ETSI GR IP6 001 V1.1.1 (2017-06)

GROUP REPORT

IPv6 Deployment in the Enterprise

Disclaimer

The present document has been produced and approved by the IPv6 Integration (IP6) ETSI Industry Specification Group (ISG)
and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference

DGR/IP6-0001

Keywords

Internet, IoT, IPv6, Mobile, NFV, SDN

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Copyright Notification

Contents

Intellectual Property Rights

Essential patents

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (https://ipr.etsi.org/).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Trademarks

The present document may include trademarks and/or tradenames which are asserted and/or registered by their owners. ETSI claims no ownership of these except for any which are indicated as being the property of ETSI, and conveys no right to use or reproduce any trademark and/or tradename. Mention of those trademarks in the present document does not constitute an endorsement by ETSI of products, services or organizations associated with those trademarks.

Foreword

This Group Report (GR) has been produced by ETSI Industry Specification Group (ISG) IPv6 Integration (IP6).

Modal verbs terminology

In the present document "**should**", "**should not**", "**may**", "**need not**", "**will**", "**will not**", "**can**" and "**cannot**" are to be interpreted as described in clause 3.2 of the ETSI Drafting Rules (Verbal forms for the expression of provisions).

"**must**" and "**must not**" are **NOT** allowed in ETSI deliverables except when used in direct citation.

Executive summary

The present document outlines the motivation for the deployment of IPv6 within enterprises, the objectives, the benefits, the risks, the challenges, the technology guidelines, the different choices that arise when designing IPv6-only or dual-stack enterprise network, step-by-step process, the addressing plan, and the milestones.

With the growing number of users and the proliferation of smart devices and things, IPv4 address space exhaustion is a major Information and Communications Technology (ICT) issue. The current IPv4-based Internet can no longer sustain the explosive growth of ICT. Any organization that relies on the Internet to any extent need to be prepared, need to support IPv6. The move to IPv6 is inevitable, there is no alternative plan at this time. IPv6 is the cornerstone of our connected society. IPv6 offers important business and technical advantages. While all/most existing IPv4-based infrastructures will continue to work after the last IPv4 address is issued, enterprises may be tempted to put off transitioning to IPv6 till some later date. However postponing the inevitable can put an enterprise at a competitive disadvantage. The present document provides guidelines and recommendations on IPv6 deployment in the enterprise.

1 Scope

The present document outlines the motivation for the deployment of IPv6 in enterprises, the objectives, the benefits, the risks, the challenges, the technology guidelines, the different choices that arise when designing IPv6-only or dual-stack enterprise network, step-by-step process, the addressing plan, and the milestones.

2 References

2.1 Normative references

Normative references are not applicable in the present document.

2.2 Informative references

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] IETF RFC 1918: "Address Allocation for Private Internets".

[i.2] IETF RFC 3493: "Basic Socket Interface Extensions for IPv6".

[i.3] IETF RFC 4193: "Unique Local IPv6 Unicast Addresses".

[i.4] IETF RFC 3531: "A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block".

[i.5] IETF RFC 3542: "Advanced Sockets Application Program Interface (API) for IPv6".

[i.6] IETF RFC 4193: "Unique Local IPv6 Unicast Addresses".

[i.7] IETF RFC 4380: "Teredo: Tunneling IPv6 over UDP through Network Address Translations".

[i.8] IETF RFC 5375: "The Locator/ID Separation Protocol (LISP)".

[i.9] IETF RFC 6146: "Stateful NAT64: - Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers".

[i.10] IETF RFC 6296: "IPv6 to IPv6 Network Prefix Translation".

[i.11] IETF RFC 6555: "Happy Eyeballs: Success with Dual-Stack Hosts".

[i.12] IETF RFC 7239: "Forwarded HTTP Extension".

[i.13] RIPE NCC: "Requirements for IPv6 in ICT Equipment".

NOTE: Available at https://www.ripe.net/publications/docs/ripe-554.

[i.14] IETF ID draft-troan-homenet-sadr-01: IPv6 Multihoming with Source Address Dependent Routing (SADR)".

[i.15] DomainKeys Identified Mail.

NOTE: Available at https://en.wikipedia.org/wiki/DomainKeys_Identified_Mail.

[i.16] IETF RFC 6145: "IP/ICMP Translation Algorithm".

3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AAAA	DNS resource records (RRs) IPv6 address record
ACL	Access Control List
AFT	Address Family Translation
ARP	Address Resolution Protocol
ALG	Application-aware Logic
AV	Audio/Video
BCP	Best Current Practices
BGP	Border Gateway Protocol
BIND	Berkeley Internet Name Domain
BYOD	Bring your own device
CDN	Content Delivery Network
CG-NAT	Carrier-grade NAT (Network Address Translation)
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
ERP	Enterprise Resource Planning
FQDN	Fully qualified domain name
GR	Group Report
HTTP	HyperText Transfer Protocol
ICMP	Internet Control Message Protocol
ICT	Information and Communications Technology
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IoT	Internet of Things
IP	Internet Protocol
IPAM	IP Address Management
IPS	Intrusion Prevention Systems
IPsec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISATAP	Intra-Site Automatic Tunnel Addressing Protocol
ISC	Internet Systems Consortium
ISG	Industry Specification Group
IS-IS	Intermediate System to intermediate System
ISP	Internet Service Provider
LISP	Locator/Identifier Separation Protocol
LLMNR	Link-Local Multicast Name Resolution Protocol
LOB	Line of Business
MAC	Media Access Conrol address
MIB	Management information base
MiM	Man-in-the-Middle
MTA	Mail Transfer Agents
NAT	Network Address Translation
ND	Neighbour Discovery
NPTv6	Network Prefix Translation
OS	Operating System
OSX	Mac OS X™

NOTE 1: Unix-based graphical operating system developed and marketed by Apple Inc.

NOTE 2: Mac OS X™ is a trademark of Apple Inc., registered in the U.S. and other countries.

