

The Internet Protocol *Journal*

September 2007

Volume 10, Number 3

*A Quarterly Technical Publication for
Internet and Intranet Professionals*

FROM THE EDITOR

In This Issue

From the Editor	1
Secure Multivendor Networks.....	2
IPv4 Address Depletion	18
IPv4 Address Consumption ..	22
Awkward /8 Assignments	29
Book Review	32
Call for Papers.....	35

For the last 10 or so years I have been involved with the organization of APRICOT, the *Asia Pacific Regional Internet Conference on Operational Technologies*. APRICOT has at its core a set of workshops featuring expert instructors with years of operational network experience. A recent addition to the APRICOT workshop program is a course focusing on Internet security in a multivendor environment. Our first article, written by Kunjal Trivedi from Cisco Systems, Inc., and Damien Holloway from Juniper Networks, is based on this workshop. It's not every day that you see an article co-authored by instructors from competing companies, but this is exactly the type of cooperation that is needed in order to deploy security in a multivendor network.

The rest of this issue is mostly devoted to IPv4 depletion and the transition to IPv6. The first article, by Geoff Huston, summarizes many of the concerns related to IPv4 depletion and IPv6 transition, and gives numerous pointers to further articles and documents of interest. Our second addressing-related article, by Iljitsch van Beijnum, looks more closely at the numbers relating to address allocation by the *Regional Internet Registries* (RIRs). The final article concerns some address blocks that are currently unassigned but actually in use. Leo Vegoda explains the potential problems that may arise when these blocks eventually become part of the RIR assignment pool.

We are pleased to announce a new online addition to this journal. *The Internet Protocol Forum* (IPF) available at www.ipjforum.org is designed to allow discussion of any article published in the printed edition of IPJ. In addition to article discussions, the forum will be used to provide updates and corrections, downloads, expanded versions of some articles, configuration and programming examples, and news and analysis that does not fall into our quarterly publication schedule. The IPF's editor and moderator is Geoff Huston, long-time contributor to this journal and chief scientist at APNIC. I am confident that IPF will become an important addition to IPJ, and I hope you will take the time to participate in the online discussions. Of course, you can always contact us at the usual e-mail address: ipj@cisco.com

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

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

A Standards-Based Approach for Offering a Managed Security Service in a Multivendor Network Environment

By Kunjal Trivedi, Cisco Systems and Damien Holloway, Juniper Networks

As transport becomes a commodity, service providers are seeking new revenue sources and new ways to differentiate themselves. Managed security services address a growing market because business customers are struggling to comply with regulatory requirements such as the *Payment Card Industry-Data Storage Standards* (PCI-DSS), the *Sarbanes-Oxley Act*, the *Gramm-Leach Bliley Act*, *Health Insurance Portability and Accountability Act* (HIPAA), *Directive 2002/58/EC*, and the *Asia-Pacific Economic Cooperation-Organization for Economic Cooperation and Development* (APEC-OECD) initiative on regulatory reform. Increasingly, business customers recognize that outsourcing network security is less costly than staffing with highly specialized security personnel who can provide 24-hour incident detection and response. Another incentive for outsourcing is to free existing IT resources to focus on the core business.

A standards-based approach helps service providers take best advantage of the managed security service opportunity because it increases the potential breadth and depth of the service offering. Multivendor solutions are becoming the norm when deploying services on an integrated backbone. Therefore, standards simplify deployment and management, helping control operational costs and accelerating time to market.

Service providers are experiencing a growing need for skilled engineers who understand multivendor environments—the motivation for conducting a multivendor security workshop at the 2006 *Asia Pacific Regional Internet Conference on Operational Technologies* (APRICOT 2006)^[15], held in Perth, Australia, in February 2006. During the workshop (which was repeated again at APRICOT 2007 in Bali), participants successfully deployed and tested a multivendor service environment using *IP Security* (IPsec)-based Layer 3 *Virtual Private Networks* (VPNs)^[1, 2, 3] over a *Border Gateway Protocol/Multiprotocol Label Switching* (BGP/MPLS) core^[4].

Technical Challenges

To offer managed security services, service providers need the following:

- A secure network infrastructure, including tools and techniques for risk mitigation
- Technical solutions for the customer's business needs, such as VPNs based on BGP/MPLS, IPsec, or both

- Web-based reporting tools that business customers can use to monitor the security service in accordance with *Service-Level Agreements* (SLAs). Service providers can scale cost-effectively by offering customers a secure, Web-based portal that shows open trouble tickets, security incident-handling detail, SLAs, and access reports that customers need to comply with regulations.

An effective managed security service requires tools and techniques to address the following challenges:

- *More sophisticated threats, and less time between vulnerability and exploitation:* In addition to worms and viruses, the service provider needs to protect its own and its customers' networks against *Denial-of-Service* (DoS) attacks. Today's botnets can launch thousands or even a million bots that carry out outbound DoS attacks. New varieties of worms have side effects similar to those of DoS attacks. These threats can take down the service provider infrastructure, thereby violating SLAs and eroding revenue.
- *A need for proactive rather than reactive threat response:* Many service provider security groups are stuck in reactive mode. Every network device and security system produces voluminous event logs every day, and vendors use different formats. Therefore, identifying security incidents in order to react to them can take hours or days—or not happen at all. The connection between two separate events in different parts of the network can easily escape human detection, especially when the clues are buried among tens of thousands of harmless events that took place around the same time.
- *Multivendor networks:* Network security and reporting are easier to achieve in single-vendor networks. Realistically, however, many service providers and business customers have multivendor networks, sometimes because of mergers and acquisitions. Even if the service provider itself has a single-vendor network, some of its customers will use other vendors' equipment.
- *Slow progress toward adopting IP Next-Generation Networks (IP NGNs):* When service providers complete the migration to IP NGN, they will achieve greater control, visibility, and operational efficiency. Until then, service providers will incur higher costs and labor requirements for support and migration.
- *A need to comply with industry standards from IETF and ITU:* Standards facilitate security in multivendor networks. MPLS helps ensure infrastructure security, whereas IPsec provides secure connectivity among the customer's branches and remote offices. By using industry standards, the service provider can select best-of-class products based on performance, features, or cost.

- *Scalability challenges:* The security operations center for a managed services provider cannot cost-effectively scale to process several million events for each customer. However, it can scale to process a few security-incident trouble tickets. Scalability hinges on the ability to minimize false positives. Products such as Cisco *Security Monitoring, Analysis and Response System* (MARS), IBM Micromuse, and NetIQ provide analysis and correlation of events from multiple elements in the IT infrastructure. They process events using consolidation, filtering, normalization, enrichment, correlation, and analysis techniques, and also notify IT staff about critical events.

Infrastructure Security in Multivendor Environments

Securing the service provider infrastructure requires the following common best practices:

- Point protection
- Edge protection
- Remote-triggered black-hole protection
- Source-address validation on all customer traffic
- Control-plane protection
- Total visibility into network activity

Point Protection

Before offering a managed security service, providers need to protect the backbone; security operations center or network operations center; *Authentication, Authorization, and Accounting* (AAA)^[10, 11] server; and remote-access networks. Securing individual network devices requires enforcing AAA, controlling the type of packets destined to network devices, and performing regular configuration audits to ensure that no unauthorized changes have been made. Best common practices include:

- *Protect the backbone by locking down the vty and console ports:* This protection helps prevent unauthorized access to network devices.
- *Encrypt management commands that staff send to devices:* Use of the *Secure Shell* (SSH) protocol helps prevent hackers from obtaining passwords that they could later use to compromise the network. Service providers that use out-of-band management for device configuration should also encrypt this management traffic and restrict access to authorized personnel.
- *Deploy a AAA server:* Using a AAA server is preferable to relying on local authorization on the devices themselves because it enables centralized policy control. The AAA server controls a user's access to the device, or even the specific commands that the user is authorized to execute.

It is strongly recommended that service providers use TACACS+^[13] authentication rather than *Remote Authentication Dial-In User Service* (RADIUS)^[12] authentication. With RADIUS, traffic is sent in the clear between the AAA servers and network devices using the *User Datagram Protocol* (UDP), which defeats the use of SSH to encrypt logins and passwords. Open-source implementations of TACACS+ are available.

- *Use one-time passwords (OTPs)*: To distribute one-time passwords, service providers can provide authorized users with a token card, soft token, or soft key. One-time passwords ensure that the user was authorized at the time of login, and was not an attacker who used a packet-sniffer program to intercept a password.
- *Protect the AAA infrastructure from DoS attacks*: Some service providers set up local accounts on routers and switches so that staff can log in if the AAA infrastructure is down, creating vulnerability. If the service provider does not secure management-plane access to the device, hackers can use SSH or Telnet and attempt a brute-force attack to crack the local account. The local account is often not as secure as an OTP because it is changed only once every 30 days, providing a longer window of opportunity for hackers to gain device access. It is strongly advised to not use default or easy-to-guess passwords. To prevent attacks against the AAA infrastructure, service providers should harden the infrastructure and consider placing the server behind a firewall with stateful inspection. Use *Access Control Lists* (ACLs), which are packet filters, on the firewall to restrict traffic between the AAA server and network devices only. Also be sure to distribute the AAA servers so that they do not create a single point of failure.
- *Regularly audit device configurations*: Frequently, the first indications of an attack, often unnoticed, are unauthorized commands executed on routers that change the configuration. An easy way to monitor configurations is using RANCID (*Really Awesome New Cisco Config Differ*)^[14], a UNIX or Linux freeware tool that logs into each of the devices in the device table file, runs various show commands, processes the output, and sends e-mail messages reporting any differences from the previous collection to staff. RANCID works with routers from Cisco and other vendors. Another tool for auditing device configurations, the *Router Audit Toolkit* (RAT) assigns security scores to ACLs and other security best practices to show the relative security of routers.

