Title: **Deliverable D2.2** **IPv6 addressing plans**	**Document Version:** 1.0

Project Number: 297239	**Project Acronym:** GEN6	**Project Title:** Governments ENabled with IPv6

Contractual Delivery Date: 30/01/2012	**Actual Delivery Date:** 30/01/2012	**Deliverable Type* - Security**:** R – PU

* Type:	P - Prototype, R - Report, D - Demonstrator, O - Other
** Security Class:	PU- Public, PP – Restricted to other programme participants (including the Commission), RE – Restricted to a group defined by the consortium (including the Commission), CO – Confidential, only for members of the consortium (including the Commission)

Responsible and Editor/Author: Jordi Palet Martínez	**Organization:** Consulintel	**Contributing WP:** WP2

Authors (organisations):

Alvaro Vives (Consulintel), Antonio Skarmeta (UMU), Carsten Schmoll (FHG), Anastasios Zafeiropoulos (GRNET), Konstantinos Koumoutsos (CTI), Martin Krengel (Citkomm)

Abstract:

This deliverable covers the IPv6 addressing plan issues to be taken into account in the public administration networks, studying and suggesting possible addressing plans.

Keywords:

IPv6, Governments, Network Topology, IPv6 Address, IPv6 Addressing Plan.

Revision History

The following table describes the main changes done in this document since its creation.

Revision	Date	Description	Author (Organization)
v0.1	01/07/2012	Document creation	Jordi Palet (Consulintel)
v0.2	21/01/2013	Added content	Alvaro Vives (Consulintel)
v0.3	24/01/2013	Document revision	Alvaro Vives (Consulintel)
v0.4	26/01/2013	Added contribution	Antonio Skarmeta (UMU)
v0.5	28/01/2013	Added content and Turkish Pilot contribution	Alvaro Vives (Consulintel)
v0.6	30/1/2013	Added contribution related to the Greek pilot	Anastasios Zafeiropoulos (GRNET) & Konstantinos Koumoutsos (CTI)
v0.7	05/02/2013	Added contribution	Martin Krengel (Citkomm)
v1.0	06/02/2013	Document revision	Alvaro Vives (Consulintel)

Disclaimer

The GEN6 project (number 261584) is co-funded by the European Commission under the ICT Policy Support Programme (PSP) as part of the Competitiveness and Innovation framework Programme (CIP). This document contains material that is the copyright of certain GEN6 partners and the EC, and that may be shared, reproduced or copied "as is", following the Creative Commons "Attribution-NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-NC 3.0) licence. Consequently, you're free to share (copy, distribute, transmit) this work, but you need to respect the attribution (respecting the project and authors names, organizations, logos and including the project web site URL "http://www.gen6.eu"), for non-commercial use only, and without any alteration, transformation or build upon this work.

The information herein does not necessarily express the opinion of the EC. The EC is not responsible for any use that might be made of data appearing herein. The GEN6 partners do not warrant that the information contained herein is capable of use, or that use of the information is free from risk, and so do not accept liability for loss or damage suffered by any person using this information.

Executive Summary

This deliverable covers the IPv6 addressing plan issues to be taken into account in the public administration networks, studying and suggesting possible addressing plans.

Issues to be taken into account will be showed and some specific examples are included at the end of this document, showing real addressing plans in public administration networks.

Table of Contents

Figure Index

Table Index

1. INTRODUCTION

One of the first things any IPv6 deployment will have to deal with is the addressing plan for the considered network. In our case, in the scope of public administration networks, there is no exception, an addressing plan is needed.

In an efficient IPv6 addressing plan, the IPv6 addressing ranges are grouped effectively and logically. This has several advantages, including:

- Security policies are easier to implement, such as the configuration of access lists and firewalls
- Addresses are easier to trace: the address contains information about the use type or location where the address is in use
- An efficient addressing plan is scalable: it can be expanded, for example, to include new locations or use types
- An efficient IPv6 addressing plan also enables more efficient network management

General issue to be taken into account when creating an IPv6 addressing plan will be described to end this document with some real examples of addressing plans made in different countries by public administrations that already have implemented IPv6 on their network or are planning to do it in the short term.