PA	Provider aggregatable
PI	Provider independent
RDNSS	Recursive DNS Server

RIPE	Réseaux IP Européens
RIR	Regional Internet Registry
RR	Resource Record
SADR	Source Address Dependent Routing
SIEM	Security Incident and Event Management
SLAAC	Stateless Address Auto Configuration
SLB	Server Load Balancer
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SSDP	Simple Service Discovery Protocol
SSL	Secure Sockets Layer
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
ULA	Unique Local IPv6 Unicast Addresses
VLAN	Virtual Local Area Network
VPN	Virtual Private Network
WAN	Wide Area Network

4 IPv6-enabled Enterprise

4.1 Introduction

When IPv4 emerged as the standard Internet protocol in the 1980s for the "Internet", the address space - some four billion IP addresses - seemed more than adequate. Today it is clearly no longer the case. The world has moved from IP enabled to IP dependent. As of 2015 more than 83 percent of the world's population no longer has access to a "public Internet address." In North America, Latin America, Asia, and Europe, the IPv4 address pool is already entirely depleted. With the growing number of users and the proliferation of smart devices and things, IPv4 address space exhaustion is a major Information and Communications Technology (ICT) issue. The current IPv4-based Internet can no longer sustain the explosive growth of ICT. Any organization that relies on the Internet to any extent needs to be prepared to support IPv6.

IPv6 needs to be adopted globally across all parts of the Internet ecosystem. Global service providers and mobile operators are already deploying IPv6 in order to keep the Internet growing. IPv6 is allowing them to continue to grow their businesses and deliver the services that all of today's e-commerce is based on.

The move to IPv6 is inevitable, there is no alternative plan at this time**.** IPv6 is the cornerstone of our connected society. IPv6 offers important business and technical advantages. Among them: higher performance, enhanced mobility, automated management, built-in multicasting for multimedia applications, enhanced security, simplified administration and many more. Indeed, many enterprises have already determined that IPv6 is a much better networking tool than its predecessor, and it offers much greater capabilities for future technologies developed for networking platforms.

Because all/most existing IPv4-based infrastructures will continue to work after the last IPv4 address is issued, enterprises may be tempted to put off transitioning to IPv6 till some later date. Postponing the inevitable, however, can put an enterprise at a competitive disadvantage. As more and more customers operate in an IPv6 world, companies supporting only IPv4 risk being shut out of high-growth markets because they are unable to reach, or be reached by, these customers. The fallacy in this position is that maintaining IPv4-only communications can put enterprises at a competitive disadvantage. Seamless, pervasive connectivity is an integral part of doing business today.

IPv6 enables the enterprise and the global Internet to keep growing in a secure and open manner, and to scale toward the demand of new applications and, literally, billions of connected devices, while streamlining operations and provisioning. Enterprises deploying IPv6-enabled services are in a better position to capture the market changes, be more competitive, increase their growth potential, and provide for improved business continuity.

Of particular note is the fact that regardless of an organization's IPv6 preference, its customer base is currently deploying it. As consumers and customers adopt IPv6, enterprises will need to ensure that their own technology assets align with how their customers prefer to do business. As the shift to online shopping has already shown, customers will continue to buy from companies that do business on their terms. Meeting the needs of the customer base (many of which already have an IPv6-connected device) should be an imperative for all organizations.

Today, enterprises should already have assessed their position toward IPv6 adoption, understood its challenges and opportunities, and drafted their own requirements and plans accordingly.

4.2 Guiding principles

There are many scenarios/variations to enable IPv6 within the enterprise, there is no "one size fits all" answer. The present document does not attempt to provide guidance for all possible networking situations. Enterprise network architects should each take the responsibility of choosing the best solution for their own case.

Enterprises should understand the impact of the IPv6 Internet on their services. Next, they should assess their own situations and requirements as early as possible, if they have not yet done so. This assessment includes network, security, and business applications. Enterprise priorities could be:

- IPv6 is strategic in order to achieve business continuity.

- IPv6 has its own value outside of IPv4 address exhaustion.

In each case, requirements demand a production-grade deployment within a two- to three-year horizon.

It is expected that for most enterprises adding IPv6 connectivity, the Internet presence will be the highest priority because enterprises should offer services and content to their IPv6 customers over the Internet. This is also the easiest part.

Deploying IPv6 across the core is also a relatively easy task and will provide the network staff with a solid working knowledge of the IPv6 protocol.

Providing IPv6 access to all internal users and applications is probably the next highest priority, especially if the enterprises move to the cloud computing paradigm or to web services.

The last step will probably be adding IPv6 in the intranet and in the data centre applications. This will take longer, as there is a clear impact on business applications.

Finally there are three core principles to shape IPv6 network development:

- The first principle is to maintain a standards-based approach and avoid proprietary technologies. Embracing open standards allows to use best-of-breed products for whatever needs arise.

  A network build on open standards will have interoperability with the broadest base of users and partners.

- The second principle is planning. Without proper planning, the transition will be plagued with rework and missteps.

- The third principle is repeatability of the network design across the network (a standard set of requirements, and a standard solution design to meet the enterprise network requirements).

  Such standardization increases efficiency in building out the network, and also makes it easier to troubleshoot the network in each location.

4.3 IPv6 Transition strategies

4.3.1 Setting the scene

The vast majority of devices, laptops, desktops, operating systems, switches, routers, content providers, carriers, and Internet service providers (ISPs) support native IPv6 today at no extra cost, which makes it possible to deploy network based on IPv6. However, enterprises typically purchased and configured their network to support IPv4 traffic only. And while most equipment can be software enabled, some may need to be replaced to add support for IPv6. Furthermore even if the capabilities to operate in a dual network configuration exist, additional planning steps and architecture design will most likely be required. Similarly management systems and security systems that can support both environments are necessary. Enterprises should also verify that applications they use can operate correctly and are IPv6-enabled.

Because IPv4 and IPv6 will coexist for some time, a phased deployment is recommended to minimize the impact of the transition and keep costs manageable. Recognizing that IPv4 and IPv6 will run parallel to each other for the foreseeable future, IETF established three standard transition mechanisms: Each of these techniques has advantages and trade-offs. The optimal solution will depend on a variety of factors, including the enterprise's current environment and long-term goals. It may also encompass all three transition mechanisms. It is important to understand that one method does not fit all.