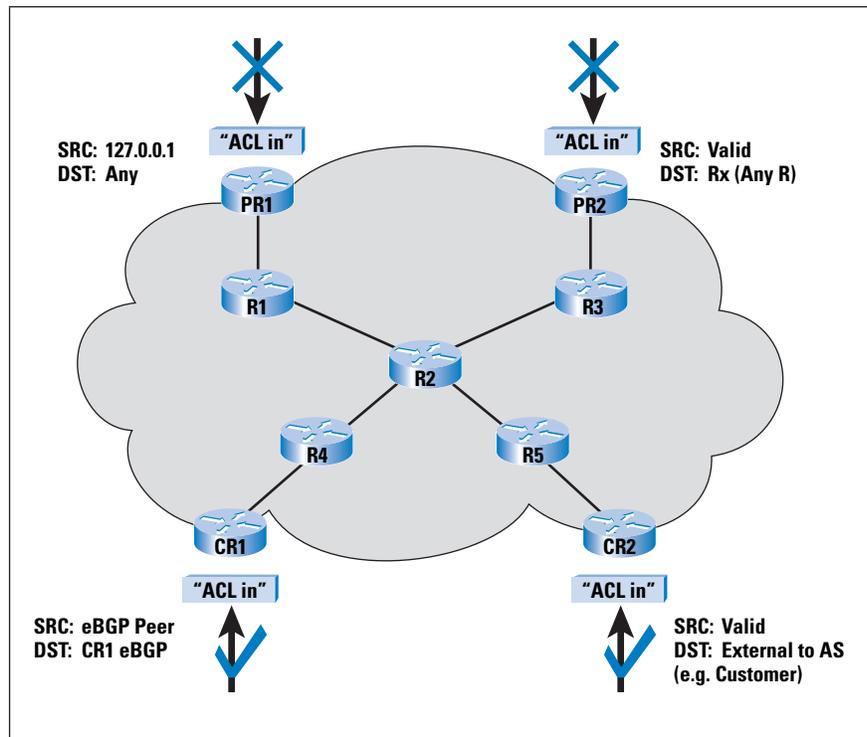
Traditionally, service providers enforced policy at the process level, using vty ACLs, *Simple Network Management Protocol* (SNMP) ACLs, and others. Some service providers used ingress ACLs when possible. Today, it is far preferable to stop DoS traffic at ingress points: the peer edge, downstream and upstream routers, colocated network devices, and the customer access edge, enabling central policy enforcement and more granular protection schemes.

In addition, many network devices at the network edge have hardware acceleration, which provides far more robust resistance to attack than the process level.

Edge Protection

In many service provider networks, each core router is individually secured but still accessible to outsiders using SNMP or Telnet. Now service providers can supplement individual router protection with infrastructure protection that prevents undesired traffic from ever touching the infrastructure.

Figure 1: Protecting the Network Edge



The following steps help protect the network edge (Figure 1):

1. Classify the required protocols that are sourced from outside the *Autonomous System* (AS) access core routers, such as *external BGP* (eBGP) peering, *Generic Routing Encapsulation* (GRE)^[5], and IPsec. (Examples of nonrequired protocols are SNMP and Telnet.) Classification can be performed using a classification packet filter or Cisco *NetFlow* telemetry. The classification packet filter comprises a series of permit statements that provide insight into required protocols. Gradually narrow down the list, keeping in mind that very few protocols need access to infrastructure equipment, and even fewer are sourced from outside the autonomous system. Summarize the IP address space as much as possible, for simpler and shorter ACLs. Be cautious: just because certain types of traffic appear in a classification packet filter or NetFlow telemetry data does not mean they should be permitted to pass through to the routers.

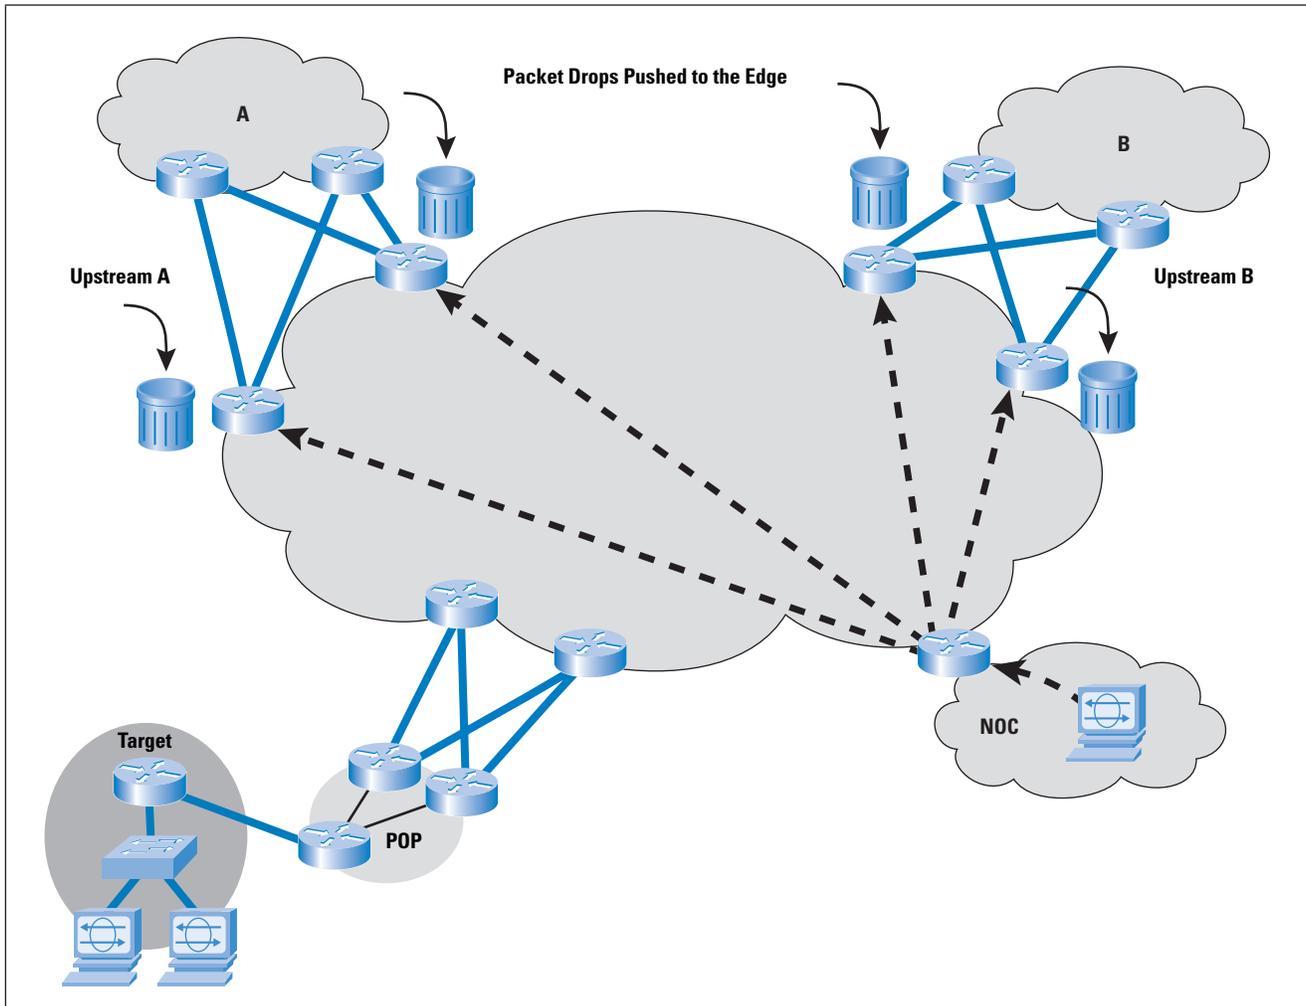
2. Begin filtering. Use an infrastructure packet filter to permit only the required protocols to access infrastructure-only address blocks, denying all other protocols. It is important to monitor the packet filter entry counters, because a high volume of hits, whether or not a protocol has been identified as required, might signal an attack. To permit transit traffic, use the following as the final line of the *Infrastructure ACL* (iACL): **permit ip any any**, protecting the core network with a basic iACL that admits only the required protocols. Note that iACLs also provide antispoof filtering by denying access to the space from external sources, denying the RFC 1918 space^[6], and denying multicast source addresses. RFC 3330^[7] defines special-use IPv4 addressing.
3. Further protect the core by identifying legitimate source addresses for the required protocols, such as external BGP peers and tunnel endpoints.
4. Deploy destination filters when possible.

Infrastructure packet filters at the edge of the network protecting the infrastructure are an effective first layer of defense. Service providers need additional forms of infrastructure protection for their older routers that do not support infrastructure packet filters and for packets that cannot be filtered with infrastructure packet filters.

Remote Triggered Black Hole Filtering

Remote Triggered Black Hole Filtering (RTBH) is among the most effective reaction and mitigation tools for DoS, *Distributed DoS* (DDoS), and backscatter tracebacks. It enables service providers to quickly drop DoS traffic at the network edge (Figure 2). Rather than sending commands to every router to drop DoS or other problem traffic, the service provider can deploy a trigger router that uses BGP to signal all other routers—just as fast as iBGP can update the network. In destination-based RTBH, all traffic headed to the destination under attack is dropped—the good traffic as well as the bad. In source-based RTBH, traffic from all or certain sources are blocked. The advantage of sourced-based RTBH is that service providers can whitelist certain addresses, such as the *Network Operations Center* (NOC) or route-name servers, so that they can continue providing services.

Figure 2: DoS Packets Dropped at the Network Edge



Source Address Validation on all Customer Traffic

Source address validation, defined in *Best Current Practices* (BCP) 38^[8], prevents service provider customers from spoofing traffic—that is, sending IP packets out to the Internet with a source address other than the address allocated to them by the service provider. Best practices from BCP 38 are to filter as close to the edge as possible, filter precisely, and filter both the source and destination address when possible.

