Guidelines and Process: IPv6 for Public Administrations in Europe

Part of a Study on Implementation of the ISA2 Programme Action 2016.10 - IPv6 Framework for European Governments – SMART 2016/0099

Erion
Plum Consulting
Internet Policy Advisors
Synetergy
iDate

9 November 2018

Plum Consulting, London
T: +44(20) 7047 1919, **www.plumconsulting.co.uk**

Table of Contents

Executive Summary

This document provides IPv6 address planning guidance for public administrations. It is intended to provide a framework that public administrations can use to learn the key differences between IPv6 and IPv4 addressing, design an IPv6 address structure, obtain IPv6 address space, deploy IPv6 addresses and manage IPv6 addresses. In addition, this guide also provides useful background information on the reasons public administrations should adopt IPv6 and how to plan for an IPv6 deployment project.

The guide is split into the following main sections:

- Drivers for adopting IPv6
- An introduction to IPv6 addressing basics
- An overview of IPv6 deployment planning
- Guidance on design, obtaining, implementing and managing IPv6 addressing

This guide is focussed on those aspects of IPv6 addressing and IPv6 deployment that relate to public administrations in Europe.

The IPv6 addressing planning sections of this guide begin by looking at some important principles that all organisation should consider when designing an IPv6 addressing plan.

This is followed by a detailed consideration of the options available to public administrations in Europe for obtaining IPv6 address space and the different types of address space that can be obtained. In this section, we provide clear guidance on how public administrations should chose and obtain IPv6 address space. In Europe, IPv6 address space can be obtained from an upstream provider or from the regional internet registrar for Europe: RIPE. Two types of space are available; Provider Independent (PI) space and Provider Aggregatable (PA) space. We discuss which type a public administration might use and how to calculate the size of space that they require.

Next, we show how a public administration can structure the IPv6 address that they have obtained into an IPv6 addressing plan. We consider the various options and provide best practice design principles for constructing an address structure. This is followed by several examples and a brief case study to illustrate different types of address structures that can be used in public administrations.

Finally, we look at how IPv6 addresses can be managed and recommend the use of automated tools such as IP Address Management (IPAM) to facilitate address management processes.

1 Drivers for Adopting IPv6 in Public Administrations

Whilst this document is focussed on IPv6 addressing, it is still important to understand the drivers for adopting IPv6, particularly those that are relevant to public administrations. These will have an influence on the process of obtaining, structuring and deploying IPv6 addresses.

The primary driver for IPv6 is the exhaustion of the IPv4 address space. This was the reason that IPv6 development began several decades ago in the early 1990s. Even by then it had become clear that the IPv4 address space was going to be inadequate for the growing public internet. Indeed, had it not been for the introduction of aggressive address conservation techniques, the IPv4 internet would have run out of addresses decades ago.

Today, the exhaustion of the IPv4 address space and the adoption of IPv4 address conservation techniques are having a direct impact on the growth, functionality, operation and security of the public internet.

This is the reason why large parts of the global Internet are now IPv6 enabled. This includes almost 100% of transit carriers and the majority of the world's largest Internet Content Providers, such as, Google, Akamai, LinkedIn, YouTube and Facebook. Today, if you have both IPv6 and IPv4, you will find that most of your public internet traffic is carried by IPv6 leaving only the minority to be carried by IPv4.

Here is a summary of some of the reasons why IPv6 is being adopted by many organisations:

- The on-going deterioration of the legacy IPv4 internet
    - Impact of Carrier Grade NAT (CGN) (and NAT44)
    - Impact of routing fragmentation
    - Impact of address squatting
- The exhaustion of their stock of public IPv4 addresses
- The exhaustion of their internal RFC1918 private IPv4 address space
- To support deploying the Internet of Things (IoT)
- Due to restrictions in certain marketplaces (e.g. Apple App Store mandating IPv6-only)
- The requirements of peer-to-peer applications (to overcome NAT and CGN issues)
- Issues with cybersecurity, legal intercept and analytics
- IPv6 is the current standard for the Internet Protocol
- IPv6 forms the basis for key technologies (e.g. mobile – 4G/5G)
- Performance and operational benefits

Public administrations have additional reasons for adopting IPv6. Not only do they rely on the internet as much as any other organisation, they also need to consider supporting the internet as a driver for economic growth and addressing the regulatory and competition issues arising from the shortage of IPv4 addresses.

Finally, public administrations should be seeking to promote best practice by example through the adoption of IPv6 within their own networks and systems.

2 IPv6 Addressing Basics

Fundamental to all IPv6 deployments is an effective IPv6 addressing plan. To create an IPv6 addressing plan, it essential that you both understand IPv6 addresses and how they are used.

2.1 How IPv6 Addresses Differ from IPv4 Addresses

IP addresses are used to identify the source and destination of network traffic on the global internet and on internal networks. IPv6 addresses are 128 bits long. This is in contrast with IPv4 addresses which are 32 bits long.

It is important to appreciate that IPv6 addresses differ from IPv4 addresses in many ways, some of which result in a profound difference in how they are used. This section provides a brief introduction to the basics of IPv6 addressing.

- There are an unimaginably huge number of IPv6 addresses
    - 2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456
    - There are also an unimaginably huge number of IPv6 addresses available in a typical single IPv6 subnet (64 bits are usually used for addresses within a subnet).
        - 2^{64} = 18,446,744,073,709,551,616
        - A single subnet contains enough addresses for over a billion addresses per person on the planet. Over 18 quintillion addresses, or over 4 billion times the total IPv4 address space within a single IPv6 subnet.
- IPv6 addresses have lifetimes
    - These are the preferred and valid lifetimes. These can change with time.
- IPv6 addresses have scope
    - Scope defines the parts of the network that an address is valid in.
- A single IPv6 enabled interface can have many IPv6 addresses
    - It is usual for an IPv6 enabled interface to have at least two addresses; common to have three addresses and legal to have a very large number of addresses.
- Addresses can change with time
    - For example, privacy addresses are ephemeral and may, for example, change every twenty-four hours.
- No broadcast addresses
    - In IPv6, multicast addresses replace the role of broadcast addresses in legacy IPv4.
- Multicast addresses are very important
    - IPv6 makes a widespread use of multicast in many of its core protocols.
- There is no equivalent to RFC1918 private addresses

- IPv6 has no equivalent to RFC1918 private address space. The use of any form of private address space and/or stateful Network Address Translation (NAT) is explicitly discouraged. Unique Local Addresses (ULAs) have some similarities with RFC1918 space but they are not equivalent.

- There are many IPv6 address structures and formats
  - IPv6 addresses are more complex than IPv4 addresses.

2.2 Representing IPv6 Addresses

When you use IPv6 addresses you represent them in a textual format designed to make it easier for humans to read and enter. The address is written using hexadecimal digits separated by colons (see Figure 1).

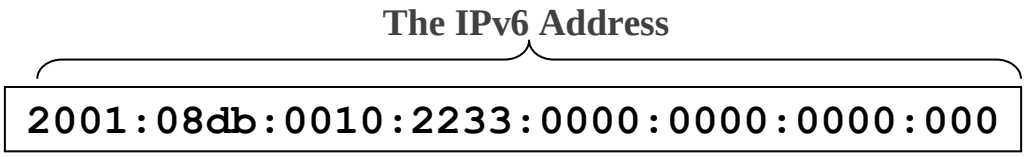

Figure 1 IPv6 Address Textual Representation

To make the address easier to read, colons are used to indicate where there is a sixteen-bit boundary. The standard for the textual representation of IPv6 addresses allows for a large variation in formats. This is because it is legal to:

- Omit leading zeros in any sixteen-bit block. (A sixteen-bit block "0001" becomes "1".)
- Write one section, between any two colons, that contains only zeros, as a double colon (::).
- Represent the bottom thirty-two bits in dotted-decimal notation when it contains an embedded IPv4 address. (For example, the IPv4-mapped address, ::ffff:192.168.1.1)

Therefore, the address in Figure 1 can be shortened to that in Figure 2.

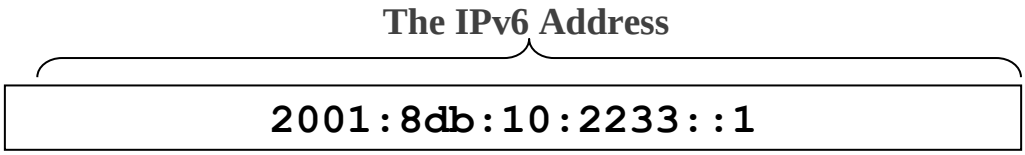

Figure 2 An IPv6 Address with Zeros Omitted

Optionally, the notation allows you to specify the interface that the address is to be used on (the percent notation) and the length of the prefix (the left-hand most-significant bits) (see Figure 3). The address prefix is the set of bits that are significant in address. This is used group addresses by prefix, for example into a subnet or in routing. Most IPv6 subnets have a 64-bit prefix leaving 64 bits to uniquely identify hosts (called nodes in the IPv6 world) within a subnet.