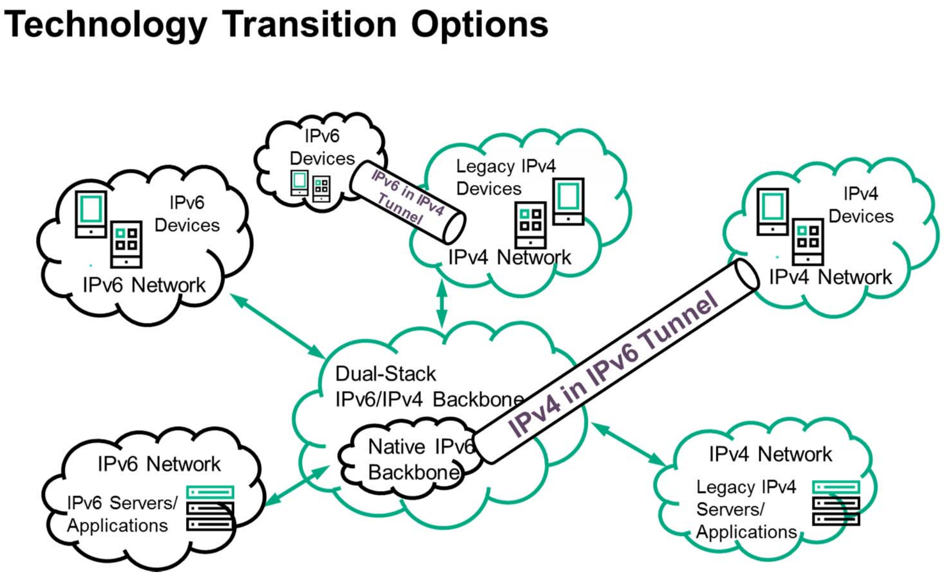

Figure 1: Technology Transition Options

In planning the IPv6 initiative, the following three key families of mechanisms enable the transition:

- **Dual-stack** - Provides support for both protocols on the same device to allow for communications with both IPv4-only and IPv6-only nodes.

  This mechanism is the most versatile. A dual stack transition strategy enables a very smooth transition and will likely remain a part of the worldwide Internet infrastructure for years to come, until IPv4 is fully retired.

- **Tunnelling** - Encapsulates IPv6 packets in IPv4 headers (or vice versa).

  Tunnelling enables the network team to create islands of IPv6 or IPv4 capabilities, and in the short term, to connect them over the existing IPv4 network. Tunnelling enables networks in transition to take advantage of IPv6 services while remaining connected to the IPv4 world.

- **Translation** - Between IPv4 and IPv6.

Enterprises should also implement a procurement policy to IPv6-ready the backbone so that IPv6 can be turned on without having to do a fork and replace the physical hardware.

The main IPv6 transition deployment models that are being discussed are listed below:

- **IPv4 only:** delays the introduction of IPv6 to a later date and remains an all-IPv4 network. Long term, it is expected that this migration strategy will lead to problems and increased costs. Due to the increase in traffic there will be an increased demand for IP addresses and the usage of NAT in the carrier's network, denoted as Carrier Grade Network Address Translation (CG-NAT). In particular, all traffic to and from the Internet will have to pass CG-NAT. Furthermore, growth in bandwidth demand can only be handled with increased CG-NAT capacity, which has a higher cost and single point of failure.

- **Coexistence of IPv4 and IPv6:** requires the use of a dual-stack, introducing IPv6 in the network next to IPv4. Please note however dual-stack networks are more complex to deploy, operate, and manage. Furthermore, this option also requires an address management solution for both IPv4 and IPv6 addresses.

- **IPv6 only:** introduces IPv6 in the network and removes IPv4 completely. This approach can provide benefits because IPv6-only networks are simpler to deploy, operate, and manage. Moreover, an address management solution is required only for IPv6 addresses. However, the problem with this approach is that many devices, websites, and applications still only work on IPv4. When moving to an IPv6-only network may lead to differences in network quality. That is why NAT64 with DNS64 should be offered in addition to offering IPv6 only.

4.3.2 Dual-Stack

Most enterprises will initially prefer the dual-stack model. All users, applications, and network equipment will be given address space from both IPv4 and IPv6. It is then up to the user's device to select which protocol version to use. Most common operating systems prefer IPv6 when a functional path is available.

The operating systems have specific checks in place to ensure quick and reliable connections when operating in dual-stack mode, based on IETF RFC 6555 "Happy Eyeballs" [i.11]. IETF RFC 6555 [i.11] tries to ensure that (where it is reliably available) IPv6 is usually preferred over IPv4. That is to say without IETF RFC 6555 [i.11] the first response back is preferred by default. Simply stated, parallel DNS queries are launched for IPv4 and IPv6 addresses for any website, and the first response back is preferred. IETF RFC 6555 [i.11] works with a slight bias toward IPv6 by giving the AAAA query a slight head start over the query for the legacy protocol address.

It is important to note that the dual-stack model will not be sustainable in the future when the available IPv4 space has been completely exhausted.

Dual-stack model also doubles the amount traffic at layer 2 for service discovery protocols like LLMNR and SSDP, as well as ARP and ND. This can have negative impact on CPU of some older network hardware. Additionally it is important to validate there is sufficient memory to store both IPv4 and multiple IPv6 addresses.

IPv6 traffic on Wi-Fi network requires additional consideration. IPv6 ND traffic and other IPv6 control plane traffic using multicast as the delivery mechanism may require optimizations over Wi-Fi network. Multicast over Wi-Fi network uses lowest management rates and mechanisms such as clients power save could impact timely delivery of multicast traffic over Wi-Fi. Typically Wi-Fi vendors convert multicast traffic to unicast to ensure robust and reliable delivery of multicast traffic. This also ensures that traffic is delivered to clients at the highest data rates possible and since the traffic is unicast, the standard power save mechanism for delivering multicast traffic does not apply.

IPv6 over Wi-Fi may require optimization techniques such as these to ensure reliable delivery of IPv6 control traffic.

4.3.3 IPv6 only

IPv6-only is the end goal of transitioning to IPv6. There have been numerous recent examples of large entities migrating to IPv6 after dual-stack transitions. Typical motivations are to avoid costs of translation equipment, reduce the cost of running a dual-stack infrastructure, reduce the attack surface to only one protocol, and simplify troubleshooting. It has deployed an IPv6-only infrastructure within its data centre.

The public-facing side of its Internet presence/WAN edge still presents a dual-stack interface to the global Internet, but it has completely removed IPv4 from the internal data centre infrastructure.

IPv6 only with NAT64/DNS64 is easy to implement and works well in a Client/Server model, there are challenges with peer-to-peer applications and IPv4 only clients. Application validation for enterprise peer-to-peer applications should be included in the scorecard.

4.3.4 Tunnelling

In 2015, native or dual-stack IPv6 deployment is possible and should be used for security and performance reasons. Tunnels should in general be avoided at all costs.

4.3.5 Enterprise Network Segments

The typical enterprise network has been built on a three-layer model, defining access, distribution and core as those layers, integrating the Internet edge, and providing maximum scalability. Smaller enterprises may have collapsed the core and distribution layers and combined that with the access layer. The adoption of IPv6 does not change those models and should be planned for in a similar fashion. The key segments of the overall network architecture are discussed laterlater.