Before starting with the IPv6 addressing plan specific issues, following there is a general introduction to IPv6 addresses types and formats, which are a prerequisite for properly creating and addressing plan.

1.1 IPv6 addresses

The Internet protocol version 6 (IPv6) addresses are 128 bits long, four times the size of an IPv4 address. There are defined different types (subclasses) of IPv6 addresses. These mostly differ in their intended scope (local/global, unicast/multicast) and are recognizable by their highest order bits (see table 1-1).

IPv6 addresses which are used in IP packets that traverse the Internet are from the range of “global unicast addresses”. These are made up from a global routing prefix, a subnet identifier, and a 64-bit interface identifier. The uppermost bits make up the global routing identifier. The remaining bits between the routing prefix and the interface identifier can be used to structure the numbering of the internal networks of an organization. The lowest 64 bits identify a communication endpoint within a given subnet and are called the interface identifier. The following figure illustrates this:

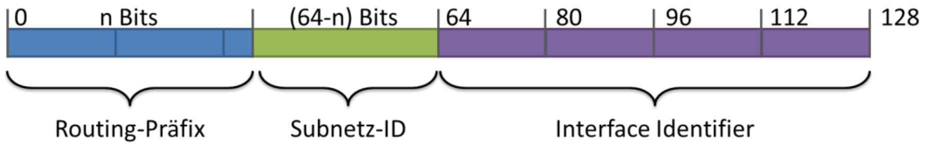

Figure 1-1: Parts of a globally routable IPv6 address

Each 128-bits IPv6 address belongs to exactly one interface (physical or logical/virtual). However each interface may have assigned multiple IPv6 addresses, usually from different address "scopes", e.g. link-local, global, mobile IP. For details on this see [RFC4291] and [RFC5952].

These global unicast addresses (GUA) can be used by clients (workstations, servers, IP-telephones etc.) to access services across multiple subnets and possibly the Internet. This works because GUA are globally routable addresses. For local communication a different IPv6 address type called unique local address (ULA) can be used. These addresses are routable too, but not globally, i.e. across the Internet. These and other address types are explained in the following text.

1.1.1 Address types

Among all the possible 128-bit IPv6 addresses, several distinct types are defined, each for different uses (local/global, unicast/multicast). The type of an IPv6 address can be recognized from the highest order bits, as follows:

Address Type	Binary Prefix	IPv6 Notation
Unspecified	00…0 (128 Bits)	::/128
Loopback interface (localhost)	00…1 (128 Bits)	::1/128
Multicast	1111 1111	ff00::/8
Link-local Unicast	1111111010	fe80::/10
Unique Local Address (ULA)	11111100 11111101	fc00::/8 fd00::/8
Global Unicast Address (GUA)	All other addresses	

Table 1-1: IPv6 Address Types

Note: IPv6 anycast addresses are part of the IPv6 unicast address space (link local or global) and are syntactically undistinguishable from IPv6 unicast addresses.

Link-Local Unicast

Link-local IPv6 addresses are defined in [RFC4291]. They are used for local communication inside an IP subnet or across point to point (PPP) connections. IPv6 routers do not forward any IP packets with link-local addresses in their headers (source or destination). For IPv6 they are denoted starting with the fe80::/64 prefix. IPv6 link-local addresses are assigned by each operating system for each IP interface in an auto configuration manner. Each IPv6 interface must at least have a link-local IP address assigned.

Multicast

IPv6 multicast addresses are used to reach a group of recipients using only one IPv6 destination address (more specifically: to reach a group of network interfaces). IP packets with a multicast destination address are sent to all interfaces which have this address assigned. Note that one interface may be a member of multiple multicast groups; in that case it would have more than one IPv6 multicast address assigned to it.

IPv6 multicast addresses are setup starting with the FF00::/8 prefix, plus 4 bits of flags, 4 bits of scope and a 112 bits long GroupID. As scope the standard defines: Interface local, link local, admin local, site local, organization local and global scope. See also [RFC4007].

Multicast support is of central importance for IPv6 (at least in the LAN environment), since numerous basic functions of the protocol are handled using multicast – and broadcast has been abandoned in IPv6. For these basic functions well-known IPv6 multicast addresses were defined in the standard, for example to reach "all routers in a subnet" or "all NTP servers inside my organization".

