Infoblox
CONTROL YOUR NETWORK

WHITEPAPER

The 7 Deadly Traps of IPv6 Deployment – and How to Avoid Them

The 7 Deadly Traps of IPv6 Deployment - and How to Avoid Them

By all accounts, 2011 is a significant year for IT. We will look back and remember that it was the year that we finally exhausted all available public IPv4 addresses[i]. For the present, however, we need to face the fact that this is the year that the clock has begun to tick on your adoption of IPv6.

Like it or not, the transition to IPv6 will happen in your organization. How that transition should occur is a question we sought to answer as we spoke with engineers and IT professionals from across the globe who have already deployed IPv6. In those conversations and in our research, we discerned a rational action plan for deploying IPv6, as well as specific threats and challenges any practitioner will face when rolling out the new addressing scheme.

Most practitioners agreed that organizations should begin the transition by running IPv4 and IPv6 in parallel for the foreseeable future, starting at the perimeter of the network, then the core, and only then the end-nodes. As deployments begin to impact internal networks, care must be paid to vulnerabilities with existing protocols and adoption plans by enterprise equipment. There are many indications that the sheer task of managing the IPv6 address space will challenge many organizations. The expansion to 128-bit addressing will be impossible to manage by spreadsheet alone, and will simply overwhelm and ultimately break some IP address management databases.

IPv4 Addresses Have Run Out

The pool of unallocated addresses has simply run out, and the Regional Internet Registries (RIRs) will exhaust their local pools by the end of 2012, starting as soon as the end of 2011.

Today, all new Internet services will continue to require IPv4 addresses to access the IPv4 part of the Internet, but the registries over time will no longer freely supply those IPv4 addresses. Some IPv4 addresses will be available from service providers, but those supplies will not last long. Consequently, the pressure to find another 200 million will continue to mount. By 2015, at least 17% of the Internet will use IPv6, and more than a quarter of new Internet users (28%) will be running the protocol[ii].

Numerous approaches have been used to extend the life of IP addresses. The two most popular — the use of public Network Address Translation (NAT) and the recovery of unused public addresses — face specific challenges. Early Internet pioneers, such as the government, IBM, AT&T and MIT, were granted large blocks of IP address. Today, many of those addresses remain unused, and there's been a call to move those addresses back into the public domain. That effort is underway and may extend of the availability of IPv4 addresses, but it cannot satisfy the need for a billion new IPv4 addresses over the next five years.

A longer-term possibility is to extend the use of NAT into carrier networks. Cascading NAT or multiple translation points could, in theory, extend the life of IPv4, but that solution comes with numerous problems. For one, the use of NAT introduces significant architectural complexity into the Internet that threatens to break peer-to-peer applications and other basic functionality. What's more, NAT devices need to retain state information on each session. The more sessions, the more resources required — and the more complex and costly the NAT. Adding NAT into the core of carrier networks is likely to become an incredibly expensive and complex undertaking.

Even then, NAT scalability will be constrained by the limitations of today's IPv4 stacks. NAT devices share IPv4 addresses by assigning port numbers to each unique session. With 65,536 ports per address, port limitations are rarely an issue with small networks. However, as networks grow, those limitations quickly become a design consideration. If a single IPv4 address is being shared among 2,000 customers, that leaves about 30 sessions per user. But modern applications run numerous sessions in parallel, far exceeding that limit. According to research conducted by Nippon Telephone and Telegraph (NTT)[iii] and presented to the IETF, a Yahoo page may create as many as 20 sessions, a Google image se arch will create 30 to 60 sessions, and if you view a YouTube video you're using 90 sessions. On average, users typically require 500 sessions, significantly limiting the ability for NAT to support the millions of users that will need to traverse them.

Reasons Why IPv6 Should Concern Your Organization Now.

Common wisdom is that since organizations run private IPv4 addressing schemes today, they can ignore the transition to IPv6 in the foreseeable future. In reality, it's much more complex. IT decisions are impacted not only by technical concerns, but also by longer-term operational and business concerns. This combination will drive many organizations to adopt IPv6 sooner than might be expected.

1. Immediate Need - For equipment and service providers addressing the public sector, the requirement to adopt IPv6 in their equipment and services is immediate. The U.S. Federal Government has an active mandate to adopt the technology. By September 30, 2012, all agencies across the U.S. Federal Government must deploy IPv6 on their public-facing Internet presence. By September 30, 2014, all agencies must upgrade their entire internal infrastructure to leverage the benefits of IPv6. Agencies must also designate an IPv6 Transition Manager and only procure equipment and services that are IPv6-compatible.[iv]

2. Application Requirements - In addition, a new generation of applications and services simply won't function using the IPv4 protocol because of the lack of addresses. These emerging solutions radically expand the notion of Internet application to devices and appliances not commonly found on the Internet today. Sensors, appliance-based controls, power management (smart grid), 4G wireless / LTE — all will remain unconnected without the plethora of addresses provided by IPv6. As these applications are introduced, the network effect will take hold, further driving IPv6 adoption.

