of the ACM Alan M. Turing Award, sometimes called the "Nobel Prize of Computer Science," in 2004 for their work on the Internet protocols. In November 2005, President George Bush awarded Cerf and Kahn the Presidential Medal of Freedom for their work. The medal is the highest civilian award given by the United States to its citizens.

Prior to rejoining MCI in 1994, Cerf was vice president of the Corporation for National Research Initiatives (CNRI). As vice president of MCI Digital Information Services from 1982 to 1986, he led the engineering of MCI Mail, the first commercial e-mail service to be connected to the Internet.

During his tenure from 1976 to 1982 with the U.S. Department of Defense Advanced Research Projects Agency (DARPA), Cerf played a key role leading the development of Internet and Internet-related packet-data and security technologies.

Vint was seated on the ICANN Board of Directors at the 1999 annual meeting, having been selected by the Protocol Supporting Organization. He was then selected by the nominating committee for a term on the board of directors that ran from June 2003 through the 2004 annual meeting. At the end of that term, he was selected by the 2004 nominating committee to an additional term, which ran from the end of the 2004 annual meeting through the conclusion of the ICANN annual meeting in 2007. He served as founding president of the Internet Society from 1992 to 1995, and in 1999 served a term as chairman of the board. In addition, Cerf is honorary chairman of the IPv6 Forum, dedicated to raising awareness and speeding introduction of the new Internet Protocol. Cerf served as a member of the U.S. Presidential Information Technology Advisory Committee (PITAC) from 1997 to 2001 and serves on several national, state, and industry committees focused on cyber security. Cerf sits on the board of directors for the Endowment for Excellence in Education, Avanex Corporation, and the ClearSight Systems Corporation. Cerf is a Fellow of the IEEE, ACM, and American Association for the Advancement of Science, the American Academy of Arts and Sciences, the International Engineering Consortium, the Computer History Museum, and the National Academy of Engineering.

Cerf is a recipient of numerous awards and commendations in connection with his work on the Internet, including the Marconi Fellowship, Charles Stark Draper Award of the National Academy of Engineering, the Prince of Asturias Award for science and technology, the National Medal of Science from Tunisia, the Alexander Graham Bell Award presented by the Alexander Graham Bell Association for the Deaf, the NEC Computer and Communications Prize, the Silver Medal of the International Telecommunications Union, the IEEE Alexander Graham Bell Medal, the IEEE Koji Kobayashi Award, the ACM Software and Systems Award, the ACM SIGCOMM Award, the Computer and Communications Industries Association Industry Legend Award, installation in the Inventors Hall of Fame, the Yuri Rubinsky Web Award, the Kilby Award, the Yankee Group/Interop/Network World Lifetime Achievement Award, the George R. Stibitz Award, the Werner Wolter Award, the Andrew Saks Engineering Award, the IEEE Third Millennium Medal, the Computerworld/Smithsonian Leadership Award, the J.D. Edwards Leadership Award for Collaboration, the World Institute on Disability Annual Award, and the Library of Congress Bicentennial Living Legend medal. In December 1994, People magazine identified Cerf as one of that year's "25 Most Intriguing People."

In addition to his work on behalf of MCI and the Internet, Cerf has served as a technical advisor to production for the "Gene Roddenberry's Earth: Final Conflict" television series and made a special guest appearance on the program in May 1998. Cerf has appeared on television programs NextWave with Leonard Nimoy and on World Business Review with Alexander Haig and Caspar Weinberger. He is also a distinguished visiting scientist at the Jet Propulsion Laboratory, where he is working on the design of an interplanetary Internet.

Cerf holds a Bachelor of Science degree in Mathematics from Stanford University and Master of Science and Ph.D. degrees in Computer Science from UCLA. He also holds honorary doctorate degrees from the Swiss Federal Institute of Technology (ETH), Zurich; Luleå University of Technology, Sweden; University of the Balearic Islands, Palma; Capitol College, Maryland; Gettysburg College, Pennsylvania; George Mason University, Virginia; Rovira i Virgili University, Tarragona, Spain; Rensselaer Polytechnic Institute, Troy, New York; the University of Twente, Enschede, The Netherlands; Brooklyn Polytechnic; and the Beijing University of Posts and Telecommunications.

Cerf's personal interests include fine wine, gourmet cooking, and science fiction. Cerf and his wife Sigrid were married in 1966 and have two sons, David and Bennett.

