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

Plum Consulting
iDate
Internet Policy Advisors
Synetergy
Erion

21 March 2018

Contents

Executive Summary

This document provides a resource for IPv6 address planning for public administrations. The intent is to provide an introduction in three parts: preparation, design and maintenance. In preparation, we will provide the resources needed to understand the elements of IPv6 addressing and how public administrations can muster support for IPv6 deployment projects. This part of the effort will also provide the tools needed for overall planning for the IPv6 deployment.

In the design component, we will provide a blueprint for IPv6 subnetting and address planning. The final part of this component is the development of a comprehensive IPv6 addressing plan appropriate to the public administration's requirements. We also examine a case study from an existing public administration deployment of IPv6.

Finally, in the maintenance component, we focus on the tools used for IPAM, managing change and growth in the IPv6 network and the details of keeping the IPv6 resources reachable and discoverable.

In Europe there are a variety of mechanisms to get IPv6 addresses: obtaining them from an upstream provider, becoming part of a larger organisation or consortium that acquires IPv6 addresses for a group of administrations, and acquiring IPv6 address on one's own by becoming a local internet registry. In Europe, organisations have access to different kinds of IPv6 address space depending on their business model and the way they intend to use that address space: IPv6 allocations; IPv6 PI assignments; and IPv6 assignments or allocations from an upstream LIR.

A public administration with a direct IPv6 allocation has the option to further assign or allocate address space from the allocation to their network infrastructure or to others for use in their own networks.

The minimum IPv6 allocation size in the European region is a /32 (4.3 billion subnets). According to the policies in effect in Europe, RIPE NCC members can request a larger IPv6 allocation (up to a /29) without having to provide additional information about how they plan to use the extra space.

A public administration can also request IPv6 address space from their upstream provider (often an ISP), if the provider is a RIPE NCC member. In this case, the public administration would be considered to be an End User.

Finally, an administration can request a Provider Independent IPv6 prefix for their internal network. In this case, the assignment cannot be further delegated. In Europe the default size for a Provider Independent assignment in a /48, or 65,536 possible subnets.

One of the fundamental requirements for moving public administrations to IPv6 is the ability to acquire needed address space. This document describes the mechanism used to begin application to become an LIR and what to expect during the process of becoming an LIR.

IPv6 Addressing Basics

Fundamental to all IPv6 deployments is an effective addressing plan for the public administration network. To build an IPv6 addressing plan, a small amount of background is required.

IPv6 Addresses

When the public Internet reached mass popularity in the 1990's, the dominant addressing and packetisation protocol was IPv4. IPv4 is simply a means to address an envelope of bits that contain information. The addresses can represent the source from which the packet was sent, and also the destination to which the packet is being sent. Addressing in the older IPv4 protocol was made up of 32 ones-or-zeroes – 32-bits. 32-bits of address space left room for more than 4 billion devices to be connected to the Internet[1].

Even in the 1990's, it was recognised that 32-bits would not be enough space to address all the devices that might connect to the Internet. The result was the development of IPv4's successor: IPv6.

The Internet protocol version 6 (IPv6)[2] has addresses that are 128 bits instead of IPv4's 32 bits. Since each bit effectively doubles the amount of address space available for use, the quadrupling of the space represents an enormous increase in the number of devices that can be connected. In IPv6, the address component for the local area network is 64-bits and the network identification uses the other 64-bits. For a sense of scale, this means that:

2^{64} = 18,446,744,073,709,551,616

18 quintillion is a very large number of devices to have on a local area network. Yet, IPv6 can easily support networks of that scale.

Representing IPv6 Addresses

Since working with IPv6 addresses means that you need to write them down, it's natural to wonder how one represents an IPv6 address. It would certainly be possible to write an IPv6 address as a string of 128 ones-and-zeroes, but that would not be easy to write, nor would it be possible to remember. Instead, IPv6 is represented as eight groupings of 16 bits, where each grouping is made up of four hexadecimal values. Colons separate each group. Here is an example:

2002:06de:0010:2233:0000:0000:0000:0001

Addresses are further simplified using these rules:

1. Leading zeroes in any grouping can be dropped. Our example becomes: 2002:6de:10:2233:0:0:0:1
2. A double colon can replace one grouping of all zeros. Our example now becomes: 2002:6de:10:2233::1

[1] Precisely 4,294,967,296.

[2] See https://tools.ietf.org/html/rfc8200

IPv6 addresses are grouped using the binary value of the address. This grouping is carried out using a "prefix." Prefixes are all addresses that start with the same series of bits, similar to an area code for phone numbers. The length of the identical series is noted after the address, separated by a forward slash. For instance, the prefix

2001:db8::/32

represents all the addresses from:

2001:0db8:0000:0000:0000:0000:0000:0000

through

2001:0db8:ffff:ffff:ffff:ffff:ffff:ffff

Basic Structure

The fundamental structure of the IPv6 address is the division of the address into three parts:

- A *global routing prefix*; a variable, n, bits that uniquely identify the network;
- A *subnet identifier*; a variable, 64 minus n, bits that identify the subnet; and,
- An *interface identifier*; 64 bits of the address that identifies the host or computer on the local network.

Every 128-bit IPv6 address represents exactly one interface in the network (physical or logical/virtual). However, each interface may have multiple IPv6 addresses assigned to it. These are often from different address "scopes" such as link-local, global or mobile IP[3].

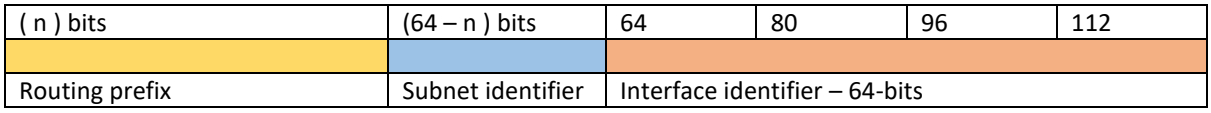

Structure of a globally routable IPv6 Address

IPv6 Address Types

Essentially, IPv6 has three types of addresses: *unicast, multicast and anycast*.

IPv6 Address Type	Type of Communication
Unicast	One-to-One
Multicast	One-to-Many
Anycast	One-to-All

IPv6 Address Types

Unicast addresses are assigned to interfaces and allow packets to be sent specifically to that interface. It's typical to think of computers or "hosts," to be configured with an address. This address then allows two computers to communicate with each other by using their corresponding unicast addresses. While that is a convenient approach to thinking of addressing on the Internet, in fact a computer could have multiple interfaces to the network, each one uniquely identified by its own unicast address.

IPv6 has some special unicast subtypes:

[3] See https://tools.ietf.org/html/rfc7346 for details.

Loopback address ::1/128	The loopback address is something carried over from IPv4 (i.e. 127.0.0.1) it is the address used for an IPv6 node to send packets to itself.
Link-Local unicast addresses fe80::/10	These are IP addresses reserved to be used on a local link. They play an essential role in Neighbor Discovery and auto-address configuration and management. A Link-Local address is automatically configured on an interface at the moment an IPv6 stack is activated. IPv6 routers do not forward any IP packets with link-local addresses in their headers (either source or destination).
Global unicast addresses[4] 2000::/3	These are the public, globally routable, allocations that each public administration will be assigned. The assignments come from either another part of the public administration, the ISP providing connectivity to the public administration, or, in Europe, the RIPE RIR. The GUA allocations come in two varieties: provider independent (PI) and provider assigned (PA). Provider independent allocations are portable and can be used and routed with any ISP. They are assigned by an RIR directly to the public administration via a Local Internet Registry. PA allocations are assigned by the ISP and must be returned if switching to a new ISP. This requires network renumbering.