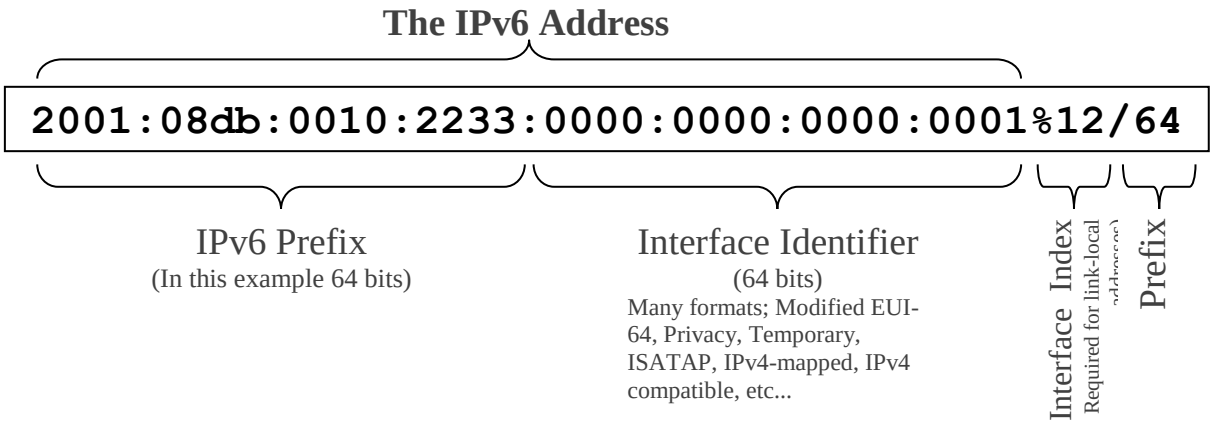

Figure 3 IPv6 Address Textual Representation with Interface Index and Prefix Length

In the examples above, the subnet prefix, specifying the subnet is written as 2001:8db:10:2233::/64.

Because IPv6 addresses can be written in many ways, a canonical form of an IPv6 address has been defined[1]. Canonical addresses are all lowercase, they suppress all leading zeros, they represent a single empty sixteen-bit field as “0” and they shorten consecutive empty fields to ::. For full details see the RFC[1].

2.3 IPv6 Address Types

There are three main categories of IPv6 addresses:

- Unicast: Represents one interface on one node
- Multicast: Represents a group of interfaces on one or more nodes
- Anycast: A group of interfaces where packets are delivered to the "nearest" interface

Unlike IPv4, interfaces on IPv6 nodes are often identified by multiple unicast addresses, sometimes with different scopes. This is a significant change which has implications for the design, management, software development and security of IP addresses.

Beyond these three categories of addresses, IPv6 addresses are further grouped by the first few bits of the address. These first few bits of an address (the prefix) are used to specify the address type and purpose. Some common and important IPv6 addresses and prefixes are listed in Table 1.

[1] See RFC5952

Address or Prefix	Description
`::/128`	Unspecified address
`::/8`	Reserved prefix
`::1`	Loopback address
`::ffff:0:0/96`	IPv4-mapped IPv6 prefix
`fe80::/10`	Link-local prefix
`fc00::/7`	Unique-local address prefix (ULA)
`ff00::/8`	Multicast prefix
`2002::/16`	6to4 prefix
`2001::/32`	Teredo prefix
`64:ff0b::/96`	Well-known prefix
`2000::/3`	Aggregatable globally assigned unicast (GUA) prefix (Public address space)
`fec0::/10`	Site-local prefix **Deprecated**
`3ffe::/16`	6bone prefix **Deprecated**
`::/96`	Use for IPv4-compatible IPv6 addresses **Deprecated**

Table 1 Examples of Assigned IPv6 Addresses and Prefixes

A full list of the registered prefixes[2] is maintained by the Internet Assigned Number Authority (IANA).

2.4 Unicast IPv6 Addresses

IPv6 has several different types of unicast addresses, some of which are shown in Table 1. The main two are:

- Global Unicast Addresses (GUAs) that are used to communicate across subnets, networks and the global internet. These are routable on the global internet. (2000::/3)
- Link-local addresses that can only be used within a subnet. These are not routable on the global internet, they are used extensively by the core IPv6 protocols. (fe80::/10)

2.4.1 IPv6 Global Unicast Addresses (GUAs)

GUAs are often referred to as public IPv6 addresses. In contrast to IPv4, IPv6 uses public addresses everywhere, even on internal networks. As a result, IPv6 has no need for Network Address Translation (NAT) and its associated proxies, gateways and traversal techniques. This is a significant change from IPv4 and has a large impact on the design of IPv6 address structures and IPv6 networks.

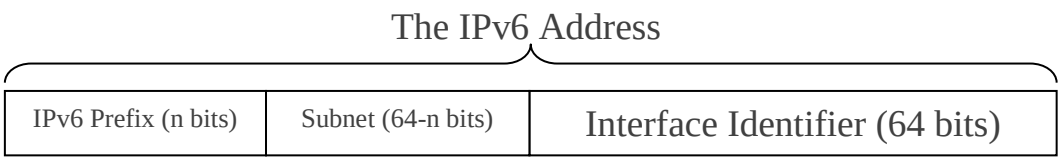

Figure 4 The Common Global Unicast Address Structure

[2] See https://www.iana.org/assignments/ipv6-unicast-address-assignments/ipv6-unicast-address-assignments.xhtml

GUAs are usually split into three main components:

1. A global prefix provided by a local internet provider (LIR) or regional internet registry (RIR)
2. A subnet specified by the organisation to identify a network
3. An interface identifier (IID) that is usually unique to the end node.

The prefix (consisting of the global prefix and the subnet identifier) is usually 64 bits long. The IID is almost always 64 bits long too. Whilst different prefix and IID lengths are possible several key IPv6 protocols require them to be 64 bits. Therefore, most subnets have a prefix of 64 bits.

2.4.2 IPv6 Link-Local Addresses

Link-local addresses are essential for the operation of many core IPv6 protocols. They can be used as any other unicast address, except that they are not routed and can only communicate between nodes within the same subnet. Every IPv6-enabled interface must have at least one link-local address, although it can have many other addresses as well.

Link-local addresses are unaffected by any changes in global prefixes.

```
fe80::020c:2384:1fe6:efab%eth0
```

Figure 5 Example of a Link-local Addresses

Since link-local addresses all have the prefix fe80::/10 it is impossible to differentiate link-local addresses by prefix or determine which interface they should be used on by their prefix. Therefore, it is normal when using link-local addresses to specify the interface as in the example in Figure 1.

2.5 IPv6 Anycast Addresses

IPv6 anycast addresses are IPv6 unicast addresses (usually GUAs) that have been assigned to more than one interface on one or more nodes. The only difference between anycast and unicast addresses is that they are used in more than one place. It is then up to routing and neighbor discovery to determine which interface or node a datagram is delivered to. Anycast addresses are particularly useful for load-balancing and resilience. For example, it is common for large DNS providers to give their DNS servers anycast addresses that are assigned to geographically distributed DNS servers. This improves performance and reliability.

2.6 IPv6 Multicast Addresses

Multicast is very important in IPv6. It used in many core IPv6 protocols. It has completely replaced the use of broadcast addresses in IPv4. All IPv6 multicast addresses begin with the prefix ff00::/8. The next four bits are used for flags. The four bits after the flag bits define the scope of the address. Common scopes are link-local, site and global. There are many more multicast scopes.

2.7 IPv6 Address Interface Identifiers

There are many formats for the final 64 bits of an IPv6 address. A public administration will need to choose which formats to use and where to use them. Historically the two most common types were manually assigned and modified-EUI64 IIDs. Modified EUI-64 addresses are IIDs that are constructed from an existing IEEE datalink address such as a MAC address.

Today, the recommended best practice is to use both semantically opaque IPv6 addresses[3] and privacy addresses[4]. However, different platforms support different sets of IID types and often they have chosen different default IID types.

[3] See RFC7217

[4] See RFC4941

3 Assigning IPv6 Addresses to Nodes

After the brief introduction to IPv6 addressing in the previous section, it is now time to look at how to assign addresses hosts and nodes. There are many ways to assign addresses in IPv6. The three main ways to assign IPv6 addresses to nodes in IPv6 are:

- Manual address configuration
- Automatic configuration using Stateless Address Autoconfiguration (SLAAC)
- Automatic configuration using Dynamic Host Configuration Protocol version 6 (DHCPv6)

3.1 Link-Local Address Configuration

As noted in the previous section, every interface on a node that has IPv6 turned on automatically configures a link-local address. Usually this address is configured using SLAAC. However, in some instances the link-local address may be configured manually.

3.2 Manual Address Configuration

Manual address configuration is often used for servers, routers, switches, firewalls, and any other network resources where addresses are unlikely to change over time. Static addresses are often used by network devices, servers and services to ensure that they can be reached at a consistent address.

3.3 Automatic Configuration using Stateless Address Autoconfiguration (SLAAC)

SLAAC (Stateless Address Autoconfiguration) is a mechanism that allows IPv6 nodes to generate their own addresses and to configure other basic network settings. A node uses SLAAC to configure an interface's link-local address. SLAAC can also be used to configure other addresses, including global IPv6 addresses (GUAs) and other network settings.

IPv6 routers can provide the configuration information that nodes need to configure network settings and non-link-local addresses. IPv6 routers do this by sending out ICMPv6 router advertisement messages that contain the necessary configuration information options.

IPv6 router advertisements do not usually assign specific addresses, instead they provide prefixes for the local subnet (usually 64 bits long) that the node can combine with a unique interface identifier (IID) to create a non-link-local address.

Router advertisements can be used to configure the node's default routers and other network parameters such as, static routes, maximum transmission unit (MTU) and DNS.

In contrast to IPv4, in IPv6, SLAAC is used to trigger the use of DHCPv6. Therefore, SLAAC is necessary even if DHCPv6 is used to assign addresses. Furthermore, DHCPv6 lacks the default router option that is found in IPv4. When DHCPv6 is used, the default router must still be configured using SLAAC.

SLAAC presents a few challenges when designing IPv6 networks. By default, router advertisements are not authenticated. Therefore, steps should be taken to protect the network against attacks that use fake router advertisements (for example using RA-guard or an equivalent technology).