E-mail: **vint@google.com**

Remembering Itojun: The IPv6 Samurai

by Bob Hinden, Nokia

"Itojun" (Dr. Junichiro Hagino) passed away on October 29, 2007. He was 37 years old. Memorial events were held in Tokyo in November and in Vancouver at the IETF meeting in December.

Itojun was an active participant in the IETF and a member of the IAB from 2003 to 2005. He worked as a Senior Researcher at the *Internet Initiative Japan* (IIJ) and was a member of the board of the *Widely Integrated Distributed Environment* (WIDE) project. He was a strong supporter of open standards development and open software, working as a core researcher at the KAME project, a joint effort of six companies in Japan to provide a free stack of IPv6, IPsec, and Mobile IPv6 for BSD variants, from 1998 to 2006.

Itojun was totally dedicated to the development and deployment of IPv6. Most of his work was centered around building a much larger worldwide Internet based on IPv6. He was simply the "IPv6 Samurai."

Photographer: Diane Bruce

Quotes from Internet Colleagues

Steve Deering: "Those of us who got to know Itojun through his work in the Internet Engineering Task Force have lost a dear friend and much-admired colleague. From the day he arrived at his first IETF meeting, he won the respect of all in the way most honored by Internet engineers: by helping to build consensus based on running code. Moreover, he provided the best possible example of collaboration, generosity, and leadership, making not only extraordinary technological contributions but also many friends and a better world. His untimely passing is a huge loss to all who knew him, and to all those who will never have that chance."

Randy Bush: "An open heart, a big soul, and very kind and patient. A very special person. He wrote a lot of great code and got great joy from doing so."

Marc Blanchet: "Itojun adopted the Samurai's philosophy in his life: *Bushido,* which consists of values such as Honesty, Justice, Courtesy, Heroic Courage, Honor, Compassion, Sincerity, Duty, and Loyalty. Very difficult to achieve, he encompassed all these. Moreover, he was always available to help, anyone, without judging. His intelligence, his competency, and his dedication has inspired a generation of network engineers for the project he took as a mission: IPv6. Many computers in the world now run his code. My family always enjoyed meeting Itojun. He was always interested in sharing his knowledge with my children, even with the French-to-Japanese-through-English language barrier. Itojun, it was an honor to know you and to meet you. You will always be a source of inspiration to me, to my family, and to many network engineers in the world. We miss you."

Rod Van Meter: "I didn't know Itojun very well; I met him for the first time about five years ago at an IPv6 meeting in the Silicon Valley, once or twice in between, and then spent three days at the WIDE Camp this past September co-supervising (with Bill Manning, Brad Huffaker, and Kenji Saito) a group of students trying to establish long-term goals for WIDE in the area of naming. Itojun was gentle but insistent with students, a good mentor. That was the last time I saw him. Go in peace, Itojun."

Joel Jaeggli: "He cared more for the people who were going to use the code and the product of his and our labor than anyone would have had a right to expect. The Itojun that I know, our friend, has been taken from us, but we'll be the beneficiary of the fact that he cared, for decades."

Itojun IPv6 Fund

Itojun's family has expressed sincere appreciation to all who attended the memorial and funeral services. His family has set up a memorial fund in Itojun's name under the directorship of the IETF/Internet Society. The fund will be used to award an R&D grant to a person who has contributed to the deployment and further advancement of IPv6. ISOC has set up an e-mail address to accept commitments for the *Itojun IPv6 Fund*. The address is: **itojun-fund@isoc.org**

The procedure for making contributions is being developed; if you wish to contribute now, please send a note to the e-mail address describing the amount you want to contribute (and in what currency), and ISOC will collect the funds.

ROBERT HINDEN is a Nokia Fellow at Nokia and is located in Mountain View, California, USA. He has been involved in the Internet since it was a research project at ARPA. He developed one of the first TCP/IP implementations and his team at Bolt, Beranek, and Newman, Inc. built and operated the routers that formed the early Internet backbone. He was co-recipient of the 2008 IEEE Internet Award "For pioneering work in the development of the first Internet routers." He has been active in the IETF since 1985 and is the author of 35 RFCs. He was recently appointed to a position on the IETF Administrative Oversight Committee (IAOC) and co-chairs the IPv6 Maintenance (6man) working group. Prior to this he served on the Internet Architecture Board (IAB), was Area Director for Routing in the Internet Engineering Steering group from 1987 to 1994, and chaired the IPv6, Virtual Router Redundancy Protocol, Simple Internet Protocol Plus, IP over ATM, and Open Routing working groups. Hinden is also a member of the RFC Editorial Board. He holds a B.S.E.E. and a M.S. in Computer Science from Union College, Schenectady, New York. E-mail: **bob.hinden@nokia.com**