Every access technology has antispoofing mechanisms derived from BCP 38:

- Packet filters
- Dynamic packet filters that are provisioned to be AAA profiles; when a customer signs in with RADIUS, a packet filter is set up for the customer
- *Unicast Reverse Path Forwarding* (URPF)
- Cable-Source Verify and packet cable multimedia (cable)
- IP Source Verify and DHCP Snooping (Metro Ethernet)

To gain operational confidence in BCP 38, service providers can take a phased approach—for example, implementing it first on one port, then on a line card, then on an entire router, and then on multiple routers.

Control-Plane Protection

Protecting the infrastructure control plane helps prevent an attacker from taking down a BGP session and thereby causing denial of service. The exploits a service provider needs to prevent include saturating the receive-path queues so that BGP times out, saturating the link so that the link protocols time out, dropping the *Transmission Control Protocol* (TCP) session, and dropping the *Interior Gateway Protocol* (IGP), which causes a recursive loop-up failure.

Following are techniques for control-plane protection.

- *Generalized Time-to-Live (TTL) Security Mechanism (GTSM)*: This technique protects BGP peers from multihop attacks. Routers are configured to transmit their packets with a TTL of 255, and to reject all packets with a TTL lower than 254 or 253. Therefore, a device that is not connected between the routers cannot generate packets that either router will accept.
- *Configuring routing authentication*: The *Message Digest Algorithm 5* (MD5) peer authentication feature instructs the router to certify the authenticity of its neighbors and the integrity of route updates. MD5 peer authentication can also prevent malformed packets from tearing down a peering session, and unauthorized devices from transmitting routing information. Be aware that MD5 peer authentication does not protect the router if an attacker compromises the router and begins generating bogus routing updates. Although it is not a panacea, MD5 peer authentication does raise the level of protection.
- *Customer ingress prefix filtering*: Prefix hijacking is an exploit in which a service provider customer announces an address space that belongs to another customer. The remedy is customer ingress prefix filtering, which enables service providers to accept only those customer prefixes that have been assigned or allocated to their downstream customers. For example, if a downstream customer has a **220.50.0.0/20** block, customers can announce this block only to their peers, and upstream peers accept this prefix only. Service providers can apply ingress prefix filtering to and from customers, peers, and upstream routers.

Visibility into Network Activity

To gain visibility into the network for early detection of security incidents, service providers can use open-source tools to analyze flow-based telemetry data, which is retrieved from routers and switches. Open-source tools for visibility into security incidents include RRDTool, FlowScan, Stager, and NTOP *Remote Monitoring* (RMON).

These tools provide information such as packets per second, bits per second, and traffic types. For example, RRDTool shows the number of *Domain Name System* (DNS) queries per second, according to record type. A spike in *Mail Exchange* (MX) Record queries might indicate that a customer's router has been compromised and is being used as a spam proxy. Similarly, a sharp increase in round-trip-time latency might indicate a DoS attack.

MPLS Security in a Multivendor Environment

In addition to securing the infrastructure, managed security service providers need to secure packets as they travel from one customer-edge router to another—regardless of the equipment the customer uses at the edge. Layer 3 VPNs meet this need. RFC 4364, which replaced RFC 2547bis, defines a BGP/MPLS IP VPN that creates multiple virtual routers on a single physical router: one virtual router for each customer.

In BGP/MPLS VPNs, *Customer Edge* (CE) routers send their routes to the *Service Provider Edge* (PE) routers. Customer edge routers at different sites do not peer with each other, and the customer's routing algorithms are not aware of the overlay. Data packets are tunneled through the backbone so that the core routers do not need to know the VPN routes. BGP/MPLS IP VPNs support either full mesh or partial mesh, although full mesh is more cost-effective.

A unique advantage of BGP/MPLS VPNs is that two service provider customers with overlapping IP addresses can connect across the service provider backbone. The router distinguishes between traffic from different companies by examining the label at the beginning of the packet, and then instantly forwards the traffic based on the *Label Switching Path* (LSP) that has been established for each customer's VPN. Eliminating the need to look at the packet in depth enables faster forwarding. That is, the service provider core does not impose any latency as packets pass between the provider edge routers.

IPsec Security in a Multivendor Environment

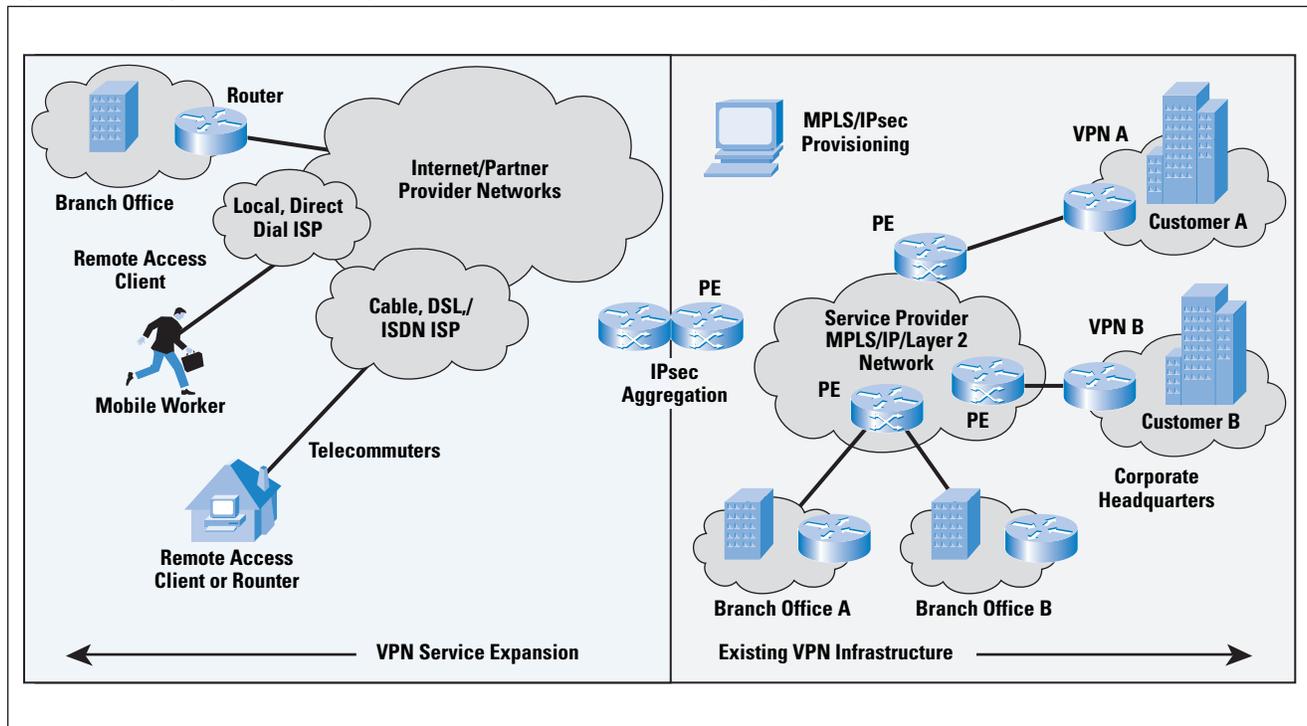
In addition to or instead of deploying a BGP/MPLS IP VPN, the service provider can extend its service to other partner provider networks using IPsec. The options are to use MPLS alone, IPsec alone, or a combination (Figure 3 on page 12). A retail customer that needs to comply with PCI-DSS, for example, needs IPsec or *Secure Sockets Layer* (SSL) encryption for payment card transaction data as part of its managed security service.

Table 1 summarizes the process based on the option the service provider selects. In the table, VPNA refers to one customer's VPN on a router that hosts VPNs for multiple customers.

Table 1: Comparing Packet Flow in IPsec VPNs, BGP/MPLS VPNs, and Combination VPNs

IPsec	BGP/MPLS VPN	BGP/MPLS VPN and IPsec
<ol style="list-style-type: none"> 1. Host A in site 1 of VPNA sends packets to host B in site 2 of VPNA. 2. Routers A and B negotiate an Internet Key Exchange (IKE) [9] phase-one session in aggressive or main mode to establish a secure and authenticated channel between peers. 3. Routers A and B negotiate an IKE phase-two session to establish security associations on behalf of IPsec services. 4. Information is exchanged securely through an IPsec tunnel. 5. The tunnel is terminated. 	<ol style="list-style-type: none"> 1. Host A in site 1 of VPNA sends packets to host B in site 2 of VPNA. 2. Packet arrives on a VPN Route-Forwarding (VRF) VPNA interface on the PE1 router. 3. The PE1 router performs an IP lookup, determines the label stack and the outgoing core-facing interface, and forwards the packet to the MPLS core. 4. The packet is label-switched at each hop in the core until it reaches the penultimate hop router. At this point, the top label is popped before the packet is forwarded to the egress provider edge router. 5. The egress PE2 router performs a MPLS lookup and determines that it should remove the label before forwarding the packet to host B in site 2. 6. Router B in site 2 receives a regular IP packet and forwards it to host B. 	<ol style="list-style-type: none"> 1. Router A in site 1 and the associated PE1 router negotiate an IKE phase-one session in aggressive or main mode to negotiate a secure and authenticated channel between peers. 2. Router A and the PE1 router negotiate an IKE phase-two session to establish security associations on behalf of IPsec services so that information is exchanged securely through an IPsec tunnel. 3. Host A in site 1 of VPNA sends packets to host B in site 2 of VPNA. 4. The PE1 router, which is enabled with VRF-aware IPsec, creates a direct association through the IPsec tunnel that connects site 1 and the corresponding VRF ID (VPNA) on the provider edge router over the Internet. 5. Encrypted traffic arrives on an Internet-facing interface on the provider edge router A, which terminates the IPsec tunnel, decrypts the incoming packet, and forwards the plaintext packet to the VRF VPNA for further processing. 6. The PE1 router performs an IP lookup, determines the label stack and the outgoing core-facing interface, and forwards the packet to the MPLS core. 7. The packet is label-switched at each hop in the core until it reaches the penultimate hop router. At this point, the top label is popped before the packet is forwarded to the egress provider edge router. 8. The egress PE2 router performs a MPLS lookup and determines that it should remove the label before forwarding the packet to host B in site 2. Router B in site 2 receives a regular IP packet and forwards it to host B. 9. If site 2 is also reachable over the Internet and the egress PE2 router is enabled with VRF-aware IPsec, the packet is encrypted and sent to site 2 across the Internet over an IPsec tunnel. 10. Router B in site 2 terminates the IPsec tunnel, performs a regular IP lookup, and forwards the packet to host B.