There are several ways that a node can create a unique IID. Current best practice is to use both semantically opaque IPv6 addresses[5] and privacy addresses[6]. Privacy addresses use pseudo random IIDs that change with time. A node configures a privacy address in addition to its existing link-local address and a static non-local address (for example a semantically opaque IPv6 address). Privacy addresses introduce additional management challenges that need to be considered when deciding whether to use them. This needs to be taken into consideration when designing an IPv6 network.

3.4 Automatic Configuration Using Dynamic Host Configuration Protocol Version 6 (DHCPv6)

Dynamic Host Configuration Protocol version 6 (DHCPv6) is the new version of DHCP for IPv6. DHCPv6 can be used to assign static or dynamic addresses and configure many other parameters. There are several major, and some subtle, differences between DHCPv6 and DHCPv4. These include:

- DHCPv6 is enabled by flags in router advertisements7 (SLAAC is required to enable DHCPv6).
- DHCPv6 cannot configure the default router (SLAAC is used for this).
- The DHCPv6 protocol has many differences from the IPv4 DHCP protocol
- DHCPv6 is not required to provide DNS information as this can be configured using SLAAC8.
- DHCPv6 is necessary to configure parameters that cannot be configured by SLAAC.
- DHCPv6 has three modes; stateful, stateless and prefix delegation (DHCPv6-PD).

DHCPv6 is useful in aiding address management. However, it cannot provide a definitive list of configured addresses on a network. This is because an IPv6 node is not forced to use DHCPv6 to configure its addresses even if DHCPv6 is configured.

The three modes of DHCPv6 are summarised below:

- Stateful DHCPv6 is like DHCP in IPv4. It can be used to assign static or dynamic addresses as well as network configuration options.
- Stateless DHCPv6 does not assign any addresses. Instead SLAAC is used to configure addresses and DHCPv6 is used only to provide additional configuration options that are not available in SLAAC.
- DHCPv6-PD is usually used by service providers to delegate a prefix to their customer's network.

[5] See RFC7217

[6] See RFC4941

[7] The flags that configure DHCPv6 can lead different behaviours on some platforms due to ambiguities in the standards. This is beyond the scope of this document and is rarely a problem in operational networks.

[8] See RFC8106

Even in networks that use DHCPv6 to configure addresses a node may lease a long-lived address and a short-lived temporary address. So, as with SLAAC, DHCPv6 can also assign temporary addresses that are identical to the privacy addresses mentioned above. Temporary addresses have sort lifetimes of typically a day. In contrast with networks that use SLAAC for address configuration, in a network using DHCPv6, there is a central record of the assignment of privacy addresses. However, this is only for those nodes that choose to interact correctly with DHCPv6.

4 What are the Differences in IPv6 Addressing in European Public Administrations?

European public administrations vary significantly in how they obtain, structure and deploy IPv6 address space. These variations arise from the wide range of differences in how public administrations are organised.

Important differences that influence IPv6 addressing in public administrations include:

- Having a centralised IT function verses having a decentralised IT function
- Having a federal structure verses a non-federal structure
- Having departments organised by function verses being organised by geography
- Differences in the legal and contractual frameworks.

In the section Obtaining Global IPv6 Addresses, we will look at how public administrations differ in the way that they ***obtain*** IPv6 addresses. In the section IPv6 , we look at how that space can be ***structured***.

5 Overview of Planning IPv6 Deployments in Public Administrations

IPv6 address planning is a crucial part of any IPv6 deployment. This section sets the context for IPv6 address planning by providing a brief overview of planning for IPv6 deployments in public administrations. These topics are covered in more detail in the materials from the workshops that were given as a part of this project[9].

Public administrations face many challenges in justifying, planning and implementing IPv6 deployments. Furthermore, there are a wide range of differences between how individual public administration's approach IPv6 deployment. This section focuses on the IPv6 planning principles that are applicable to most, if not all public administrations.

Before looking at planning IPv6 deployments it is important to cover two key areas that can be challenges for public administration. These are:

- Justifying the IPv6 deployment
- Obtaining management or ministerial support
- Budgeting for the IPv6 deployment

In both cases it is important to remember that IPv6 deployment does not have to be on-hold until it is justified or until budget is assigned to it. With suitable policies in place, IPv6 can often be deployed as a part of on-going network maintenance, upgrades and deployments. Indeed, to some extent, you will have already deployed IPv6. For example, all modern operating systems are IPv6 enabled by default. Therefore most, if not all, of your desktop and server platforms will already be IPv6 capable and most are likely to be using IPv6 by default now.

5.1 Justifying the IPv6 Deployment

Whilst some preparation for IPv6 deployment should already be taking place as a natural part of your on-going IT operations. Fully deploying IPv6 is best achieved through a centralised project with common standards and goals. Therefore, it is crucial to be able to justify IPv6 deployment and to obtain the appropriate ministerial support.

The exact nature of an IPv6 deployment plan (and indeed if a centralised project is possible) will depend on the departmental and regional structure of the public administration.

We have already covered many of the reasons for deploying IPv6 that are relevant to public administrations in an earlier section; Drivers for Adopting IPv6 in Public Administrations. So, we will not repeat these here. It is important that when you are seeking justification that you remember that there are many areas where IPv6 can be deployed without justification. The example mentioned in the previous section of desktop and server platforms is one such case.

Whilst a single coordinated deployment project is best practice, this approach sometimes gives the wrong impression. A single all-encompassing project may imply that IPv6 is a bigger step than it is or that it is not the natural evolution that it is. So, presenting IPv6 as a single all-encompassing project can sometimes be a help and sometimes it can be a hinderance.

[9] The workshop materials can be found at, http://ipv6gov.eu/.

Also, you should consider that there can be differences in the justification for different areas of deployments. For example, public verses internal deployments. Sometimes, for example, deploying IPv6-only networks in a greenfield site (such as a new data centre) can lead to cost savings and greater operational efficiency.

No one justifies the continued use of IPv4. You should seek parity for IPv6. It is reasonable to ask, why are we continuing to invest in legacy IPv4?

You should also engage your suppliers in the justification process. There are many suppliers who are keen to support IPv6 adoption. They are often eager to help public administrations deploy IPv6.

Finally, you should obtain executive and/or ministerial support for IPv6 deployment. Who and how you obtain this will depend on the structure of the public administration. A clear mandate for IPv6 is extremely important.

5.2 Obtaining Management or Ministerial Support

Any IPv6 deployment needs some level of executive, management, ministerial or related support and approval. Support and approval are necessary to obtain the capital and operational budget assignments for the deployment effort.

Experience shows that it can be difficult to acquire this support and approval. IT managers, for instance, have been ambivalent about IPv6 for many years. In public administrations it can be difficult to translate the benefits of IPv6 adoption into tangible business cases that make sense to public policymakers. It's often the case in public administrations that early-stage IPv6 adoption tasks, such as first steps, are taken with no formal project definition, additional funding or explicit approval and support. In these cases, an IT department often simply absorbs the IPv6 project into existing operational cycles and budget. However, this does not scale well to full, strategic IPv6 adoption efforts. Instead, IPv6 adoption should be justified as a unique project, one that requires dedicated planning and resources.

Depending on how an administration is organised, the executive buy-in may be on a departmental or regional basis. This allows an IPv6 adoption activity to proceed in a separate business unit or department without having to make a mass migration necessary immediately. In some units of governments, this can amount to an unfunded mandate by a department head or minister to "go figure out what we need to do to implement IPv6." Even without budget, this is at least one way to have management buy-in for early IPv6 adoption.

5.3 Budgeting for the IPv6 Deployment

Budgeting is the most common barrier mentioned by member states to the adoption of IPv6. Public administrations tend to view IPv6 as an additional cost for something that they already have: the internet. They often find it hard to understand or appreciate the drivers for adopting IPv6.

Budget constraints can affect many parts of an IPv6 deployment project including, staffing, capital costs, operational costs, planning, design, project management, readiness audits and staff training.

Therefore, addressing the budgeting challenges is key for the success of an IPv6 project. Towards this end we suggest several actions that can assist with limiting or overcoming some of the budgeting issues. First, here are three IPv6 budgeting principles:

- Integrate IPv6 into your normal on-going budgets. Include IPv6 in everything that you do regardless of whether you have an on-going IPv6 project. For example, in equipment refreshes, software development and new purchases.
- You don't have to do everything at once. IPv6 is designed to be deployed piecemeal, in a wide variety of different ways. This allows you to target easy-wins and avoid a single large budget item.

As much as possible decouple IPv6 from budget.

There are also several strategic approaches to IPv6 budgeting that you can use. These include:

- Implement strategic IPv6 policies to minimise unnecessary expenditure
- Carefully limit the scope of your IPv6 deployment project
- Include IPv6 by default in other non-IPv6-specific network projects
- Embed IPv6 in to everything that you would normally do with IPv4
- Utilise zero-cost options whenever possible
- Make sure that greenfield deployments are IPv6 from day-one
- Some key project activities should already form a part of other budgets; for example, you should already have a security policy that includes IPv6
- Only specifically budget for key essential IPv6 deployment activities

One of the most effective policies is to mandate IPv6 readiness in purchasing. You should set administration wide purchasing policies that:

- Mandate that all purchases must be IPv6-ready and capable of IPv6-only operation
- Mandate that all software development is IPv6-ready and capable of IPv6-only operation
- Mandate that all ICT job descriptions that include a knowledge of networking and particularly IP specifically require IPv6

It is important for a public administration to set these policies regardless of support for an IPv6 deployment project as they avoid wasteful investment in legacy technologies that will require additional spending in the future. Not doing so is a bad stewardship of public funds.