Anycast addresses

Anycast addresses in IPv6 [RFC4291] are used to reach one recipient out of a group of communication partners, more specifically out of a group of interfaces. IPv6 packets destined to an anycast address will be routed so that one selected interface out of a group of interfaces with this address will be reached. It is the task of the IT infrastructure to make sure that one of the communication partners with the intended IPv6 address will be reached. Often the goal here is to select the “closest”, “fastest”, or “cheapest” system in the group. This principle can be used to create a kind of load-balancing or fail-over behavior, for example to reach any working DNS server inside a domain where multiple of those do exist.

A single interface may belong to multiple anycast groups by having more than one anycast IPv6 address assigned.

IPv6 anycast addresses can be taken from either the link local unicast or global unicast address space. They are not assigned a dedicated IPv6 prefix, and are therefore not distinguishable from other unicast addresses.

Global unicast address (GUA)

IPv6 global unicast addresses ([RFC4291], section 2.5.4) are used for end-to-end IP communication between two interfaces, usually on two different hosts. Since GUA are globally routable, these two hosts can be situated in different IP subnets, also in the case where a communication between the two hosts crosses the global Internet (assuming that firewall rules and local policies permit the communication). GUA are made up from a global route prefix, a subnet ID and a 64 bit interface identifier, as shown in figure 1-1.

Unique local address (ULA)

Unique local addresses [RFC4193] are unicast addresses for local IP communication. They can be routed across subnets, but not across the global Internet. Still, ULA are globally unique with a very high probability, and therefore could be used to interconnect different sites using tunnels. Also, they can be easily filtered or blocked at administrative borders due to their distinctive prefix. Their main use is the application for local IP communication at one administrative site. Also ULA are not affected by renumbering in case that the global prefix of a site changes, e.g. when choosing a different Internet service provider (ISP) for that site.

The disadvantage of ULA is that they break the IP end-to-end principle between sites, which prohibits their use for certain protocols and applications. For “normal” tasks like Internet access ULA can still be used on clients when a local http proxy is in place to “translate” between the ULA of the clients and the global addresses of requested servers on the Internet (or at external, 3rd party sites).

From the point of view of an application that is programmed to perform IPv6 communication ULA and GUA can be treated in the same way. However from an administrative and security point of view some servers and services may only be accessible with either address type of these two.

IPv6 embedded IPv4 address

IPv6 addresses with an embedded IPv4 address are defined in [RFC4291]. Two types of embedded addresses are distinguished:

- IPv4-compatible IPv6 address
- IPv4-mapped IPv6 address

The first address type is deprecated and need not be supported anymore by devices, clients, and applications.

The IPv4-mapped IPv6 addresses however are in active use for some transition techniques, for example when deploying NAT64 or 6PE. Here they represent IPv4 addresses by using 128 bit IPv6 addresses, in which the lowest 32 bits store the original IPv4 address. Their format is:

0000:0000:0000:0000:0000:FFFF:<IPv4-Address>

See also [RFC4038] for further background information.

1.1.2 IPv6 interface addresses

For the assignment of the 64 bit interface identifier addresses inside an organization the same logical rules apply as in an IPv4 environment:

- Servers must have fixed IP addresses and therefore fixed interface identifiers
- Server IP addresses should be registered in a site's DNS server (in AAAA records for IPv6)
- Client systems may be dynamically assigned an IPv6 address from a pool of addresses (alternatively with IPv6 clients may use IP auto configuration to obtain a valid IPv6 address)

For hosts, as an example of administration of IPv6 addresses, could look like this:

Host	IPv4 Address	IPv6 Address	Remarks
File Server	192.168.22.9	<prefix>:1022::9	Fixed interface identifier ::9
Time Server	193.193.98.3	<prefix>:F098::3	Fixed interface identifier ::3
Client / Workstation	10.0.20.147	<prefix>:0010::3000	Dynamically chosen from an address pool (here: 3000-3FFF)

Table 1-2: Example IPv4/IPv6 addresses for different kinds of end systems

Further on it needs to be decided how the previously selected IPv6 addresses are technically configured to the hosts. For this job there exist several techniques:

- Static configuration directly on the hosts
- Use of stateful DHCPv6
- Use of auto configuration (SLAAC) with router advertisements (RA)

In addition an IP address management system (IPAM) can be used to maintain the set of configured IP addresses centrally. The use of an IPAM system is recommended to reduce the overall configuration overhead. It also allows easier, central logging and reduces errors compared to manually typed IPv6 addresses.