4.4 Enterprise Design Considerations - Building a cross functional team

Planning and developing an IPv6 strategy will involve multiple phases and should include the creation of governance cross-functional team of IT professionals, technical business owners, and an assigned project manager.

The team should meet regularly to discuss the progress and address any outstanding issues. Training and education are critical factors toward success. Among the first steps are:

- Determining an architectural approach
- Developing an execution strategy
- Assessing existing IT components for IPv6 readiness
- IT readiness communication to end users

Governance activities related to IPv6: It is important that key people from various areas in an organization understand the importance of IPv6 and the associated governance elements. There should be strategic oversight of development, resource allocation and purchasing that covers at least these key areas of the IT organizations:

- Executive and Risk management
- Application development
- Desktop and server build
- Information security
- Architecture (for both infrastructure and applications)
- Networking

4.5 Preparation and Assessment Phase

An enterprise needs to assess its network, applications, hosts, and critical infrastructure to determine its IPv6 readiness and identify challenges that might occur during the transition. This step will provide insight and visibility, enabling the team to proactively manage and budget resources with timelines for later testing, trailing, and deployment phases. Most current hardware and software are already IPv6 ready.

Table 1: Example Application and Service testing matrix to help access IPv6 readiness

Applications <--> Devices	LOB applications	Web portals	Cloud content	AV applications
Windows OS				
MAC OS X™				
Smart phones				
Tablets				
IoT				

What is most important is for the enterprise to ensure that all new procurements undertaken automatically require IPv6 support. In this regard, an extremely helpful baseline has been created (Requirements for IPv6 in ICT Equipment - RIPE 554 [i.13])

Table 2: Example Scorecard

Dependencies	Product Vendor	Category	Milestone	Finish Date	Team ownership	Notes
Windows Server	Microsoft	datacenter	upgrade to windows server		datacenter engineering	
Windows Client	Microsoft	Campus edge	Upgrade to Windows 10		systems engineering	
Network Config mgmt	vendor	network tools	Upgrade to support IPv6 or roadmap		Tools engineering team	
Network performance mgmt	vendor	network tools	Upgrade to support IPv6 or roadmap		Tools engineering team	
IP database	vendor	network tools	integrate IPv6		IP services team	
Active Directory IP sync	Microsoft	Identify mgmt	automate IPv6 subnet assocation to AD		systems engineering	
Firewall Policy manager	vendor	network hardware	Integrate or update policy manager to support IPv6		security	
Netflow	vendor	network tools	All devices to Netflow v9, collector upgrade to support IPv6		Tools engineering team	
DHCP	Microsoft	IP infastructure	DHCPv6 enabled		IP services team	
Naming services	Microsoft	IP infastruture	naming services support IPv6		IP services team	
Active Directory	Microsoft	Identify mgmt	IPv6 support AD infastructure		IP services team	
Wireless controllers	vendor	network hardware	Enable IPv6 on wireless controllers		backbone engineering	
Load balancer	vendor	datacenter	enable IPv6 on load balancers		datacenter engineering	
Firewall	vendor	Network security	enable IPv6 cabability and performance		Security	
Wan Optimizer	vendor	network hardware	no plan of record		backbone engineering	
Backbone router	vendor	network hardware	validate IPv6 current approved code or upgrade lifecycle		backbone engineering	
Building router	vendor	network hardware	validate IPv6 current approved code or upgrade lifecycle		backbone engineering	
WAN router	vendor	network hardware	validate IPv6 current approved code or upgrade lifecycle		backbone engineering	
O365	Microsoft	Application	validate IPv6 support or support roadmap		systems engineering	
Skype	Microsoft	Application	validate IPv6 support or support roadmap		systems engineering	

4.6 IPv6 address plan

4.6.1 Setting the scene

An enterprise needs to determine what type of IPv6 address space to use, where it will get it from, and how much it will need. One of the first choices a network designer needs to make is the type of addresses to be used in the network core. Should the network use provider-independent global addresses, "private" addresses (either IETF RFC 1918 [i.1] addresses or unique-local addresses) or something else? A related choice is whether or not to use only link-local addresses on certain links.

A smaller enterprise with a single ISP may opt to use IPv6 address space allocated from its provider (this is known as Provider Aggregable, or PA, and is typically given at no extra cost), as for IPv4.

For much larger enterprises (typically multihomed to multiple providers) PA space will not be practical. They need to apply directly to their RIR for what is known as a provider-independent (PI) allocation. This type of allocation comes with an annual operational cost and should be routable across multiple providers, as is the case with most existing legacy address allocations:

- PI provider independent - Globally-unique IPv4 or IPv6 addresses obtained directly from an address registry. An organization which has such addresses is considered to have "its own" address space.
- PA (provider aggregable) - Globally-unique IPv4 or IPv6 addresses obtained from an upstream provider. Such addresses will need to be returned if the relationship with the upstream provider ceases.
- Private - Either IETF RFC 1918 [i.1] IPv4 addresses or unique-local IPv6 addresses IETF RFC 4193 [i.3].

The IPv6 prefix should satisfy the following addressing requirements (IETF RFC 3531 [i.4] provides some schema guidelines):

- The corporate network. This includes services such as wireless. An effort should also be made to accommodate the existing IPv6 implementation with the minimum of change.
- Guest services, wired and wireless.
- Future mobile devices.
- Miscellaneous devices such as security cameras, etc.
- Multicast.

NOTE: There are no translation mechanisms for IPv4/IPv6 multicast.

- Internet.

There are also address space management requirements:

- The address space should be easily aggregable. That is, fragmentation should be kept to a minimum.
- If the current IPv4 schema attaches semantic to an address, it is possible to extend this facility to IPv6. It is debatable that is the right thing to do

As with the early days of IPv4, IPv6 will see change as adoption matures and new devices and ways of addressing those devices evolve. The IPv6 address plan should be flexible to change with automation in place to facilitate re-ip. Fortunately, IPv6 was designed to support multiple IP addresses on an interface, allowing a new address to be assigned, while maintaining existing connections with original address.

4.6.2 To use or not to use ULA (Unique Local IPv6 Unicast Addresses)

IPv6 has a concept comparable in certain respects to IPv4 private addressing (IETF RFC 1918 [i.1]), known as unique local addressing (ULA). A specific prefix (FC00:/7) has been set aside for private internal use within an administrative domain. Use of this space has numerous documented side effects and is generally recommended only for use in networks that will never need a direct, end-to-end connection with a device outside the company's network.