Some IPv6 deployment is zero cost. You should take every opportunity to make use of many zero-cost options for deploying IPv6. You should also emphasize to management that there are many aspects of IPv6 deployment that are cost free. For example:

- Modern operating systems are IPv6 by default and are IPv6 enabled by default
- Windows has an IPv6 stack that provides legacy support for IPv4
- Equipment refreshes will often be IPv6-ready even if you didn't mandate it

Furthermore, some services provide IPv6 at no additional cost. For example, several cloud and Content Delivery Network (CDN) operators will IPv6 enable your services at the edge for no additional cost (often all this requires is checking a tick box – in some instances the default is for this is to be enabled)

There are some costs that are difficult to avoid. So even though a potentially significant portion of an IPv6 deployment may not need an explicit budget, some elements will.

Finally, it is worth noting that whilst an IPv6 deployment usually increases costs, there are areas where sometimes savings can be made. For example, the cost of purchasing additional IPv4 address space can be reduced by freeing up IPv4 addresses through the deployment of IPv6. IPv6 can also lead to savings in certain infrastructure areas, such as VPN concentrators, removal of NAT/CGN and the operational overheads of complex private IPv4 addresses. In the longer term, the removal of IPv4 and deployment of IPv6-only networks will simplify network operations. In some cases, it is even possible to generate income from the deployment of IPv6. For example, some administrations have been able to free up IPv4 address space that they have then sold on the IPv4 address transfer market.

5.4 Planning the IPv6 Deployment

Most IPv6 deployment projects have a common set of key deployment activities. These are interrelated and are often carried out in parallel. An overview of a typical deployment project is shown in Figure 6. This shows the key deployment activities of an IPv6 project that are listed here:

- IPv6 project planning activities:
    - Define the project goals and scope
    - Plan an IPv6 Awareness and Education Programme
    - Define an IPv6 Deployment Strategy
    - Create an IPv6 Address Strategy
    - Develop an IPv6 Security Strategy
    - Plan the phases of the IPv6 Deployment
- Key IPv6 project activities:
    - Obtain IPv6 Address Space
    - Initiate an IPv6 Awareness and Education Programme
    - Carry out a Readiness Audit and/or Pilots/Trials

This guide is focused on the IPv6 address elements of these activities.

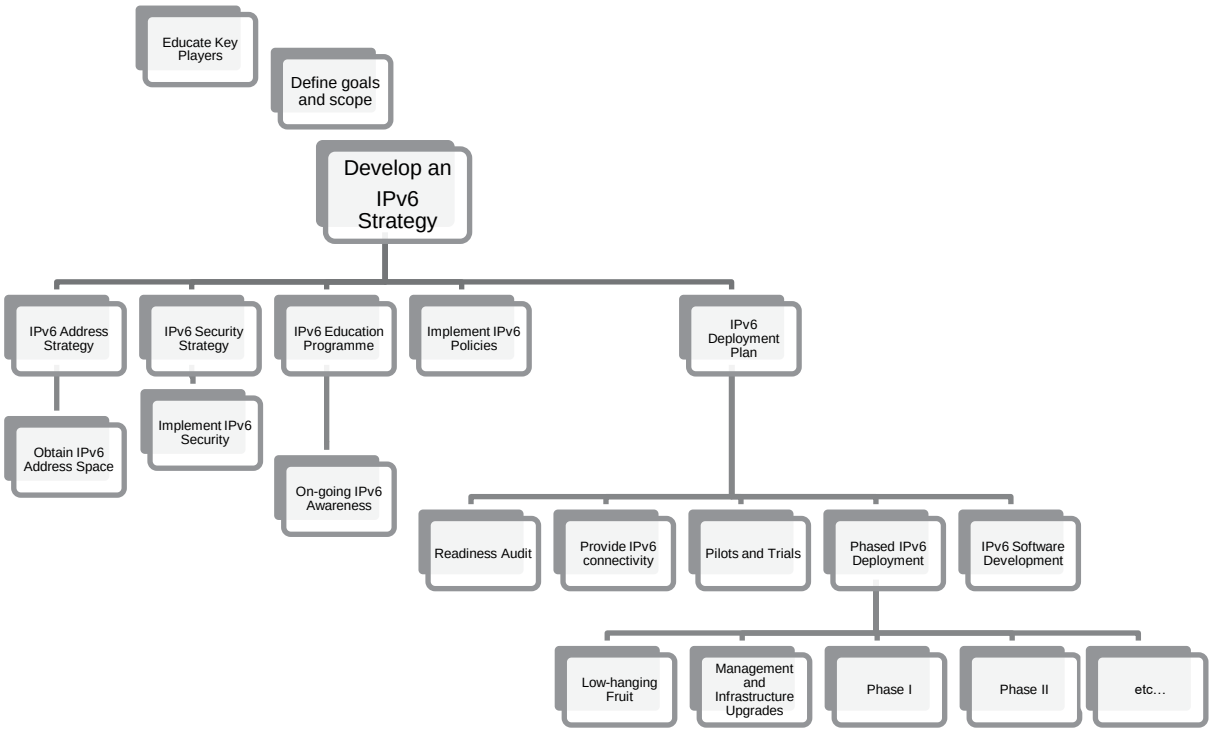

Figure 6 Overview of the IPv6 Deployment Process[10]

Two early activities that you should prioritise are the training of key staff and the setting of project goals and aims. Training for key project staff is essential to ensure that they make valid decisions and so that they can communicate knowledgably about IPv6 with other parties. Setting the goals and aims is crucial to the project's success.

The following sections provide a brief overview of some of the elements of the IPv6 deployment process.

5.5 Defining the IPv6 Deployment Project Goals and Scope

The long-term goal for an IPv6 deployment project should be to migrate to an IPv6-only network. There will be short-term and intermediate goals, such as deploying dual-stack and enabling IPv6 at the edge.

Ideally the scope should ensure that all areas impacted by a protocol and addressing change will be addressed. This includes applications, information, computing platforms, networking, infrastructure services, processes, standards, security, governance, buildings, sites, transport, communications, media, human resource, etc.

[10] Based on Erion's IPv6 Deployment Process and used with permission.

5.6 IPv6 Awareness and Training

IPv6 education is pivotal to the success of an IPv6 deployment project. One of the most significant risks in an IPv6 deployment project is that of legacy-IPv4 thinking. This is because technology professionals, including those with extensive networking experience, often do not appreciate the extent or significance of the differences between IPv4 and IPv6. For example, it is not unusual for IPv4 best practice to be the opposite of IPv6 best practice.

Therefore, IPv6 education is critical to an IPv6 deployment project. So that staff can both learn IPv6 protocols and technologies and so that staff can change their thinking from the IPv4 world to the IPv6 world.

You need to ensure that all staff are competent in IPv6 and that they are working to an IPv6 agenda not a legacy IPv4 agenda. We have found that all organisations underestimate the IPv6 training that they require.

IPv6 training can take many forms, from awareness through self-learning to formal, instructor-led, technical training[11].

In addition, your IT hiring policies should mandate that new hires of technical staff can demonstrate that they already have IPv6 skills and experience.

It is especially important that key players in the project have a good understanding of IPv6 and are capable of being advocates of IPv6 adoption.

5.7 Strategic Planning

In addition to creating an IPv6 addressing plan, as described later in this document, you should create at least two other strategic documents; an IPv6 deployment strategy and an IPv6 security strategy. The rest of this section describes some of the elements of the IPv6 deployment strategy. An IPv6 security strategy is important because IPv6 is already enabled on most platforms by default. This means that your network is already vulnerable to many IPv6 vulnerabilities. Ideally an organisation should have implemented IPv6 security at least a decade ago. If you have not yet done so, you should take measures to deploy IPv6 security.

5.8 Creating an IPv6 Task Force or Stakeholder Group

Creating an IPv6 Task Force or a Stakeholder Group can be an effective way of coordinating support for an IPv6 deployment. They can bring together government and industry players and create awareness, generate support and assist with the case for IPv6. In addition, they can also bring together key players that are necessary for success such as service providers

Some public administrations have created national or public administration IPv6 Task Forces. Others have provided support for existing IPv6 Task Forces. It has been noted that even a single meeting can have a significant effect on the deployment of IPv6 within an administration.

[11] For example, see the range of course at www.ipv6training.com.

5.9 Obtaining IPv6 Connectivity and Transit

Native IPv6 transit is a key prerequisite for an IPv6 deployment. It is something that can be obtained prior to an internal deployment of IPv6. Whilst IPv6 can be deployed without native IPv6 connectivity this does not provide the necessary scalability, manageability, reliability and performance necessary for a government or enterprise deployment. Therefore, non-native connectivity should be avoided. Most transit providers today are IPv6 enabled.

In addition to transit connectivity, you will need native IPv6 connectivity from your service providers. When seeking IPv6 connectivity from service providers you should beware of immaturity or inexperience within your service providers. You should also beware of service providers that provide IPv6 as a second-class IPv6 service to IPv4.

5.10 IPv6 Readiness Audit and Gap Analysis

You need to estimate the gap between what you have and what you need to deploy IPv6. You need to identify potential problem areas that may impact the IPv6 deployment. Some organisations carry out a comprehensive audit and others only look at specific areas of interest.

A comprehensive audit is not always necessary for two reasons:

- Modern hardware, operation systems and software are generally IPv6 ready.
- You may be able to push responsibility for ensuring that products are IPv6 ready onto the suppliers. If they claim that their products are IPv6 ready they should be prepared to ensure that they are.