3. Loss of critical Web statistics - To serve those users on IPv6, enterprises can't simply address them through IPv4-to-IPv6 translation. Even if the IPv6 devices can reach your website though the use of translation techniques, the enterprises that deliver Internet services will lose key customer information. Various marketing metrics delivered by software tools, such as identification of visitor information, country, and organization, are based on Web server logs that will require native IPv6.

4. Operational risks - Operationally, businesses need to consider the potential risks that are introduced by the depletion of IPv4 addresses and what this will mean to the organization. What plan does your organization have in place to migrate its public-facing Web sites and applications to IPv6? The reduction in available IPv4 address will increasingly become a significant risk to the organization. How long will it be before a Sarbanes Oxley and SAS 70 auditors start requiring disclosure of IPv4 exhaustion as a business continuity risk?

Stages of IPv6 Deployment

It is recommended that the organization adopt a multi-step staged deployment of IPv6:

Stage 1: Implement IPv6 on your external facing Internet presence – Given that there will be a significant introduction of IPv6-only clients in the coming months and years, namely new mobile devices, organizations should support IPv6 on external-facing Web sites, mail servers and other applications. There are two choices on how to achieve this: either run an IPv4-to-IPv6 translation technology, or run both IPv4 and IPv6 protocol stacks on Internet-facing infrastructure.

Translating between the IPv4 protocol and the IPv6 protocol introduces a number of operational concerns. All traffic will need to traverse the protocol translation (PT) device, creating a potential performance bottleneck. Availability is also of concern, and will require redundancy and other high availability and performance features in the PT device. As such, protocol translation introduces significant architectural complexity into the network design.

A better alternative is dual-stack architecture, where the Internet-facing infrastructure runs both IPv4 and IPv6 stacks. With this approach, there is no single point of failure or performance bottleneck with which the IT organization must contend. There are, of course, dual-routing domains that need to be monitored and managed appropriately.

Stage 2: Migrate the core backbone and WAN to dual stack– The organization should consider deploying IPv6 internally on switches and routers by the end of 2012 to achieve parity with industry direction. This means that one must switch over to a dual-stack deployment of all internal switching and routing within 2011, leaving enough time for testing and analysis for IPv6 endpoint deployment which should begin in 2012. WAN migration, though, is more than just evaluating the routing infrastructure. Services must be IPv6-compatible, and, if not, the carrier must articulate an IPv6 migration plan for making them IPv6-compatible. WAN optimization equipment, firewalls, and the infrastructure and security components impacting the WAN must also be IPv6-compatible

Stage 3: Migrate the Intranet to IPv6 - With the routing and switching (core backbone) infrastructure in place, organizations should enable local IPv6 access to the Intranet. Using an IPv6-to-IPv4 protocol transition device can also accomplish this goal. However, for the reasons enumerated above, the preferred deployment route is a dual-stack design.

Stage 4: Implement IPv6 Internet access – The prevalence of IPv6-enabled sites is growing. Back in 2008, just a little more than three percent of Autonomous Systems (AS) announced IPv6 routes. In 2010 that number reached nearly eight percent. Content on the IPv6 network is also starting to grow quickly. Google has announced that YouTube will be IPv6-compatible. Netflix has already demonstrated IPv6 access. Also, eBay and Facebook plan to deploy IPv6-enabled sites. While we can assume that most content will be IPv4-compatible for the foreseeable future, some services will be better suited for IPv6 access, such as connectivity to native IPv6 smart phones and devices.

Stage 5: Enable native IPv6 access to the end client – With the core IPv6 infrastructure in place, organizations should then push native IPv6 access to the edge of their networks. This inflection point requires significant changes in how IPv6 addresses are assigned and managed, which we'll discuss later.

At the very least, though, if we assume that one amortizes the cost of equipment for three to five years, all new equipment purchased today by organizations must be IPv6-compatible. Regardless of how one evaluates the exact cutover to IPv6, for most organizations all the projections anticipate a transition to a significant amount of IPv6 within the next three years. Furthermore, as IPv6 clients grow, a network effect will occur that will likely accelerate the adoption of this new protocol.

The 7 Deadly Traps of IPv6 Deployment - And How to Avoid Them

What should organizations watch for as they develop their approach to IPv6? Clearly, some things are well known. A thorough audit is needed of all equipment and software. Staff education about IPv6 is essential, and networking policies need to be reassessed or put into place. However, in speaking with practitioners that have deployed IPv6, we found seven important lessons they learned that you should put into your cutover plan to make your migration go smoothly

1. Review how you will configure and track IP addresses – The most fundamental change in IPv6 is the impact this new protocol has on IP address assignment and management. With the shift from 32-bit IPv4 addressing to 128-bit IPv6 addressing, organizations must reconsider how they assign and track their IP addresses. While manual tools, such as spreadsheet programs, may have been useful in tracking IPv4-addressed networks, they become impossible to use with IPv6. The sheer length of IPv6 addresses makes human address management impractical, if not impossible. Organizations should move to automated IP Address Management (IPAM) tools that are fully IPv6-compatible, such as IPAM products from Infoblox.