2. Addressing Plan Considerations

There is no template to be used for all the IPv6 addressing plans, so each network will need some previous study of its topology, services, IPv4 addressing plan, internal procedures and organization, and foreseen changes to create a full custom IPv6 addressing plan.

There are different options to be taken into account when creating an addressing plan for a public administration, like in any other network, and a subset of them should be used.

When redesigning IP addressing schemes, you can allocate according to your need. Such logical addressing plans have the potential to simplify network operations and service offerings as well as network management and troubleshooting.

The addressing plan must take the following into consideration:

- Prefix aggregation: The large IPv6 addresses can lead to large routing tables unless network designers actively pursue aggregation.
- Network growth: It is important to design the address infrastructure to take network growth into account
- Use of Unique Local Addresses [RFC4193]: In IPv6 the address space is just for the one network and is globally unique. This private address space can be used to address devices and services that do not need to connect to the Internet.

IPv6 addresses use 128 bits to be codified [RFC4291], what provides a huge amount of prefixes, subnets and IPv6 addresses. We will not enter into details about IPv6 address notation and sub-netting but will cover some high level issues to be taken into account and that can be used in the IPv6 addressing plan.

Once the prefix to be used in our network is received from the RIR, in case we are a LIR, or from our ISP, we have to divide it to fit our needs and ease its management.

Different considerations can be taken into account about how to manage sub-netting in the network design. Below we present different approaches.

2.1 By-service sub-netting

We can divide the prefix we have received in smaller prefixes, assigning prefixes to each service being used within our network.

As an example, let's say we receive a /32 prefix (2001:db8::/32) and we identify ten different

services or uses for our addresses. In this example we need four bits ($2^4 = 16$) to create a field within our addressing architecture that can be used to identify the service to which the prefix is assigned.

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

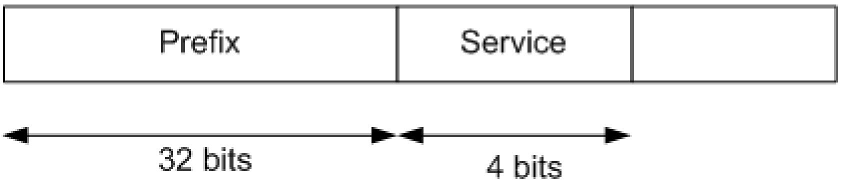

Figure 2-1: By-service sub-netting

The following table shows some detail of this example.

Prefix	Service / Use	Comments
2001:db8:0000::/36	Internal networks, p-to-p links between routers, loopbacks, etc.	
2001:db8:1000::/36	Management network	
2001:db8:2000::/36	Datacenters	
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	Ministries	
2001:db8:8000::/36	National Army and Defense Networks	
2001:db8:9000::/36	Police	
2001:db8:a000::/36	Emergency Services	
2001:db8:b000::/36	Free	Available for future growth
2001:db8:c000::/36	National Infrastructures	
2001:db8:d000::/36	Free	Available for future growth
2001:db8:e000::/36	Free	Free for future use or new services
2001:db8:f000::/36	Free	Free for future use or new services

Table 2-1: By-service sub-netting

Note that, as we have 16 possible prefixes and just 10 identified services, we have some free prefixes that can be reserved for future growth in the needs of some services.

2.2 Geographical sub-netting

We can divide the prefix we have received in smaller prefixes, assigning prefixes to each geographical region our network expands over.

As an example, let's say we receive a /32 prefix (2001:db8::/32) and we identify ten different regions for our addresses. In this example we need four bits ($2^4 = 16$) to create a field within our addressing architecture that can be used to identify the geographical region to which the prefix is assigned.