Book Review

Network Routing

Network Routing: Algorithms, Protocols, and Architectures, by Deepankar Medhi and Karthikeyan Ramasamy, Morgan Kaufmann Publishers, ISBN-13:978-0120885886, 2007,
`http://www.NetworkRouting.net`

Routing is a fundamental architectural component of any network, and in this book the authors examine in detail the routing technologies of the Internet and the *Public Switched Telephone Network* (PSTN).

Organization

The book is divided into five parts, with an additional advanced section provided on CDROM. The first part examines the fundamentals of routing technology, looking in detail at the basic approaches of distance-vector and link-state routing. The second part looks at the routing protocols used in the Internet today, as well as Traffic Engineering. The third part addresses routing in the PSTN, examining the SS7 signaling protocol and the overall architecture of the PSTN. The next part explores the internal architecture of routers, address-lookup algorithms, and packet-classification techniques. Finally, the authors consider topics encompassed in the so-called "Next-Generation Network," including *Quality-of-Service Routing*, *Multiprotocol Label Switching* (MPLS), and *Voice over IP* (VoIP). The advanced-topic section includes a more detailed examination of packet-switching approaches, scheduling, and conditioning. This book is positioned as a graduate-level text, and each chapter is accompanied by exercises that review the material.

The book covers a broad range of material: each topic has been the subject of entire books. The level of detail in the book varies considerably. In some instances, such as in the area of IP Traffic Engineering, it presents a highly detailed mathematical analysis of aspects of the topic, whereas in other instances, such as in the treatment of the *Border Gateway Protocol* (BGP), the material appears to be obviously condensed. I was expecting a little more use of algorithms to illustrate routing concepts, and found at times the mathematical analysis to be unhelpful in terms of understanding the underlying problem space being described.

Comparison

In this area of Internet routing, any publication is inevitably compared to Radia Perlmann's book *Interconnections: Bridges, Routers, Switches and Internetworking Protocols,* and this book is no exception. To my mind it falls a little short of this rather demanding standard. Radia spends some time discussing the underlying rationale as to why a particular technology was devised for a given problem space, and also discusses the strengths and limitations of the technology in various areas of application.

In *Network Routing* the authors limit their approach to a description of the technology by looking at packet payloads and protocol interactions and numerous deployment scenarios that illustrate the features of this particular routing technology. The consideration of choices made in the development of the protocol, and the consequent implications of such design choices, are missing in such a treatment, and the reader is often left wondering why a routing protocol has chosen to support certain functions but not others.

I found this to be a very ambitious book, because it appears to position itself both as a reference publication on routing technologies and architecture and also on the description of routing protocols, while at the same time wanting to encompass the role of a course text. This goal could have been attainable if the book had chosen a tighter focus, but the all-encompassing approach that led to the inclusion of considerations of the PSTN topics makes the outcome less than fully satisfying.

Recommended

However, the book manages to bring together the basic topics in routing in both the Internet and the PSTN, and it not only includes a good description of the routing technologies in use today, but also looks at some of the advanced topics in routing today. I found the major strength of the book in its role as a graduate-course text, where there is sufficient description of the topic to lead into further reading of current research papers and more-detailed technical material. Although the book has some shortcomings, I'd certainly recommend it as a suitable addition to the shelves of any professional in the area of Internet routing technologies and architecture.

—Geoff Huston
gih@apnic.net

The Author Responds:

I thank Geoff Huston for writing a well-thought-out review; in general, this review is fair. This book was certainly an ambitious project. I wanted to do it as I've investigated various routing protocols for almost two decades—and many people I talked to thought that it would be useful to have such a book. In fact, Dave Clark, when he read the original book proposal, wrote "It is ambitious—there may be issues of how much depth they can get on all these topics in one book," but felt that "...the approach is distinctive and very valuable. So I support the idea." As can be observed from the book, the depth on different topics remained a major trade-off we pondered without making the book go over 1000 pages (with 140 pages on CD-ROM it came pretty close).