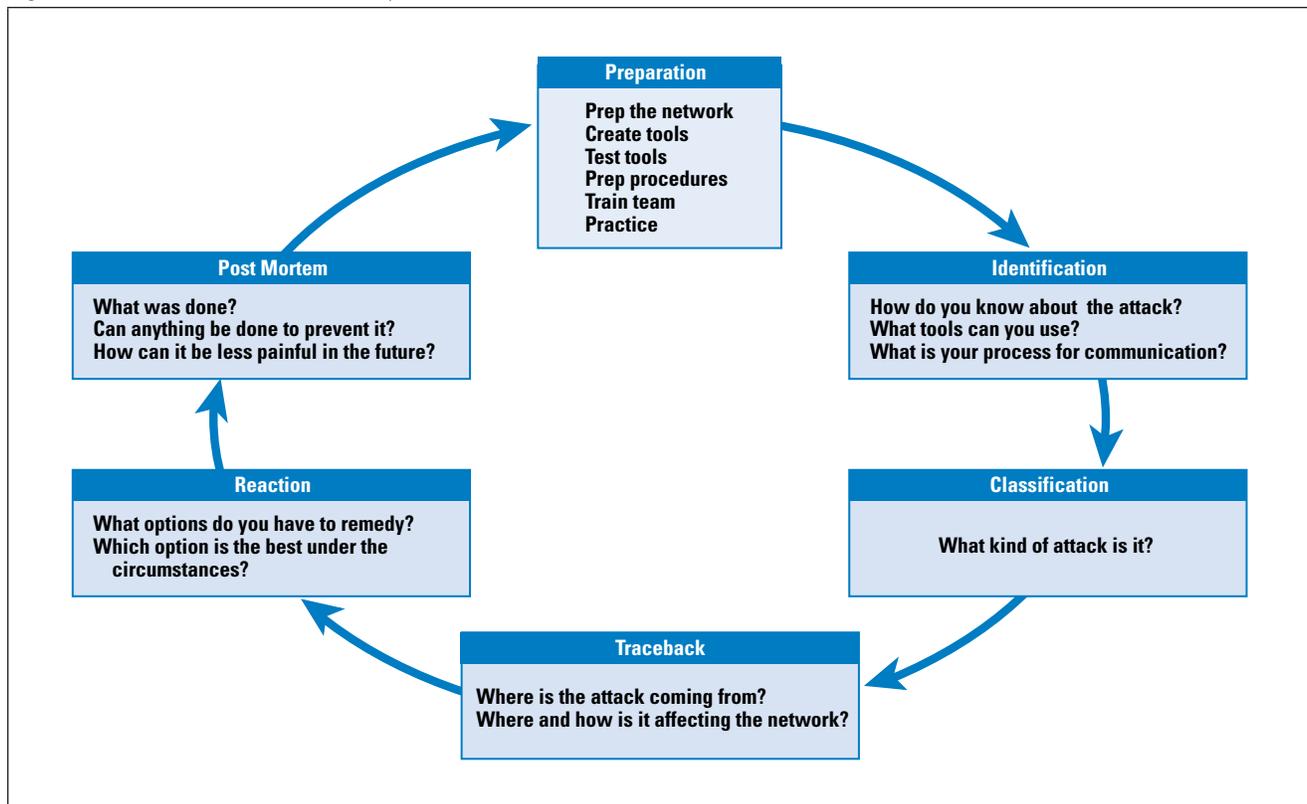
Figure 3: Managed IP VPN Security Services: IP MPLS, IPsec, or MPLS plus IPsec



Six-Step Methodology

Service providers can detect and mitigate attacks on the infrastructure using a six-step incident-response methodology (Figure 4).

Figure 4: Six Phases of Incident Response



- *Preparation:* The service provider needs to prepare the network, acquire the needed tools, develop and document a security plan, implement security procedures, and train NOC staff to use tools and procedures. *It is vital that security be a practice; the first time that the NOC staff follows its incident-response procedures should not be during an actual attack.*
- *Identification:* Unfortunately, service providers sometimes learn about a security incident from their customers. It is far better to be able to identify the threat before it becomes a problem, using NetFlow telemetry data and analysis tools, for example.
- *Classification:* The service provider needs to be able to quickly assess the nature of the threat and its scope: single customer, multiple customers, or entire infrastructure.
- *Traceback:* After classifying the threat, the IT staff needs to identify the point of ingress: peer, upstream server, downstream server, or compromised network device in the data center.
- *Reaction:* Following classification and traceback, the IT team applies the tools and processes needed to mitigate the attack. Success requires visibility into the network and well-defined procedures. Adherence to standard operating procedures helps prevent the service provider from inadvertently making the problem worse.
- *Post-mortem:* After the incident, the security team should analyze the root causes and integrate new insights into the security incident-handling procedures for use during the next incident.

Real-Life Observations About Interoperability from the APRICOT Workshops

Cisco and Juniper conducted a multivendor security workshop at APRICOT 2006 in Perth, Australia, and again at APRICOT 2007 in Bali, Indonesia. The workshops were offered in response to the fact that service providers often deploy a multivendor network for reasons ranging from financial to political.

Hands-on workshops were conducted in a lab using 12 routers running the Cisco IOS Software and another 12 running JUNOS software. Topics included:

- Password protection
- Packet filtering at the network edge
- Protecting the control plane
- Securing routing protocols
- Network monitoring techniques: NetFlow, syslog, SNMP, and *Network Time Protocol* (NTP)
- BGP MPLS Layer 3 VPNs
- IPsec VPNs

The goal of the workshops was to achieve a working configuration that interoperated with JUNOS and the Cisco IOS Software, resulting in consistent technology implementation, as well as common security policy enforcement. The workshops underscored the fact that interoperability is not automatic—even among standards-based network products. The reason is that standards bodies such as IETF, ITU, IEEE, and others define some aspects of protocols but leave others to vendor discretion. Standards do define *protocol format*, which is a syntactical structure identifying bit-field definition, length, and more. They also define *protocol behavior*, which specifies when actions occur, such as sending Hello and Keepalive timer probes and handling retransmission and reset packets. For purposes of analogy, a spoken language such as English is like a protocol format, and polite conversation conventions, such as beginning with a greeting and concluding with goodbye, is like a protocol behavior.

What standards do *not* cover are vendor-specific internal implementations, such as software coding techniques, hardware acceleration for performance, *command-line interface* (CLI) structure, and so on. Therefore, even though the APRICOT workshops involved deploying standards-based technology such as BGP-based MPLS VPNs and IPsec, vendor-specific differences had to be accounted for in the workshop materials and were noticed by participants. Following are examples noted at the APRICOT workshop:

- *Label Distribution Protocol*: With BGP MPLS VPN, JUNOS and Cisco IOS Software did not interoperate in their default configurations. However, routers from the same vendor did establish *Label Distribution Protocol* (LDP) sessions. The explanation, which participants found by troubleshooting with debug commands and referring to the manual, is that Cisco IOS Software uses the *Tag Distribution Protocol* (TDP) by default, whereas JUNOS uses LDP. After the Cisco IOS Software was changed to use LDP, the BGP-based MPLS VPN configuration succeeded.
- *IPsec tunnel establishment*: To simplify IPsec configuration, the workshop employed a *Graphical User Interface* (GUI) that prompted the user to choose source and destination IP addresses for the tunnel endpoints, a shared key, and the prefixes that defined the “interesting” traffic that was to use the IPsec tunnel. On the first attempt, the IPsec tunnel was not established. Workshop participants used the CLI to determine the problem, which was that the default encryption being negotiated was incompatible. The root cause for this mismatched encryption standard was that some routers were using an export version of software and needed an upgrade to support a higher encryption standard. Furthermore, even with common encryption capabilities, the two operating systems used different criteria to identify the interesting traffic that would be encrypted. Using the GUI, JUNOS defined interesting traffic as sourced from “ANY” network and destined to **192.168.1.0/24**.

In contrast, the Cisco IOS Software defined interesting traffic as sourced from **10.1.1.0/24** and destined to **192.168.1.0/24**. Following a discussion about whether the JUNOS default was too permissive or the Cisco IOS Software default was too restrictive, workshop participants agreed to disallow traffic that did not require encryption in the IPsec tunnel. The consensus was that the customer's security policy would provide a more conclusive answer to how permissive the policy should be, and that it was reasonable to require use of the CLI to tweak the configuration because the GUI performed most of the more difficult parts of the configuration on both platforms.

- *Loopback interface cost with Open Shortest Path First (OSPF):* During the OSPF deployment, participants noticed that the OSPF cost associated with interfaces was the same for each vendor. The OSPF cost is based upon a reference bandwidth of 100 Mbps. However, the loopback interfaces had different values: a default OSPF cost of 1 for the Cisco IOS Software and 0 for JUNOS. It is advisable to change one of the defaults to make them the same.

Although these subtle differences in protocols are documented by the vendors, service provider operational teams often have little time to research them. Therefore, it can be valuable for them to participate in multivendor hands-on workshops. Anecdotal evidence suggests that operators who are comfortable with multiple vendors understand the protocols, helping them design networks that can support new, revenue-generating services.

It is hoped that events such as the APRICOT workshops will help build a community of professionals who can add value for their employers, each other, and the broader Internet community. The result will be a secure and trusted networking environment that people and industry can rely on and use to connect in new and innovative ways.