IPv6 Unicast Subtypes

In IPv6, multicast addresses are assigned to, and identify, groups of interfaces. IPv6 packets with a multicast destination address are sent to all interfaces which have this address assigned. Multicast is important in public administration environments because some of the fundamental management features of IPv6 are done through multicast (for instance, the ability to send a message to reach all routers in a subnet). The multicast addresses always have all ones in the first eight high-order bits of the address:

ff00:/8

The next four bits set flags that help describe the scope of the multicast communication. In practice only the last of the four bits is set[5]:

Scope	Name
0	Reserved
1	Interface-local scope
2	Link-local scope
3	Realm-local scope
4	Admin-local scope
5	Site-local scope
6-7	Unassigned
8	Organization-local scope
9-d	Unassigned
e	Global scope

[4] See https://tools.ietf.org/html/rfc3587.

[5] For further details about the definitions and use of multicast scope in IPv6, please see https://tools.ietf.org/html/rfc7346.

f	Reserved

Multicast Scope for IPv6 Addresses

An anycast address is, like the multicast address above, assigned to multiple, different interfaces, but packets are delivered to the closest interface to the sender as measured by routing metrics. Thus, an anycast address can be any IPv6 unicast address. They do not have a reserved IPv6 prefix and thus come from the link-local or global-unicast addresses.

A typical use for anycast addresses is when large network operators configure their wide-area DNS networks. Name servers maintained by the operators of those networks are configured with anycast addresses, and end-user DNS queries are then routed to the closest name server. This tends to provide the end user with better performance with regard to the DNS query. The goal is often to provide access to the "closest" or "fastest" available server amongst a group of servers providing a similar service.

Host Addressing Assignment

With a fundamental understanding of how IPv6 addressing works, it is now time to understand how hosts and nodes in the network get these addresses assigned to them. There are three major ways to do host address assignment in IPv6:

1. Direct, static addressing on the host or node itself;
2. Dynamic addressing from a pool of addresses using DHCPv6; and.
3. Self-assignment of addresses on a local link

Direct/Static Addressing

Static addressing is often used for servers, routers, switches, firewalls, and any other network resources where address assignments shouldn't (or, are unlikely to) change over time. Often, in these cases, the static address is mapped to a name in the DNS. As a result, much of the network's management and content infrastructure receives static IPv6 addresses configured on the device itself, and then mapped in an organization's DNS server via DNS AAAA records for IPv6.

DHCPv6 Addressing

IPv6 hosts can generate their own IP addresses, unlike IPv4 hosts. They do this using Stateless Address Autoconfiguration, or SLAAC, for short. However, in some circumstances, a network manager may choose to assign IPv6 addresses from a pool, just as they are in IPv6. The tool used for this task is called DHCPv6.

One of the main reasons to use DHCPv6 is that the host can request information other than an IP address or route (as in autoconfiguration). DHCPv6 can be used to provide this information, even though it is not being used to configure IP addresses. DHCPv6 is not necessary for configuring hosts with the addresses of DNS servers, because they can be configured using IPv6's Neighbor Discovery Protocol, which is also the mechanism for stateless autoconfiguration.

Self-Assignment of Addresses

SLAAC is a mechanism that allows hosts to generate their own addresses. Default router information is provided to the local network by ICMPv6 router advertisements. As a result, when a router supports IPv6 SLAAC, it can gather all the information it needs for self-configuration on the network.

There are important considerations for public administrations who consider using SLAAC for network IPv6 assignments. First, because SLAAC does not provide any authentication mechanism and allows

a host to connect to the network and communicate with other nodes, the protocol is not recommended when security is required or preferred.

Another issue with SLAAC arises when privacy extensions are enabled on a host. A privacy extension allows the host to randomise the interface ID portion of the IPv6 address on a SLAAC assigned address. This increases privacy for the host, but it eliminates the ability to trace particular packets to particular hosts. Privacy extensions are enabled by default in many operating systems and may need to be disabled in the public administration setting if strict tracking and control of the hosts is required.

Why is IPv6 Different for Public Administrations in Europe?

We will see, in the section below titled "Getting IPv6 Addresses," that one of the things that makes public administrations different is the way they ***acquire*** IPv6 addresses.

However, once acquired, the way that public administrations are organized can influence the design and planning of the ***use*** of the address space. In particular, the organisation of the public administration may affect the design of the subnetting of the address space. In a centralised, federated organisation it might be natural to divide the space by the departments or organisations that are implementing the IPv6 services. In decentralised organisations, it might be more natural to divide the space by the region or locality that is being served.

In a later section, we will also see that public administrations have different approaches to the ***deployment*** of IPv6 technology: often in terms of budgets, planning cycles or horizons and infrastructure maintenance.

Planning the Public Administration Deployment

In a Public Administration setting, there are many challenges to setting out a business case for IPv6 adoption. As an example, although public administration IT environments use many common standards, hardware, software and best practices, it is still difficult to provide a consistent and compelling bottom-line benefit across units of government. Indeed, in many governments the approach to serving the public and performing regulation functions are varied according to subject and stakeholder – seldom does "one size fit all."

Also, part of the problem is budgetary. In an environment where government budgets are challenged, a project that delivers on changes to infrastructure is harder to justify. It has little immediate evident impact. In fact, the business case for IPv6 is based on managing the risks and costs to an administration's IT practice that increase over time in the absence of an IPv6 adoption initiative.

Another part of the puzzle is finding the responsible party for leading the IPv6 initiative. Planning the initiative is crucial to its success, and without a clear mandate for an individual responsible party, the planning is at risk.

The number of devices, especially mobile and Internet of Things devices, being connected means that the exhausted IPv4 address space is unable to meet the requirements of the future. Those who remain using IPv4 face the unprecedented operational nightmare in attempting to persist in using IPv4 private addressing and carrier-grade NAT technologies to support new users and devices. Only IPv6 has sufficient addressing capacity for the future.

Once an administration has committed to migrating to IPv6, the question is: how to plan for that migration. Many organisations have succeeded by breaking the plan into phases. In this section we will describe a simple plan that involves three phases for adoption:

- Preparation and research
- Support for external stakeholders
- Internal adoption

Preparation and Research

Preparation and research gives the administration a low-risk mechanism to assess support for the project, the costs and risks to the project, and a plan for implementation. Typically, the preparation and research phase of an IPv6 migration consists of some typical steps:

- Getting management or ministerial buy-in
- IPv6 training
- Creating an IPv6 Task Force or Stakeholder Group
- IPv6 Connectivity and Transit
- IPv6 Network / Hardware / Software Audit
- Working with Key 3rd Parties
- Obtaining an IPv6 Address Allocation

Getting management or ministerial buy-in

Simply stated, any IPv6 deployment plan needs some level of executive, management, ministerial or related support and approval. Support is added to approval because the approval needs to have some awareness of what the project goals are and the metrics to be used for evaluating success. Support and approval should lead to capital and operational budget assignments for whatever size deployment effort is put in place. The budget also needs to consider the need for training, additional staff, hardware, software and other infrastructure support.

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 administration settings 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. This does not scale 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 – a silo – 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.

