Internet Engineering Task Force (IETF)
Request for Comments: 7381
Category: Informational
ISSN: 2070-1721

K. Chittimaneni
Dropbox, Inc.
T. Chown
University of Southampton
L. Howard
Time Warner Cable
V. Kuarsingh
Dyn, Inc.
Y. Pouffary
Hewlett Packard
E. Vyncke
Cisco Systems

October 2014

Enterprise IPv6 Deployment Guidelines

Abstract

Enterprise network administrators worldwide are in various stages of preparing for or deploying IPv6 into their networks. The administrators face different challenges than operators of Internet access providers and have reasons for different priorities. The overall problem for many administrators will be to offer Internet-facing services over IPv6 while continuing to support IPv4, and while introducing IPv6 access within the enterprise IT network. The overall transition will take most networks from an IPv4-only environment to a dual-stack network environment and eventually an IPv6-only operating mode. This document helps provide a framework for enterprise network architects or administrators who may be faced with many of these challenges as they consider their IPv6 support strategies.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7381.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

An enterprise network is defined in [RFC4057] as a network that has multiple internal links, one or more router connections to one or more providers, and is actively managed by a network operations entity (the "administrator", whether a single person or a department of administrators). Administrators generally support an internal network, consisting of users' workstations; personal computers; mobile devices; other computing devices and related peripherals; a server network, consisting of accounting and business application servers; and an external network, consisting of Internet-accessible services such as web servers, email servers, VPN systems, and customer applications. This document is intended as guidance for enterprise network architects and administrators in planning their IPv6 deployments.

The business reasons for spending time, effort, and money on IPv6 will be unique to each enterprise. The most common drivers are due to the fact that when Internet service providers, including mobile wireless carriers, run out of IPv4 addresses, they will provide native IPv6 and non-native IPv4. The non-native IPv4 service may be NAT64, NAT444, Dual-Stack Lite (DS-Lite), Mapping of Address and Port using Translation (MAP-T), Mapping of Address and Port using Encapsulation (MAP-E), or other transition technologies. Compared to tunneled or translated service, native traffic typically performs better and more reliably than non-native. For example, for client networks trying to reach enterprise networks, the IPv6 experience will be better than the transitional IPv4 if the enterprise deploys IPv6 in its public-facing services. The native IPv6 network path should also be simpler to manage and, if necessary, troubleshoot. Further, enterprises doing business in growing parts of the world may find IPv6 growing faster there, where again potential new customers, employees, and partners are using IPv6. It is thus in the enterprise's interest to deploy native IPv6 at the very least in its public-facing services but ultimately across the majority or all of its scope.

The text in this document provides specific guidance for enterprise networks and complements other related work in the IETF, including [IPv6-DESIGN] and [RFC5375].

1.1. Enterprise Assumptions

For the purpose of this document, we assume the following:

- The administrator is considering deploying IPv6 (but see Section 1.2 below).
- The administrator has existing IPv4 networks and devices that will continue to operate and be supported.
- The administrator will want to minimize the level of disruption to the users and services by minimizing the number of technologies and functions that are needed to mediate any given application. In other words, provide native IP wherever possible.

Based on these assumptions, an administrator will want to use technologies that minimize the number of flows being tunneled, translated, or intercepted at any given time. The administrator will choose transition technologies or strategies that both allow most traffic to be native and manage non-native traffic. This will allow the administrator to minimize the cost of IPv6 transition technologies by containing the number and scale of transition systems.

Tunnels used for IPv6/IPv4 transition are expected as near-/mid-term mechanisms, while IPv6 tunneling will be used for many long-term operational purposes such as security, routing control, mobility, multihoming, traffic engineering, etc. We refer to the former class of tunnels as "transition tunnels".

1.2. IPv4-Only Considerations

As described in [RFC6302], administrators should take certain steps even if they are not considering IPv6. Specifically, Internet-facing servers should log the source port number, timestamp (from a reliable source), and the transport protocol. This will allow investigation of malefactors behind address-sharing technologies such as NAT444, MAP, or DS Lite. Such logs should be protected for integrity, safeguarded for privacy, and periodically purged within applicable regulations for log retention.

Other IPv6 considerations may impact ostensibly IPv4-only networks, e.g., [RFC6104] describes the rogue IPv6 Router Advertisement (RA) problem, which may cause problems in IPv4-only networks where IPv6 is enabled in end systems on that network. Further discussion of the security implications of IPv6 in IPv4-only networks can be found in [RFC7123].

1.3. Reasons for a Phased Approach

Given the challenges of transitioning user workstations, corporate systems, and Internet-facing servers, a phased approach allows incremental deployment of IPv6, based on the administrator's own determination of priorities. This document outlines suggested phases: a Preparation and Assessment Phase, an Internal Phase, and an External Phase. The Preparation Phase is highly recommended to all administrators, as it will save errors and complexity in later phases. Each administrator must decide whether to begin with an External Phase (enabling IPv6 for Internet-facing systems, as recommended in [RFC5211]) or an Internal Phase (enabling IPv6 for internal interconnections first).

Each scenario is likely to be different to some extent, but we can highlight some considerations:

- In many cases, customers outside the network will have IPv6 before the internal enterprise network. For these customers, IPv6 may well perform better, especially for certain applications, than translated or tunneled IPv4, so the administrator may want to prioritize the External Phase such that those customers have the simplest and most robust connectivity to the enterprise, or at least its external-facing elements.

- Employees who access internal systems by VPN may find that their ISPs provide translated IPv4, which does not support the required VPN protocols. In these cases, the administrator may want to prioritize the External Phase and any other remotely accessible internal systems. It is worth noting that a number of emerging VPN solutions provide dual-stack connectivity; thus, a VPN service may be useful for employees in IPv4-only access networks to access IPv6 resources in the enterprise network (much like many public tunnel broker services, but specifically for the enterprise). Some security considerations are described in [RFC7359].

- Internet-facing servers cannot be managed over IPv6 unless the management systems are IPv6 capable. These might be Network Management Systems (NMS), monitoring systems, or just remote management desktops. Thus, in some cases, the Internet-facing systems are dependent on IPv6-capable internal networks. However, dual-stack Internet-facing systems can still be managed over IPv4.

- Virtual Machines (VMs) may enable a faster rollout once initial system deployment is complete. Management of VMs over IPv6 is still dependent on the management software supporting IPv6.

- IPv6 is enabled by default on all modern operating systems, so it may be more urgent to manage and have visibility on the internal traffic. It is important to manage IPv6 for security purposes, even in an ostensibly IPv4-only network, as described in [RFC7123].

- In many cases, the corporate accounting, payroll, human resource, and other internal systems may only need to be reachable from the internal network, so they may be a lower priority. As enterprises require their vendors to support IPv6, more internal applications will support IPv6 by default, and it can be expected that eventually new applications will only support IPv6. The inventory, as described in Section 2.2, will help determine the systems' readiness, as well as the readiness of the supporting network elements and security, which will be a consideration in prioritization of these corporate systems.

- Some large organizations (even when using private IPv4 addresses [RFC1918]) are facing IPv4 address exhaustion because of the internal network growth (for example, the vast number of VMs) or because of the acquisition of other companies that often raise private IPv4 address overlapping issues.

- IPv6 restores end-to-end transparency even for internal applications (of course security policies must still be enforced). When two organizations or networks merge [RFC6879], the unique addressing of IPv6 can make the merger much easier and faster. A merger may, therefore, prioritize IPv6 for the affected systems.

These considerations are in conflict; each administrator must prioritize according to their company's conditions. It is worth noting that the reasons given in "A Large Corporate User's View of IPng", described in [RFC1687], for reluctance to deploy have largely been satisfied or overcome in the intervening years.