Summary

Managed security services represent a growing revenue opportunity for service providers. Most service providers operate in a multivendor environment, either because of mergers and acquisitions or because their customers use other vendors' equipment. Therefore, a standards-based approach positions providers to capitalize on the managed security service opportunity. Providers can secure their infrastructure in a multivendor environment by following best practices for point protection, edge protection, RTBH protection, source-address validation, control-plane protection, and total visibility into network activity.

References

- [1] S. Kent and R. Atkinson, "Security Architecture for the Internet Protocol," RFC 2401, November 1998.
- [2] S. Kent and R. Atkinson, "IP Authentication Header," RFC 2402, November 1998.
- [3] S. Kent and R. Atkinson, "IP Encapsulating Security Payload (ESP)," RFC 2406, November 1998.
- [4] E. Rosen and Y. Rekhter, "BGP/MPLS VPNs," RFC 2547, March 1999.
- [5] S. Hanks, T. Li, D. Farinacci, and P. Traina, "Generic Routing Encapsulation (GRE)," RFC 1701, October 1994.
- [6] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, and E. Lear, "Address Allocation for Private Internets," RFC 1918, February 1996.
- [7] Internet Assigned Numbers Authority (IANA), "Special-Use IPv4 Addresses," RFC 3300, September 2002.
- [8] F. Baker and P. Savola, "Ingress Filtering for Multihomed Networks," RFC 3704, March 2004.
- [9] D. Harkins and D. Carrel, "The Internet Key Exchange (IKE)," RFC 2409, November 1998.
- [10] Convery, S., "Network Authentication, Authorization, and Accounting – Part One: Concepts, Elements, and Approaches," *The Internet Protocol Journal*, Volume 10, No. 1, March 2007.
- [11] Convery, S., "Network Authentication, Authorization, and Accounting – Part Two: Protocols, Applications, and the Future of AAA," *The Internet Protocol Journal*, Volume 10, No. 2, June 2007.
- [12] Rigney et. al., "Remote Authentication Dial- In User Service (RADIUS)," RFC 2865 (Obsoletes RFC 2138, and RFC 2058), June 2000.
- [13] Carrel et al., "The TACACS+ Protocol Version 1.78," Internet Draft, Work in Progress, **draft-grant-tacacs-02.txt**, January 1997.
- [14] <http://www.shrubbery.net/rancid/>
- [15] <http://www.apricot.net>

KUNJAL TRIVEDI joined Cisco in 1999 as a consulting engineer initially and then worked in product management covering Cisco IOS Software infrastructure security. Currently, he is helping Cisco shape a Managed Security Services marketing vision and strategy. A widely respected networking security expert, Kunjal presents infrastructure security, IP Security, and Managed Security topics at Cisco Networkers events as well as at conferences such as APRICOT. Kunjal has a Bachelor of Engineering degree with honors in electrical and electronics engineering from University of Wales, College of Cardiff, and a Master of Science degree in Artificial Intelligence from Cranfield Institute of Technology, UK. He holds CISSP and CCIE designations in routing and switching as well as security. Recently, he published a book titled *[Read Me First]: Building or Buying VPNs*; Kunjal has been awarded Chartered Engineer status by Institute of Engineering and Technology. He can be reached at kunjal@cisco.com

DAMIEN HOLLOWAY joined Juniper Networks in 2004 as an Instructing Engineer. He contributes to the development of the Juniper Technical Certification Program and custom delivery of training in the Asia Pacific region. Previously he was a consulting engineer and provided design, installation, and training to providers in Australia and the United States. Damien has presented a wide variety of topics relevant to customers, including backbone design, application acceleration, and *Broadband Remote Access Server* edge design, to audiences, including APRICOT and SANOG. Damien has a Bachelor of Electrical Engineering and Bachelor of Science from University of Sydney, Australia. He is a CCIE expert in routing and switching and JNCIE-M, JNCIP-E, and CISSP. He can be reached at holloway@juniper.net

Kunjal Trivedi (left) and Damien Holloway (center) share a joke with workshop students at APRICOT 2007



Kunjal with APRICOT 2007 workshop attendees



IPv4 Address Depletion and Transition to IPv6

by Geoff Huston, APNIC

At the recent APNIC meeting in New Delhi, the subject of IPv4, IPv6, and transition mechanisms was highlighted in the plenary session^[1]. This article briefly summarizes that session and the underlying parameters in IPv4 address depletion and the transition to IPv6.

IPv4 Status

As of September 2007 we have some 18 percent of the unallocated IPv4 address pool remaining with the *Internet Assigned Numbers Authority* (IANA), and 68 percent has already been allocated to the *Regional Internet Registries* (RIRs) and through the RIRs to *Internet Service Providers* (ISPs) and end users. The remaining 14 percent of the IPv4 address space is reserved for private use, multicast, and special purposes. Another way of looking at this situation is that we have exhausted four-fifths of the unallocated address pool in IPv4, and one-fifth remains for future use. It has taken more than two decades of Internet growth to expend this initial four-fifths of the address space, so why shouldn't it take a further decade to consume what remains?

At this point the various predictive models come into play, because the history of the Internet has not been a uniformly steady model. The Internet began in the 1980s very quietly; the first round of explosive growth in demand was in the early 1990s as the Internet was adopted by the academic and research sector. At the time, the address architecture used a model where class A networks (or a /8) were extremely large, the class B networks (/16) were also too large, and the class C networks (/24) were too small for most campuses. The general use of class B address blocks was an uncomfortable compromise between consuming too much address space and consuming too many routing slots through address fragmentation. The subsequent shift to a classless address architecture in the early 1990s significantly reduced the levels of IPv4 address consumption for the next decade. However, over the past five years the demand levels for addresses have been accelerating again. Extensive mass-market broadband deployment, the demand for public non-*Network Address Translation* (NAT) addresses for applications such as *Voice over IP* (VoIP), and continuing real cost reductions in technology that has now brought the Internet to large populations in developing economies all contribute to an accelerating IPv4 address consumption rate.

Various approaches to modeling this address consumption predict that the IANA unallocated address pool will be fully depleted sometime in 2010 or 2011^[2, 3, 4, 5].

Transitioning to IPv6

The obvious question is “What then?”, and the commonly assumed answer to that question is one that the *Internet Engineering Task Force* (IETF) started developing almost 15 years ago, namely a shift to use a new version of the Internet Protocol: what we now know as IP Version 6, or IPv6. But if IPv6 really is the answer to this problem of IPv4 unallocated address-pool depletion, then we appear to be leaving the transition process quite late. The uptake of IPv6 in the public Internet remains extremely small as compared to IPv4^[6]. If we really have to have IPv6 universally deployed by the time we fully exhaust the unallocated IPv4 address pools, then this objective appears to be unattainable during the 24 months we have to complete this work. The more likely scenario we face is that we will not have IPv6 fully deployed in the remaining time, implying a need to be more inventive about IPv4 in the coming years, as well as inspecting more closely the reason why IPv6 has failed to excite much reaction on the part of the industry to date.

We need to consider both IPv4 and IPv6 when looking at these problems with transition because of an underlying limitation in technology: *IPv6 is not “backward-compatible” with IPv4*. An IPv6 host cannot directly communicate with an IPv4 host. The IETF worked on ways to achieve this through intermediaries, such as a protocol to translate NATs^[7], but this approach has recently been declared “historic” because of technical and operational difficulties^[8]. That decision leaves few alternatives. If a host wants to talk to the IPv4 world, it cannot rely on clever protocol translating intermediaries somewhere, and it needs to have a local IPv4 protocol stack, a local IPv4 address, and a local IPv4 network and IPv4 transit. And to speak to IPv6 hosts, IPv6 has the same set of prerequisites as IPv4. This approach to transition through replication of the entire network protocol infrastructure is termed “Dual Stack.” The corollary of Dual Stack is continued demand for IPv4 addresses to address the entire Internet for as long as this transition takes. The apparent contradiction here is that we do not appear to have sufficient IPv4 addresses in the unallocated address pools to sustain this Dual Stack approach to transition for the extended time periods that we anticipate this process to take.

What Can We Expect?

So we can expect that IPv4 addresses will continue to be in demand well beyond any anticipated date of exhaustion of the unallocated address pool, because in the Dual Stack transition environment all new and expanding network deployments need IPv4 service access and addresses. But the address distribution process will no longer be directly managed through address allocation policies after the allocation pool is exhausted.

Ideas that have been aired in address policy forums include encouraging NAT deployment in IPv4, expanding the private use of IPv4 address space to include the last remaining “reserved-for-future-use” address block, various policies relating to rationing the remaining IPv4 address space, increased efforts of address reclamation, the recognition of address transfers, and the use of markets to support address distribution.

Of course the questions here are about how long we need to continue to rely on IPv4, how such new forms of address distribution would affect existing notions of fairness and efficiency of use, and whether this effect would imply escalation of cost or some large-scale effect on the routing system.

On the other hand, is IPv6 really ready to assume the role of the underpinning of the global Internet? One view is that although the transition to a universal deployment of IPv6 is inevitable, numerous immediate concerns have impeded IPv6 adoption, including the lack of backward compatibility and the absence of simple, useful, and scalable translation or transition mechanisms^[9]. So far the business case for IPv6 has not been compelling, and it appears to be far easier for ISPs and their customers to continue along the path of IPv4 and NATs.

When we contemplate this transition, we also need to be mindful of what we need to preserve across this transition, including the functions and integrity of the Internet as a service platform, the functions of existing applications, the viability of routing, the capability to sustain continued growth, and the integrity of the network infrastructure.

It appears that what could be useful right now is clear and coherent information about the situation and current choices, and analyzing the implications of various options. When looking at such concerns of significant change, we need to appreciate both the limitations and the strengths of the Internet as a global deregulated industry and we need, above all else, to preserve a single coherent networked outcome. Perhaps this topic is far broader than purely technical, and when we examine it from a perspective that embraces economic considerations, business imperatives, and public policy objectives, we need to understand the broader context in which these processes of change are progressing^[10].

It is likely that some disruptive aspects of this transition will affect the entire industry, and this transition will probably be neither transparent nor costless.