IPv6 Training

IPv6 is often unfamiliar to both technical and other staff in public administration settings. Since IP addressing changes impact everything in a network, the uncertainty regarding IPv6 can lead to fear

and doubt regarding its implementation. Formal training is a good mechanism for technology transfer and technology transfer is one of the key ways to eliminate fear and doubt associated with IPv6.

IPv6 training comes in many forms, but an extremely effective one is formal training provided on-site or at a training centre. Europe has many such resources. There are two approaches to using this kind of formal training. The first is to build the knowledge base of individuals within the administration through systematic delivery of knowledge and skills to those individuals. The second is to use the training as a springboard to having the individuals turn around and deliver the knowledge and skills, themselves, into the rest of the organisation.

There are also a large number of independent resources for independent study. These can range from systematic certificate preparation courses to independent research in other books, RFCs, blogs or white papers. Europe's Regional Internet Registry, RIPE, provides significant training resources on IPv6 to its members and the public.

There are also national and regional IPv6 conferences and forums throughout Europe.

Creating an IPv6 Task Force or Stakeholder Group

An effective way of building support within a public administration for the IPv6 transition effort is to identify stakeholders and allies in relevant departments, ministries and governmental units who have an interest in the IPv6 initiative. These allies can spread awareness, and possibly support, for the initiative.

In some case, administrations have created an IPv6 Task Force. It often includes interested or involved stakeholders from principal organisations within the administration. Even a single meeting of such a group can spread awareness of the IPv6 adoption initiative and a level of organisational interest. This can help build both internal and management support. Also, this can avoid the organisational inertia and lack of participation that might lead to difficulties for the success of the project.

IPv6 Connectivity and Transit

Public administrations in Europe get their IP transit from a variety of service providers. It can be a traditional commercial ISP, the government itself can be the provider, or the administration can be in a cooperative arrangement with other units of government. In all of these cases, an increasingly larger number of transit providers support IPv6 connectivity. Experience shows, however, that the same providers who often have rich experience as an IPv4 service provider, might not be as able as an IPv6 provider.

An administration coming to an IPv6 initiative is looking for dual-stack connectivity over an existing IPv4 connection. ISPs that are relying on tunnelling to provide IPv6 services are not ready to provide government-grade IPv6 support. Assessing your connectivity options is one of the most important steps in the matrix of preparation and research.

IPv6 Network / Hardware / Software Audit

IPv6 migration requires considerable planning and an assessment of the impact of migration on the underlying network and services. For many administrations, a critical step is an assessment of the "readiness" of the components of the network for the IPv6 transition.

An IPv6 audit is an inventory and assessment of all network assets to determine the degree to which they support IPv6. Any legacy hardware or software that is dependent on the legacy IPv4 network must be identified and assessed.

This includes a huge variety of network infrastructure, for instance:

Routers/Switches	Web Servers	Desktops
Laptops	Operating Systems	Mobile and Tables Devices
Optical gear	Collaboration tools	Route server software
NATs	Security Tokens	Middleware
Cloud based services	Wireless Access Points	Web and Email content filters
Firewalls	Database servers	PBXs and VoIP systems
Sensors/sensor networks	Mail servers	CRM systems
Embedded or custom systems	POS terminals	Load balancers

The list is long, and this may not be exhaustive.

The marketplace has software that helps identify, discover and assess IT assets and then generate reports. There are also online services that provide certification for certain kinds of network equipment[6]. There are also third-party services that provide consultancy services to assist in the inventory/audit/assessment process.

The European RIR, RIPE, has an excellent description of what standards each unit of equipment should support. If testing network equipment at Layer 2 or Layer 3, there are clear specifications for which standards are mandatory to support and which standards are optional. An extract of the RIPE standards list is presented in *Appendix A: RIPE Requirements for IPv6 Compatibility*.

Working with Key 3rd Parties

One of the essential steps for many administrations will be working with 3rd parties who provide software, custom programs, databases, security infrastructure or other hardware and software. In these cases, the custom products will not be under the control of the local authority. As a result, the administration will have to identify the components that take advantage of IP and then perform, in cooperation with the vendor, an assessment of IPv6 readiness.

This is an area in which testing is optimal, but sometimes a challenge. For some custom programs, services, and hardware, there is no lab environment available to do IPv6 conformance testing.

Obtaining an IPv6 Address Allocation

In Europe, assignments and allocations of IPv6 addresses are governed by the RIPE Regional Internet Registry. To request IPv6 address space, you need to have a contractual agreement with the RIPE NCC. To directly request an IPv6 allocation, you need to be a member of the RIPE NCC. To request an IPv6 PI or IPv6 IXP assignment, you can also become a member of the RIPE NCC or you can request an independent assignment via a sponsoring LIR.

For smaller units of government or administrations, it is natural to obtain IPv6 address space from their upstream provider – either an ISP or a consortia/cooperative that provides transit for public administrations in the member state.

[6] For instance the IPv6 Ready Logo program at https://www.ipv6ready.org

Support for External Stakeholders

The first of the three basic phases for adoption was the preparation and research phase. The next is support for external stakeholders to use IPv6 to gain access to the administration's services. The most basic of these services is to have the primary website for the public administration be available over the IPv6 Internet. For services that are delivered through a third party, this could be as easy as checking a box on a web-hosting (or, content delivery network) service portal. On the other hand, if the administration runs its own web services, the provisioning of those servers as IPv6-compatible would be far more complex.

The goal of external stakeholder support for IPv6 is to ensure that content from IPv6 hosts does not have to be translated to and from IPv4. While this does not ensure that the IPv6 experience will be improved, it does ensure that it isn't degraded by unnecessary translation between IPv6 and IPv4.

This is important for external stakeholders because of the mechanisms at the end user's computer. In the early days of IPv6 transition, standards developers decided that if a host has both IPv4 and IPv6 connectivity, it should prefer the IPv6 connection.[7] For reasons beyond the scope of this paper, this broke some implementations of end-user IPv6 access and a change of approach was taken. In today's Internet the operating system or browser simply tests both the IPv4 and IPv6 connection simultaneously and picks the one that works. This has led to greater adoption of IPv6 in end-user devices including tablets, mobile devices and laptops. In these cases, the operating system almost always ships with this methodology in place and IPv6 gets used for many major websites as a result – without any intervention from the end-user.

Internal Adoption

The last of the three phases of IPv6 deployment for public administrations is the deployment of IPv6 on internal networks. Many units of government are slower at deploying IPv6 internally than they are providing resources for external stakeholders on IPv6. Given that all modern operating systems have IPv6 enabled by default, IPv6 is effectively running on any network when it hasn't been explicitly disabled.

Still many administrations do not make the internal transition for reason of budget, business case or lack of technical skills/knowledge. Other administrations rely on IPv4's private address space architecture for internal networks. Still others have little operational best practice to leverage from other organisations due to slow overall adoption.

The workshop component of this project seeks to provide case studies for success for internal adoption of IPv6 that can be shared across national boundaries.

IPv6 Subnetting

IPv6 addresses don't have a fixed structure, like the older class A/B/C system originally used with IPv4. Instead, IPv6 subnets usually are /64 prefixes. Other subnet sizes are possible but interfere with mechanisms such as stateless address autoconfiguration. As a result, even very small subnets, such as a point-to-point link, use the same size IPv6 address block as very large subnets, such as a large office or enterprise networks containing a number of Ethernet switches.