You need to specify what you are going to audit and what you are looking for. A typical gap analysis should cover:

- IPv6 support and functionality in network infrastructure, nodes, services, security, applications and management systems
- Resources to support the addition of IPv6 (e.g. memory)
- Management tools
- Expertise (see education/training)

Several public profiles exist for IPv6. These are not always kept up to date, but they do provide a useful starting point:

- Requirements for IPv6 in ICT Equipment (RIPE-554)
- IPv6 Ready Logo Program of the IPv6 Forum
- A Profile for IPv6 in the US Government (V1, NIST, 2008)
- Department of Defence Unified Capabilities Requirements 2013 (UCR 2013)
- IPv6 Node Requirements (RFC6434)
- IPv6 Profile (www.bmi.bund.de)

Finally, for key systems and network infrastructure it is prudent to carry out a pilot deployment to ensure that everything works as expected. This is also useful for staff to gain experience with IPv6.

5.11 The IPv6 Deployment Plan

IPv6 has been designed so that it is very flexible in the order and how it can be deployed. This provides public administrations with a great deal of choice in how they deploy IPv6 and in what sequence. Discussing all the options in detail is beyond the scope of this document. However, here are a couple of examples of the types of things that should be considered.

At the edge, it may be possible to IPv6-enable public facing services that are hosted externally very quickly. Some Content Deliver Networks (CDNs) and cloud platforms provide the option of IPv6-enabling a service simply by checking a tick box (some have this enabled by default, e.g. Akamai).

Internally, as noted in a previous section, you will need IPv6 connectivity to the public internet before enabling anything else. Once you have that connectivity you have a wide range of options in how you enable IPv6 to the rest if the network. It is usual to enable core network services first and then enable data centres or offices. Some organisations have chosen to begin with specific areas of the network, for example Wi-Fi guest access, before deploying more widely.

6 General IPv6 Addressing Principles

When designing and deploying IPv6 addressing there are some general principles that you should adopt. We have split these into high-level design principles and low-level design principles. The high-level principles are:

- Education. It is crucial to ensure that staff are competent to design, implement and manage IPv6 addresses. Legacy IPv4 knowledge is not enough.
- Longevity. Ensure that your design avoids the need for renumbering and has capacity for growth. Aim to avoid having to obtain additional address space in the future. Obtain the shortest prefix or prefixes that you can justify.
- Flexibility. Geographic boundaries can change, organisational structures can change, and department responsibilities can move. Create a design that can accommodate change.
- Simplicity. Avoid the temptation to create a complex structure.
- Manageability. Consider operational and security issues.

The low-level design principles deal with the detail of the addressing structure. They should include:

- Routing efficiency. Build in opportunities for route aggregation. Create a hierarchical structure.
- If Possible, use PI Space. Provider independent address space (PI) avoids the renumbering risks associated with using provider aggregatable (PA) space owned by a provider.
- Align Major Address Boundaries on Nibbles. Nibbles (four bits) are the smallest unit represented by a hexadecimal digit. Aligning major boundaries on nibble boundaries makes an addressing scheme easier to read and facilitates reverse DNS lookups which are delegated on nibble boundaries.
- Avoid Encoding Decimal Numbers as Hexadecimal Numbers. Instead train your staff to use hexadecimal. It is confusing and wasteful to represent decimal numbers as hexadecimal numbers. Furthermore, it is explicitly discouraged in current standards.
- Avoid Encoding IPv4 Addresses in IPv6 Addresses. This artificially links IPv6 with legacy IPv4 and sends the message that IPv4 is the primary IP protocol. Furthermore, it may leave a long-term archaism that is difficult to make sense of when networks are moved to IPv6-only operation.

Your biggest challenge here is staff mind-set. Technical staff must be familiar with IPv6 addresses.

7 Obtaining Global IPv6 Addresses Space

In Europe, there are three ways a public administration can acquire public IPv6 address space:

- Direct from the regional internet registry (RIR), which in Europe is RIPE.
- Indirectly through a local internet registry (LIR), which is an organisation that can make sub-allocations such as a service provider (ISP).

By becoming (or already being) a local internet registry (LIR).

Public administrations may use one or more of the above methods. The best choice of method for a public administration will depend on several factors, including:

- Pre-existing arrangements.
- IPv6 specific requirements.
- Whether the responsibility for IP address space is centralised or decentralised.
- Whether the public administration is itself a local internet registry (LIR)
- Whether the public administration will need to sub-allocate address space to departments or regional administrations.
- Whether the public administration has out-sources IP connectivity and address space provision to a third party.

An Internet Registry (IR) is an organisation that is responsible for distributing IP address space to its members or customers and for registering those distributions. A Local Internet Registry (LIR) is an IR that primarily assigns address space to the users of the network services that it provides. LIRs are often, but not always, ISPs whose customers are primarily End Users and possibly, other ISPs.

Throughout Europe and the Middle East, IPv6 address allocation rules are set by the membership of RIPE NCC[12].

7.1 Types of Address Space (Provider Aggregatable vs. Provider Independent)

Currently there are two types of address space that can be obtained by a public administration. These are provider aggregatable (PA) space and provider independent (PI) space. The main differences between PA and PI space are:

- PI space is independent of a service provider. A public administration may use PI space across multiple providers. PI can be kept even if they change providers.
- PA space is specific to a service provider. A public administration with PA space must use the service provider whose space it is. If they change their service provider then the address space will have to change, resulting in an address renumbering exercise. (This is very undesirable.)

[12] https://www.ripe.net/publications/docs/ripe-707

There are other differences too. These include, who the contract is with, the cost, who is responsible for advertising the address space, whether the space can be sub-allocated to other organisations and whether multi-homing is possible.

7.2 Options for Obtaining IPv6 Address Space for Public Administrations

A public administration may obtain PI space directly from a RIR (RIPE in Europe) or a sponsoring LIR such as their ISP. PA space can be obtained directly from an ISP.

If a public administration becomes an LIR then it will be allocated both PA space for its customers and PI space for its own internal use. If a public administration is an LIR then it can make sub allocations of address space to other organisations from its PA space.

Generally, using PI IPv6 address space is recommended for any medium to large organisation. In most cases public administrations will want to obtain PI space to ensure that they are not tied to one service provider and to minimise the risk of having to renumber if they change providers. PI space can be obtained directly from RIPE by becoming a member of RIPE or through a sponsoring LIR such as a service provider (ISP).

Usually it is advantageous to obtain a single prefix for the public administration and create a single IPv6 address schema that allocates portions of the space to different regions, departments or functions. However, in some public administrations, autonomy between departments, regions and functions may lead to different prefixes being obtained for multiple entities. For example, in some public administrations a large department, such as the military, may seek to obtain its own space.

In public administrations with centralised responsibility for IP addressing, either PI or PA space may be used. If departments or regions are to some extent legally or otherwise autonomous, the centralise function may wish to act as an LIR and sub-allocation portions of its PA space to other departments or regions.

Finally, in the case that responsibility for IP addressing is out-sourced, it is possible that the service provider will allocate portions of its PA space to the public administration.

The many differences between public administrations mean that the best choice between PI and PA space and the method of obtaining the address space will vary.

7.3 Size of the Initial Address Space Allocation (Prefix Length)

Since it is important to avoid renumbering and ensure that there is enough space for future growth and changes in addressing structures it is crucial to obtain the largest address space possible. RIRs, such as RIPE, have a goal to conserve address space. In contrast, public administrations should be more concerned about ensuring that they have enough address space for their current and future needs. The IPv6 address space is immense, therefore there should be no reason why a public administration should obtain insufficient address space.

However, RIPE's address allocation policies set limits on how address can be used (for example who to can be sub allocated to and where it can be sub allocated to) and how much space is allocated.

If a public administration is a LIR, and has a plan to make sub-allocations to other organizations (for instance, a national government making sub-allocations to regional and local units of government) in the following twenty-four months, the administration is entitled to request an initial allocation of /32 of PA space. The administration can request a /29 without any further documentation. If the public administration is an LIR they can also request a PI allocation for their own internal use.

If the administration desires an initial allocation of PA space that is larger than a /29, the request must be supported by an addressing plan that supports the larger request. RIPE's decision on the allocation size will be based on the number of users, the extent of the organisation's infrastructure, the hierarchical and geographical structuring of the organisation, the segmentation of infrastructure for security and the planned longevity of the allocation.

If an administration requests PI space, then the default allocation is a /48. If the administration requires more space than that provided by the default /48 prefix, then they must provide documentation to justify the shorter prefix.

7.4 The Public Administration as LIR

In Europe, if a public administration wants to be able to allocate IPv6 address space to other organisations then it should choose to become a LIR by becoming a member of the RIPE NCC. RIPE members can request IPv6 address space, make assignments to other organisations and are eligible for a one-time allocation of a very small (/22) amount of IPv4 address space. For public administrations, any legal entity can become a member of RIPE NCC.

Application to become a member requires the full legal name and registered legal address of the public administration's organisation, and a digital copy of official company registration papers. It also requires the organisation's billing address (if different from the legal address), the email address that invoices will be sent to, and your VAT number (if your organisation is VAT-registered). Finally, names and email addresses of people in the organisation for administrative or technical questions, and an email address where users can contact the organisation in case of network abuse are required.

Once RIPE NCC receives the application form[13], it validates the copies of the documents provided during the application. When the verification is complete, RIPE NCC sends the organisation an invoice by electronic mail. At the same time, RIPE NCC also sends a contract (service agreement), another copy of the invoice and a "Statement of Lawful Presentation" by courier.

Once RIPE NCC receives the payment for the invoice and the signed contract, it activates the LIR account. At this point, the public administration has become an LIR. RIPE sends an introductory email with all the information needed for a new member. The applicant will then be able to log in to the LIR Portal[14], the secure web area for RIPE NCC members to manage everything related to their membership and the Internet number resources they hold, and start requesting Internet number resources – including IPv6 address space[15].