2. Preparation and Assessment Phase

2.1. Program Planning

Since enabling IPv6 is a change to the most fundamental Internet Protocol, and since there are so many interdependencies, having a professional project manager organize the work is highly recommended. In addition, an executive sponsor should be involved in determining the goals of enabling IPv6 (which will establish the order of the phases) and should receive regular updates.

It may be necessary to complete the Preparation Phase before determining whether to prioritize the Internal or External Phase,

since needs and readiness assessments are part of that phase. For a large enterprise, it may take several iterations to really understand the level of effort required. Depending on the required schedule, it may be useful to roll IPv6 projects into other architectural upgrades -- this can be an excellent way to improve the network and reduce costs. However, by increasing the scope of projects, the schedule is often affected. For instance, a major systems upgrade may take a year to complete, where just patching existing systems may take only a few months.

The deployment of IPv6 will not generally stop all other technology work. Once IPv6 has been identified as an important initiative, all projects, both new and in progress, will need to be reviewed to ensure IPv6 support.

It is normal for assessments to continue in some areas while execution of the project begins in other areas. This is fine, as long as recommendations in other parts of this document are considered, especially regarding security (for instance, one should not deploy IPv6 on a system before security has been evaluated).

2.2. Inventory Phase

To comprehend the scope of the Inventory Phase, we recommend dividing the problem space in two: network infrastructure readiness and applications readiness.

2.2.1. Network Infrastructure Readiness Assessment

The goal of this assessment is to identify the level of IPv6 readiness of network equipment. This will identify the effort required to move to an infrastructure that supports IPv6 with the same functional service capabilities as the existing IPv4 network. This may also require a feature comparison and gap analysis between IPv4 and IPv6 functionality on the network equipment and software. IPv6 support will require testing; features often work differently in vendors' labs than production networks. Some devices and software will require IPv4 support for IPv6 to work.

The inventory will show which network devices are already capable, which devices can be made IPv6 ready with a code/firmware upgrade, and which devices will need to be replaced. The data collection consists of a network discovery to gain an understanding of the topology and inventory network infrastructure equipment and code versions with information gathered from static files and IP address management, DNS, and DHCP tools.

Since IPv6 might already be present in the environment, through default configurations or VPNs, an infrastructure assessment (at minimum) is essential to evaluate potential security risks.

2.2.2. Application Readiness Assessment

Just like network equipment, application software needs to support IPv6. This includes OS, firmware, middleware, and applications (including internally developed applications). Vendors will typically handle IPv6 enablement of off-the-shelf products, but often enterprises need to request this support from vendors. For internally developed applications, it is the responsibility of the enterprise to enable them for IPv6. Analyzing how a given application communicates over the network will dictate the steps required to support IPv6. Applications should avoid instructions specific to a given IP address family. Any applications that use APIs, such as the C language, that expose the IP version specifically, need to be modified to also work with IPv6.

There are two ways to IPv6-enable applications. The first approach is to have separate logic for IPv4 and IPv6, thus leaving the IPv4 code path mainly untouched. This approach causes the least disruption to the existing IPv4 logic flow, but introduces more complexity, since the application now has to deal with two logic loops with complex race conditions and error recovery mechanisms between these two logic loops. The second approach is to create a combined IPv4/IPv6 logic, which ensures operation regardless of the IP version used on the network. Knowing whether a given implementation will use IPv4 or IPv6 in a given deployment is a matter of some art; see Source Address Selection [RFC6724] and Happy Eyeballs [RFC6555]. It is generally recommended that the application developer use industry IPv6-porting tools to locate the code that needs to be updated. Some discussion of IPv6 application porting issues can be found in [RFC4038].

2.2.3. Importance of Readiness Validation and Testing

Lastly, IPv6 introduces a completely new way of addressing endpoints, which can have ramifications at the network layer all the way up to the applications. So to minimize disruption during the transition phase, we recommend complete functionality, scalability, and security testing to understand how IPv6 impacts the services and networking infrastructure.

2.3. Training

Many organizations falter in IPv6 deployment because of a perceived training gap. Training is important for those who work with addresses regularly, as with anyone whose work is changing. Better knowledge of the reasons IPv6 is being deployed will help inform the assessment of who needs training and what training they need.

2.4. Security Policy

It is obvious that IPv6 networks should be deployed in a secure way. The industry has learned a lot about network security with IPv4, so network operators should leverage this knowledge and expertise when deploying IPv6. IPv6 is not so different than IPv4: it is a connectionless network protocol using the same lower-layer service and delivering the same service to the upper layer. Therefore, the security issues and mitigation techniques are mostly identical with the same exceptions that are described further.

2.4.1. IPv6 Is No More Secure Than IPv4

Some people believe that IPv6 is inherently more secure than IPv4 because it is new. Nothing can be more wrong. Indeed, being a new protocol means that bugs in the implementations have yet to be discovered and fixed and that few people have the operational security expertise needed to operate securely an IPv6 network. This lack of operational expertise is the biggest threat when deploying IPv6: the importance of training is to be stressed again.

One security myth is that, thanks to its huge address space, a network cannot be scanned by enumerating all IPv6 addresses in a /64 LAN; hence, a malevolent person cannot find a victim. [RFC5157] describes some alternate techniques to find potential targets on a network, for example, enumerating all DNS names in a zone. Additional advice in this area is also given in [HOST-SCANNING].

Another security myth is that IPv6 is more secure because it mandates the use of IPsec everywhere. While the original IPv6 specifications may have implied this, [RFC6434] clearly states that IPsec support is not mandatory. Moreover, if all the intra-enterprise traffic is encrypted, both malefactors and security tools that rely on payload inspection (Intrusion Prevention System (IPS), firewall, Access Control List (ACL), IP Flow Information Export (IPFIX) ([RFC7011] and [RFC7012]), etc.) will be affected. Therefore, IPsec is as useful in IPv6 as in IPv4 (for example, to establish a VPN overlay over a non-trusted network or to reserve for some specific applications).

The last security myth is that amplification attacks (such as [SMURF]) do not exist in IPv6 because there is no more broadcast. Alas, this is not true as ICMP error (in some cases) or information messages can be generated by routers and hosts when forwarding or receiving a multicast message (see Section 2.4 of [RFC4443]). Therefore, the generation and the forwarding rate of ICMPv6 messages must be limited as in IPv4.

It should be noted that in a dual-stack network, the security implementation for both IPv4 and IPv6 needs to be considered, in addition to security considerations related to the interaction of (and transition between) the two, while they coexist.

2.4.2. Similarities between IPv6 and IPv4 Security

As mentioned earlier, IPv6 is quite similar to IPv4; therefore, several attacks apply for both protocol families, including:

- Application layer attacks: such as cross-site scripting or SQL injection
- Rogue device: such as a rogue Wi-Fi access point
- Flooding and all traffic-based denial of services (including the use of control plane policing for IPv6 traffic: see [RFC6192])

A specific case of congruence is IPv6 Unique Local Addresses (ULAs) [RFC4193] and IPv4 private addressing [RFC1918], which do not provide any security by 'magic'. In both cases, the edge router must apply strict filters to block those private addresses from entering and, just as importantly, leaving the network. This filtering can be done by the enterprise or by the ISP, but the cautious administrator will prefer to do it in the enterprise.

IPv6 addresses can be spoofed as easily as IPv4 addresses, and there are packets with bogon IPv6 addresses (see [CYMRU]). Anti-bogon filtering must be done in the data and routing planes. It can be done by the enterprise or by the ISP, or both, but again the cautious administrator will prefer to do it in the enterprise.

2.4.3. Specific Security Issues for IPv6

Even if IPv6 is similar to IPv4, there are some differences that create some IPv6-only vulnerabilities or issues. We give examples of such differences in this section.