[7] See https://datatracker.ietf.org/doc/rfc6724/

Principles of Address Planning

- Properly sized initial allocations
- Sparse assignments of subnets
- Hierarchical organisation of subnets
- Adherence to 4-bit boundaries
- Uniform subnetting and aggregation

There is no single template for every IPv6 addressing plan. However, almost every IPv6 addressing plan follows the principles above. The goals include stable routing table size through aggregation, ensuring support for growth through sparse assignment of subnets, and the general use of Global Unicast Addresses as the default address space for devices and services in the administration.

Elsewhere in this report, the details of acquiring an IPv6 prefix in Europe are provided.

In its essentials, the addressing plan is a plan to divide the available space to meet both short-term needs and support long-term management of the IPv6 space.

Unlike IPv4 addressing plans which relied on host counts and resulted in a bottom-up approach to address planning, the vast amount of address space in IPv6 allows us to use a top-down, organisational approach. The main steps to building the plan is to obtain a properly sized initial allocation of IPv6, then identify the organizing attributes and hierarchy of the network or governmental agency (e.g. geographical location, organizational role or other), and then assign blocks of subnets within each attribute or hierarchical level to satisfy their immediate and long-term subnet requirements.

Two common organisational features used to define and assign subnets in IPv6 are location and function. Location is not just based on physical geography but could be an administrative division in a country. Function can be any logical or administrative entity and are often associated with some specific set of stakeholders (e.g. licensees, regulated entities, tourists, citizens), hosts (e.g. mobile devices), server types or roles. The benefit of using either function or location comes from the ability to generalise and keep consistent the approach across multiple sites form ease of operations and management. This remains true even as growth takes place of the site prefix changes.

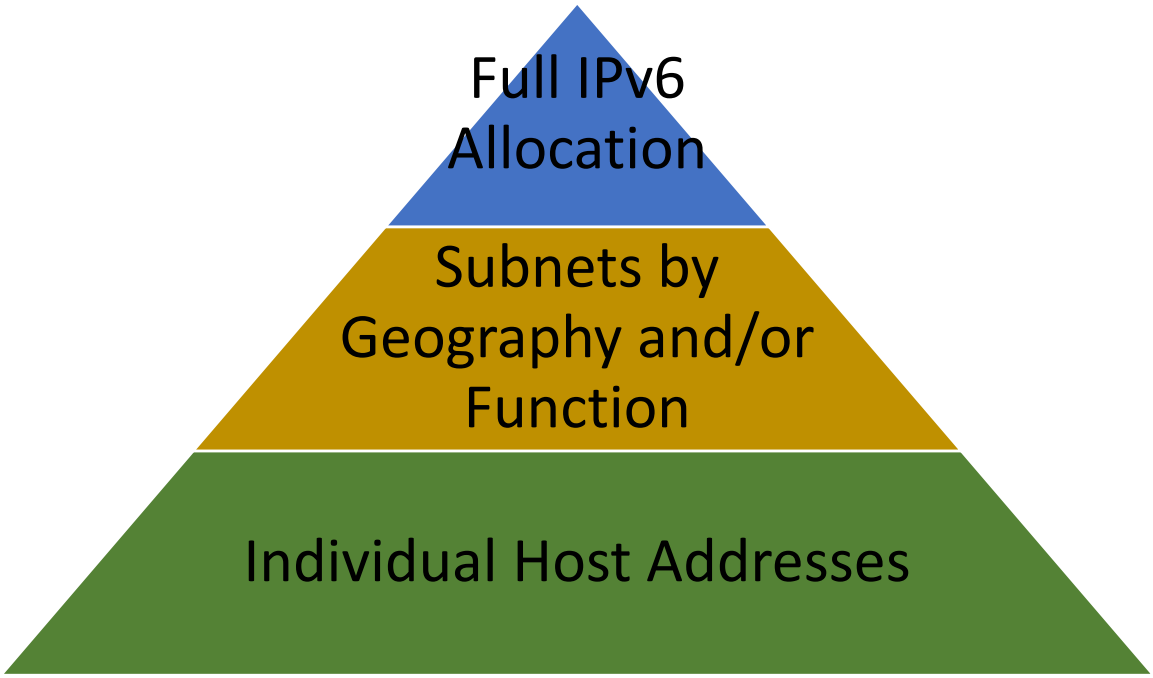

In the examples in this paper, we assume that the public administration has become an LIR and a RIPE member. LIRs get at least a /32, so we've used that as our example for building an IPv6 addressing plan.

Thus, the first 32 bits of the address are allocated to the public administration and specific to it (I.e. assigned by RIPE). The last 64 bits are used within each subnet. Thus, an IPv6 addressing plan is about how to use bits 32 to 63.

Geographical Subnetting Example

Here's an example of geographical subnetting. Suppose we acquire a /32 prefix (2001:db8::/32) and we are building subnets for the provinces of The Netherlands:

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 we would need for a subnet we would total the number of geographic areas, add a few to accommodate for growth and consult the following chart to see the number of bits needed:

Bits	Regions or Functions
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

Since there are eleven regions we need four bits ($2^4 = 16$) to accommodate the regions plus a little more for growth.

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

The following table is an example of how the subnets for the geographic regions might be laid out.

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	

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.

It is easy to imagine a similar subnetting plan based not on geographic region, but on function instead. The addressing plan would be similar:

- Establish the list of organisational functions to be supported;
- Ensure that you have 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; and,
- Then, sparsely distribute the functions across the available subnets created by full initial prefix plus functions.

Once again, imagine we acquire a /32 prefix (2001:db8::/32) and we are building subnets for a collection of national networks:

1. Management networks
2. Educational networks
3. National Health System networks
4. National Army and Defense 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 a little more for growth.

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

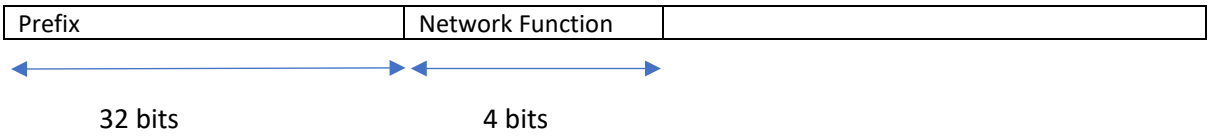

The following table is an example of how the subnets for the national networks might be laid out.

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	

However, this is a relatively simplistic model for building an IPv6 addressing plan for public administrations. It would work well in administrations where there was little functional or organizational hierarchy, but what about in more complex organisations where there was a need to represent both function and geographic distribution. Such an addressing plan is certainly possible but only a little more complex.

One choice to be made is whether to use location or function first. If location is the primary subnet, each location is given a subnet which is further broken down by function. This has the advantage of keeping all regional traffic on the same subnet, and this might have the side-effect of optimising aggregation in the routing table. However, applying access control and other administrative rules which are based on use or function is made more difficult. Using location as the primary subnet happens in cases where each location is customising their own address plan.

In the other case, a public administration might group the functions as the primary subnet and then divide these subnets by region. The advantage of grouping the functions in this way is the ease in which access control and other security policies can be applied to the subnets. For example, many firewall policies are based on the type of use and not on the location of the network.

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.

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

The following table shows some of the detail of how the addressing plan would be laid out for a network where function was captured first and then 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
	. . .	. . .