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

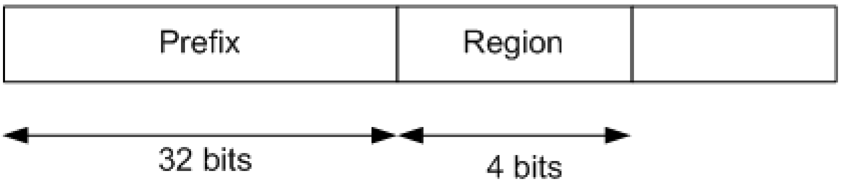

Figure 2-2: Geographical sub-netting

The following table shows some detail of this example.

Prefix	Use	Comments
2001:db8:0000::/36	Region 1	
2001:db8:1000::/36	Free	Available for future growth
2001:db8:2000::/36	Region 2	
2001:db8:3000::/36	Free	Available for future growth
2001:db8:4000::/36	Region 3	
2001:db8:5000::/36	Free	Available for future growth
2001:db8:6000::/36	Region 4	
2001:db8:7000::/36	Region 5	
2001:db8:8000::/36	Region 6	
2001:db8:9000::/36	Region 7	
2001:db8:a000::/36	Region 8	
2001:db8:b000::/36	Region9	
2001:db8:c000::/36	Region 10	
2001:db8:d000::/36	Free	Free for future use or new regions
2001:db8:e000::/36	Free	Free for future use or new regions
2001:db8:f000::/36	Free	Free for future use or new regions

Table 2-2: Geographical sub-netting

Note that, as we have 16 possible prefixes and just 10 indentified regions we have some free prefixes that can be reserved for future growth in the needs of some regions with potential of growth.

2.3 Combined approach

We can combine the two approaches seen above, identifying services and geographical regions and creating one field for each in our addressing architecture.

As an example, let's say we receive a /32 prefix (2001:db8::/32) and we identify ten different regions and ten services for our addresses. In this example we need four bits ($2^4 = 16$) to create a field for each one within our addressing architecture.

The following figures illustrates the address structure of the described example using the two options, using first a service field and then dividing by region, or using first a region field and then dividing each one in smaller prefixes for each service.

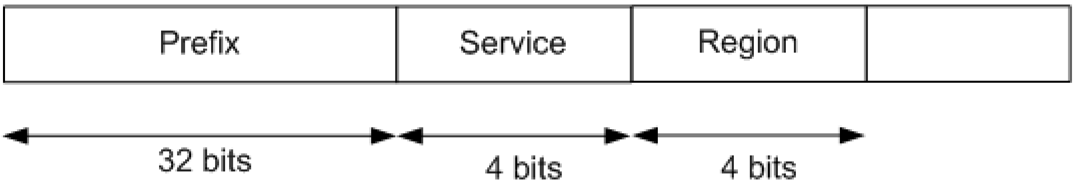

Figure 2-3: Combined approach: Service and Region

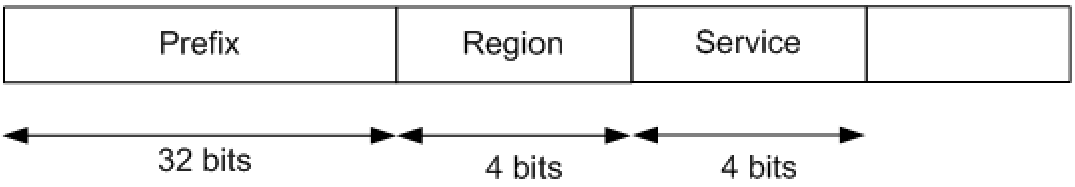

Figure 2-4: Combined approach: Region and Service

The following table shows some detail of this example, using the service and region Scheme.

Service	Region	Prefix	Comments
Internal networks, p-to-p links between routers, loopbacks, etc.	Region 1	2001:db8:0000::/40	
	Free	2001:db8:0100::/40	
	Region 2	2001:db8:0200::/40	
	...	...	
Management network	Region 1	2001:db8:1000::/40	
	Free	2001:db8:1100::/40	
	Region 2	2001:db8:1200::/40	
	...	...	
Datacenters	Region 1	2001:db8:2000::/40	
	Free	2001:db8:2100::/40	
	Region 2	2001:db8:2200::/40	
	...	...	
Educational networks	Region 1	2001:db8:3000::/40	
	Free	2001:db8:3100::/40	
	Region 2	2001:db8:3200::/40	
	...	...	
Free		2001:db8:4000::/36	Available for future growth
National Health System Networks	Region 1	2001:db8:5000::/40	
	Free	2001:db8:5100::/40	
	Region 2	2001:db8:5200::/40	
	...	...	
Free		2001:db8:6000::/36	Available for future growth
Ministries	Region 1	2001:db8:7000::/40	
	Free	2001:db8:7100::/40	
	Region 2	2001:db8:7200::/40	
	...	...	
National Army and	Region 1	2001:db8:8000::/40	