One of the key challenges of IPv4 private addresses has been mostly mitigated with IPv6 ULA: the particular issue of uniqueness. With IPv4 all companies shared a limited amount of private space, and it was not uncommon for two merging companies to have overlapping address space.

However, using ULA would require careful consideration of application proxies, translation devices, merger strategies, application layer gateways and new application development. As with IETF RFC 1918 [i.1], there are no magic properties of ULA that guarantee complete isolation of the Internet. In other words, Access Control Lists (ACLs) and routing need to be explicitly used to provide isolation.

4.6.3 Understanding the differences between IPv4 and IPv6 addressing schemes.

IPv4 uses a variable subnet mask that is based on the maximum number of expected hosts on a subnet. For example, if a network is to support 250 hosts, a subnet mask of at least /24 would be used for IPv4. The address allocation (from a service provider) is therefore based on the aggregate number of hosts expected across a network.

IPv6 uses a fixed subnet mask of /64 on broadcast media such as Ethernet and Wi-Fi. The most significant 64 bits are used for routing, and the least significant 64 bits are used for the node identification. With the subnet mask fixed, the number of deployed hosts in a network becomes irrelevant. Hence the allocation from a service provider would be based roughly on the total number of expected networks within an enterprise.

The common rule for a service provider is to allocate a /48 to enterprises, leaving 16 bits to the enterprise for enumerating all of the /64 networks (several enterprises simply use the VLAN ID to fill those 16 bits). This gives the typical enterprise 65 536 subnets (which for the vast majority should be more than enough). Only very large multisite, multibuilding enterprises may need to request a larger address block from their service provider or RIR.

4.6.4 Address Management

IPv6 provides a mechanism for automatic host address configuration, called stateless address auto configuration (SLAAC). By default, SLAAC addresses are configured to change daily to enhance the privacy of the user. This behaviour may affect an enterprise security audit.

Alternate techniques for address configuration are static configuration or the use of Dynamic Host Configuration Protocol (DHCP) for IPv6. The deployment of DHCPv6 is very similar to the typical enterprise deployment of DHCP for IPv4. As part of the design phase, an enterprise will have to decide which address strategy to use (SLAAC vs. DHCPv6). It is likely that an enterprise will deploy both, though not typically on the same access layer segment.

NOTE: SLAAC does not provide DNS server address so stateless DHCPv6 (DHCPv6 provides DHCP option 23 DNS recursive name sever list) is required. Additionally, the network infrastructure will require features such as RDNSS to support Android BYOD devices.

For example large corporation envisions DHCPv6 as the typical secure deployment for most enterprises, while SLAAC may be useful in IoT and BYOD or Guest network.

4.7 Routing considerations

Choice of IGP will depend on current IGP in use, experience and network tooling automation. The IGP should provide the flexibility to support both a converged (single topology), or separate (multi-topology) forwarding paths in addition to a transition mechanism to ease migration from one to the other. A consideration with separate IPv4 and IPv6 network is that in most cases the clients are dual-stack enabled and by default prefer IPv6 over IPv4. When the service they connect to is IPv6 enabled, all dual-stack enabled, will shift immediately to the IPv6 path. This will be transparent to the end user and can complicate troubleshooting if the IPv6 path is less optimal to the IPv4 path. This is not an issue if the path is the same (i.e. Single topology).

One benefit of IS-IS, in dual stack environment, is that it provides single device level view of both IPv6 and IPv4 with a single command:

EXAMPLE: sh isis database detail.

```
6-P4#sh isis database detail
IS-IS Level-2 Link State Database:
LSPID LSP Seq Num LSP Checksum LSP Holdtime ATT/P/OL
5-ASBR1.00-00 0x0000840A 0x7C60 841 0/0/0
 Area Address: 49.0000
 NLPID: 0xCC 0x8E
 Hostname: 5-ASBR1
 IP Address: 10.100.0.5
 IPv6 Address: 2001:4898:100::5
 Metric: 10 IS-Extended 9-ASBR2.05
 Metric: 10 IS-Extended 2-P1.03
 Metric: 10 IP 10.100.0.5/32
 Metric: 10 IP 10.100.1.0/30
 Metric: 10 IP 10.100.1.4/30
 Metric: 10 IPv6 2001:4898:100:2::/64
 Metric: 10 IPv6 2001:4898:100:1::/64
 Metric: 10 IPv6 2001:4898:100::5/128
```

4.8 IPv6 Data Centre

The enterprise data centre is defined as the different data centres of an enterprise that are located within the enterprise network and managed by the enterprise. It contains all the servers, applications, and data storage accessed by Internet users, partners, and internal users. The different user types (external vs. internal) and data served are logically isolated, although they are often physically collocated.

The steps to enable an IPv6-only data centre may include:

- IPv4-only: No translation, load balancers inline, services firewalled.
- Dual-stack front end: Translation on the front end (hard to move forward, may work with NAT).
- Dual-stack servers: Current recommended approach (requires dual everything-policy, quality of service, security).
- IPv6-only data centre: Stateless translation for internal users (reduces operating cost, enables quicker innovation).

The type of IPv4 address space used in the data centre is another consideration. For the enterprise data centre using public IPv4 addresses, address exhaustion will be an immediate issue because access to new IPv4 address space may not be possible. Moving to IPv6 is clearly the right way forward. For the enterprise using private address space, changes such as mergers, application development, and gateway deployment will be barriers to success at some point in the future.

Deploying IPv6 in the data centre requires two main steps:

- **Network deployment:** Because of the higher performance required in the data centre, all networking devices are required to have the same performance level for IPv6 as for IPv4, not only for routing but also for convergence, high availability, security inspection, and so on. Other points that could be sensitive are the load balancers, SSL acceleration devices, and network management tools.
- **Application deployment:** While large application vendors are aggressively moving to IPv6, this is not the case for all application vendors or open-source applications. If a server runs an IPv4-Only application or its code has IPv4 literals embedded in the application, a migration strategy may be needed. This could be accomplished by using the legacy protocol for the life of the application or by updating or even rewriting the code. There are two primary RFCs to assist an application developer with this task: IETF RFC 3493 [i.2] for open socket calls and IETF RFC 3542 [i.5] for raw socket calls.

- **Software development guidelines for IPv6:** All types of internal software development should be aware of the impact that IPv6 will have on their development philosophy. This can include development of client/server applications, web development, and even script and tool development. Guidelines can include things such as accommodating both IPv4 and IPv6 addresses, or the use of IP version-agnostic APIs for network connectivity.