2. Review your DNS architecture – In later stages, when planning for IPv6 on internal networks, organizations will also need to assess the IPv6 readiness of the rest of the IP management infrastructure. Naturally, if you use Dynamic Host Configuration Protocol (DHCP) for address assignment, you need to ensure your DHCP server is IPv6-compliant. But DHCP is only part of the requirements for supporting an IPv6 endpoint. IPv6 endpoint support also requires configuring IPv6 DNS domain support, DNS server addresses, network time server addresses, and more — all of which is related to the DHCP server.

   It is critical that the IT organization run a modern DNS infrastructure that is equipped to deliver both IPv4 (A Records) and IPv6 (AAAA Records). Both DNS and DHCP must also be interoperability-tested to ensure compatibility between the systems.

3. Review your security and maintenance policies – Every organization needs to appreciate and prepare for the likelihood they will need to update their maintenance and security policies when they implement IPv6. The vulnerabilities of the IPv4 stack are well known since they have run wild for so many years, but there simply has not been the same level of experience with IPv6. As a consequence, thorough threat assessment of the new protocol is needed. Organizations, in turn, need to review their security posture as they deploy the new protocols.

4. Inventory your current network infrastructure – It should go without saying that an IPv6 migration can only begin once an organization understands what's actually deployed within their current IPv4-based network. IT must conduct a thorough analysis of its network infrastructure and how traffic is routed in order to deploy IPv6. As you implement IPv6 in each subnet, you need to make sure your path to the backbone is fully enabled for the new protocol or you will have a broken link.

5. Review your application compatibility – Organizations can't assume that their network applications will continue to function as expected on an IPv6 network. They should be tested before a switchover. The changes in the IPv6 stack imply that new TCP layer 4 protocols (TCP6 and UDP6) will be deployed, and that could impact some applications.

   Take VoIP, for example: the popular Asterisk PBX only became compatible with IPv6 in the fall of 2010, according to Timothy Winters, a senior manager at the University of New Hampshire Interoperability Laboratory (UNH-IOL), one of two organizations currently accredited by the National Institute of Standards and Technology (NIST), to perform the U.S. Government IPv6 (USGv6) compliance testing.

6. Update your backend tools – The process of managing and troubleshooting an IPv6 network will require a new set of tools, or at the very least a revision of old ones. At both the administrative and the maintenance levels, organizations need to recheck their existing toolkits to be sure those tools are IPv6-compatible. For example, the sheer length of IPv6 addresses poses problems for some databases that are unable to store IPv6 address. Many IT managers who have neglected to check before converting have reported that analyzers and the other monitoring tools tend not be IPv6-compatible.

7. Monitor network performance – IPv6 introduces changes that may impact network performance. Header size is doubled from that of IPv4, increasing to 40 bytes. In applications that rely on small packet size, there will be a noticeable impact on application performance. SIP, for example, uses small packets, on average about 1000 bytes in length. The header increase will add about 2 percent on to the packet size — not a huge increase but enough to impact extreme cases. Practically, research has shown similar differences in IPv6 performance.[v]

   Performance is also an issue with IPv6 networking equipment. While most system vendors have some sort of IPv6 implementation strategy, the performance of systems running IPv6 protocols is likely to be suspect. The IPv6 protocol is generally supported in the firmware of many network vendors, but it is not yet optimized in silicon, notes Jerry Johnson CIO of the Pacific Northwest Laboratory. Firmware migration is an initial step, but significant deployments of IPv6 internally will require a hardware upgrade in much of the networking infrastructure to maintain the level of performance customers and employees expect today.

The Unforeseen Deadly Trap — A Bonus Warning

When troubleshooting an elusive IT problem, often the last culprit that comes to mind is one that frequently causes headaches and frustration — SPAM. Hackers use this annoyance to slow traffic, clog pipes, infuriate users, and generally interrupt business-as-usual as often as they can. The problem will only grow larger with IPv6, at least in the near term. So we offer this bonus alert, the eighth deadly trap to avoid in migrating to IPv6

8. SPAM tools need to be reinvented – SPAM-blocking solutions have relied heavily on the use of DNS blacklists or blocklists, known as DNSBLs — but DNSBLs will not work effectively in IPv6. With IPv4, the limited supply of IP addresses means that hosts will have at most a few hundred addresses to screen for, so listing and blocking individual addresses is trivial. With IPv6, however, the large number of addresses will enable hackers to allocate thousands of addresses to a server and simply jump from address to address for each new SPAM message.