In practice, the process of applying for membership, due diligence, and approval can take as little as two weeks. Usually, however, the process takes longer because of the administrative process of getting a contract signed by a public authority, making a payment to RIPE NCC and waiting for the

[13] https://my.ripe.net/#/public/membership

[14] https://lirportal.ripe.net/

[15] https://www.ripe.net/manage-ips-and-asns/ipv6/request-ipv6

due diligence process to complete. For many public administrations, the process can be between six to eight weeks from start to finish.

8 IPv6 Address Planning for Public Administrations

IPv6 subnet prefixes are usually 64 bits[16] long, leaving 64 bits for the IID. This means that the bits between the public global address prefix and the IID are available for creating an addressing structure (see Figure 7). This section of the guide is about how you can structure these bits.

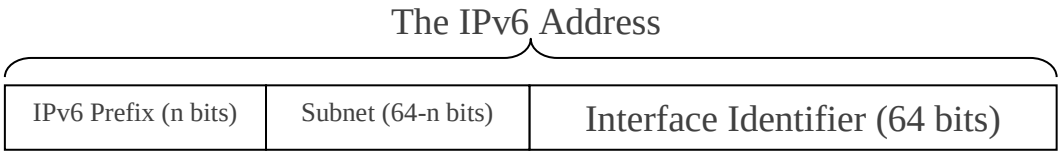

Figure 7 The Common Global Unicast Address Structure

There is no single template for an IPv6 address structure. Most organisations create a structure that encapsulates some or all the following factors:

- Network topology
- Geographical location
- Organisations structure
- Functional structure
- Security boundaries

The primary function of IPv6 prefixes, including subnet prefixes is to enable the efficient routing of datagrams to the correct destination network. For reasons of efficiency (route aggregation) and performance (minimisation of hops) it is important that the factors used in the IPv6 address structure are useful and relevant to the structure of the network and organisation. You should avoid trying to encode too many factors into the prefix. This increases complexity and often wastes space through duplication. For example, in a public administration that has a division of departments both by their location and function, including both in the addressing structure may be redundant.

For these and other reasons a common hierarchical address structure for public administrations is one that is built upon geographical regions. Geographical regions that may be included in an addressing hierarchy are, country (usually one and therefore often unnecessary), regions or municipalities, cities and sites. Underneath the regional structure it is common to have a functional structure, for example within an individual site (see Figure 8).

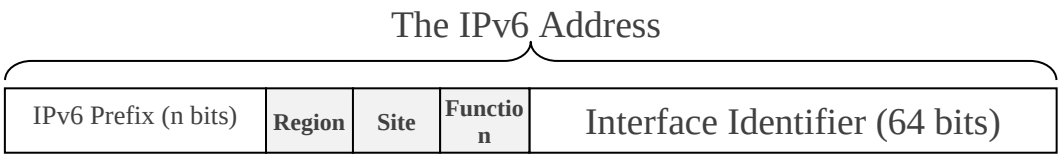

Figure 8 Common Elements of an IPv6 Address Structure

[16] There are several exceptions to the default /64 subnet prefix length. For example, the current standard for router interlinks is /127. See RFC6164.

Whilst the example structure in Figure 8 is common, the reverse structure is also possible, particularly in public administrations where the organisation is split by function rather than region. In this case the function bits would come before the geographical bits.

Some organisations encode security boundaries within their IPv6 addressing structure. This splits the network into separate areas with different security levels. Using bits in this fashion eases the filtering of traffic by firewalls, however it comes at the cost of the consumption of prefix bits and potentially greater complexity.

8.1 General IPv6 Address Planning Principles

When designing your IPv6 address structure you should refer to the general IPv6 addressing principles covered in the section General IPv6 Addressing Principles. The principles that are particularly relevant to address planning are:

- Ensure that initial allocations are large enough for current and future requirements.
- Build in flexibility by the sparse assignment of subnets and keeping space for expansion.
- Keep it simple and avoid the temptation to create a complex structure.
- Create a hierarchical structure to increase routing efficiency.
- Align Major Address Boundaries on Nibbles (4-bit boundaries).
- Avoid Encoding Decimal Numbers as Hexadecimal Numbers.
- Avoid Encoding IPv4 Addresses in IPv6 Addresses.
- Consider operational and security issues in the design.

8.2 Sizing and Creating a Hierarchical IPv6 Address Structure

In sharp contrast to IPv4 network design, the enormous space available in IPv6 means that the main priority is not the conservation of address space. Instead, the focus is on a design that best fits the organisation's needs and one that makes the best use of the organisation's IT infrastructure. Towards this end, it is crucial that the sizing of each part of the structure is enough both for current requirements and for requirements into the foreseeable future.

Sometimes it is appropriate to have alternative structures for different parts of the network. For example:

- The structure used in data centres may be different from that used within offices
- The structure may need to include provision for a legacy infrastructure within it. For example, a structure based on VLAN numbers
- Router interlinks may also be different from other subnets. It is common to use /127 prefixes for router interlinks. Therefore, it is usually possible to place all router interlinks into a single /64 within a location, function or even an organisation.

You should include in your structure provision for future network technologies that are not yet deployed. For example, you should include space for the Internet of Things (IoT).

Once you have identified the organisation attributes that you will use in your IPv6 address structure (for example, organisational geographical location or function) you then need to determine the current size of each attribute and the maximum possible size that might be required in the future. You can then assign an appropriate number of bits (or nibbles) to each organisational attribute.

Whilst it is usually essential and beneficial to align boundaries on nibbles. Occasionally it is unnecessary and potentially wasteful. One nibble (4 bits) gives sixteen options. If, for example, you only have two options then one or two bits would be enough. If that part of the address hierarchy does not require DNS reverse lookup delegation and if it can be merged with other bits in the hierarchy then using part of a nibble may be a better solution.

When assigning prefixes to parts of the hierarchy it useful to place geographically, departmentally or functionally "near" or "similar" regions numerically near to each other. This can allow for further route aggregation in the future.

Note that there are some common "break points" used in IPv6 prefixes. These are /32, /48, /56 and /64.

Bits	Subdivisions
1	2
2	3 or 4
3	5 – 8
4	9 – 16
5	17 – 32
6	33 – 64
7	65 – 128
8	129 – 256
9	257 – 512
10	513 – 1024
11	1025 - 2048
12	2049 - 4096

Table 2 Number of Subdivisions by Number of Bits (Nibble boundaries are highlighted)

Table 2 illustrates the number of subdivisions available for a specific number of bits. Notice that the nibble boundaries (multiples of 4 bits) are highlighted.

Therefore, if you have thirty-five geographic regions that you wish to include in your address hierarchy then you would need a minimum of 6 bits. If you need this boundary to align with nibble boundaries, then you would need a minimum of 8 bits.

Six bits gives you 64 options, nearly double the number that you need. The extra space could be used for future expansion. Eight bits gives you 256 options. This is significantly more than the 35 options that you currently need to represent each geographic region.

An important point to note here is that some parts of the address structure are unlikely to change (countries don't often restructure their regions) whilst other parts are likely to change (for example the number of departments). Therefore, space for growth is best allocated where it is likely to be needed rather than where it is not.

8.3 Sparse Address Allocation

When designing an address structure care should be taken to accommodate growth and the potential need for more space in some parts of the hierarchy. Ideally the initial space allocated would be sufficiently large to accommodate all future requirements. In the real-world, requirements change, and sometimes additional space will be required.

Adding additional space to a part of the hierarchy from a non-continuous block of address space fragments the address space. Two non-continuous blocks of address space allocated to one subnet will require two routes in the routing table and makes it difficult to merge the two blocks. This creates additional complexity and increases the number of routes required for subnets.

This problem can be avoided if address space is allocated sparsely. That is, address blocks are allocated with free space between them. One or more block is left free immediately after the allocated block. If additional space is required, a free block adjacent to the existing space can be added. Since the blocks are adjacent, only one route is required to reach both blocks. Therefore, the address space is not fragmented.

8.4 Prefix Visibility

When deciding on the length of the prefix that will be assigned to parts of a public administration it is important to consider if any part of the public administration will be advertising its own space, say by peering with a service provider. If this is the case, then the length of the prefix assigned to that part of the public administration should be sufficiently large to not be filtered in the public internet.

As a rough guide, prefixes longer than a /48 are usually filtered. Prefixes of length /48 should not be filtered, but they are in some parts of the public internet. Therefore, it is normal to advertise at least a /48. Bear in mind that the public administration's main prefix should also be being advertised, so even if these longer prefixes are filtered the address space should still be reachable.

9 IPv6 Address Planning Examples

In this section we use the Netherlands as an example of how a public administration IPv6 address structure can be created. We assume that the public administration has become an LIR by joining RIPE. As an LIR the public administration will receive at least a /32 prefix (PA space) and optionally some PI space. We have used this prefix length in our examples of building an IPv6 addressing plan. (Note that whether the organisations use PA or PI space they may use a /32 prefix or a different length prefix.)

The first 32 bits of the address are allocated to the public administration by RIPE. The last 64 bits are the IID and are used within each subnet to identify nodes. Thus, an IPv6 addressing plan is about how to use the remaining unspecified bits in the prefix (in this case the twelve bits from bit 33 to bit 64).

9.1 Example of Geographical Subnetting - The Netherlands

Here's an example of geographical subnetting. We have acquired a /32 prefix (2001:db8::/32) and we are using this to create an addressing structure for the provinces of The Netherlands. The Netherlands has eleven provinces. These are:

1. Groningen
2. Friesland
3. Drenthe
4. Overijssel
5. Gelderland
6. Flevoland
7. Utrecht
8. North Holland
9. South Holland
10. Zeeland
11. North Brabant

To determine how many bits are required for the provinces you count the number of provinces and consult Table 2 to see the number of bits needed.