Privacy extension addresses [RFC4941] are usually used to protect individual privacy by periodically changing the interface identifier

part of the IPv6 address to avoid tracking a host by its otherwise always identical and unique 64-bit Extended Unique Identifier (EUI-64) based on Media Access Control (MAC). While this presents a real advantage on the Internet, moderated by the fact that the prefix part remains the same, it complicates the task of following an audit trail when a security officer or network operator wants to trace back a log entry to a host in their network because when the tracing is done, the searched IPv6 address could have disappeared from the network. Therefore, the use of privacy extension addresses usually requires additional monitoring and logging of the binding of the IPv6 address to a data-link layer address (see also the monitoring section in [IPv6-SECURITY], Section 2.5). Some early enterprise deployments have taken the approach of using tools that harvest IP/MAC address mappings from switch and router devices to provide address accountability; this approach has been shown to work, though it can involve gathering significantly more address data than in equivalent IPv4 networks. An alternative is to try to prevent the use of privacy extension addresses by enforcing the use of DHCPv6, such that hosts only get addresses assigned by a DHCPv6 server. This can be done by configuring routers to set the M bit in RAs, combined with all advertised prefixes being included without the A bit set (to prevent the use of stateless autoconfiguration). Of course, this technique requires that all hosts support stateful DHCPv6. It is important to note that not all operating systems exhibit the same behavior when processing RAs with the M bit set. The varying OS behavior is related to the lack of prescriptive definition around the A, M, and O bits within the Neighbor Discovery Protocol (NDP). [DHCPv6-SLAAC-PROBLEM] provides a much more detailed analysis on the interaction of the M bit and DHCPv6.

Extension headers complicate the task of stateless packet filters such as ACLs. If ACLs are used to enforce a security policy, then the enterprise must verify whether its ACLs (but also stateful firewalls) are able to process extension headers (this means understand them enough to parse them to find the upper-layer payloads) and to block unwanted extension headers (e.g., to implement [RFC5095]). This topic is discussed further in [RFC7045].

Fragmentation is different in IPv6 because it is done only by the source host and never during a forwarding operation. This means that ICMPv6 packet-too-big messages must be allowed to pass through the network and not be filtered [RFC4890]. Fragments can also be used to evade some security mechanisms such as RA-Guard [RFC6105]. See also [RFC5722] and [RFC7113].

One of the biggest differences between IPv4 and IPv6 is the introduction of NDP [RFC4861], which includes a variety of important IPv6 protocol functions, including those provided in IPv4 by the

Address Resolution Protocol (ARP) [RFC0826]. NDP runs over ICMPv6 (which as stated above means that security policies must allow some ICMPv6 messages to pass, as described in RFC 4890), but has the same lack of security as, for example, ARP, in that there is no inherent message authentication. While Secure Neighbor Discovery (SEND) [RFC3971] and Cryptographically Generated Addresses (CGAs) [RFC3972] have been defined, they are not widely implemented). The threat model for RAs within the NDP suite is similar to that of DHCPv4 (and DHCPv6), in that a rogue host could be either a rogue router or a rogue DHCP server. An IPv4 network can be made more secure with the help of DHCPv4 snooping in edge switches, and likewise RA snooping can improve IPv6 network security (in IPv4-only networks as well). Thus, enterprises using such techniques for IPv4 should use the equivalent techniques for IPv6, including RA-Guard [RFC6105] and all work in progress from the Source Address Validation Improvement (SAVI) WG, e.g., [RFC6959], which is similar to the protection given by dynamic ARP monitoring in IPv4. Other DoS vulnerabilities are related to NDP cache exhaustion, and mitigation techniques can be found in ([RFC6583]).

As stated previously, running a dual-stack network doubles the attack exposure as a malevolent person has now two attack vectors: IPv4 and IPv6. This simply means that all routers and hosts operating in a dual-stack environment with both protocol families enabled (even if by default) must have a congruent security policy for both protocol versions. For example, permit TCP ports 80 and 443 to all web servers and deny all other ports to the same servers must be implemented both for IPv4 and IPv6. It is thus important that the tools available to administrators readily support such behavior.

2.5. Routing

An important design choice to be made is what IGP is to use inside the network. A variety of IGPs (IS-IS, OSPFv3, and Routing Information Protocol Next Generation (RIPng)) support IPv6 today, and picking one over the other is a design choice that will be dictated mostly by existing operational policies in an enterprise network. As mentioned earlier, it would be beneficial to maintain operational parity between IPv4 and IPv6; therefore, it might make sense to continue using the same protocol family that is being used for IPv4. For example, in a network using OSPFv2 for IPv4, it might make sense to use OSPFv3 for IPv6. It is important to note that although OSPFv3 is similar to OSPFv2, they are not the same. On the other hand, some organizations may chose to run different routing protocols for different IP versions. For example, one may chose to run OSPFv2 for IPv4 and IS-IS for IPv6. An important design question to consider here is whether to support one IGP or two different IGPs in the longer term. [IPv6-DESIGN] presents advice on the design choices

that arise when considering IGPs and discusses the advantages and disadvantages to different approaches in detail.

2.6. Address Plan

The most common problem encountered in IPv6 networking is in applying the same principles of conservation that are so important in IPv4. IPv6 addresses do not need to be assigned conservatively. In fact, a single, larger allocation is considered more conservative than multiple non-contiguous small blocks because a single block occupies only a single entry in a routing table. The advice in [RFC5375] is still sound and is recommended to the reader. If considering ULAs, give careful thought to how well it is supported, especially in multiple address and multicast scenarios, and assess the strength of the requirement for ULA. [ULA-USAGE] provides much more detailed analysis and recommendations on the usage of ULAs.

The enterprise administrator will want to evaluate whether the enterprise will request address space from a Local Internet Registry (LIR) such as an ISP; a Regional Internet Registry (RIR) such as AfriNIC, APNIC, ARIN, LACNIC, or RIPE-NCC; or a National Internet Registry (NIR) operated in some countries. The normal allocation is Provider-Aggregated (PA) address space from the enterprise's ISP, but use of PA space implies renumbering when changing providers. Instead, an enterprise may request Provider-Independent (PI) space; this may involve an additional fee, but the enterprise may then be better able to be multihomed using that prefix and will avoid a renumbering process when changing ISPs (though it should be noted that renumbering caused by outgrowing the space, merger, or other internal reason would still not be avoided with PI space).

The type of address selected (PI vs. PA) should be congruent with the routing needs of the enterprise. The selection of address type will determine if an operator will need to apply new routing techniques and may limit future flexibility. There is no right answer, but the needs of the External Phase may affect what address type is selected.

Each network location or site will need a prefix assignment. Depending on the type of site/location, various prefix sizes may be used. In general, historical guidance suggests that each site should get at least a /48, as documented in RFC 5375 and [RFC6177]. In addition to allowing for simple planning, this can allow a site to use its prefix for local connectivity, should the need arise, and if the local ISP supports it.

When assigning addresses to end systems, the enterprise may use manually configured addresses (common on servers) or Stateless Address Autoconfiguration (SLAAC) or DHCPv6 for client systems.

Early IPv6 enterprise deployments have used SLAAC both for its simplicity and the time DHCPv6 has taken to mature. However, DHCPv6 is now very mature; thus, workstations managed by an enterprise may use stateful DHCPv6 for addressing on corporate LAN segments. DHCPv6 allows for the additional configuration options often employed by enterprise administrators, and by using stateful DHCPv6, administrators correlating system logs know which system had which address at any given time. Such an accountability model is familiar from IPv4 management, though DHCPv6 hosts are identified by a DHCP Unique Identifier (DUID) rather than a MAC address. For equivalent accountability with SLAAC (and potentially privacy addresses), a monitoring system that harvests IP/MAC mappings from switch and router equipment could be used.