There were a few "design" decisions I deliberately made in organizing and writing this book. One of them was based on years of teaching and interacting with industry folks: I decided to divide materials broadly on "how and why" away from "what;" this approach is somewhat surprising, but people's learning style seems to fall into these two categories (certainly there are overlaps). Therefore, details on "how and why" of different protocols went into Chapter 3 (and for algorithms into Chapter 2), while details on "what" went with chapters on specific protocols such as OSPF or BGP. Similarly, I also separated out the topic of "how" routing in the global Internet works and is organized (such as public exchange points) from the chapter on BGP. Secondly, we separated math parts from non-math parts—this way, those who are interested, for example, in detailed Traffic Engineering modeling can read the relevant chapters. Others may skip them and read just the first couple of overview sections; it should be noted that math-oriented chapters are generally organized from simple concepts to difficult concepts. Thirdly, we covered address lookup, packet filtering and classification, and router architectures separately because they can be read independently; Karthik brought his wealth of experience in writing these chapters.

I want to take this opportunity to respond to a few of Geoff's comments:

1. "...expecting a little more use of algorithms to illustrate routing concepts." I suspect that Geoff didn't think that Chapters 2 and 3 covered enough, although these chapters included details illustrative of distance-vector protocols, link-state protocols, path-vector protocols (and their pitfalls), and so on. As stated previously, by design of the book, illustrative examples of routing concepts were separated from specific protocols so that readers can read different portions of the material according to their interests. As an indicator to the reader, each chapter starts with a brief "reading guideline" (which is a unique feature of this book) that states how the material is organized and its relation to other chapters or sections in the rest of the book.
2. "The consideration of choices made in the development of the protocol, and the consequent implications of such design choices are missing in such a treatment." We did indeed cover these aspects in many instances. For example, the book covers why, for I-BGP scalability, the route reflector or the confederation approach are needed; why route flap damping was developed; why ROUTE-REFRESH was added; what MPLS was trying to solve that IP-only couldn't do at that time; the need for age with Sequence Number field in link-state protocols; what led to the development of dynamic call routing from hierarchical routing in the PSTN, and so on.

That said, I did not include certain discussions because some choices on protocols have been based on personality clashes and "camps;" I felt that this is not easily explainable in many instances—trying to do so would require quite a bit of discussion, and could potentially divert from the main focus of the book. For example, I explained why the route reflector or confederation approach was needed for I-BGP scalability, but I didn't discuss why both route reflector and confederation approaches were developed simultaneously when both convey the same idea conceptually.

3. "... the all-encompassing approach that led to the inclusion of considerations of the PSTN topics into this book make the outcome less than fully satisfying." I included routing in the PSTN because of its historic context, and particularly to make readers aware of the evolution from hierarchical routing to dynamic routing and recent changes in routing due to Local Number Portability—these lessons are important ones to learn for anyone interested in routing or designing future routing protocols. Secondly, many concepts in MPLS/GMPLS have parallels in the PSTN, thus certain aspects in MPLS/GMPLS are easily explainable if a reader is familiar with PSTN details. We therefore felt it was appropriate to include all this material in one place. Furthermore, control- and data-path separation in GMPLS is strikingly similar to separation of signaling in PSTN through SS7 from actual voice communication. Thus, lessons learned from failure propagation from SS7 to voice paths are relevant lessons to be aware of for anyone involved in deploying GMPLS-based networks. Lastly, to discuss VoIP routing, it is critical to tie into PSTN because in the real operational environment PSTN-Internet interworking for VoIP routing is expected to remain prevalent for years to come.

Finally, the "barrier to entry" in learning about routing is very high, especially for entry-level professionals—I've attempted to position the book as both a text and a reference for professionals. Thus, I very much appreciated Geoff's concluding comment "... as a suitable addition to the shelves of any professional in the area of Internet routing technologies and architecture."

—Deep Medhi
dmedhi@umkc.edu

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

Nii Quaynor Receives 2007 Postel Service Award

The *Internet Society* (ISOC) has awarded pioneering Internet engineer Nii Quaynor the prestigious *Jonathan B. Postel Service Award* for 2007 for his leadership in advancing Internet technology in Africa and galvanizing technologists to improve Internet access and capabilities throughout the continent. ISOC presented the award, including a $20,000 [USD] honorarium, during the 70th meeting of the *Internet Engineering Task Force* (IETF) in Vancouver, BC, Canada.