In its essentials, the IPv6 addressing plan methodology is defining and assigning blocks of subnets based on the structure of the network or the organisation. It may be easier for administrative and security (or other) reasons to group by function first and then by region. It is also easy to imagine further subdivisions that create a hierarchy – for instance, having a separate subnet per building in each location. Having a /32 to work with initially (which is what RIPE gives, by default, to an LIR) means that the designer of the addressing plan can use many layers of hierarchy.

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[8] that documents the German IPv4 addressing plan.

At a high level, the German IPv4 addressing plan looks like this:

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

From the point of view of a local administration, each gets a /48 prefix to work with. The 16 bits available for local subnets are theirs to organise as their network requirements dictate. 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 Center or Large-Scale public administration
- A home office and very small administrations

8

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

Getting IPv6 Addresses

In Europe, there are three ways to acquire IPv6 addresses:

- To be a part of an organisation that already has an allocation of IPv6 addresses;
- To be a customer or consumer of an organisation that makes sub-allocations, for instance, an ISP; or,
- To be a Local Internet Registry (LIR).

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, IPv6 address allocation rules are set by the membership of RIPE NCC[9].

Initial Allocation

If a public administration is a LIR, or 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. In fact, the administration can request address space of the size /29 without any further documentation.

If the administration desires an initial allocation 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.

Future Allocation

For public administrations that have used their initial allocation, RIPE has a set of policies that govern how to request subsequent allocations.

Subsequent allocations take place in two cases:

- Where the initial allocation has been utilized past a certain threshold; or,
- When new needs are documented that could not be met by the initial allocation.

In the first case, measurement of the utilisation of the initial allocation is done via a metric called the HD-ratio. The HD-ratio is documented in RFC 3194[10] and is effectively:

$$\text{HD} = \frac{\log(number\ of\ allocated\ objects)}{\log(maximum\ number\ of\ allocatable\ objects)}$$

An HD-Ratio value of 0.94 is set as the threshold for indicating that an initial allocation has been effectively used and as justifying the allocation of additional address space.

[9] https://www.ripe.net/publications/docs/ripe-699
[10] https://tools.ietf.org/html/rfc3194

For a given prefix size, the following table shows the number of /56 prefixes available, the number of prefixes that would have to be in use and the raw utilisation percentage that would qualify for a subsequent allocation.

Prefix	Total /56s	/56s HD 0.94	Util %
10	70368744177664	10388121308479	14.76
11	35184372088832	5414630391777	15.39
12	17592186044416	2822283395519	16.04
13	8796093022208	1471066903609	16.72
14	4398046511104	766768439460	17.43
15	2199023255552	399664922315	18.17
16	1099511627776	208318498661	18.95
17	549755813888	108582451102	19.75
18	274877906944	56596743751	20.59
19	137438953472	29500083768	21.46
20	68719476736	15376413635	22.38
21	34359738368	8014692369	23.33
22	17179869184	4177521189	24.32
23	8589934592	2177461403	25.35
24	4294967296	1134964479	26.43
25	2147483648	591580804	27.55
26	1073741824	308351367	28.72
27	536870912	160722871	29.94
28	268435456	83774045	31.21
29	134217728	43665787	32.53
30	67108864	22760044	33.92
31	33554432	11863283	35.36

32	16777216	6183533	36.86

If the threshold for subsequent allocation is met, it is immediately eligible to obtain an additional allocation that results in a doubling of the address space allocated to it. Where possible, the allocation will be made from an adjacent address block, meaning that its existing allocation is extended by one bit to the left. The idea behind this is to meet the goal of limiting the size of the routing table by having the organisation able to announce a single, contiguous prefix, rather than two separate prefixes.

If an organisation needs more address space for reasons other than utilisation, it must provide documentation justifying its new requirements. The allocation made will be based on the documentation provided, just as in the case of an initial allocation.

The Public Administration as LIR

In Europe, if a public administration wants to have direct control of an allocation of IPv6 address space, it should choose to become a member of the RIPE NCC. RIPE members can request IPv6 address space, make assignments to other organisations and also are eligible for a one-time allocation of a very small (/22) amount of IPv4 address space. For public administrations, any legal entity is allowed to 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[11], 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[12], 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[13].

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

[11] https://my.ripe.net/#/public/membership
[12] https://lirportal.ripe.net/
[13] https://www.ripe.net/manage-ips-and-asns/ipv6/request-ipv6

IPv6 Address Maintenance

IP Address Management in Public Administrations

IP Address Management (IPAM) refers to the management of allocation, administration, reporting and tracking of IP addresses – including IPv6 addresses. Enterprises generally deploy systems and processes that interact with the DNS (naming) and DHCP (IP address allocation) infrastructure in order to provide IPAM capabilities.

In small organisations, in both the public and private sector, IT departments use manual processes, spreadsheets or other home-grown tools for IP address management. The simple addition of a printer to the network might involve several steps and hours or days of elapsed time as new people are involved in the process of allocating and assigning an address.

With IPv6 the problem of IP address management is made more difficult by the expanded address space and the longer literal representations of the addresses. This is not the only difference: most devices on a network will have both link-local and a global unicast address. The resulting complexity – and the interaction between the address and the DNS and DHCP - is why organisations move away from manual systems to automated approaches.

Automated IPAM systems usually have some essential features 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.
- **Measuring Consumption of IPv6 Address Resources.** An IPAM system also gives the ability to measure and track the consumption of available IPv6 addresses – measured in subnets rather than hosts. This becomes a useful tool for capacity planning and network management.
- **Reporting.** The ability of the IPAM tool to generate reports of the IPv6 network to assist in planning and maintenance of the network.
- **IPv6 Address Plan Design.** IPAM systems provide a tool to help build 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 its addresses via an LIR rather than getting them from an upstream ISP, the unit of government 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 RIPE'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.**

IPAM and DHCPv6

In IPv6, autoconfiguration can take place for hosts using stateful DHCPv6, a combination of DHCPv6 and SLAAC, and SLAAC only. When using stateful DHCPv6, the host receives an IPv6 address from the server. In stateless DHCPv6 the host uses SLAAC to generate its own address while DHCP

provides options and other information. With either approach to DHCPv6, the server can provide optional information including DNS servers and search domain information.

Unlike DHCPv4, DHCPv6 relies on router advertisements from one or more routers on the local link to provide default gateway information.

DHCPv6 is complicated by the fact that the router advertisement flags have to be properly configured depending on the auto-address configuration method in use. This configuration is further complicated when a combination of autoconfiguration options are being used on the same local link. Because of this, many network designers use only a single auto-addressing method per link as a way to minimise operational challenges related to router configurations.

Many public administration networks will choose to run DHCPv6 instead of SLAAC because it provides tighter control over access to network resources. If a router is properly configured for SLAAC, this style of auto-configuration allows any host to connect to the network and autoconfigure a usable address.

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

IPAM and DNS

From an IPv6 addressing point of view, the DNS has two main uses:

- Forward mapping of domain names to IPv6 addresses; and
- Reverse mapping of IPv6 addresses to domain names

Since the IPv6 addresses are long and unwieldy, it is extremely useful to automate the process of configuring these resource records in the DNS. An IPAM solution will provide the tools to forward-map an authoritative zone or reverse-map an authoritative zone for IPv6. The advantage to using the IPAM solution is that the resource records are automatically generated, and name servers (including primary and secondaries) can be automatically updated.