A common deployment consideration for any enterprise network is how to get host DNS records updated. Commonly, either the host will send DNS updates or the DHCP server will update records. If there is sufficient trust between the hosts and the DNS server, the hosts may update (and the enterprise may use SLAAC for addressing). Otherwise, the DHCPv6 server can be configured to update the DNS server. Note that an enterprise network with this more controlled environment will need to disable SLAAC on network segments and force end hosts to use DHCPv6 only.

In the data center or server room, assume a /64 per VLAN. This applies even if each individual system is on a separate VLAN. In a /48 assignment, typical for a site, there are then still 65,535 /64 blocks. Some administrators reserve a /64 but configure a small subnet, such as /112, /126, or /127, to prevent rogue devices from attaching and getting numbers; an alternative is to monitor traffic for surprising addresses or Neighbor Discovery (ND) tables for new entries. Addresses are either configured manually on the server or reserved on a DHCPv6 server, which may also synchronize forward and reverse DNS (though see [RFC6866] for considerations on static addressing). SLAAC is not recommended for servers because of the need to synchronize RA timers with DNS Times to Live (TTLs) so that the DNS entry expires at the same time as the address.

All user access networks should be a /64. Point-to-point links where NDP is not used may also utilize a /127 (see [RFC6164]).

Plan to aggregate at every layer of network hierarchy. There is no need for variable length subnet mask (VLSM) [RFC1817] in IPv6, and addressing plans based on conservation of addresses are shortsighted. Use of prefixes longer then /64 on network segments will break common IPv6 functions such as SLAAC [RFC4862]. Where multiple VLANs or other Layer 2 domains converge, allow some room for expansion. Renumbering due to outgrowing the network plan is a nuisance, so

allow room within it. Generally, plan to grow to about twice the current size that can be accommodated; where rapid growth is planned, allow for twice that growth. Also, if DNS (or reverse DNS) authority may be delegated to others in the enterprise, assignments need to be on nibble boundaries (that is, on a multiple of 4 bits, such as /64, /60, /56, ..., /48, /44), to ensure that delegated zones align with assigned prefixes.

If using ULAs, it is important to note that AAAA and PTR records for ULAs are not recommended to be installed in the global DNS. Similarly, reverse (address-to-name) queries for ULA must not be sent to name servers outside of the organization, due to the load that such queries would create for the authoritative name servers for the ip6.arpa zone. For more details, please refer to Section 4.4 of [RFC4193].

Enterprise networks are increasingly including virtual networks where a single, physical node may host many virtualized addressable devices. It is imperative that the addressing plans assigned to these virtual networks and devices be consistent and non-overlapping with the addresses assigned to real networks and nodes. For example, a virtual network established within an isolated lab environment may, at a later time, become attached to the production enterprise network.

2.7. Tools Assessment

Enterprises will often have a number of operational tools and support systems that are used to provision, monitor, manage, and diagnose the network and systems within their environment. These tools and systems will need to be assessed for compatibility with IPv6. The compatibility may be related to the addressing and connectivity of various devices as well as IPv6 awareness of the tools and processing logic.

The tools within the organization fall into two general categories: those that focus on managing the network and those that are focused on managing systems and applications on the network. In either instance, the tools will run on platforms that may or may not be capable of operating in an IPv6 network. This lack in functionality may be related to operating system version or based on some hardware constraint. Those systems that are found to be incapable of utilizing an IPv6 connection, or which are dependent on an IPv4 stack, may need to be replaced or upgraded.

In addition to devices working on an IPv6 network natively, or via a transition tunnel, many tools and support systems may require additional software updates to be IPv6 aware or even a hardware

upgrade (usually for additional memory, IPv6 addresses are larger and for a while, IPv4 and IPv6 addresses will coexist in the tool). This awareness may include the ability to manage IPv6 elements and/or applications in addition to the ability to store and utilize IPv6 addresses.

Considerations when assessing the tools and support systems may include the fact that IPv6 addresses are significantly larger than IPv4, requiring data stores to support the increased size. Such issues are among those discussed in [RFC5952]. Many organizations may also run dual-stack networks; therefore, the tools need to not only support IPv6 operation but may also need to support the monitoring, management, and intersection with both IPv6 and IPv4 simultaneously. It is important to note that managing IPv6 is not just constrained to using large IPv6 addresses, but also that IPv6 interfaces and nodes are likely to use two or more addresses as part of normal operation. Updating management systems to deal with these additional nuances will likely consume time and considerable effort.

For networking systems, like node management systems, it is not always necessary to support local IPv6 addressing and connectivity. Operations such as SNMP MIB polling can occur over IPv4 transport while seeking responses related to IPv6 information. Where this may seem advantageous to some, it should be noted that without local IPv6 connectivity, the management system may not be able to perform all expected functions -- such as reachability and service checks.

Organizations should be aware that changes to older IPv4-only SNMP MIB specifications have been made by the IETF and are related to legacy operation in [RFC2096] and [RFC2011]. Updated specifications are now available in [RFC4292] and [RFC4293] that modified the older MIB framework to be IP protocol agnostic, supporting both IPv4 and IPv6. Polling systems will need to be upgraded to support these updates as well as the end stations, which are polled.

3. External Phase

The External Phase for enterprise IPv6 adoption covers topics that deal with how an organization connects its infrastructure to the external world. These external connections may be toward the Internet at large or to other networks. The External Phase covers connectivity, security and monitoring of various elements, and outward-facing or accessible services.

3.1. Connectivity

The enterprise will need to work with one or more service providers to gain connectivity to the Internet or transport service infrastructure such as a BGP/MPLS IP VPN as described in [RFC4364] and [RFC4659]. One significant factor that will guide how an organization may need to communicate with the outside world will involve the use of PI and/or PA IPv6 space.

Enterprises should be aware that, depending on which address type they selected (PI vs. PA) in their planning phase, they may need to implement new routing functions and/or behaviors to support their connectivity to the ISP. In the case of PI, the upstream ISP may offer options to route the prefix (typically a /48) on the enterprise's behalf and update the relevant routing databases. Otherwise, the enterprise may need to perform this task on their own and use BGP to inject the prefix into the global BGP system.

Note that the rules set by the RIRs for an enterprise acquiring PI address space have changed over time. For example, in the European region, the RIPE-NCC no longer requires an enterprise to be multihomed to be eligible for an IPv6 PI allocation. Requests can be made directly or via a LIR. It is possible that the rules may change again and may vary between RIRs.

When seeking IPv6 connectivity to a service provider, native IPv6 connectivity is preferred since it provides the most robust and efficient form of connectivity. If native IPv6 connectivity is not possible due to technical or business limitations, the enterprise may utilize readily available transition tunnel IPv6 connectivity. There are IPv6 transit providers that provide robust tunneled IPv6 connectivity that can operate over IPv4 networks. It is important to understand the transition-tunnel mechanism used and to consider that it will have higher latency than native IPv4 or IPv6, and may have other problems, e.g., related to MTUs.

It is important to evaluate MTU considerations when adding IPv6 to an existing IPv4 network. It is generally desirable to have the IPv6 and IPv4 MTU congruent to simplify operations (so the two address families behave similarly, that is, as expected). If the enterprise uses transition tunnels inside or externally for IPv6 connectivity, then modification of the MTU on hosts/routers may be needed as mid-stream fragmentation is no longer supported in IPv6. It is preferred that Path MTU Discovery (pMTUD) be used to optimize the MTU, so erroneous filtering of the related ICMPv6 message types should be monitored. Adjusting the MTU may be the only option if undesirable upstream ICMPv6 filtering cannot be removed.