"Dr. Quaynor has selflessly pioneered Internet development and expansion throughout Africa for nearly two decades, enabling profound advances in information access, education, healthcare and commerce for African countries and their citizens," said ISOC president Lynn St. Amour. "Today, Dr. Quaynor continues to champion not just technological advances but also African involvement in Internet standards, processes and deployments, discussion on Internet policies and regulations, and ensuring African interests are well-represented globally. He has shaped a community of Africans who share his vision and reflect the dedication shown by Jon Postel."

"I am humbled by the award and what Jon Postel represents to our community in Africa. Jon Postel's efforts and the global view he maintained on the operation of the *Domain Name System* and the numbering services assured that Africa would share in the Internet growth and early. I thank the Internet Society for the recognition and am very pleased to be associated with Jon's memorial," said Dr. Nii Quaynor. "We will work to develop more African engineers to meet the fast network growth needs of the region, being a late starter, and to join the technical policy processes. Our overall objective is to strengthen education and research in network technologies in Africa."

The annual ISOC award is named after Dr. Jonathan B. Postel to commemorate his extraordinary stewardship exercised throughout his thirty-year career in networking. Between 1971 and 1998, Postel managed, nurtured and transformed the RFC series of notes, which encompasses the technical specifications and recommendations for the Internet and was created by Steve Crocker in 1969 as a part of his work on the ARPANET, the forerunner of today's Internet. Postel was a founding member of the Internet Architecture Board and the first individual member of the Internet Society, where he also served as a trustee until his untimely death.

Dr. Quaynor is chairman of *Network Computer Systems* (NCS) Ghana.COM and a professor of computer science at University of Cape-Coast, Ghana. He is also the convener of the *African Network Operators Group* (AfNOG), a network technology transfer institution since 2000 and the founding chairman of AfriNIC, the African numbers registry.

Dr. Quaynor began his pioneering Internet work in Africa in 1993 when he returned to his home country of Ghana to establish the first Internet Service operated by NCS in West Africa. At NCS, he and his team worked on the early development of the Internet in Africa. Today, there are more than 43 million Internet users in Africa.

Prior to NCS, Dr. Quaynor worked with Digital Equipment Corporation in the United States from 1977 till 1992. In 1979, he established the Computer Science department at the University of Cape Coast, Ghana. Dr. Quaynor graduated from Dartmouth College in 1972 with B.A (Engineering Science) and received a Ph.D. (Computer Science) in distributed systems in 1977 from State University of New York at Stony Brook.

The Jonathan B. Postel Service Award was established by the Internet Society to honor those who, like Postel, have made outstanding contributions in service to the data communications community. The award is focused on sustained and substantial technical contributions, service to the community, and leadership. With respect to leadership, the nominating committee places particular emphasis on candidates who have supported and enabled others in addition to their own specific actions.

Previous recipients of the Postel Award include Jon himself (posthumously and accepted by his mother), Scott Bradner, Daniel Karrenberg, Stephen Wolff, Peter Kirstein, Phill Gross, Jun Murai, Bob Braden, and Joyce K. Reynolds. The award consists of an engraved crystal globe and $20,000 [USD]. This year's award is sponsored in part by Afilias Global Registry Services. For more information about ISOC, please visit: **`www.isoc.org`**

Steps Taken for Multilingual Internet

The *Internet Corporation for Assigned Names and Numbers* (ICANN), the *International Telecommunication Union* (ITU), and the *United Nations Educational, Scientific and Cultural Organization* (UNESCO) will collaborate on global efforts to forge universal standards towards building a multilingual cyberspace. The three agencies organized a workshop on this subject during the second *Internet Governance Forum* (IGF) which took place in Rio de Janeiro, Brazil from 12 to 15 November 2007.

The Internet is a key factor in developing a more inclusive and development-oriented information society, which stresses plurality and diversity instead of global uniformity. Multilingualism is a key concept to ensure cultural diversity and participation for all linguistic groups in cyberspace. There is growing concern that hundreds of local languages may be sidestepped, albeit unintentionally in the radical expansion of Internet communication and information. The *World Summit on the Information Society* (WSIS) recognized the importance attached to linguistic diversity and local content, with UNESCO given the responsibility to coordinate implementation of the *Summit Action Line*.