References

- [1] APNIC 24 Plenary Session: “The Future of IPv4,” September 2007. <http://www.apnic.net/meetings/24/program/plenaries/apnic/>
- [2] Geoff Huston, “The IPv4 Report.” <http://ipv4.potaroo.net>
- [3] Tony Hain, “IPv4 Address Pool.” <http://www.tndh.net/~tony/ietf/ipv4-pool-combined-view.pdf>
- [4] Tony Hain, “A Pragmatic Report on IPv4 Address Space Consumption,” *The Internet Protocol Journal*, Vol. 8, No. 3, September 2005. http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_8-3/ipv4.html
- [5] K.C. Claffy, CAIDA, “ ‘Apocalypse Then’: IPv4 Address Space Depletion,” Presentation to ARIN XVI, October 2005. http://www.arin.net/meetings/minutes/ARIN_XVI/PDF/wednesday/claffy_ipv4_roundtable.pdf
- [6] Geoff Huston, “IPv6 / IPv4 Comparison Metrics.” <http://bgp.potaroo.net/v6/v6rpt.html>
- [7] G. Tsirtsis and P. Srisuresh, “Network Address Translation – Protocol Translation (NAT-PT),” RFC 2766, February 2000.
- [8] C. Aoun and E. Davies, “Reasons to Move the Network Address Translator – Protocol Translator (NAT-PT) to Historic Status.” RFC 4966, July 2007.
- [9] Randy Bush, “IPv6 Operational Reality,” APNIC 24 Plenary Presentation, September 2007. <http://www.apnic.net/meetings/24/program/plenaries/apnic/presentations/bush-ipv6-op-reality.pdf>
- [10] Geoff Huston, “IPv4 Exhaustion,” APNIC 24 Plenary Presentation, September 2007. <http://www.apnic.net/meetings/24/program/plenaries/apnic/presentations/huston-ipv4-exhaustion.pdf>

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

IPv4 Address Space: 2.46 Billion Down, 1.25 Billion to Go

by Iljitsch van Beijnum

In September 2005, *The Internet Protocol Journal* published an article about the IPv4 address space consumption^[1]. At that time, projections done by Geoff Huston and Tony Hain varied widely, because the number of /8 address blocks in use had gone up sharply in early 2005. So what has happened since then, and what can we expect for the not-too-distant future?

Address Assignment and Allocation

The *Internet Assigned Numbers Authority* (IANA, part of the *Internet Corporation for Assigned Names and Numbers* [ICANN]) has authority over the IPv4 address space. In the past, IANA gave out address blocks directly to end users, but now IANA distributes address space in the form of /8 blocks, each holding 24 bits worth of address space, or 16,777,216 addresses, to five *Regional Internet Registries* (RIRs). There are a few exceptions, but AfriNIC^[2] gives out address space in Africa; APNIC^[3] in the Asia-Pacific region; ARIN^[4] in North America; LACNIC^[5] in Latin America and the Caribbean; and RIPE NCC^[6] in Europe, the former Soviet Union, and the Middle East. These RIRs sometimes assign address space to end users, but mostly allocate it to *Internet Service Providers* (ISPs), who then assign it to their customers, meaning that there are two pools of available address space: the global pool of /8 blocks that IANA has not delegated to anyone^[7], and the address space held by the RIRs that they have not given out yet. The article in the September 2005 issue of *The Internet Protocol Journal*^[1] looked at the depletion of the IANA global pool, whereas this article mostly looks at the amounts of address space given out by the RIRs, providing a more granular view. The RIRs publish daily reports of their address assignments and allocations on their respective FTP servers. According to these reports as downloaded on January 1, 2007, the amounts of address space shown in Table 1 were given out over the past seven years.

Table 1: Address Space Allocated 2000–2006 [January 2007 data]

	2000	2001	2002	2003	2004	2005	2006
AfriNIC	0.56	0.39	0.26	0.22	0.51	1.03	2.72
APNIC	20.94	28.83	27.03	33.05	42.89	53.86	51.78
ARIN	30.83	28.55	21.08	22.32	34.26	47.57	38.94
LACNIC	0.88	1.61	0.65	2.62	3.77	10.97	11.50
RIPE NCC	24.79	25.36	19.84	29.61	47.49	62.09	56.53
Total	78.00	84.73	68.87	87.82	128.92	175.52	161.48

However, if we compare these totals to the totals seen on January 1, 2006, we see some differences (Table 2).

Table 2: Address Space Allocated 2000–2006 [January 2006 data]

	2000	2001	2002	2003	2004	2005	2006
Total	78.35	88.95	68.93	87.77	128.45	165.45	–

For the years 2000 to 2002, the number of addresses registered as given out is slightly lower, as seen in the January 1, 2007 data compared to the January 1, 2006 data—a result that is to be expected because address space given out in that year that is no longer used is returned. However, for the later years, and especially for 2005, there is a retroactive *increase* in the number of addresses given out. The reason: When ARIN suspects an address space user may come back for more space relatively soon, it takes a larger block than requested, and then fulfills the request from part of that block and keeps the rest in reserve. So an organization requesting a /16 may get the first half of a /15. When that organization then requests another /16 one or two years later, ARIN gives the organization the second half of the /15. ARIN subsequently records this as a /15 given out on the date when the original /16 was requested.

For instance, ARIN’s January 1, 2006, data shows that a block of 12.6 million addresses was given out within **73.0.0.0/8** block:

arin|US|ipv4|73.0.0.0|12582912|20050419|allocated

In the January 1, 2007, data, this number had changed to 13.6 million addresses:

arin|US|ipv4|73.0.0.0|13631488|20050419|allocated

This change means that simply looking at the registration date does not provide very good information. It also does not account for address space given out in earlier years that is returned. An alternative approach is to count the amount of address space given out based on the RIR records published on a certain date (Table 3).

Table 3: RIR Records for Address Space Allocation

	IANA (/8)		RIRs (millions)			Total	
	Delegated	Free	Received	Delegated	Free	Free	Delta
Jan. 1, 2004	133	88	1509.95	1245.63	264.32	1740.71	
Jan. 1, 2005	142	79	1660.95	1351.66	309.30	1634.69	106.02
Jan. 1, 2006	155	66	1879.05	1517.74	361.31	1468.61	166.08
Jan. 1, 2007	166	55	2063.60	1685.69	377.90	1300.65	167.96
May 1, 2007	172	49	2181.04	1754.68	426.36	1248.44	52.21

(Note that block **7.0.0.0/8** shows up as unused in the IANA global pool and is counted as available in the table, but this block is in fact used by the U.S. Department of Defense.)

The jump in address consumption between 2004 (106 million) and 2005 (166 million) is even more dramatic in this light, while consumption numbers of 2005 and 2006 (168 million) are now almost identical. The figure for the first four months of 2007 seems rather modest at 52 million addresses, but the reason lies in the fact that Bolt, Beranek and Newman returned **46.0.0.0/8** to IANA in April. So the number of addresses given out from January to April was 69 million, a rate that puts the RIRs on track to give out more than 200 million addresses in 2007.

The size of address blocks given has been increasing steadily. Table 4 shows the number of requests for a certain range of block sizes: equal or higher than the first, lower than the second value (2005 and earlier values from the January 1, 2006 data, 2006 values from the January 1, 2007 data).

Table 4: Number of Requests for Ranges of Block Sizes

	2000	2001	2002	2003	2004	2005	2006
< 1,000	326	474	547	745	1022	1309	1526
1,000 – 8,000	652	1176	897	1009	1516	1891	2338
8,000 – 64k	1440	868	822	1014	1100	1039	1133
64k – 500k	354	262	163	215	404	309	409
500k – 2M	19	39	29	46	61	60	56
> 2M	3	5	5	6	7	18	13

The number of blocks in the two smallest categories has increased rapidly, but not as fast as the number of blocks in the largest category, in relative numbers. However, the increase in large blocks has a very dramatic effect whereas the small blocks are insignificant, when looking at the millions of addresses involved (Table 5).

Table 5: Millions of Addresses Given Out

	2000	2001	2002	2003	2004	2005	2006
< 1,000	0.10	0.16	0.18	0.25	0.35	0.44	0.52
1,000 – 8,000	2.42	4.47	3.23	3.45	4.49	5.07	6.10
8,000 – 64k	18.79	12.81	11.35	14.00	15.99	15.46	17.17
64k – 500k	35.98	32.19	20.28	25.51	42.01	34.23	49.64
500k – 2M	12.68	24.64	21.30	31.98	44.63	41.63	46.64
> 2M	8.39	14.68	12.58	12.58	20.97	68.62	41.42

The increase in the 2M+ blocks was solely responsible for the high number of addresses given out in 2005. In 2006, there was growth in all categories except the 2M+ one (even the 500k – 2M category increased in number of addresses if not in number of blocks). When the 2M+ blocks are taken out of the equation, 2005 had a total of 96.83 million addresses (January 1, 2006) and 2006 had 119.06 million given out, even without fully correcting for the ARIN reporting particularities. Apparently there is still an underlying upward trend.

Figure 1 shows the amounts of address space given out by IANA and by the RIRs every year from 1994 to 2006.