3.2. Security

The most important part of security for external IPv6 deployment is filtering and monitoring. Filtering can be done by stateless ACLs or a stateful firewall. The security policies must be consistent for IPv4 and IPv6 (or else the attacker will use the less-protected protocol stack), except that certain ICMPv6 messages must be allowed through and to the filtering device (see [RFC4890]):

- Packet Too Big: essential to allow Path MTU discovery to work
- Parameter Problem
- Time Exceeded

In addition, NDP messages (including Neighbor Solicitation, RAs, etc.) are required for local hosts.

It could also be safer to block all fragments where the transport layer header is not in the first fragment to avoid attacks as described in [RFC5722]. Some filtering devices allow this filtering. Ingress filters and firewalls should follow [RFC5095] in handling routing extension header type 0, dropping the packet and sending ICMPv6 Parameter Problem, unless Segments Left = 0 (in which case, ignore the header).

If an IPS is used for IPv4 traffic, then an IPS should also be used for IPv6 traffic. In general, make sure IPv6 security is at least as good as IPv4. This also includes all email content protection (anti-spam, content filtering, data leakage prevention, etc.).

The edge router must also implement anti-spoofing techniques based on [RFC2827] (also known as BCP 38).

In order to protect the networking devices, it is advised to implement control plane policing as per [RFC6192].

The potential NDP cache exhaustion attack (see [RFC6583]) can be mitigated by two techniques:

- Good NDP implementation with memory utilization limits as well as rate limiters and prioritization of requests.
- Or, as the external deployment usually involves just a couple of exposed statically configured IPv6 addresses (virtual addresses of web, email, and DNS servers), then it is straightforward to build an ingress ACL allowing traffic for those addresses and denying traffic to any other addresses. This actually prevents the attack

as a packet for a random destination will be dropped and will never trigger a neighbor resolution.

3.3. Monitoring

Monitoring the use of the Internet connectivity should be done for IPv6 as it is done for IPv4. This includes the use of IPFIX [RFC7012] to report abnormal traffic patterns (such as port scanning, SYN flooding, and related IP source addresses) from monitoring tools and evaluating data read from SNMP MIBs [RFC4293] (some of which also enable the detection of abnormal bandwidth utilization) and syslogs (finding server and system errors). Where NetFlow is used, Version 9 is required for IPv6 support. Monitoring systems should be able to examine IPv6 traffic, use IPv6 for connectivity, and record IPv6 addresses, and any log parsing tools and reporting need to support IPv6. Some of this data can be sensitive (including personally identifiable information) and care in securing it should be taken, with periodic purges. Integrity protection on logs and sources of log data is also important to detect unusual behavior (misconfigurations or attacks). Logs may be used in investigations, which depend on trustworthy data sources (tamper resistant).

In addition, monitoring of external services (such as web sites) should be made address specific, so that people are notified when either the IPv4 or IPv6 version of a site fails.

3.4. Servers and Applications

The path to the servers accessed from the Internet usually involves security devices (firewall and IPS), server load balancing (SLB), and real physical servers. The latter stage is also multi-tiered for scalability and security between presentation and data storage. The ideal transition is to enable native dual stack on all devices; but as part of the phased approach, operators have used the following techniques with success:

- Use a network device to apply NAT64 and basically translate an inbound TCP connection (or any other transport protocol) over IPv6 into a TCP connection over IPv4. This is the easiest to deploy as the path is mostly unchanged, but it hides all IPv6 remote users behind a single IPv4 address, which leads to several audit trail and security issues (see [RFC6302]).

- Use the server load balancer, which acts as an application proxy to do this translation. Compared to the NAT64, it has the potential benefit of going through the security devices as native IPv6 (so more audit and trace abilities) and is also able to insert an HTTP X-Forward-For header that contains the remote IPv6

address. The latter feature allows for logging and rate limiting on the real servers based on the IPV6 address even if those servers run only IPv4.

In either of these cases, care should be taken to secure logs for privacy reasons and to periodically purge them.

3.5. Network Prefix Translation for IPv6

Network Prefix Translation for IPv6, or NPTv6 as described in [RFC6296], provides a framework to utilize prefix ranges within the internal network that are separate (address independent) from the assigned prefix from the upstream provider or registry. As mentioned above, while NPTv6 has potential use cases in IPv6 networks, the implications of its deployment need to be fully understood, particularly where any applications might embed IPv6 addresses in their payloads.

Use of NPTv6 can be chosen independently from how addresses are assigned and routed within the internal network, how prefixes are routed towards the Internet, or whether PA or PI addresses are used.

4. Internal Phase

This phase deals with the delivery of IPv6 to the internal user-facing side of the Information Technology (IT) infrastructure, which comprises various components such as network devices (routers, switches, etc.), end-user devices and peripherals (workstations, printers, etc.), and internal corporate systems.

An important design paradigm to consider during this phase is "dual stack when you can, tunnel when you must". Dual stacking allows a more robust, production-quality IPv6 network than is typically facilitated by internal use of transition tunnels that are harder to troubleshoot and support, and that may introduce scalability and performance issues. Of course, tunnels may still be used in production networks, but their use needs to be carefully considered, e.g., where the transition tunnel may be run through a security or filtering device. Tunnels do also provide a means to experiment with IPv6 and gain some operational experience with the protocol. [RFC4213] describes various transition mechanisms in more detail. [RFC6964] suggests operational guidance when using Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) tunnels [RFC5214], though we would recommend use of dual stack wherever possible.

4.1. Security

IPv6 must be deployed in a secure way. This means that all existing IPv4 security policies must be extended to support IPv6; IPv6 security policies will be the IPv6 equivalent of the existing IPv4 ones (taking into account the difference for ICMPv6 [RFC4890]). As in IPv4, security policies for IPv6 will be enforced by firewalls, ACL, IPS, VPN, and so on.

Privacy extension addresses [RFC4941] raise a challenge for an audit trail as explained in Section 2.4.3 of this document. The enterprise may choose to attempt to enforce use of DHCPv6 or deploy monitoring tools that harvest accountability data from switches and routers (thus making the assumption that devices may use any addresses inside the network).

One major issue is threats against ND. This means, for example, that the internal network at the access layer (where hosts connect to the network over wired or wireless) should implement RA-Guard [RFC6105] and the techniques being specified by the SAVI WG [RFC6959]; see also Section 2.4.3 of this document for more information.

4.2. Network Infrastructure

The typical enterprise network infrastructure comprises a combination of the following network elements -- wired access switches, wireless access points, and routers (although it is fairly common to find hardware that collapses switching and routing functionality into a single device). Basic wired access switches and access points operate only at the physical and link layers and don’t really have any special IPv6 considerations other than being able to support IPv6 addresses themselves for management purposes. In many instances, these devices possess a lot more intelligence than simply switching packets. For example, some of these devices help assist with link-layer security by incorporating features such as ARP inspection and DHCP snooping, or they may help limit where multicast floods by using IGMP (or, in the case of IPv6, Multicast Listener Discovery (MLD)) snooping.

Another important consideration in enterprise networks is first-hop router redundancy. This directly ties into network reachability from an end host’s point of view. IPv6 ND [RFC4861] provides a node with the capability to maintain a list of available routers on the link, in order to be able to switch to a backup path should the primary be unreachable. By default, ND will detect a router failure in 38 seconds and cycle onto the next default router listed in its cache. While this feature provides a basic level of first-hop router redundancy, most enterprise IPv4 networks are designed to fail over

much faster. Although this delay can be improved by adjusting the default timers, care must be taken to protect against transient failures and to account for increased traffic on the link. Another option in which to provide robust first-hop redundancy is to use the Virtual Router Redundancy Protocol Version 3 (VRRPv3) for IPv6 [RFC5798]. This protocol provides a much faster switchover to an alternate default router than default ND parameters. Using VRRPv3, a backup router can take over for a failed default router in around three seconds (using VRRPv3 default parameters). This is done without any interaction with the hosts and a minimum amount of VRRP traffic.