Public Administrations, IPv6 and Network Growth and Change

The addressing strategy outlined in this paper 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 part of foundation 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 has to 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 may 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 providers. If the public administration is not an LIR and does not have Provider Independent 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. With an enormous available address space, renumbering is never necessary because of insufficient host addresses. In addition, each node in the network can have as many IPv6 addresses configured from as many prefixes as needed.

Renumbering in a public administration is essentially a reconfiguration plan and procedure in which an existing, in-use address prefix (or, perhaps a set of prefixes) is replaced by a new, yet-unused address prefix.

One common strategy for renumbering is to start by adding addresses from the new prefix to network infrastructure like routers, switches and the links that connect them. IPv6 has been designed to allow nodes to have multiple addresses and subnets on single interfaces. As a result, this step, adding the new prefixes to infrastructure, can often be accomplished without impacting production traffic using the old prefix.

A next step is to configure the hosts so that they begin using the new prefix. It's possible to do this using SLAAC, but DHCPv6 is much simpler. In public administration settings where DHCPv6 is in use, a unicast "RECONFIGURE" message from the DNCPv6 server to the hosts will cause the hosts to change their addresses. Dynamic DNS will then automatically update the hosts AAAA and PTR DNS resource records. If an IPAM solution is in place, it is relatively easy to break the renumbering up into blocks of devices (for instance, floor-by-floor) to limit the operational strain and isolate any problems that might arise. In fact, in the iETF's IPv6 Enterprise Numbering Scenarios, Considerations and Methods document, it says: "It is recommended that the site have an automatic and systematic procedure for updating/synchronizing its DNS records, including both forward and reverse mapping. In order to simplify the operational procedure, the network architect should combine the forward and reverse DNS updates in a single procedure. A manual on-demand updating model does not scale and increases the chance of errors.[14]"

Once the hosts are configured with a new IPv6 address, the internal logic at the host will determine which of the two addresses are used for new sessions[15]. Once everything is stable, the old prefixes and addresses can be removed from hosts, network infrastructure and the DNS.

Public Administrations, IPv6 and Reachability

Whether a public administration is organized as a single entity, or if it is a complex hierarchy of inter-related national, regional and local organisations, the IPv6 address plan is intended to make those addresses reachable. Small public administrations may be using Provider Aggregatable addresses from an ISP that provides IPv6 transit. In those cases, it is the ISP that ensures reachability.

For the public administration who has worked to become an LIR at RIPE NCC and acquired an IPv6 PI allocation, the situation is different. Working with the ISP or provider of transit, the public administration has available all the routing protocols that are familiar from IPv4. Selecting the right one is often a matter of cooperation between the public administration and its Internet Service Provider. Often, operational continuity with an existing IPv4 network is the overriding factor in

[14] https://tools.ietf.org/html/rfc6879
[15] https://tools.ietf.org/html/rfc6724

selecting a routing protocol. However, using a different routing protocol or dual-topology mode may provide better isolation and fault tolerance for both the IPv4 and IPv6 networks.

Appendix A: RIPE Requirements for IPv6 Compatibility

Requirements for "host" equipment

Mandatory support:

- IPv6 Basic specification [RFC2460] *
- IPv6 Addressing Architecture [RFC4291] *
- Default Address Selection [RFC3484]
- Unique Local IPv6 Unicast Addresses (ULA) [RFC4193]
- ICMPv6 [RFC4443] *
- DHCPv6 client [RFC3315] *
- SLAAC [RFC4862] *
- Path MTU Discovery [RFC1981] *
- Neighbor Discovery [RFC4861] *
- If support for tunneling and dual stack is required, the device must support Basic Transition Mechanisms for IPv6 Hosts and Routers [RFC4213]
- If support for mobile IPv6 is required, the device must support “MIPv6” [RFC6275, RFC5555] and “Mobile IPv6 Operation With IKEv2 and the Revised IPsec Architecture” [RFC4877]
- DNS protocol extensions for incorporating IPv6 DNS resource records [RFC3596]
- DNS message extension mechanism [RFC2671]
- DNS message size requirements [RFC3226]
- Deprecation of Type 0 Routing Headers in IPv6 [RFC5095] *

Optional support:

- IPv6 Router Advertisement Options for DNS Configuration [RFC6106]
- Extended ICMP for multi-part messages [RFC4884]
- SeND [RFC3971]
- SLAAC Privacy Extensions [RFC4941]
- Stateless DHCPv6 [RFC3736] *
- DS (Traffic class) [RFC2474, RFC3140]
- Cryptographically Generated Addresses [RFC3972]
- IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] *
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]
- Multicast Listener Discovery version 2 [RFC3810] *
- Packetisation Layer Path MTU Discovery [RFC4821]
- IPv6 Host-to-Router Load Sharing [RFC4311]
- Default Router Preferences and More-Specific Routes [RFC4191]

Requirements for consumer grade "Layer 2 switch" equipment

Optional support (management)

- MLDv2 snooping [RFC4541]
- IPv6 Basic specification [RFC2460] *
- IPv6 Addressing Architecture [RFC4291] *

- Default Address Selection [RFC3484]
- ICMPv6 [RFC4443] *
- SLAAC [RFC4862] *
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]

Requirements for enterprise/ISP grade "Layer 2 switch" equipment

Mandatory support:

- MLDv2 snooping [RFC4541]
- DHCPv6 filtering [RFC3315]
- Router Advertisement (RA) filtering [RFC4862]
- Dynamic "IPv6 Neighbor solicitation/advertisement" inspection [RFC4861]
- Neighbor Unreachability Detection [NUD, RFC4861] filtering
- Duplicate Address Detection [DAD, RFC4429] snooping and filtering.[2]

Optional support (management):

- IPv6 Basic specification [RFC2460] *
- IPv6 Addressing Architecture [RFC4291] *
- Default Address Selection [RFC3484]
- ICMPv6 [RFC4443] *
- SLAAC [RFC4862] *
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]
- IPv6 Routing Header [RFC2460, Next Header value 43] filtering *
- Deprecation of Type 0 Routing Headers in IPv6 [RFC5095] *
- UPnP filtering

Requirements for "router or Layer 3 switch" equipment

Mandatory support:

- IPv6 Basic specification [RFC2460] *
- IPv6 Addressing Architecture [RFC4291] *
- Default Address Selection [RFC3484]
- Unique Local IPv6 Unicast Addresses (ULA) [RFC4193]
- ICMPv6 [RFC4443] *
- SLAAC [RFC4862] *
- MLDv2 snooping [RFC4541]
- Multicast Listener Discovery version 2 [RFC3810] *
- Router-Alert option [RFC2711]
- Path MTU Discovery [RFC1981] *
- Neighbor Discovery [RFC4861] *
- Deprecation of Type 0 Routing Headers in IPv6 [RFC5095] *