Figure 1: IPv4 Address Space Given Out from 1994 to 2006

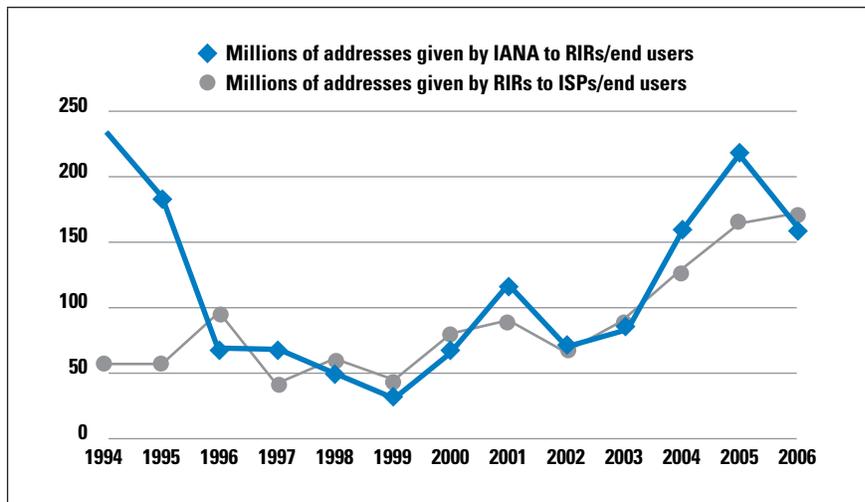
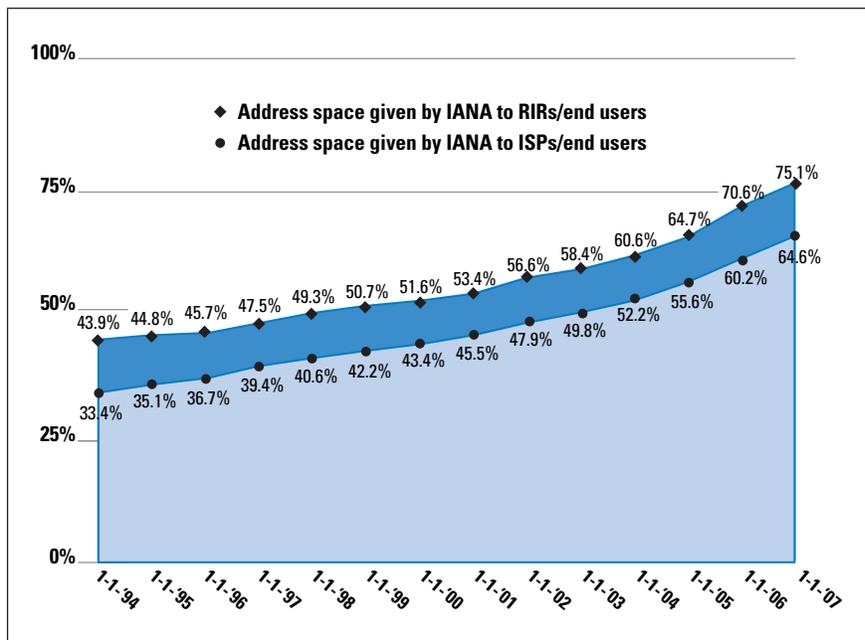


Figure 2 shows the amounts of address space marked as “in use” by IANA and by the RIRs. The difference between the two numbers is what the RIRs hold in order to satisfy day-to-day address space requests. This amount is usually two years’ worth of address space.

Figure 2: IPv4 Address Space in Use from 1994 to 2006



Depletion

The exact moment when the IPv4 address space will be depleted depends on numerous factors. Since 1997, the three-year period with the largest growth in yearly address use was from 2003 to 2005 relative to the 2002 figure: a factor 2.4, or 34 percent per year. If this growth repeats itself in the next three years, we will be out of IPv4 addresses in the second half of 2010.

Interestingly, the period with the lowest growth also includes the year 2003: from 2001 to 2003 relative to 2000. In 2003, 12 percent more addresses were given out than in 2000, for an average increase in yearly use of 4 percent. If this is the new trend for the coming years, we can expect to run out of IPv4 addresses in mid-2013. There is of course no reason to assume that future IP address use will conform to patterns seen in earlier years, but we really have nothing else to base our projections upon.

So anyone expecting to obtain new IPv4 address space more than three years from now is taking a big risk. With the IPv4 reserves visibly diminishing each year, the question is: What can we, as a community, do to make the IPv4 address depletion as painless as possible? IPv4 addresses are useful only if the people who need them can obtain them, meaning that using up addresses unnecessarily fast or locking up the still-available reserves are both suboptimal solutions. It has been suggested that turning IPv4 address space into a tradable commodity would allow a free market to form, aiding the efficient distribution of address space from those who have it to those who need it.

This scenario has several problems. First, when supply is limited and demand is high, prices rise and hoarding becomes lucrative. So the effect of making address space tradable could be a reduction of available address space rather than an increase. And certainly, as trading IPv4 space becomes more likely, holders of large address blocks will be less inclined to return them. Finally, more than half of the IPv4 address space in use is held by organizations in the United States, whereas the developing world has comparatively little address space. The prospect of having to buy address space from American companies that got the space for free is not likely to be popular in the rest of the world.

Address Reclamation a Solution?

There are two large classes of potentially reclaimable address space: the class E reserved space (**240.0.0.0 – 255.255.255.255**) and the class A blocks given out directly to end users by IANA. The class E space has 268 million addresses and would give us in the order of 18 months worth of IPv4 address use. However, many TCP/IP stacks, such as the one in Windows, do not accept addresses from class E space and will not even communicate with correspondents holding those addresses. It is probably too late now to change this behavior on the installed base before the address space would be needed. There are currently 42 class A blocks and another two /8s from class C space listed as given out to end users—738 million addresses. The U.S. government uses about 10 of those blocks; 21 of them are not present in the *Border Gateway Protocol* (BGP) routing table.

Although harsh judgments about the need for so much address space are easily made from the outside without having all the pertinent information, it seems reasonable to try to reclaim some of this space. I would consider getting back half of this space a big success, but that would give us only 2 years worth of additional address space. There are also 645 million addresses of older class B assignments, but reclaiming those will be extremely difficult because nearly 8,000 individual assignments are involved. Reclaiming a class B block is probably not much easier than reclaiming a class A block, but the amount of address space returned is less than half a percent.

Planning for the End Game

So what should we do? In my opinion: promote predictability. The situation where we run out of IPv4 address space much faster than expected would be very harmful as organizations struggle to adjust to the new circumstances. On the other hand, if the IPv4 space unexpectedly lasts longer, people may be disinclined to believe space is really running out and then would be unprepared when it does. Artificially delaying running out of IPv4 address space also prolongs the situation in which it is difficult to get IPv4 space, but not enough people feel the pain to initiate IPv6 deployment. One solution worthy of consideration would be to impose a worldwide moratorium on the change of IPv4 address allocation and assignment policies after a certain date to aid this predictability. If some kind of encouraged or forced reclamation of older class A blocks is desired, this process should be instigated sooner rather than later, both for the sake of predictability and because it gives the address holders involved time to reorganize their networks. Another small but useful step would be to limit the size of address blocks given out. This scenario would be like the agreement between the RIRs and IANA that the RIRs will receive two /8s at a time in the future. The situation where a single /9 or /8 allocation constitutes 5 or even 10 percent of the address space given out in that year makes adequate predictions extremely difficult, and also runs the risk that a good part of the address block in question will never be used as circumstances change. Limiting individual allocations to a /11 or /12 would be better, even if it requires the requesting organization to come back for more address space several times per year.

Finally, it seems prudent for all organizations using public IPv4 address space to start planning for the moment that they themselves, or third parties that they communicate with over the public Internet, can no longer obtain additional IPv4 address space.

References

- [1] Tony Hain, "A Pragmatic Report on IPv4 Address Space Consumption," *The Internet Protocol Journal*, Volume 8, No. 3, September 2005.
- [2] AfriNIC, <http://www.afrinic.net>
- [3] APNIC, <http://www.apnic.net>
- [4] ARIN, <http://www.arin.net>
- [5] LACNIC, [http://www/lacnic.net](http://www.lacnic.net)
- [6] RIPE NCC, <http://ripe.net>
- [7] IANA Internet Protocol v4 Address Space
<http://www.iana.org/assignments/ipv4-address-space>

ILJITSCH VAN BEIJNUM holds a Bachelor of Information and Communication Technology degree from the Haagse Hogeschool in The Hague, Netherlands. In 1995, he found himself in the emerging Internet Service Provider business. There he learned about system administration, IP networking, and especially routing. After first starting a small ISP with four others and working as a senior network engineer for UUNET Netherlands, he became a freelance consultant in 2000. Not long after that, he started contributing to the IETF Multihoming in IPv6 working group. He wrote the book *BGP: Building Reliable Networks with the Border Gateway Protocol*, ISBN 0-596-00254-8, published by O'Reilly in 2002, and *Running IPv6*, ISBN 1590595270, published by Apress in 2005. E-mail: iljitsch@muada.com

Used but Unallocated: Potentially Awkward /8 Assignments

by Leo Vegoda, ICANN

IPv4 has proven to be exceedingly popular, so it should be no surprise that the time is rapidly approaching when the last /8 block will be allocated and the *Internet Assigned Numbers Authority's* (IANA's) free pool of address space will be empty. At the time of writing, Geoff Huston of the *Asia Pacific Network Information Centre* (APNIC) is projecting^[1] the IANA free pool will run out in mid-2010. Unfortunately, it is possible that some of these remaining /8s may cause problems for enterprise and *Internet Service Provider* (ISP) network operators when they are put back into use. These blocks are not the /8s that have been returned to IANA by the original registrants; they are previously unassigned address blocks.

Concerns

There are many concerns about the IANA free pool depletion, but one of them seems particularly straightforward to identify and fix. Many organizations have chosen to use unregistered IPv4 addresses in their internal networks and, in some cases, network equipment or software providers have chosen to use unregistered IPv4 addresses in their products or services. In many cases the choice to use these addresses was made because the network operators did not want the administrative burden of requesting a registered block of addresses from a *Regional Internet Registry* (RIR)^[2, 11]. In other cases they may not have realized that RFC 1918^[3] set aside three blocks of address space for private networks, so they just picked what they believed to be an unused block, or their needs exceeded the RFC 1918 set-aside blocks. Other organizations used the default address range suggested by their equipment vendor, or supplied in example documentation, when configuring *Network Address Translation* (NAT) devices. Regardless of the reason, these uses of unregistered addresses will conflict with routed addresses when the /8s in question are eventually assigned to ISPs or enterprise users.