Last but not least, one of the most important design choices to make while deploying IPv6 on the internal network is whether to use SLAAC [RFC4862], the Dynamic Host Configuration Protocol for IPv6 (DHCPv6) [RFC3315], or a combination thereof. Each option has advantages and disadvantages, and the choice will ultimately depend on the operational policies that guide each enterprise's network design. For example, if an enterprise is looking for ease of use, rapid deployment, and less administrative overhead, then SLAAC makes more sense for workstations. Manual or DHCPv6 assignments are still needed for servers, as described in the Address Plan and External Phase sections of this document; see Sections 2.6 and 3, respectively. However, if the operational policies call for precise control over IP address assignment for auditing, then DHCPv6 may be preferred. DHCPv6 also allows you to tie into DNS systems for host entry updates and gives you the ability to send other options and information to clients. It is worth noting that in general operation, RAs are still needed in DHCPv6 networks, as there is no DHCPv6 Default Gateway option. Similarly, DHCPv6 is needed in RA networks for other configuration information, e.g., NTP servers or, in the absence of support for DNS resolvers in RAs [RFC6106], DNS resolver information.

4.3. End-User Devices

Most operating systems (OSes) that are loaded on workstations and laptops in a typical enterprise support IPv6 today. However, there are various out-of-the-box nuances that one should be mindful about. For example, the default behavior of OSes vary; some may have IPv6 turned off by default, some may only have certain features such as privacy extensions to IPv6 addresses (RFC 4941) turned off, while others have IPv6 fully enabled. Further, even when IPv6 is enabled, the choice of which address is used may be subject to source address selection (RFC 6724) and Happy Eyeballs (RFC 6555). Therefore, it is advised that enterprises investigate the default behavior of their installed OS base and account for it during the Inventory Phases of their IPv6 preparations. Furthermore, some OSes may have some

transition tunneling mechanisms turned on by default, and in such cases, it is recommended to administratively shut down such interfaces unless required.

It is important to note that it is recommended that IPv6 be deployed at the network and system infrastructure level before it is rolled out to end-user devices; ensure IPv6 is running and routed on the wire, and secure and correctly monitored, before exposing IPv6 to end users.

Smartphones and tablets are significant IPv6-capable platforms, depending on the support of the carrier's data network.

IPv6 support for peripherals varies. Much like servers, printers are generally configured with a static address (or DHCP reservation) so clients can discover them reliably.

4.4. Corporate Systems

No IPv6 deployment will be successful without ensuring that all the corporate systems that an enterprise uses as part of its IT infrastructure support IPv6. Examples of such systems include, but are not limited to, email, video conferencing, telephony (VoIP), DNS, RADIUS, etc. All these systems must have their own detailed IPv6 rollout plan in conjunction with the network IPv6 rollout. It is important to note that DNS is one of the main anchors in an enterprise deployment, since most end hosts decide whether or not to use IPv6 depending on the presence of IPv6 AAAA records in a reply to a DNS query. It is recommended that system administrators selectively turn on AAAA records for various systems as and when they are IPv6 enabled; care must be taken though to ensure all services running on that host name are IPv6 enabled before adding the AAAA record. Care with web proxies is advised; a mismatch in the level of IPv6 support between the client, proxy, and server can cause communication problems. All monitoring and reporting tools across the enterprise will need to be modified to support IPv6.

5. IPv6 Only

Early IPv6 enterprise deployments have generally taken a dual-stack approach to enabling IPv6, i.e., the existing IPv4 services have not been turned off. Although IPv4 and IPv6 networks will coexist for a long time, the long-term enterprise network roadmap should include steps to simplify engineering and operations by deprecating IPv4 from the dual-stack network. In some extreme cases, deploying dual-stack networks may not even be a viable option for very large enterprises due to the address space described in RFC 1918 not being large enough to support the network's growth. In such cases, deploying IPv6-only

networks might be the only choice available to sustain network growth. In other cases, there may be elements of an otherwise dual-stack network that may be run in IPv6 only.

If nodes in the network don't need to talk to an IPv4-only node, then deploying IPv6-only networks should be straightforward. However, most nodes will need to communicate with some IPv4-only nodes; an IPv6-only node may, therefore, require a translation mechanism. As [RFC6144] points out, it is important to look at address translation as a transition strategy towards running an IPv6-only network.

There are various stateless and stateful IPv4/IPv6 translation methods available today that help IPv6-to-IPv4 communication. RFC 6144 provides a framework for IPv4/IPv6 translation and describes in detail various scenarios in which such translation mechanisms could be used. [RFC6145] describes stateless address translation. In this mode, a specific IPv6 address range will represent IPv4 systems (IPv4-converted addresses), and the IPv6 systems have addresses (IPv4-translatable addresses) that can be algorithmically mapped to a subset of the service provider's IPv4 addresses. NAT64 [RFC6146] describes stateful address translation. As the name suggests, the translation state is maintained between IPv4 address/port pairs and IPv6 address/port pairs, enabling IPv6 systems to open sessions with IPv4 systems. DNS64 [RFC6147] describes a mechanism for synthesizing AAAA resource records (RRs) from A RRs. Together, RFCs 6146 and RFC 6147 provide a viable method for an IPv6-only client to initiate communications to an IPv4-only server. As described in Enterprise Assumptions, Section 1.1, the administrator will usually want most traffic or flows to be native and only translate as needed.

The address translation mechanisms for the stateless and stateful translations are defined in [RFC6052]. It is important to note that both of these mechanisms have limitations as to which protocols they support. For example, RFC 6146 only defines how stateful NAT64 translates unicast packets carrying TCP, UDP, and ICMP traffic only. The classic problems of IPv4 NAT also apply, e.g., handling IP literals in application payloads. The ultimate choice of which translation mechanism to chose will be dictated mostly by existing operational policies pertaining to application support, logging requirements, etc.

There is additional work being done in the area of address translation to enhance and/or optimize current mechanisms. For example, [DIVI] describes limitations with the current stateless translation, such as IPv4 address sharing and application layer gateway (ALG) problems, and presents the concept and implementation of dual-stateless IPv4/IPv6 translation (dIVI) to address those issues.

It is worth noting that for IPv6-only access networks that use technologies such as NAT64, the more content providers (and enterprises) that make their content available over IPv6, the less the requirement to apply NAT64 to traffic leaving the access network. This particular point is important for enterprises that may start their IPv6 deployment well into the global IPv6 transition. As time progresses, and given the current growth in availability of IPv6 content, IPv6-only operation using NAT64 to manage some flows will become less expensive to run versus the traditional NAT44 deployments since only IPv6-to-IPv4 flows need translation. [RFC6883] provides guidance and suggestions for Internet Content Providers and Application Service Providers in this context.

Enterprises should also be aware that networks may be subject to future convergence with other networks (i.e., mergers, acquisitions, etc.). An enterprise considering IPv6-only operation may need to be aware that additional transition technologies and/or connectivity strategies may be required depending on the level of IPv6 readiness and deployment in the merging networking.

6. Considerations for Specific Enterprises

6.1. Content Delivery Networks

Some guidance for Internet Content and Application Service Providers can be found in [RFC6883], which includes a dedicated section on Content Delivery Networks (CDNs). An enterprise that relies on a CDN to deliver a 'better' e-commerce experience needs to ensure that their CDN provider also supports IPv4/IPv6 traffic selection so that they can ensure 'best' access to the content. A CDN could enable external IPv6 content delivery even if the enterprise provides that content over IPv4.