- If a dynamic interior gateway protocol (IGP) is requested, then RIPng [RFC2080], OSPF-v3 [RFC5340] or IS-IS [RFC5308] must be supported. The contracting authority shall specify the required protocol.
- If OSPF-v3 is requested, the equipment must comply with "Authentication/Confidentiality for OSPF-v3" [RFC4552]
- If BGP4 protocol is requested, the equipment must comply with RFC4271, RFC1772, RFC4760, RFC1997, RFC3392 and RFC2545
- Support for QoS [RFC2474, RFC3140]
- If support for tunneling and dual stack is required, the device must support Basic Transition Mechanisms for IPv6 Hosts and Routers [RFC4213]
- If support for tunneling and dual stack is required, the device must support Generic Packet Tunneling and IPv6 [RFC2473]
- If 6PE is requested, the equipment must support "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)" [RFC4798]
- If mobile IPv6 is requested, the equipment must support MIPv6 [RFC6275, RFC5555] and "Mobile IPv6 Operation With IKEv2 and the Revised IPsec Architecture" [RFC4877]
- If the IS-IS routing protocol is requested the equipment must support "M-ISIS: Multi-Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)" [RFC5120]
- If MPLS functionality (for example, BGP-free core, MPLS TE, MPLS FRR) is requested, the PE-routers and route reflectors must support "Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)" [RFC4798]
- If Layer 3 VPN functionality is requested, the PE-routers and route reflectors must support "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN" [RFC4659]
- If MPLS Traffic Engineering is used in combination with IS-IS routing protocol, the equipment must support "M-ISIS: Multi-Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)" [RFC5120]

Optional support:

- IPv6 Router Advertisement Options for DNS Configuration [RFC6106]
- DHCPv6 client/server/relay [RFC3315] *
- Extended ICMP for multi-part messages [RFC4884]
- SeND [RFC3971]
- SLAAC Privacy Extensions [RFC4941]
- Stateless DHCPv6 [RFC3736] *
- DHCPv6 PD [RFC3633] *
- Route Refresh for BGP-4 Capabilities [RFC2918]
- BGP Extended Communities Attribute [RFC4360]
- (QOS) Assured Forwarding [RFC2597]
- (QOS) Expedited Forwarding [RFC3246]
- Generic Routing Encapsulation [RFC2784]
- Cryptographically Generated Addresses [RFC3972]
- IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] *
- Using IPsec to Secure IPv6-in-IPv4 tunnels [RFC4891]
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]

- DNS protocol extensions for incorporating IPv6 DNS resource records [RFC3596]
- DNS message extension mechanism [RFC2671]
- DNS message size Requirements [RFC3226]
- 127-bit IPv6 Prefixes on Inter-Router Links [RFC6164]
- Packetisation Layer Path MTU Discovery [RFC4821]
- IPv6 Host-to-Router Load Sharing [RFC4311]
- Default Router Preferences and More-Specific Routes [RFC4191]

Requirements for "network security equipment"

Equipment in this section is divided into three subgroups:

- Firewall (FW)
- Intrusion prevention device (IPS)
- Application firewall (APFW)

For every mandatory standard the applicable subgroups are specified in parentheses at the end of the line.

Mandatory support:

- IPv6 Basic specification [RFC2460] (FW, IPS, APFW) *
- IPv6 Addressing Architecture [RFC4291] (FW, IPS, APFW)
- Default Address Selection [RFC3484] (FW, IPS, APFW)
- ICMPv6 [RFC4443] (FW, IPS, APFW) *
- SLAAC [RFC4862] (FW, IPS) *
- Deprecation of Type 0 Routing Headers in IPv6 [RFC5095] *
- Inspecting IPv6-in-IPv4 protocol-41 traffic, which is specified in: Basic Transition Mechanisms for IPv6 Hosts and Routers [RFC4213] (IPS)
- Router-Alert option [RFC2711] (FW, IPS)
- Path MTU Discovery [RFC1981] (FW, IPS, APFW) *
- Neighbor Discovery [RFC4861] (FW, IPS, APFW) *
- If the request is for the BGP4 protocol, the equipment must comply with RFC4271, RFC1772, RFC4760 and RFC2545 (FW, IPS, APFW)
- If the request is for a dynamic internal gateway protocol (IGP), then the required RIPng [RFC2080], OSPF-v3 [RFC5340] or IS-IS [RFC5308] must be supported. The contracting authority shall specify the required protocol. (FW, IPS, APFW)
- If OSPF-v3 is requested, the device must support "Authentication/Confidentiality for OSPFv3" [RFC4552] (FW, IPS, APFW)
- Support for QoS [RFC2474, RFC3140] (FW, APFW)
- If tunneling is required, the device must support Basic Transition Mechanisms for IPv6 Hosts and Routers [RFC4213] (FW)

A Network Security Device is often placed where a Layer 2 switch or a router/Layer 3 switch would otherwise be placed. Depending on this placement those requirements should be included.

Functionality and features that are supported over IPv4 should be comparable with the functionality supported over IPv6. For example, if an intrusion prevention system is capable

of operating over IPv4 in Layer 2 and Layer 3 mode, then it should also offer this functionality over IPv6. Or if a firewall is running in a cluster capable of synchronising IPv4 sessions between all members of a cluster, then this must also be possible with IPv6 sessions.

Optional support:

- IPv6 Router Advertisement Options for DNS Configuration [RFC6106]
- DHCPv6 client/server/relay [RFC3315] *
- Extended ICMP for Multipart Messages [RFC4884]
- SeND [RFC3971]
- SLAAC Privacy Extensions [RFC4941]
- Stateless DHCPv6 [RFC3736] *
- DHCPv6 PD [RFC3633] *
- BGP Communities Attribute [RFC1997]
- BGP Capabilities Advertisement WITH-4 [RFC3392]
- (QOS) Assured Forwarding [RFC2597]
- (QOS) Expedited Forwarding [RFC3246]
- Unique Local IPv6 Unicast Addresses (ULA) [RFC4193]
- Cryptographically Generated Addresses [RFC3972]
- IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] *
- Using IPsec to Secure IPv6-in-IPv4 Tunnels [RFC4891] (FW)
- OSPF-v3 [RFC5340]
- Authentication/Confidentiality for OSPF-v3 [RFC4552]
- Generic Packet Tunneling and IPv6 [RFC2473]
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]
- DNS extensions to support IPv6 [RFC3596]
- DNS message extension mechanism [RFC2671]
- DNS message size requirements [RFC3226]
- Using IPSec to Secure IPv6-in-IPv4 Tunnels [RFC4891]
- Multicast Listener Discovery version 2 [RFC3810] *
- MLDv2 snooping [RFC4541] (when in L2 or passthrough mode) *
- Packetisation Layer Path MTU Discovery [RFC4821]
- IPv6 Configuration in Internet Key Exchange Protocol Version 2 (IKEv2) [RFC5739]
- IPv6 Host-to-Router Load Sharing [RFC4311]
- Default Router Preferences and More-Specific Routes [RFC4191]

Requirements for CPE equipment

Mandatory support:

- RFC6204 (Basic Requirements for IPv6 Customer Edge Routers) *

Optional support:

- IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] *

- If support for mobile IPv6 is required, the device needs to comply to "MIPv6" [RFC6275, RFC5555] and "Mobile IPv6 Operation With IKEv2 and the Revised IPsec Architecture" [RFC4877]
- Extended ICMP for multi-part messages [RFC4884]
- SeND [RFC3971]
- SLAAC Privacy Extensions [RFC4941]
- DS (Traffic class) [RFC2474, RFC3140]
- Cryptographically Generated Addresses [RFC3972]
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]
- Multicast Listener Discovery version 2 [RFC3810] *
- Packetisation Layer Path MTU Discovery [RFC4821]
- IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) [RFC5969]
- Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion [RFC6333] If support this then also must support Dynamic Host Configuration protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite [RFC6334]
- The A+P Approach to the IPv4 Address Shortage [RFC6346]
- IPv6 Configuration in Internet Key Exchange Protocol Version 2 (IKEv2) [RFC5739]
- IPv6 Host-to-Router Load Sharing [RFC4311]
- Default Router Preferences and More-Specific Routes [RFC4191]