- **IPv6 application compatibility considerations:** It is very important to understand how to seek out potential IPv6 incompatibilities within applications, tools, and scripts in order to avoid unnecessary surprises as IPv6 is deployed. An IPv6 application compatibility test plan, will help with this topic, but this should be supplemented with an IPv6 remediation plan that has committed key stakeholders.

The trends toward virtualization and orchestration (due to their prolific and dynamic nature) may also require the enterprise to establish and deploy IPv6 services as the increasing need for the addresses for virtual machines further causes the exhaustion of IPv4 address space at an even greater rate. Moving data from one data centre to the next for the purpose of replication or disaster recovery will require that the enterprise deploy data centre interconnect technologies that support IPv6.

4.9 Building an IPv6 Internet Presence

4.9.1 Setting the scene

An enterprise Internet presence usually consists of the services the enterprise offers to its partners, customers and the Internet community: web servers, email, remote access VPN and DNS. Enabling IPv6 on those services makes the enterprise present on both the IPv4 and IPv6 Internet. Certain operational support systems and network operations procedures needs to be updated and become IPv6 aware.

4.9.2 The network edge

When IPv6 is enabled, it will be necessary to define how IPv6 will be allowed (or not allowed) to traverse the network edge, if this has not already been done. Specifically, it should be considered whether connection to the IPv6 Internet will be allowed. Many organizations opt to block all IPv6 traffic, including native IPv6, Teredo (on port UDP/3544) [i.7], and ISATAP/6to4 (on IPv4 protocol #41), until such time as they are prepared to offer a full solution for IPv6 Internet connectivity to computers on their intranets.

4.9.3 Web

In order to get an IPv6 web presence, it is usually enough to implement IPv6 on the front end of all web servers. There is no initial need to upgrade any back-end database or back-end server, as those servers are never accessed directly from the Internet. There are multiple ways of adding IPv6 connectivity to a web server farm. A few of these methods use address family translation (AFT):

- **Adding native IPv6 to existing web servers:** Configure IPv6 on the web server itself. Most modern web servers have supported IPv6 for several years. This is the clean and efficient way to do it. Some applications or scripts running on the web servers may need some code changes, particularly if they use, manipulate, or store remote IP addresses of their clients.

- **Adding a set of standalone native IPv6 web servers:** Configure standalone web servers separately from the IPv4 infrastructure. This approach has the benefit of reducing dependencies on other components, perhaps even allowing selection of different hosting providers for IPv4 and IPv6.

- **Server load balancers (SLB):** Load balancers are able to have clients connecting over IPv6 while the physical servers still run IPv4. They do this by translating back and forth between the two address families (IPv4 and IPv6). This is probably the easiest way to add IPv6 to the web servers. Without a specific configuration, some information is lost in the web servers' logs because all IPv6 clients will appear as a single IPv4 address. However, with IETF RFC 7239 [i.12] a Forwarded: header (well-known practice of injecting an X- identified.

- **Reverse web proxies:** If reverse proxies are used (for example, to enforce some security policies), they similarly can be used to perform AFT, with the same caveat as for load balancers.

- **Network Address Translation (NAT64):** The Internet Engineering Task Force (IETF) has specified AFT to be done in network devices when the connection is initiated from an IPv6-only host to an IPv4-Only server. The specification includes both stateful and stateless translation methods. An enterprise may desire to use stateless NAT64 in front of an IPv4 web farm, where the clients connect from native IPv6.

- **Enabling IPv6 via CDN:** Nowadays, an increasing number of content delivery networks (CDNs) provide an IPv6 proxy for the enterprise on their public-facing web presence. For example, Akamai® and Cloudflare® both support IPv6 in their infrastructures today. Any customer of these CDN services can request dual-stack delivery of their content, and by proxy they become IPv6 reachable over the Internet.

4.9.4 VPN

VPN connections are one of a number of applications that are known to have challenges with traversing CG-NAT environments. VPN remote access and site to site works today with IPv6 natively. Both IP Security (IPsec) and SSL VPNs work well today and can transport both IPv4 and IPv6 connections over IPv4 or IPv6. Over the long term, enabling IPv6 at the head end will make it easier for IPv6-only clients to connect.

4.9.5 DNS

DNS is a critical piece of any Internet presence, as it is used to announce the IP addresses of the web and email servers. There are two steps to fully support IPv6 on a DNS server:

- **IPv6 information in the DNS zones:** Adding the IPv6 addresses of all public servers in the DNS database involves simply adding specific resource records (RRs) with the IPv6 address (those records are called AAAA). To facilitate debugging and operation, it is also advisable to add the reverse mapping of IPv6 addresses to fully qualified domain names (FQDNs). For dual-stack servers, there are two RRs per FQDN: one IPv4 address (type A) and one IPv6 address (type AAAA).

- **IPv6 transport of DNS information:** The DNS server accepts DNS requests over IPv6 and replies over IPv6. It's more common to have a dual-stack DNS server accepting requests and replies over IPv4 and IPv6.

These two steps are independent. In order to have an Internet IPv6 presence, only the first step is required; that is, the enterprise is required to publish the IPv6 addresses of all its Internet servers in its DNS zone information. All major DNS server implementations (including Cisco Prime™ Network Registrar, ISC DNS BIND https://www.isc.org/, and Microsoft DNS Server) have supported IPv6 for several years. This step does not need IPv6 connectivity to the DNS servers.

Table 3: Hostname to IP Address example

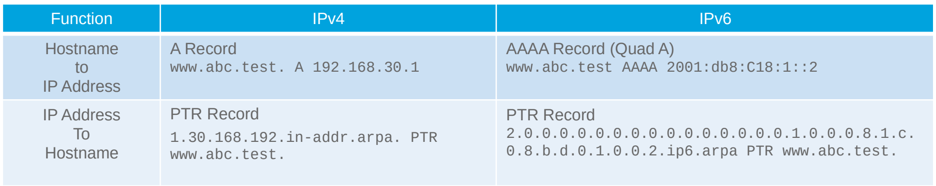

Function	IPv4	IPv6
Hostname to IP Address	A Record `www.abc.test. A 192.168.30.1`	AAAA Record (Quad A) `www.abc.test AAAA 2001:db8:C18:1::2`
IP Address To Hostname	PTR Record `1.30.168.192.in-addr.arpa. PTR www.abc.test.`	PTR Record `2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d.0.1.0.0.2.ip6.arpa PTR www.abc.test.`

4.9.6 NAT

As IPv6 has no shortage of address space, there is no reason to deploy NAT for IPv6.