6.2. Data Center Virtualization

IPv6 Data Center considerations are described in [IPv6-DC].

6.3. University Campus Networks

A number of campus networks around the world have made some initial IPv6 deployments. This has been encouraged by their National Research and Education Network (NREN) backbones, having made IPv6 available natively since the early 2000's. Universities are a natural place for IPv6 deployment to be considered at an early stage, perhaps compared to other enterprises, as they are involved by their very nature in research and education.

Campus networks can deploy IPv6 at their own pace; there is no need to deploy IPv6 across the entire enterprise from day one. Rather, specific projects can be identified for an initial deployment that are both deep enough to give the university experience but small enough to be a realistic first step. There are generally three areas in which such deployments are currently made.

In particular, those initial areas commonly approached are:

- External-facing services. Typically, the campus web presence and commonly also external-facing DNS and mail exchange (MX) services. This ensures early IPv6-only adopters elsewhere can access the campus services as simply and as robustly as possible.

- Computer science department. This is where IPv6-related research and/or teaching is most likely to occur, and where many of the next generation of network engineers are studying, so enabling some or all of the campus computer science department network is a sensible first step.

- The eduroam wireless network. Eduroam [EDUROAM] is the de facto wireless roaming system for academic networks and uses authentication based on 802.1X, which is agnostic to the IP version used (unlike web-redirection gateway systems). Making a campus' eduroam network dual stack is a very viable early step.

The general IPv6 deployment model in a campus enterprise will still follow the general principles described in this document. While the above early stage projects are commonly followed, these still require the campus to acquire IPv6 connectivity and address space from their NREN (or other provider in some parts of the world) and to enable IPv6 on the wire on at least part of the core of the campus network. This implies a requirement to have an initial address plan, and to ensure appropriate monitoring and security measures are in place, as described elsewhere in this document.

Campuses that have deployed to date do not use ULAs, nor do they use NPTv6. In general, campuses have very stable PA-based address allocations from their NRENs (or their equivalent). However, campus enterprises may consider applying for IPv6 PI; some have already done so. The discussions earlier in this text about PA vs. PI still apply.

Finally, campuses may be more likely than many other enterprises to run multicast applications, such as IP TV or live lecture or seminar streaming, so they may wish to consider support for specific IPv6 multicast functionality, e.g., the Embedded Rendezvous Point

(Embedded-RP) [RFC3956] in routers and MLDv1 and MLDv2 snooping in switches.

7. Security Considerations

This document has multiple security sections detailing with how to securely deploy an IPv6 network within an enterprise network.

8. Informative References

[CYMRU] Team CYMRU Community Services, "THE BOGON REFERENCE", Version 7, April 2012, <http://www.team-cymru.org/Services/Bogons/>.

[DHCPv6-SLAAC-PROBLEM] Liu, B. and R. Bonica, "DHCPv6/SLAAC Address Configuration Interaction Problem Statement", Work in Progress, draft-liu-bonica-dhcpv6-slaac-problem-02, September 2013.

[DIVI] Bao, C., Li, X., Zhai, Y., and W. Shang, "dIVI: Dual-Stateless IPv4/IPv6 Translation", Work in Progress, draft-xli-behave-divi-06, January 2014.

[EDUROAM] Wierenga, K., Winter, S., and T. Wolniewicz, "The eduroam architecture for network roaming", Work in Progress, draft-wierenga-ietf-eduroam-04, August 2014.

[HOST-SCANNING] Gont, F. and T. Chown, "Network Reconnaissance in IPv6 Networks", Work in Progress, draft-ietf-opsec-ipv6-host-scanning-04, June 2014.

[IPv6-DC] Lopez, D., Chen, Z., Tsou, T., Zhou, C., and A. Servin, "IPv6 Operational Guidelines for Datacenters", Work in Progress, draft-ietf-v6ops-dc-ipv6-01, February 2014.

[IPv6-DESIGN] Matthews, P. and V. Kuarsingh, "Design Choices for IPv6 Networks", Work in Progress, draft-ietf-v6ops-design-choices-02, September 2014.

[IPv6-SECURITY] Chittimaneni, K., Kaeo, M., and E. Vyncke, "Operational Security Considerations for IPv6 Networks", Work in Progress, draft-ietf-opsec-v6-04, October 2013.

[RFC0826] Plummer, D., "Ethernet Address Resolution Protocol: Or converting network protocol addresses to 48.bit Ethernet address for transmission on Ethernet hardware", STD 37, RFC 826, November 1982, <http://www.rfc-editor.org/info/rfc826>.

[RFC1687] Fleischman, E., "A Large Corporate User's View of IPng", RFC 1687, August 1994, <http://www.rfc-editor.org/info/rfc1687>.

[RFC1817] Rekhter, Y., "CIDR and Classful Routing", RFC 1817, August 1995, <http://www.rfc-editor.org/info/rfc1817>.

[RFC1918] Rekhter, Y., Moskowitz, R., Karrenberg, D., Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2011] McCloghrie, K., "SNMPv2 Management Information Base for the Internet Protocol using SMIv2", RFC 2011, November 1996, <http://www.rfc-editor.org/info/rfc2011>.

[RFC2096] Baker, F., "IP Forwarding Table MIB", RFC 2096, January 1997, <http://www.rfc-editor.org/info/rfc2096>.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000, <http://www.rfc-editor.org/info/rfc2827>.

[RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003, <http://www.rfc-editor.org/info/rfc3315>.

[RFC3956] Savola, P. and B. Haberman, "Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address", RFC 3956, November 2004, <http://www.rfc-editor.org/info/rfc3956>.

[RFC3971] Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005, <http://www.rfc-editor.org/info/rfc3971>.

[RFC3972] Aura, T., "Cryptographically Generated Addresses (CGA)", RFC 3972, March 2005, <http://www.rfc-editor.org/info/rfc3972>.

[RFC4038] Shin, M-K., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005, <http://www.rfc-editor.org/info/rfc4038>.

[RFC4057] Bound, J., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005, <http://www.rfc-editor.org/info/rfc4057>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005, <http://www.rfc-editor.org/info/rfc4193>.

[RFC4213] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005, <http://www.rfc-editor.org/info/rfc4213>.

[RFC4292] Haberman, B., "IP Forwarding Table MIB", RFC 4292, April 2006, <http://www.rfc-editor.org/info/rfc4292>.

[RFC4293] Routhier, S., "Management Information Base for the Internet Protocol (IP)", RFC 4293, April 2006, <http://www.rfc-editor.org/info/rfc4293>.

[RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006, <http://www.rfc-editor.org/info/rfc4364>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006, <http://www.rfc-editor.org/info/rfc4443>.

[RFC4659] De Clercq, J., Ooms, D., Carugi, M., and F. Le Faucheur, "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, September 2006, <http://www.rfc-editor.org/info/rfc4659>.

[RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007, <http://www.rfc-editor.org/info/rfc4861>.

[RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007, <http://www.rfc-editor.org/info/rfc4862>.

[RFC4890] Davies, E. and J. Mohacsi, "Recommendations for Filtering ICMPv6 Messages in Firewalls", RFC 4890, May 2007, <http://www.rfc-editor.org/info/rfc4890>.

[RFC4941] Narten, T., Draves, R., and S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 4941, September 2007, <http://www.rfc-editor.org/info/rfc4941>.

[RFC5095] Abley, J., Savola, P., and G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", RFC 5095, December 2007, <http://www.rfc-editor.org/info/rfc5095>.

[RFC5157] Chown, T., "IPv6 Implications for Network Scanning", RFC 5157, March 2008, <http://www.rfc-editor.org/info/rfc5157>.