   Listing IPv6 ranges in the DNSBL, as is done in IPv4, isn't a feasible approach because the ranges are so vast that the caches and DNS servers will be overloaded. What's more, since DNS caches tend to keep the most recent answers around in preference to older ones, the flood of DNSBL data will force all of the other DNS information out of the cache. On most systems, DNSBLs use the same cache as all other DNS queries, so it will also increase the load on every other DNS server, re-fetching answers that were flushed out of the cache. Even if the DNSBL servers use a single DNS wildcard record to cover a large range of DNSBL entries, that won't solve the problem because DNS caches can't tell that a response was created from a wildcard, and so it keeps separate entries for each response.

   And remember, the SPAM hassle is not only yours to deal with. Make sure your SPAM protection vendors have updated their products to address these issues as well.[vi]

How Infoblox can Help

Infoblox provides robust automation solutions for DNS, DHCP and IP Address Management (IPAM), and network change and configuration management to help plan, implement and operate IPv6 networks.

Infoblox capabilities address the IPv6 migration issues discussed in this paper relating to taking inventory, visually mapping, and configuring network equipment. Infoblox will also help you optimize performance on your network and analyze your network for internal and regulatory policy compliance.

Infoblox DNS, DHCP and IP Address Management products provide a dual-stack, appliance-based infrastructure for IPv6-capable DNS delivery and visual IPAM tools for IPv6 address space allocation and management. The IPAM tools automate IPAM procedures to reduce human error associated with complex IPv6 addresses and eliminate repetitive tasks, allowing organizations to easily scale management processes across their enterprise with existing IT staff.

From a network infrastructure point of view, IPv6 will have a significant impact on the traditional tactics of managing the routers, switches and other core devices. The days of using naming conventions to predict where devices are located and how they are connected are quickly drawing to a closely. Infoblox can help organizations like yours automate the discovery, analysis and management of the network infrastructure as you migrate from IPv4 to IPv6.

Using Infoblox products, customers can reduce the risks and costs of IPv6 migration and operate a parallel IPv4 and IPv6 DNS and DHCP network service infrastructure. The table below provides a summary of key Infoblox IPv6 capabilities.
Interested to learn more? Register to see a live Infoblox weekly product demo at: www.infoblox.com/demo/

About Infoblox

Infoblox (NYSE:BLOX) helps customers control their networks. Infoblox solutions help businesses automate complex network control functions to reduce costs and increase security and uptime. Our technology enables automatic discovery, real-time configuration and change management and compliance for network infrastructure, as well as critical network control functions such as DNS, DHCP and IP Address Management (IPAM) for applications and endpoint devices. Infoblox solutions help over 6,500 enterprises and service providers in 25 countries control their networks.

[i] See IPv4 Address Report, Feb 03,2011 (http://www.potaroo.net/tools/ipv4/)
[ii] "IT Market Clock for Enterprise Networking", Gartner Group, September 1, 2010
[iii] See http://www.nttv6.jp/~miyakawa/IETF72/IETF-IAB-TECH-PLENARY-NTT-miyakawa-extended.pdf
[iv] "Memorandum to CIOs for Executive Departments and Agencies", Vivek Kundra, federal chief information officer,
[v] See http://www.ftsm.ukm.my/network/files/m%20khairil_journal%20ijcsns.pdf
[vi] "Why DNS Blacklists Don't Work for IPv6 Networks", by John Levine, IPv6 News. See full post "http://ipv6.net/News/ipv6-news.html "

Infoblox solution for IPv6 migration and management	
IPv6 Capable External DNS	• DNS for IPv6 • Dual Stack DNS Appliance
IPv6 IPAM	• Automated IP Address Management • Role based accessibility • Integrated with DNS/DHCP
Planning tools for Internal IPv6 Migration	• Current Network Equipment Inventory (with OS version running) • Current Network Topology and Connectivity • Current Subnet Inventory
Internal IPv6 Capabilities	• IPv6 IP Address Allocation, Tracking and Reclaiming • IPv6 Subnet Allocation and Tracking • Dual Stack Devices Tracking (Smart Folders) • Reduced Complexity of Dual Stack Environment & IP Address Explosion
IPv6 Network Infrastructure Management	• Automated Network Change and Configuration for IPv6 • Compliance, Policy Enforcement and Auditing

CORPORATE HEADQUARTERS:

+1.408.986.4000

+1.866.463.6256

(toll-free, U.S. and Canada)

info@infoblox.com

www.infoblox.com

EMEA HEADQUARTERS:

+32.3.259.04.30

info-emea@infoblox.com

APAC HEADQUARTERS:

+852.3793.3428

sales-apac@infoblox.com