Defense Networks	Free	2001:db8:8100::/40	
	Region 2	2001:db8:8200::/40	
	...	...	
Police	Region 1	2001:db8:9000::/40	
	Free	2001:db8:9100::/40	
	Region 2	2001:db8:9200::/40	
	...	...	
Emergency Services	Region 1	2001:db8:a000::/40	
	Free	2001:db8:a100::/40	
	Region 2	2001:db8:a200::/40	
	...	...	
Free		2001:db8:b000::/36	Available for future growth
National Infrastructures	Region 1	2001:db8:c000::/40	
	Free	2001:db8:c100::/40	
	Region 2	2001:db8:c200::/40	
	...	...	
Free		2001:db8:d000::/36	Free for future use
Free		2001:db8:e000::/36	Free for future use
Free		2001:db8:f000::/36	Free for future use

Table 2-3: Combined approach: Service and Region

In the table above there are free prefixes available for some services growth or for new services, and within each service's prefix (/36) 4 bits have been used to assign /40 prefixes to different geographical regions, again with some free prefixes for future growth or new regions to be addressed.

3. EXAMPLES

3.1 German Example

3.1.1 National government addressing scheme[1]

The Cooperation Committee on Automated Data Processing for the Federation, the Länder (states) and the municipalities “KoopA” and the State Secretary Committee “Deutschland-Online” (Decision Nr. 04 - 09/2007 / November 2007) decided that, as central authority, the Federal Government should request IPv6 address space for the entire public administration sector in Germany. This decision was accompanied by the information and communications technologies strategy “Deutschland Digital 2015” for Germany adopted by the Federal Cabinet on 10 November 2010 [http://www.bmwi.de]. The strategy makes clear that, for Federation, Länder and municipalities, the introduction of IPv6 (Internet protocol version 6) represents an essential element in the implementation of new Internet technologies within the context of modern and secure communication infrastructures. The German public administration sector’s request for Internet address space was granted by the European address allocation authority (RIPE NCC, Réseaux IP Européens Network Coordination Centre) in December 2009. In consequence the Federal government, represented by the Federal Ministry of the Interior (BMI), received a /26 IPv6 address block and is therefore considered sufficient for the entire German public administration sector. A contract establishing this allocation has been signed by the BMI and the RIPE NCC. In this process the BMI undertakes the function of a Local Internet Registry (LIR) (referred to as “de.government”) which is comparable to that of an Internet Service Provider. An IPv6 working group consisting of members from different administration levels (“IPv6 AG”), the Deutschland-Online Netz e.V. association and the BMI has established plans for structuring the address space as well as for organising and implementing the address assignment. Technical recommendations concerning the introduction of IPv6 have also been prepared.

Addresses will be made available on the basis of an address framework which describes the allocation of address space to the public administration. The address space has been divided into 64 equally sized blocks (/32). It has been established for the long term and is believed to hold sufficient reserves for the foreseeable future. In a first step, address blocks were reserved for the Länder, the Federation and the indirectly subordinate administration. Therefore the addressing structure is indirectly geographic occurring on the regions of the federal states (Länder). Beneath the Länder the structure follows different rules, but most time the municipalities are also addressed on political regions which lean on geographic structures.

[1] Extracted from [IPv6RefManGermany]

The assignment and usage of addresses is based on a description of roles and processes. The de.government LIR has the overall responsibility for the address management at the highest level. End users will be assigned address blocks by means of so-called Sub-LIRs. At the present time, it is intended that the Länder, the DOI, the NdB (Network of the Federal Government) and the Federal Ministry of Defence should act as Sub-LIRs. Additional Sub-LIRs may be established according to a clearly defined process. The Sub-LIRs are responsible for administering the address space of the respective municipalities in general.