A few examples of /8s where problems are likely to occur follow:

- 1.0.0.0/8** Widely used as private address space in large organizations whose needs exceed those provided for by RFC 1918^[4]
- 5.0.0.0/8** Used by one of numerous zero-configuration Internet applications (including the Hamachi VPN service^[5, 6])
- 42.0.0.0/8** Default range used in the NAT configuration of at least one Internet appliance (the HP Procurve 700wl^[7])

Organizations using these address ranges in products or services may experience problems when more specific Internet routes attract traffic that was meant for internal hosts, or alternatively find themselves unable to reach the legitimate users of those addresses because those addresses are being used internally. The users of unregistered networks may also find problems with reverse *Domain Name System* (DNS) resolution, depending on how their DNS servers are configured. These problems are likely to result in additional calls to helpdesks and security desks at both enterprises and ISPs, with unexpected behavior for end users that might be hard to diagnose. Users of unregistered address space may also experience problems with unexpected traffic being received at their site if they leak internal routes to the public Internet. Many ISPs have already had experience with this type of routing inconsistency as recent /8 allocations reach routing tables and bogon filters are updated.

Alternatives

There are several alternatives to using unregistered IPv4 address space:

- Use RFC 1918 IPv4 address space (no need to obtain this space from an RIR)
- Use IPv4 address space registered with an RIR
- Use IPv6 address space registered with an RIR
- Use IPv6 Unique Local Address^[8] space (no need to obtain this space from an RIR)

Obviously, all of these efforts will involve renumbering networks, a sometimes painful and time-consuming process. Those using unregistered unique IPv4 address space should look at renumbering their networks or services before the previously unallocated /8s are allocated to avoid address clashes and routing difficulties.

Additionally, vendors and documentation writers can clean up their configurations to ensure they use RFC 1918 addresses, or make it clear to their users that they must use registered addresses to avoid routing conflicts.

All RIRs provide free telephone helpdesks that can advise you on obtaining unique IPv4 or IPv6 address space. But if you want to continue using unregistered space and can transition to IPv6, the prefix selection mechanism described in RFC 4193 makes the probability of a clash a mere 1 in 550 billion. Ultimately, transitioning to IPv6 is most likely the best solution, and this approach offers an opportunity for those having to renumber parts of their network to avoid a subsequent renumbering later into IPv6.

About IANA and ICANN

IANA allocates address space to RIRs according to the global IPv4 [9] and IPv6^[10] policies. Enterprise and ISP networks need to obtain IP addresses from their upstream provider or from the appropriate RIR.

The *Internet Corporation for Assigned Names and Numbers* (ICANN) is an internationally organized, nonprofit corporation that has responsibility for *Internet Protocol* (IP) address space allocation, protocol identifier assignment, *generic* (gTLD) and *country code* (ccTLD) *Top-Level Domain* name system management, and root server system management functions. These services were originally performed under U.S. government contract by IANA and other entities. ICANN now performs the IANA function.

References

- [1] <http://www.potaroo.net/tools/ipv4/>
- [2] RFC 1174, para 2.2 states in part, “The term Internet Registry (IR) refers to the organization which has the responsibility for gathering and registering information about networks to which identifiers (network numbers, autonomous system numbers) have been assigned by the IR. An RIR does this function for its service area.”
- [3] <http://www.ietf.org/rfc/rfc1918.txt>
- [4] <http://tools.ietf.org/id/draft-hain-1918bis-01.txt>
- [5] <https://secure.logmein.com/products/hamachi/howitworks.asp>
- [6] <http://en.wikipedia.org/wiki/Hamachi>
- [7] <http://www.hp.com/rnd/support/faqs/700w1.htm>
- [8] <http://www.ietf.org/rfc/rfc4193.txt>
- [9] <http://www.icann.org/general/allocation-IPv4-rirs.html>
- [10] <http://www.icann.org/general/allocation-IPv6-rirs.htm>
- [11] 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.

LEO VEGODA holds a BA (Hons) from the University of Central England. He joined ICANN in 2006 and is the Manager, Number Resources - IANA. He has previously worked for the RIPE NCC, where he ran the Registration Services department. He can be reached at: leo.vegoda@icann.org

Book Review

Uncommon Sense *Uncommon Sense: Out of the Box Thinking for An In the Box World*, By Peter Cochrane, ISBN 1-84112-477-x, Published by Capstone, 2004, <http://www.wileyurope.com>

A series of articles published in **silicon.com** form the basis for this book, which looks at the effect that new technology has on business and its implications for society. In many ways it attacks conventional wisdom and forces a reevaluation of the effect of technology, often exposing flaws in the business logic that lead to many investments and decisions.

The book is aimed at technologists, managers, and professionals who are interested in change and progress, offering them a glimpse of the future. It is easy to read, with liberal use of figures and tables to aid understanding.

Organisation

Cochrane begins by looking at the communication of ideas, particularly fairly complex and novel concepts. He notes the lack of agreement on the major concerns of the future and bemoans the handling of complex business and political topics—and the lack of engineering type rigour applied to their assessment. He suggests a much more rigorous modeling of complex business problems is required, especially of business processes, which are typically complex and inter-related, so treating them as isolated “stovepipes” is inappropriate and error-prone. Cochrane emphasises the need for nonlinear thinking.

Cochrane’s analysis continues with an assessment of technology markets, not surprisingly beginning with the forces behind the dot-com bubble, with particular reference to the effect that the so-called new and old economies have had on each other. He suggests that short-term approaches, with their tendency to hit high-visibility symptoms and not the underlying commercial factors, are a barrier to progress. Cochrane reflects that whilst the dot-com boom is over, it is now clear that the online world has been very successful and has dragged the old world along in its wake.

The book then looks at change: considering the adoption of new technology and the impact effect of the Internet, comparing this new technology with the adoption of television. Cochrane spends a significant amount of time on both entertainment and learning. He examines topics as varied as security, the ease of movement of information across borders, and the role of specialist and general devices.

His assessment of security considers the range and rate of spread of threats and some advanced countermeasures such as biometrics. He considers the nature of change programmes and the harmful ways insensitive micromanagement can affect their progress.

Cochrane explores the role of the consumer in deciding which technical innovations survive, as exemplified by the growth of the American cable TV (CATV) market. He notes that most consumers have a fixed level of disposable income and new innovations allow them to redirect rather than increase their level of spending. Cochrane argues that this truth is reflected in the saturation within the mobile handset market and the dynamics seen between the media companies and new innovators such as Napster.

The penultimate collection of essays considers the speed of innovation. Cochrane notes that many consumers are suffering from “technology fatigue” and many products are suffering from “feature death.” Here he discusses stagnation within the mobile market and disillusionment with the *Wireless Application Protocol* (WAP), *General Packet Radio Service* (GPRS), and Bluetooth. He notes that the adoption of technology is linked to the willingness of customers to pay.

Cochrane concludes by looking at leading-edge variables, including reliability, noting that this variable goes hand-in-hand with maturity, with the *Public Switched Telephone Network* (PSTN) delivering extremely high levels of reliability and most modern IT solutions delivering considerably less. He makes this comparison a critical test of the five-nines availability claims of many new technology solutions. Cochrane looks at some more less-conventional ideas such as the replication of ant logic in IT systems and the possible future use of plasma screens and voice recognition as convenient input/output devices. He notes the increasing intelligence of devices, but also acknowledges that rapid communications and minimal hierarchy can triumph over better organised structures as demonstrated by protesters in France in 2000 and 2001.

Synopsis

Cochrane takes the reader through many contemporary technology developments and concerns and in the process invites his readers to form their own views. His mission is to “communicate the implications of what we have done, are doing and are about to do.” In 50 short articles, delivered in 233 pages, it is possible for the author to cover only a small portion of a rapidly growing field, providing sufficient detail to appeal to the technologist without losing the bigger picture. He examines the implications of new technology for society and notes that the progress we are seeing means that we have to take on the new, changing the way we manage, operate, and govern our businesses as a result.

The Author

Peter Cochrane is the ex-BT Chief Technologist, who with a group of ex-Apple Computer technologists founded Concept Labs, where he advises a range of companies across the world. He has published widely, holds B.Sc., M.Sc., Ph.D., and D.Sc. degrees from Nottingham (Trent) and Essex Universities, is an Apple Master, and is a visiting professor at London, Essex, and Southampton Universities. He is best known for his incisive and often provocative views on the United Kingdom and world telecommunications industries.

—Edward Smith, BT, UK

edward.a.smith@btinternet.com

Read Any Good Books Lately?

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

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

This publication is distributed on an “as-is” basis, without warranty of any kind either express or implied, including but not limited to the implied warranties of merchantability, fitness for a particular purpose, or non-infringement. This publication could contain technical inaccuracies or typographical errors. Later issues may modify or update information provided in this issue. Neither the publisher nor any contributor shall have any liability to any person for any loss or damage caused directly or indirectly by the information contained herein.

The Internet Protocol Journal

Ole J. Jacobsen, Editor and Publisher

Editorial Advisory Board

Dr. Vint Cerf, VP and Chief Internet Evangelist
Google Inc, USA

Dr. Jon Crowcroft, Marconi Professor of Communications Systems
University of Cambridge, England

David Farber
Distinguished Career Professor of Computer Science and Public Policy
Carnegie Mellon University, USA

Peter Löthberg, Network Architect
Stupi AB, Sweden

Dr. Jun Murai, General Chair Person, WIDE Project
Vice-President, Keio University
Professor, Faculty of Environmental Information
Keio University, Japan

Dr. Deepinder Sidhu, Professor, Computer Science &
Electrical Engineering, University of Maryland, Baltimore County
Director, Maryland Center for Telecommunications Research, USA

Pindar Wong, Chairman and President
Verifi Limited, Hong Kong

*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*

Copyright © 2007 Cisco Systems, Inc. All rights reserved. Cisco, the Cisco logo, and Cisco Systems are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. All other trademarks mentioned in this document or Website are the property of their respective owners.

Printed in the USA on recycled paper.



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
--