[RFC5211] Curran, J., "An Internet Transition Plan", RFC 5211, July 2008, <http://www.rfc-editor.org/info/rfc5211>.

[RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008, <http://www.rfc-editor.org/info/rfc5214>.

[RFC5375] Van de Velde, G., Popoviciu, C., Chown, T., Bonness, O., and C. Hahn, "IPv6 Unicast Address Assignment Considerations", RFC 5375, December 2008, <http://www.rfc-editor.org/info/rfc5375>.

[RFC5722] Krishnan, S., "Handling of Overlapping IPv6 Fragments", RFC 5722, December 2009, <http://www.rfc-editor.org/info/rfc5722>.

[RFC5798] Nadas, S., "Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6", RFC 5798, March 2010, <http://www.rfc-editor.org/info/rfc5798>.

[RFC5952] Kawamura, S. and M. Kawashima, "A Recommendation for IPv6 Address Text Representation", RFC 5952, August 2010, <http://www.rfc-editor.org/info/rfc5952>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, October 2010, <http://www.rfc-editor.org/info/rfc6052>.

[RFC6104] Chown, T. and S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement", RFC 6104, February 2011, <http://www.rfc-editor.org/info/rfc6104>.

[RFC6105] Levy-Abegnoli, E., Van de Velde, G., Popoviciu, C., and J. Mohacsi, "IPv6 Router Advertisement Guard", RFC 6105, February 2011, <http://www.rfc-editor.org/info/rfc6105>.

[RFC6106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 6106, November 2010, <http://www.rfc-editor.org/info/rfc6106>.

[RFC6144] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", RFC 6144, April 2011, <http://www.rfc-editor.org/info/rfc6144>.

[RFC6145] Li, X., Bao, C., and F. Baker, "IP/ICMP Translation Algorithm", RFC 6145, April 2011, <http://www.rfc-editor.org/info/rfc6145>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, April 2011, <http://www.rfc-editor.org/info/rfc6146>.

[RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, April 2011, <http://www.rfc-editor.org/info/rfc6147>.

[RFC6164] Kohno, M., Nitzan, B., Bush, R., Matsuzaki, Y., Colitti, L., and T. Narten, "Using 127-Bit IPv6 Prefixes on Inter-Router Links", RFC 6164, April 2011, <http://www.rfc-editor.org/info/rfc6164>.

[RFC6177] Narten, T., Huston, G., and L. Roberts, "IPv6 Address Assignment to End Sites", BCP 157, RFC 6177, March 2011, <http://www.rfc-editor.org/info/rfc6177>.

[RFC6192] Dugal, D., Pignataro, C., and R. Dunn, "Protecting the Router Control Plane", RFC 6192, March 2011, <http://www.rfc-editor.org/info/rfc6192>.

[RFC6296] Wasserman, M. and F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC 6296, June 2011, <http://www.rfc-editor.org/info/rfc6296>.

[RFC6302] Durand, A., Gashinsky, I., Lee, D., and S. Sheppard, "Logging Recommendations for Internet-Facing Servers", BCP 162, RFC 6302, June 2011, <http://www.rfc-editor.org/info/rfc6302>.

[RFC6434] Jankiewicz, E., Loughney, J., and T. Narten, "IPv6 Node Requirements", RFC 6434, December 2011, <http://www.rfc-editor.org/info/rfc6434>.

[RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", RFC 6555, April 2012, <http://www.rfc-editor.org/info/rfc6555>.

[RFC6583] Gashinsky, I., Jaeggli, J., and W. Kumari, "Operational Neighbor Discovery Problems", RFC 6583, March 2012, <http://www.rfc-editor.org/info/rfc6583>.

[RFC6724] Thaler, D., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, September 2012, <http://www.rfc-editor.org/info/rfc6724>.

[RFC6866] Carpenter, B. and S. Jiang, "Problem Statement for Renumbering IPv6 Hosts with Static Addresses in Enterprise Networks", RFC 6866, February 2013, <http://www.rfc-editor.org/info/rfc6866>.

[RFC6879] Jiang, S., Liu, B., and B. Carpenter, "IPv6 Enterprise Network Renumbering Scenarios, Considerations, and Methods", RFC 6879, February 2013, <http://www.rfc-editor.org/info/rfc6879>.

[RFC6883] Carpenter, B. and S. Jiang, "IPv6 Guidance for Internet Content Providers and Application Service Providers", RFC 6883, March 2013, <http://www.rfc-editor.org/info/rfc6883>.

[RFC6959] McPherson, D., Baker, F., and J. Halpern, "Source Address Validation Improvement (SAVI) Threat Scope", RFC 6959, May 2013, <http://www.rfc-editor.org/info/rfc6959>.

[RFC6964] Templin, F., "Operational Guidance for IPv6 Deployment in IPv4 Sites Using the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 6964, May 2013, <http://www.rfc-editor.org/rfc/rfc6964.txt>.

[RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013, <http://www.rfc-editor.org/info/rfc7011>.

[RFC7012] Claise, B. and B. Trammell, "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, September 2013, <http://www.rfc-editor.org/info/rfc7012>.

[RFC7045] Carpenter, B. and S. Jiang, "Transmission and Processing of IPv6 Extension Headers", RFC 7045, December 2013, <http://www.rfc-editor.org/info/rfc7045>.

[RFC7113] Gont, F., "Implementation Advice for IPv6 Router Advertisement Guard (RA-Guard)", RFC 7113, February 2014, <http://www.rfc-editor.org/info/rfc7113>.

[RFC7123] Gont, F. and W. Liu, "Security Implications of IPv6 on IPv4 Networks", RFC 7123, February 2014, <http://www.rfc-editor.org/info/rfc7123>.

[RFC7359] Gont, F., "Layer 3 Virtual Private Network (VPN) Tunnel Traffic Leakages in Dual-Stack Hosts/Networks", RFC 7359, August 2014, <http://www.rfc-editor.org/info/rfc7359>.

[SMURF] The Cert Division of the Software Engineering Institute, "Smurf IP Denial-of-Service Attacks", CERT Advisory CA-1998-01, March 2000, <http://www.cert.org/advisories/CA-1998-01.html>.

[ULA-USAGE] Liu, B. and S. Jiang, "Considerations of Using Unique Local Addresses", Work in Progress, draft-ietf-v6ops-ula-usage-recommendations-03, July 2014.

Acknowledgements

The authors would like to thank Robert Sparks, Steve Hanna, Tom Taylor, Brian Haberman, Stephen Farrell, Chris Grundemann, Ray Hunter, Kathleen Moriarty, Benoit Claise, Brian Carpenter, Tina Tsou, Christian Jacquenet, and Fred Templin for their substantial comments and contributions.

Authors' Addresses

Kiran K. Chittimaneni
Dropbox, Inc.
185 Berry Street, Suite 400
San Francisco, CA 94107
United States
EMail: kk@dropbox.com

Tim Chown
University of Southampton
Highfield
Southampton, Hampshire SO17 1BJ
United Kingdom
EMail: tjc@ecs.soton.ac.uk

Lee Howard
Time Warner Cable
13820 Sunrise Valley Drive
Herndon, VA 20171
United States
Phone: +1 703 345 3513
EMail: lee.howard@twcable.com

Victor Kuarsingh
Dyn, Inc.
150 Dow Street
Manchester, NH
United States
EMail: victor@jvknet.com

Yanick Pouffary
Hewlett Packard
950 Route Des Colles
Sophia-Antipolis 06901
France
EMail: Yanick.Pouffary@hp.com

Eric Vyncke
Cisco Systems
De Kleetlaan 6a
Diegem 1831
Belgium
Phone: +32 2 778 4677
EMail: evyncke@cisco.com