"The discussions at this multilingualism workshop—combined with our current evaluation of *Internationalized Domain Names* (IDNs)—are going to help ICANN keep moving toward full implementation of Internationalized Domain Names," said Dr Paul Twomey, ICANN's President and CEO. "ICANN is in the midst of the largest ever evaluation of IDNs at the top level."

Thanks to ICANN's evaluation of Internationalized Domain Names, Internet users around the globe can now access wiki pages (see **`http://idn.icann.org/`**) with the domain name **`example.test`** in the 11 test languages—Arabic, Persian, Chinese (simplified and traditional), Russian, Hindi, Greek, Korean, Yiddish, Japanese and Tamil. The wikis will allow Internet users to establish their own sub pages with their own names in their own language; one suggestion is: **`example.test/yourname`**

Domain Names, which are currently mainly limited to characters from the Latin or Roman scripts, are seen as an important element in enabling the multilingualization of the Internet, reflecting the diverse and growing language needs of all users. "ITU is fully committed to assist its membership in promoting the diversity of language scripts for domain names," said Dr Hamadoun Touré, Secretary-General of ITU. "This workshop represents an important opportunity to strengthen the need for cooperation with relevant organizations, such as UNESCO, the *World Intellectual Property Organization* (WIPO) and ICANN among others to ensure Internet use and advancement across language barriers."

The Plenipotentiary Conference of ITU, which took place in Antalya, Turkey in November 2006, recognized the need to make Internet content available in non-Latin based scripts. Internet users are more comfortable reading or browsing through texts in their own language and a multilingual Internet is essential to make it more widely accessible. The WSIS outcomes also focused on the commitment to work towards multilingualization of the Internet as part of a multilateral, transparent and democratic process involving governments and all stakeholders.

UNESCO, joined by both ITU and ICANN, seeks to convene all major stakeholders around the world towards an agreement on universal standards regarding language issues in cyberspace. Such issues are far broader than the single issue of IDNs as they extend to standards for fonts and character sets, text encoding, language implementations within major computer operating systems, content development tools, automatic translation software, and search engines across languages. Ultimately, equitable access to information can be only achieved if we resolve language barriers at the same time we build communications infrastructures and capacity building programs.

RIPE Community Resolution on IPv4 Depletion and Deployment of IPv6

During the RIPE 55 meeting in Amsterdam in October 2007, the RIPE community agreed to issue the following statement on IPv4 depletion and the deployment of IPv6:

"Growth and innovation on the Internet depends on the continued availability of IP address space. The remaining pool of unallocated IPv4 address space is likely to be fully allocated within two to four years. IPv6 provides the necessary address space for future growth. We therefore need to facilitate the wider deployment of IPv6 addresses.

While the existing IPv4 Internet will continue to function as it currently does, the deployment of IPv6 is necessary for the development of future IP networks.

The RIPE community has well-established, open and widely supported mechanisms for Internet resource management. The RIPE community is confident that its *Policy Development Process* meets and will continue to meet the needs of all Internet stakeholders through the period of IPv4 exhaustion and IPv6 deployment.

We recommend that service providers make their services available over IPv6. We urge those who will need significant new address resources to deploy IPv6. We encourage governments to play their part in the deployment of IPv6 and in particular to ensure that all citizens will be able to participate in the future information society. We urge that the widespread deployment of IPv6 be made a high priority by all stakeholders."

For more information, see: **`http://ripe.net/ripe/`**

Upcoming Events

The next *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT) will be held in Taipei, Taiwan from February 20th to 29th, 2008. As usual, this conference is co-located with an APNIC Open Policy Meeting. For more information about these events see: **`http://www.apricot2008.net/`** and **`http://www.apnic.net/meetings/25/index.html`**

The *Internet Engineering Task Force* (IETF) will meet in Philadelphia, Pennsylvania, March 9–14 and "somewhere in Europe" July 27–August 1. (The announcement of the exact location is expected soon). The final IETF meeting in 2008 will take place in Minneapolis, Minnesota, November 16–21. For more information see: **`http://www.ietf.org/meetings/0mtg-sites.txt`**

The *Internet Corporation for Assigned Names and Numbers* (ICANN) will meet in New Delhi, India, February 10–15, and in Paris, France, June 22–27. See: **`http://icann.org/meetings/`**

Call for Papers

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

Topics include, but are not limited to:

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

In addition to feature-length articles, IPJ 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