Requirements for mobile devices

Mandatory support:

- IPv6 basic specification [RFC2460] *
- Neighbor Discovery for IPv6 [RFC4861] *
- IPv6 Stateless Address Autoconfiguration [RFC4862] *
- IPv6 Addressing Architecture [RFC4291] *
- ICMPv6 [RFC4443] *
- IPv6 over PPP [RFC2472]
- Multicast Listener Discovery version 2 [RFC3810] *
- IPv6 Router Alert Option [RFC2711]
- DNS protocol extensions for incorporating IPv6 DNS resource records [RFC3596]

Optional support:

- Privacy Extensions for Stateless Address Autoconfiguration in IPv6 [RFC4941]
- Path MTU Discovery for IPv6 [RFC1981] *
- Generic Packet Tunneling for IPv6 [RFC2473]
- DHCPv6 [RFC3315] *
- Stateless DHCPv6 [RFC3736]
- DHCPv6 option for SIP servers [RFC3319]
- IPv6 Prefix Options for DHCPv6 [RFC3633]
- Prefix Exclude Option for DHCPv6-based Prefix Delegation [draft-ietf-dhc-pd-exclude]
- Default Address Selection [RFC3484]
- IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] *

- IKEv2 Mobility and Multihoming Protocol MOBIKE [RFC 4555]
- IPv6 Host-to-Router Load Sharing [RFC4311]
- Default Router Preferences and More-Specific Routes [RFC4191]

References:

- 3GPP
- Internetworking Between Public Land Mobile Network (PLMN) supporting packet based services and Packet Data Networks (PDN) [3GPP TS 29.061]
- GPRS Service Description [3GPP TS 23.060]
- General Packet Radio Service (GPRS) enhancements for Evolved Universal Terrestrial Radio Access Network (E-UTRAN) access [3GPP TS 23.401]
- Signaling flows for IP multimedia Call control based on SIP and SDP [3GPP TS 24.228]
- IP multimedia call control protocol based on SIP and SDP [3GPP TS 24.229]
- IP Based Multimedia Framework [3GPP TS 22.941]
- Architectural Requirements [3GPP TS 23.221]
- Packet domain; Mobile Stations (MS) Supporting Packet Switching Service [3GPP TS 27.060]
- IPv6 migration guidelines [3GPP TR 23.975]
- IETF
- IPv6 for Some Second and Third Generation Cellular Hosts [RFC3316]
- Recommendations for IPv6 in 3GPP Standards [RFC3314]
- IPv6 in 3rd Generation Partnership Project (3GPP) [RFC6459]

Requirements for load balancers

A load balancer distributes incoming requests and/or connections from clients to multiple servers. Load balancers will have to support several combinations of IPv4 and IPv6 connections:

- Load balancing IPv6 clients to IPv6 servers (6-to-6) **must** be supported
- Load balancing IPv6 clients to IPv4 servers (6-to-4) **must** be supported
- Load balancing IPv4 clients to IPv4 servers (4-to-4) **should** be supported
- Load balancing IPv4 clients to IPv6 servers (4-to-6) **should** be supported
- Load balancing a single external/virtual IPv4 address to a mixed set of IPv4 and IPv6 servers **should** be supported
- Load balancing a single external/virtual IPv6 address to a mixed set of IPv4 and IPv6 servers **should** be supported

If a load balancer provides Layer 7 (application level / reverse proxy, defined as 'surrogate' in section 2.2 of RFC3040) load balancing then support for the X-forwarded-for (or equivalent) header in HTTP **must** be provided in order to make the source IP address of the client visible to the servers.

Mandatory support:

- IPv6 Basic specification [RFC2460] *
- IPv6 Addressing Architecture [RFC4291] *
- Default Address Selection [RFC3484]
- Unique Local IPv6 Unicast Addresses (ULA) [RFC4193]

- ICMPv6 [RFC4443] *
- Path MTU Discovery [RFC1981] *
- Neighbor Discovery [RFC4861] *
- DNS protocol extensions for incorporating IPv6 DNS resource records [RFC3596]
- DNS message extension mechanism [RFC2671]
- DNS message size requirements [RFC3226]
- Deprecation of Type 0 Routing Headers in IPv6 [RFC5095] *

Optional support:

- IPv6 Router Advertisement Options for DNS Configuration [RFC6106]
- Extended ICMP for multi-part messages [RFC4884]
- SeND [RFC3971]
- DS (Traffic class) [RFC2474, RFC3140]
- Cryptographically Generated Addresses [RFC3972]
- SNMP protocol [RFC3411]
- SNMP capabilities [RFC3412, RFC3413, RFC3414]
- SNMP MIBs for IP [RFC4293] Forwarding [RFC4292] and DiffServ [RFC3289]
- Multicast Listener Discovery version 2 [RFC3810] *
- Packetisation Layer Path MTU Discovery [RFC4821]
- NAT64/DNS64 [RFC6146, RFC6147]
- If support for IPsec is required, the device must support IPsec/IKEv2 [RFC4301, RFC4303, RFC4302, RFC5996] * and Redirect Mechanism for the Internet Key Exchange Protocol Version 2 (IKEv2) [RFC5685]
- If support for BGP4 is required, the equipment must comply with RFC4271, RFC1772, RFC4760 and RFC2545
- If support for a dynamic internal gateway protocol (IGP) is required, the RIPng [RFC2080], OSPF-v3 [RFC5340] or IS-IS [RFC5308] must be supported. The contracting authority shall specify the required protocol.
- If OSPF-v3 is requested, the device must support "Authentication/Confidentiality for OSPFv3" [RFC4552] (FW, IPS, APFW)
- IPv6 Host-to-Router Load Sharing [RFC4311] (FW)
- Default Router Preferences and More-Specific Routes [RFC4191] (FW)

Requirements for IPv6 support in software

All software must support IPv4 and IPv6 and be able to communicate over IPv4-only, IPv6-only and dual-stack networks. If software includes network parameters in its local or remote server settings, it should also support configuration of IPv6 parameters.

All features that are offered over IPv4 must also be available over IPv6. The user should not experience any noticeable difference when software is communicating over IPv4 or IPv6, unless this is providing explicit benefit to the user.

It is strongly recommended not to use any address literals in software code, as described in “Default Address Selection for Internet Protocol version 6” [RFC3484].

Appendix B: Special Use IPv6 Addresses

Addresses	Description	IETF RFC
::1/128	Node-scoped Unicast, loopback address	4291
::/128	Node-Scoped Unicast, unspecified address	4291
::ffff:0:0/96	IPv6-mapped addresses	4291
::/96	IPv4-compatible addresses (deprecated)	4291
fe80::/10	Link-scoped Unicast	4291
fc00::/7	Unique-Local addresses (ULAs)	4193
2001:0002::/48	Reserved for IPv6 Benchmarking	5180
2001:db8::/32	Documentation prefix	3849
100::/64	Remotely Triggered Black Hole Addresses	6666
2001:10::/28	Overlay Rouf Cryptographic Hash Identifiers (ORCHID)	4843
fec0::/10	Site-Local Unicast (deprecated)	3879
ff00::/8	Multicast	4291
5f00::/8	First 6bone allocation	1897
3ffe::/16	Second 6bone allocation	2471