The removal of NAT represents a simplification, not just to an enterprise's network design, but also to application designs. The dubious security value of IPv4 NAT is easily replaced by any stateful firewall solution for IPv6 (which can, of course, be complemented by other security techniques such as Intrusion Prevention Systems (IPS)). For this reason, IPv6/IPv6 NAT has not been specified by the IETF.

NAT also breaks the end-to-end connectivity model and either breaks or complicates application deployment. This mostly applies to applications that embed IP address semantics inside the application payload, which requires the NAT gateways to implement Application-aware Logic (ALG). Another concern is how NAT impacts host mobility and how connectivity works for device roaming between the inside and outside of network domains. Therefore, the eventual removal of NAT represents a simplification not just to an enterprise's network design but also to its application designs. Audits are also more complex with NAT, as all NAT logs should be kept.

IPv6/IPv4 NAT (NAT64, which is a specific AFT technique) does have applications today. NAT64 is a technology that facilitates communications between IPv6-only and IPv4-only hosts and networks. This solution allows enterprises to accelerate IPv6 adoption and also helps with IPv4 address depletion at the same time. While NAT64 supports several translation scenarios, stateless NAT64 (IETF RFC 6145 [i.16]) and stateful NAT64 (IETF RFC 6146 [i.9]) are the two most common use cases:

- **Stateless NAT64:** Maps IPv6 addresses to IPv4 addresses and vice versa without maintaining any bindings or session state. It supports both IPv6-initiated and IPv4-intiatied communications (1-to-1 translation).

- **Stateful NAT64:** Similar to stateless NAT64 except that it creates or modifies bindings/session state while performing translation. It supports both IPv6-initiated and IPv4-initiated communications with static or manual mappings (1-to-N translation).

It should be noted that NAT64 suffers from the same operational and security issues as NAT.

4.9.7 Multihoming

Multihoming on the Internet edge of an enterprise network refers to having redundant reliable paths through one or more ISPs. The criticality of application uptime serving Internet content and e-commerce has typically facilitated the enterprise requirement for multi-path and/or multi-provider deployments. Larger enterprise environments typically solve problems such as asymmetric routing and prefix advertisement with a combination of NAT and Border Gateway Protocol (BGP) peering.

When deploying IPv6 for those environments, medium-sized enterprises using multi-provider designs can benefit from a recent technology: IETF RFC 6296 [i.10] Network Prefix Translation (NPTv6), a stateless prefix-swapping technology operating on the network topology portion of an IPv6 address while still allowing inbound access.

Another technology that not only solves multihoming but also more effectively load-balances traffic across multiple service provider links is Locator/ID Separation Protocol (LISP) IETF RFC 5375 [i.8]. LISP is a very powerful set of services and tools that reduces the operational burden of tuning BGP for load balancing and provides an extra layer of resilience at the Internet layer.

Waiting in the wings (and with a promise of far more options in future) is Source Address Dependent Routing for IPv6 (SADR) [i.14].

The approach here will allow an application/host to select the next hop from multiple options by virtue of the fact that the host itself will have multiple global addresses and prefixes allocated, and selection of the correct source address will in turn determine the correct upstream egress route.

4.9.8 Email

The sending and receiving of email messages over the Internet occurs through Simple Mail Transfer Protocol (SMTP) over TCP. Most popular mail transfer agents (MTAs) are fully capable of using IPv6. Email reputation (which is the measurement of email sending practices) also supports IPv6.

Google Gmail™ webmail service and Microsoft Office 365® are IPv6 enabled. See https://en.wikipedia.org/wiki/DomainKeys_Identified_Mail [i.15].

4.10 Cloud

As more services and operations move to the cloud, operations are affected in many ways. The dependence on the Internet becomes mission critical for both the cloud provider and the customer. From the cloud provider standpoint, being IPv6 enabled increases site reliability, as resources may be reached over two protocols that are orthogonal to each other. In addition, peering paths for IPv4 and IPv6 tend to be different, which enhances the path diversity to reach their site.

Having IPv6 allocation for the tenants also provides some of the following operational efficiencies:

- Growth: Cloud services are growing at an exponential rate. The entire IETF RFC 1918 [i.1] can easily be consumed, as it provides only 17 891 328 addresses. This sets a limit on the cloud provider's operation.
- Management and cost saving: Growing past IETF RFC 1918 [i.1] space requires NAT/CDN within the cloud provider's own network, which in turn requires more equipment and more complex network designs. This can all be avoided by deploying IPv6 to customer resources.
- Increased uptime: From the cloud customer's point of view, having the cloud provider running IPv6 also increases the reachability of the provider's resources and tools. If for some reason IPv4 routes are compromised or dropped, the provider can still be reached over IPv6. In this case, IPv6 addressing is contributing to the cloud provider's reliability/backup potential.

Criteria for selecting an IPv6 Cloud Provider

The arrival of cloud computing and cloud-based services presents enterprises with another area where the impact of IPv6 will require careful consideration. Among the areas of concern are the following:

- Is the cloud provider dual-stack capable? If not, is their roadmap for support acceptable?
- Is the IPv6 support provided with full feature parity to IPv4? If not, is their roadmap for support acceptable?
- Are there any performance limitations connected with using IPv6 (tunnels, for example)?
- Have security considerations been taken into account? In particular, are IPv6 First Hop Security controls in place where relevant?
- How is the cloud provider's IPv6 network connected upstream? Pay particular attention to transit and peering, which may be completely different from the legacy protocol.
- Are any options offered that help with transitioning to IPv6?

Because public addresses are needed in the cloud, cost pressures on IPv4-based services are already being seen. Some providers are actually offering lower rates for IPv6 services than for those that require the legacy protocol.

4.11 Management

Monitoring and management capabilities of the network:

- Does the software have the capacity to measure and "see" IPv6?
- IPv6 capability of the directory infrastructure (e.g. name resolution, address assignment, site management, authentication, logging) equivalent to IPv4?

Network management tools capable of operating an IPv6 network:

Network monitoring and management tools within the organization and IP Address Management (IPAM) tools need to be capable of working in an IPv6 environment. IPv6 networking cannot successfully be deployed if the network management toolset is not IPv6-capable.

Any tool that monitors network activity should be reviewed to make sure that it could handle the new address format. IPv6 requires that tools use the most updated MIBs in SNMP and version 9 of NetFlow, for instance. Similarly, any tools that perform packet analysis, inspection, or access control have to be reviewed.

4.12 Security

4.12.1 Setting the scene

Secure deployment and understanding of risk are key criteria for the successful deployment of IPv6 in the enterprise.