Since there are eleven regions, we need at least four bits ($2^4 = 16$) to accommodate the regions and provide some room for growth.

The following figure illustrates the address structure of the described example:

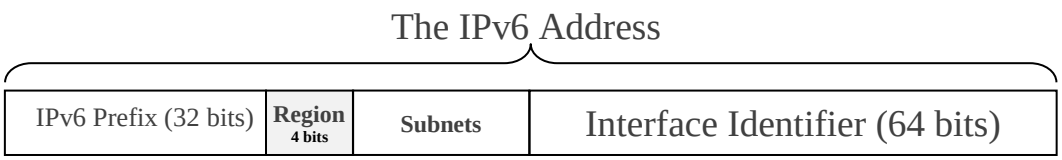

Figure 9 An Address Structure Based on Provinces

Next, we need to assign specific address blocks to specific regions. Table 3 gives an example of how this can be done. The free space should be adjacent to the regions that are most likely to require additional space.

Bits	Regions or Functions	Purpose
2001:db8:0000::/36	Free	Available for future growth
2001:db8:1000::/36	Groningen	
2001:db8:2000::/36	Friesland	
2001:db8:3000::/36	Free	Available for future growth
2001:db8:4000::/36	Drenthe	
2001:db8:5000::/36	Overijssel	
2001:db8:6000::/36	Free	Available for future growth
2001:db8:7000::/36	Gelderland	
2001:db8:8000::/36	Flevoland	
2001:db8:9000::/36	Free	Available for future growth
2001:db8:a000::/36	Utrecht	
2001:db8:b000::/36	North Holland	
2001:db8:c000::/36	Free	Available for future growth
2001:db8:d000::/36	South Holland	
2001:db8:e000::/36	Zeeland	
2001:db8:f000::/36	North Brabant	

Table 3 Example of a Geographical Addressing Structure

An approach such as this has several advantages. First, the subnets to be routed remain aggregated, saving space in the routing table. By using contiguous prefixes, we make it easy to build Access Control Lists or other network infrastructure rules. Finally, we have achieved a subnetting plan that makes it possible to support growth.

9.2 Example of Functional Subnetting - The Netherlands

A similar subnetting plan can be create that is based on function rather than on geographic region. The approach to creating the addressing plan would be similar:

- Establish the list of organisational functions to be supported
- Ensure that you have included space for growth
- Decide how many bits are needed to support the functions plus the growth

- Add this to the prefix of the initial allocation
- Sparsely assign prefixes to the functions

Once again, we use a /32 prefix (2001:db8::/32) that we have acquired from RIPE to create a structure for a collection of national networks:

1. Management networks
2. Educational networks
3. National Health System networks
4. National Army and Defence networks
5. Police networks
6. Transportation networks
7. Government ministries and departments
8. Emergency Services
9. Other National Infrastructure

Since there are nine functions, we need four bits ($2^4 = 16$) to accommodate the networks plus space for growth.

Figure 10 illustrates an address structure based on network function.

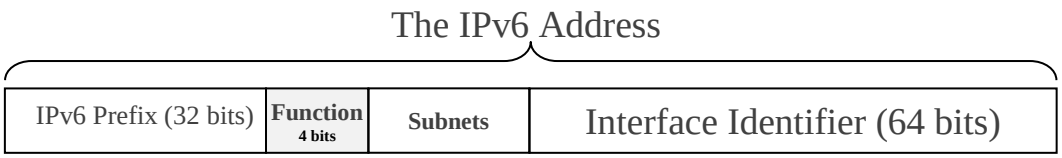

Figure 10 An Address Structure Based on Function (National Networks)

Next, we need to assign specific address blocks to specific network functions. Table 4Table 3 gives an example of how this can be done. The free space should be adjacent to the regions that are most likely to require additional space.

Bits	Regions or Functions	Purpose
2001:db8:0000::/36	Free	Available for future growth
2001:db8:1000::/36	Management networks	
2001:db8:2000::/36	Free	Available for future growth
2001:db8:3000::/36	Educational networks	
2001:db8:4000::/36	Free	Available for future growth
2001:db8:5000::/36	National Health System networks	
2001:db8:6000::/36	Free	Available for future growth
2001:db8:7000::/36	National Army and Défense nets	
2001:db8:8000::/36	Police networks	
2001:db8:9000::/36	Free	Available for future growth
2001:db8:a000::/36	Transportation networks	
2001:db8:b000::/36	Free	Available for future growth
2001:db8:c000::/36	Government ministries and departments	
2001:db8:d000::/36	Free	Available for future growth
2001:db8:e000::/36	Emergency services	
2001:db8:f000::/36	Other national infrastructure	

Table 4 Example of a Functional Addressing Structure

It is unusual for an IPv6 address structure to only include geographic or functional attributes in its hierarchy. More often public administrations must further breakdown the address space by other attributes. Often this means encoding both geographic location and function into the structure.

A key decision that you must make is whether to place geographic location or function first in the addressing hierarchy. This decision depends on the type of public administration, its organisational structure and its network infrastructure.

If geographic location comes first, it is then easier to aggregate routing by geographic location. If function comes first, then aggregation by function is easier but it may be more difficult to align the routing with the network topology.

Which attribute comes first in the prefix can also assist security, such as traffic filtering using Access Control Lists (ACLs). Some public administrations would wish to filter by region, others by function and others by both. In theory filtering on any sequence of bits anywhere in the address is possible. However, many hardware filtering platforms filter by the most-significant-bits. As a result, filtering is made easier by putting the bits that will be filtered first.

Putting the geographic location first also makes it easier for regional administrations to create their own local addressing plan or make their own modifications to a global template.

9.3 Example of Functional and Regional Subnetting - The Netherlands

Here is an example of how such a combined approach would work. The two previous examples are combined into an addressing plan that emphasises function over the region in which the function is provided.

Figure 11 illustrates an address structure based on function and geographical region.

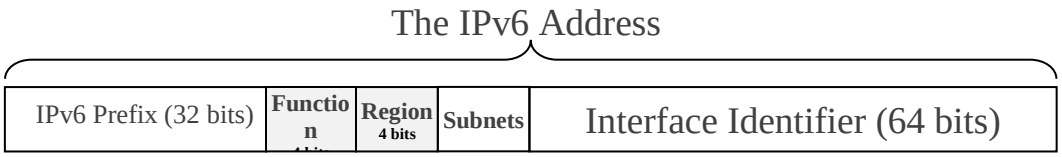

Figure 11 An Address Structure Based on Network Function and Geographic Region

Table 5 shows an example addressing plan for a network where function comes first followed by location.

Free		2001:db8:0000::/36
Management networks	Free	2001:db8:1000::/40
	Groningen	2001:db8:1100::/40
	Friesland	2001:db8:1200::/40
	Free	2001:db8:1300::/40
	Drenthe	2001:db8:1400::/40
	Overijssel	2001:db8:1500::/40
	Free	2001:db8:1600::/40
	Gelderland	2001:db8:1700::/40
	Flevoland	2001:db8:1800::/40
	Free	2001:db8:1900::/40
	Utrecht	2001:db8:1a00::/40
	North Holland	2001:db8:1b00::/40
	Free	2001:db8:1c00::/40
	South Holland	2001:db8:1d00::/40
	Zeeland	2001:db8:1e00::/40
	North Brabant	2001:db8:1f00::/40
Free		2001:db8:2000::/36
Educational networks	Free	2001:db8:3000::/40
	Groningen	2001:db8:3100::/40
	Friesland	2001:db8:3200::/40
	Free	2001:db8:3300::/40
	Drenthe	2001:db8:3400::/40
	. . .	. . .
Free		2001:db8:4000::/36
National Health System networks	Free	2001:db8:5000::/40
	Groningen	2001:db8:5100::/40
	Friesland	2001:db8:5200::/40
	Free	2001:db8:5300::/40
	Drenthe	2001:db8:5400::/40
	. . .	. . .
Free		2001:db8:6000::/36
National Army and Défense nets	Free	2001:db8:7000::/40

	Groningen	2001:db8:7100::/40
	Friesland	2001:db8:7200::/40
	Free	2001:db8:7300::/40
	Drenthe	2001:db8:7400::/40
	. . .	. . .
Police networks	Free	2001:db8:8000::/40
	Groningen	2001:db8:8100::/40
	Friesland	2001:db8:8200::/40
	Free	2001:db8:8300::/40
	Drenthe	2001:db8:8400::/40
	. . .	. . .
Free		2001:db8:9000::/36
Transportation networks	Free	2001:db8:a000::/40
	Groningen	2001:db8:a100::/40
	Friesland	2001:db8:a200::/40
	Free	2001:db8:a300::/40
	Drenthe	2001:db8:a400::/40
	. . .	. . .
Free		2001:db8:b000::/36
Government ministries and departments	Free	2001:db8:c000::/40
	Groningen	2001:db8:c100::/40
	Friesland	2001:db8:c200::/40
	Free	2001:db8:c300::/40
	Drenthe	2001:db8:c400::/40
	. . .	. . .
Free		2001:db8:d000::/36
Emergency services	Free	2001:db8:e000::/40
	Groningen	2001:db8:e100::/40
	Friesland	2001:db8:e200::/40
	Free	2001:db8:e300::/40
	Drenthe	2001:db8:e400::/40
	. . .	. . .
Other national infrastructure	Free	2001:db8:f000::/40
	Groningen	2001:db8:f100::/40
	Friesland	2001:db8:f200::/40
	Free	2001:db8:f300::/40
	Drenthe	2001:db8:f400::/40
	. . .	. . .