At a minimum parity with the legacy protocol should be expected. In fact, the majority of security concerns do not change with the introduction of IPv6. The differences occur when the protocol specifics become important. IPv6 introduces the concept of extension headers. Some types of extension headers have been deprecated and others may be blocked, depending on the policy and needs of the enterprise. Another change to the protocol occurs in the extensive use of the Internet Control Message Protocol (ICMPv6). Certain ICMPv6 message types need to be allowed through the firewall in order to provide connectivity, while other types may be blocked or allowed per policy.

There has been guidance from network operators in the form of best common practices (BCPs) related to IPv6. These BCPs include bogon filtering and anti-spoofing techniques. It is expected that an enterprise security policy will be updated to properly allow and control IPv6 traffic. Most intrusion prevention systems (IPS) have adapted to IPv6 and function similar to the legacy protocol designs. Another critical element in securely operating IPv6 is to ensure that the security incident and event management (SIEM) systems are capable of providing the forensics and correlation required by the enterprise security policy.

Of interest to enterprises about to embark on an IPv6 deployment are the challenges found when retrofitting security controls over existing IPv4 deployments. When deployments do not have all modern security controls in place, retrofitting things such as proper zoning of addressing to simplify controls and first hop security has not been trivial or without impact. As IPv6 is (in many cases) a greenfield for enterprises, architects have the ability (from the outset) to enable a secure environment, so that in the event they need to tighten access in the future, the framework is laid and the enterprise can be agile in the deployment of additional controls.

4.12.2 First Hop Security

When a device has an IPv6 stack enabled, it will automatically send router and neighbour solicitations (to find network information). A rogue device could (through either misconfiguration or malicious intent) provide that information via router advertisements. If that were to occur, there would be many possibilities of Man-in-the-Middle (MiM) attacks. Furthermore, unchecked malicious or misconfigured hosts could attempt to spoof, steal, or deny service to other hosts in the access layer domain.

Many access switches provides specific security features (for both IPv4 and IPv6 threats) to handle these types of issues and increase the security level.

5 Lessons learned: IPv6 touches everything

What has industry learned along the way? First and foremost, that introducing IPv6 is a long journey. Dual-stack network infrastructure will likely endure for many years, for as long as IPv4-only devices remain in place. In the meantime, organizations that rely on the Internet should undertake a transition now. Remember by the time the network is asked for IPv6 for competitive purposes, it will be too late.

Another key lesson is that IPv6 is not just a network challenge. Moving to IPv6 is more than a network protocol upgrade since everything in IT is connected to the network, and the network in turn touches everything in the IT environment. It requires an across-the-board, holistic-IT approach. It touches everything from server and desktop operating systems to office productivity suites, ERP platforms, email, web services, management software and security tools. IPv6 impacts the entire IT ecosystem.

One key aspect that is often underestimated by some IT professionals is the challenge of application IPv6 enablement. While providers of off-the-shelf software have been dealing with this challenge by writing applications to be IP version independent, home-grown applications are behind and will also need to be updated to accommodate IPv6. Application teams are going to have to add IPv6 support to their applications and test IPv6 in their applications. Enablement of IPv6 applications happens one application at a time.

Plan ahead with partners, users - Moreover, achieving IPv6 deployment is not entirely within any single organization's hands. Beyond internal infrastructure and applications, every organization should plan with partners and end-users. When transitioning to IPv6, inventory is in order. Interactions with partners and suppliers to make sure they are ready to move forward, is necessary.

6 Conclusions

The prospect of transitioning to IPv6 may be daunting, but everyone who relies on the Internet faces the challenge. The worst mistake, is assuming the IPv6 transition can wait. Anyone who relies exclusively on IPv4 will eventually be put at a competitive disadvantage. The legacy IPv4 based Internet can no longer grow. Without deployment and support of IPv6, it is just a matter of time before networks/businesses become isolated and unable to communicate. The transition to date has been gradual but the steep curve is starting and it is critical to be prepared. For the past 30 years the IT industry has embedded IPv4 related knowledge in all its processes, in all its infrastructure gear like network management tools, load balancers, firewalls and unfortunately, in all applications. So it will take time.

There is no single recipe for IPv6 transformation. Each enterprise is unique and depends on its unique business goals, long-term vision and constraints. It is critical to put in place a joint Business & IT Task Force. This will help ensuring a smooth path toward IPv6. A pragmatic roadmap for an IPv6 transition, while also developing clear business benefits that can be achieved through the transition, is needed.

Being intentional and methodical, while keeping all stakeholders informed and accountable, is the best way to ensure a smooth transition from IPv4, to dual protocol, and eventually to IPv6 only. Reviewing the advantages of IPv6 while planning and deploying an IPv6 network is essential. Previous IPv4 limitations which dictated a specific network topology may no longer exist with IPv6. This might allow for a more efficient network design, which over time, may lead to cost savings for network management.

Annex A: Authors & contributors

The following people have contributed to the present document:

Rapporteur:
Yanick Pouffary, Hewlett Packard Enterprise

Co-Authors

Yanick Pouffary - Hewlett Packard Enterprise

Eric Vyncke - Cisco

Justine Vick & Madhuri Kaniganti - Microsoft IT

Contributors

Eric Beauchesne - Microsoft Consulting Services

Annex B: Bibliography

IETF RFC 4852 (2007): "IPv6 Enterprise Network analysis - Layer 3 Focus".

IETF RFC 6106: "IPv6 Router Advertisement Options for DNS Configuration".

IETF RFC 6147: "DNS64 - DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers".

IETF RFC 6180 (2011): "Guidelines for Using IPv6 Transition Mechanisms during IPv6 Deployment".

IETF RFC 6342 (2011): "Mobile Networks Considerations for IPv6 Deployment".

IETF RFC 6782 (2012): "Wireline Incremental IPv6".

IETF RFC 7381 (2014): "Enterprise IPv6 Deployment Guidelines".

IETF ID draft-ietf-v6ops-design-choices-12: "Routing-Related Design Choices for IPv6 Networks".

Annex C:
Change History

Date	Version	Information about changes
V0.0.1	July 2015	Document creation
V0.0.2	January 2016	Input from HPE, Cisco and Microsoft
V0.0.3	January 2016	Input from HPE, Cisco and Microsoft
V0.0.4	February 2017	Input from HPE, Cisco and Microsoft ,Spell checks, ETSI drafting rules compliance
V0.0.5	May 2017	UK spell checking and Edits following Christine' inputs

History

Document history		
V1.1.1	June 2017	Publication