Table 5 Example of a Functional and a Geographic Addressing Structure

9.4 IPv6 Address Planning Case Study – German Ministry of the Interior

The German Federal government decided that is should request and hold IPv6 address space for the entire public administration sector in Germany. In 2009, Germany was given a /26, specifically 2a02:1000/26. The intent is to provide a common, central repository of address space for all public administration entities in Germany. The contract for the address space is between the Federal Ministry of the Interior and RIPE NCC. As a result, the Federal Ministry takes on the role of Local Internet Registry for the public sector in Germany.

The address space has been divided into 64 equally sized address blocks where each block is a /32. Using this as a foundation, the Ministry of the Interior has created a IPv6 Transition Guide[17] that documents the German IPv6 addressing plan.

The high-level German IPv6 addressing plan is shown in Table 6.

26-bits	Prefix for public administrations in Germany 2a02:1000::/26 Assigned by RIPE NCC to the German Ministry of the Interior
6-bits	Assigned by the German LIR (The Ministry of the Interior) to German federal states (Landler, the federal government's backbone network, and some in reserve for growth)
16-bits	Assigned by Sub LIRs (for instance, German federal states) to individual units and departments of state government, cities, counties, municipalities, and some in reserve for growth
16-bits	Used for local subnets in each individual unit and department of state government, cities, counties, municipalities
64-bits	Interface identifier on each local subnet

Table 6 The German IPv6 Addressing Structure

Each local administration is allocated a /48 prefix for their internal use. This prefix leaves 16 bits available for local subnets. The local administration can use these 16 bits to structure their IPv6 subnets to meet their local organisational and network requirements. The German address management scheme provides some examples of how address planning could be done in:

- A medium-sized public administration
- A small public administration
- A Data Centre or Large-Scale public administration
- A home office and very small administrations

17

http://www.bva.bund.de/DE/Organisation/Abteilungen/Abteilung_BIT/Leistungen/IT_Beratungsleistungen/IPv6/best_practice/ipv6migrationsleitfaden_EN/download/fue_migrationsleitfaden.pdf?__blob=publicationFile&v=7

10 IPv6 Address Management

10.1 IPv6 Address Management (IPAM) Systems

IP Address Management (IPAM) refers to the management of the allocation, administration, reporting and tracking of IP addresses. IPAM systems often interact with the DNS and DHCP infrastructure. Such systems are sometimes referred to as DDI systems (DNS, DHCP and IPAM).

Historical, in both the public and private sector, IT departments, particularly in small organisations have often used manual processes, including spreadsheets or other home-grown tools for IP address management.

In the case of IPv6, the increased length and complexity of IPv6 addresses makes it imperative that an organisation implements some form of IPAM. The challenges of IPv6 address management include:

- The huge address spaces
- Longer literal (textual) address representation
- Multiple addresses per node and per interface
- Addresses that change regularly (e.g. privacy addresses)

The resulting complexity and the interaction with DNS and DHCP is why organisations need to move away from manual to automated systems.

There are many IPAM systems on the market. Some include their own built-in DNS and DHCPv6 servers, whilst others integrate with your existing infrastructure. Many IPAM systems include support for IPv6.

- Display of IPv6 Address Space. This is simply a way to visualise the IPv6 address space under management, and a good IPAM system will offer more than one view. Common views are a graphical map of the network space in a list or table view. It will also provide a view of the hierarchy of the network including utilisation by subnets and sites.
- Measurement of the Consumption of IPv6 Address Resources. An IPAM system usually gives the ability to measure and track the consumption of available IPv6 addresses. In IPv6, this is measured in subnets rather than in node addresses. This is a useful tool for capacity planning and network management.
- Reporting. The ability to generate reports of the IPv6 network to assist in planning and maintenance of the network.
- IPv6 Address Plan Design. Some IPAM systems provide a tool to help create an IPv6 address plan. Application of address planning principles and best practices can be combined with the IPAM system's ability to display allocations and assignments according to organisational hierarchy, both within and between sites. This is especially useful in cases where a public administration is spread across multiple regions or cities.
- RIR Registration Updating. If a public administration has pursued obtaining its addresses through an LIR rather than from an upstream ISP, the administration will have a reporting responsibility to the RIR. The IPAM tool can assist with any changes to the registration reporting that is needed to meet the RIR's requirements.

- Policy Compliance. An IPAM system tracking address resource can make it easier to validate that new or ongoing assignments comply with existing security, SLA or routing policy requirements.
- Integration with DNS and DHCPv6.

10.2 IPAM and DHCPv6

Earlier in this document we looked at the different methods used to assign IPv6 addresses to nodes. There we saw that there are three main methods of assigning addresses, these are:

1. Manual address configuration
2. Automatic configuration using Stateless Address Autoconfiguration (SLAAC)
3. Automatic configuration using Dynamic Host Configuration Protocol version 6 (DHCPv6)

It is usual to use one of the two automatic methods of address configuration on a subnet and very unusual to use two. In any case, we noted earlier that in IPv6, SLAAC is still necessary when using DHCPv6. Flags in router advertisements enable DHCPv6 and specify whether it will only provide options or also lease addresses. Also, router advertisements are required to provide the default router as DHCPv6 does not have a default router option.

IPAM systems usually use DHCPv6 to manage IPv6 address assignment. Some can also obtain information about assigned IPv6 addresses using other means, such as features in switches and routers (for example, using Neighbor Discovery Inspection and scraping the neighbor discovery caches).

Most IPAM solutions will have mechanisms and user interfaces that help build DHCPv6 address pools and manage IPv6 prefix delegation. They also have interfaces that help report on assigned addresses and details of client leases.

10.3 IPAM and DNS

The domain name system (DNS) is one of the most important network services. Most users and applications use names to refer to nodes and services. For nodes to communicate, these names must be converted to IP addresses. DNS provides two main services that are relevant to address management. These are:

- The resolving of domain names to IPv6 addresses (forward lookups)
- Finding the IPv6 addresses for domain names (reverse lookups)

Since the IPv6 addresses are long, it is extremely useful to automate the process of configuring these resource records in DNS. Most IPAM solutions provide tools to automatically create directory entries for forward and reverse lookups (the forward and reverse zones). The advantage to using an IPAM solution to do this is that the resource records are automatically generated, and name servers (including master and slave servers) can be automatically updated as addresses are assigned or changed.

10.4 Managing IPv6 Network Growth and Change in Public Administrations

The addressing strategy outlined in this guide is not intended to be a static, unchanging plan. Instead, it has room for growth built-in. Supporting a public administration as it adapts to change, and growth is foundational part of the plan.

One of the features of sparse assignment of subnets is that there is ongoing room for growth without affecting the size of routing tables. By including room for growth and change in the original plan, the foundation is set so that the existing plan can adapt to a change in the environment, structure or technical situation. Networks grow, and shrink, along with the units of government they support. There can be organisational changes, but there can also be technological changes.

One source of change is when a public administration must adapt to a disruptive technology. For many units of government, the deployment of sensor networks and the Internet of Things is a good example of this. IPv6 is likely to be the identifier technology for these new technologies and, if so, planning for large-scale deployments of IPv6 addresses for sensor networks will be a new challenge.

Another source of change is when a public administration changes network provider. If the public administration is not an LIR and does not have PI address space, then changing providers may mean numbering into a new assignment from the new provider, and out of the old one to return it.

IPv6 was designed, in part, to reduce the necessity and frequency of network renumbering. This is particularly important due to the renewed emphasis on the use of public addresses. With the enormous address space that is available in a single IPv6 subnet (/64), renumbering within a subnet is never be necessary because of a shortage of node addresses. However, if the public administration's prefix is changed then the network using that prefix must be renumbered. Therefore, it is desirable to obtain enough address space to ensure that renumbering is not necessary.

Unlike many IPv4 networks that use private address space (RFC1918 space) renumbering IPv6 networks cannot rely on reconfiguring Network Address Translation (NAT) at the edge of the network. Instead the whole network must be renumbered, not just the edge. In IPv4, this can be extremely onerous and difficult and cannot be achieved without breaking existing network connections. In contrast, IPv6 has been designed to be renumbered. It has built-in features that make it theoretically possible to move from an old address prefix to a new address prefix with the minimum of disruption. Despite this it is still a significant undertaking and something that should be avoided if possible.

10.5 IPv6 Reachability in Public Administrations

For nodes in the public administration to communicate with nodes on the public internet the IPv6 addresses that they are using must be reachable from the global public internet. A crucial part of ensuring that the addresses assigned to nodes in the administration are reachable is to make the administration's IPv6 prefix/es visible in the public internet.

Prefixes are advertised on the global internet using the Border Gateway Protocol (BGP). Therefore, whether a public administration is a single entity or is a complex hierarchy of inter-related national, regional and local organisations, at some point, the administrations IPv6 prefix must be advertised by someone using BGP.

In the case of small public administrations that are using PA address space provided by their ISP the prefix will most likely be advertised by the ISP.

A public administration that has become an LIR and that has acquired an IPv6 prefix directly from RIPE must advertise that prefix using BGP with one or more peers (usually ISPs or transit providers).

A public administration that is not an LIR and has acquired an IPv6 prefix (PI space) from RIPE or through an ISP, must either advertise the prefix itself or contract with one or more service providers to advertise the prefix on its behalf.

A public administration will also need to implement IPv6 dynamic routing protocols for its internal network. The administration will need to consider whether to use the dynamic routing protocols that they are currently using for IPv4 or to use different protocols for IPv6. There are several factors that must be considered when making this decision. These include; technical expertise with the protocols, whether multi-topology is required (that is different topologies for IPv4 and IPv6) and whether their routers have enough resources for the chosen protocols.