Other organs of public administration may be permitted to assume the role of a Sub-LIR where the address requirement can be adequately justified.

Matters of IPv6 technology in public administration infrastructures will be dealt with at a later time and will take into consideration the results of the current pilot projects.

De.government has requested so-called global unicast addresses (GUA) from the RIPE NCC which means that they can be routed globally. Globally routable addresses from the de.government address space must be used for any comprehensive communication. In cases where the IP traffic remains within one organisational unit and does not need to be routed between different organisations, unique local addresses (ULA) may also be used. ULAs are generally not routed by a provider.

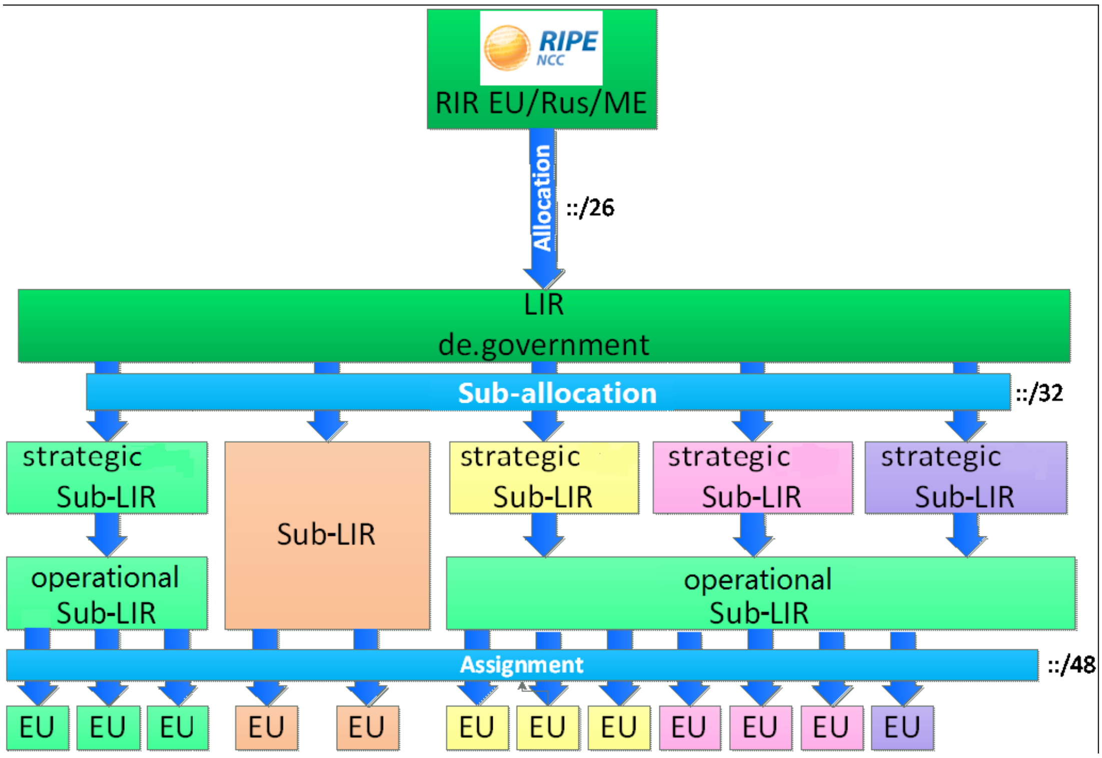

Figure 3-1: German Example - IPv6 address allocation/assignment hierarchy

De.government holds the allocation 2a02:1000::/26.

The following figure depicts an example of a German municipality which got a /48 IPv6 prefix. In this specific case, the prefix is composed of the 26 bits that RIPE NCC allocated for use by the German government as mentioned above, plus 22 bits chosen by the German LIR and the Sub-LIR which is responsible for the specific municipality. This results in a hierarchical IP address structure which is used for all public administrations in the IPv6 address space.

Since an IPv6 subnet with clients (servers, work stations, and other networked devices) in it always has a /64 prefix length, there are 16 bits left (green part in figure XXX) that can be used by the municipality for structuring the IPv6 address space of its own networks. The following text gives suggestions how these 16 bits (or 8 bits, in the case of a /56 bit long prefix) can be assigned and used systematically.

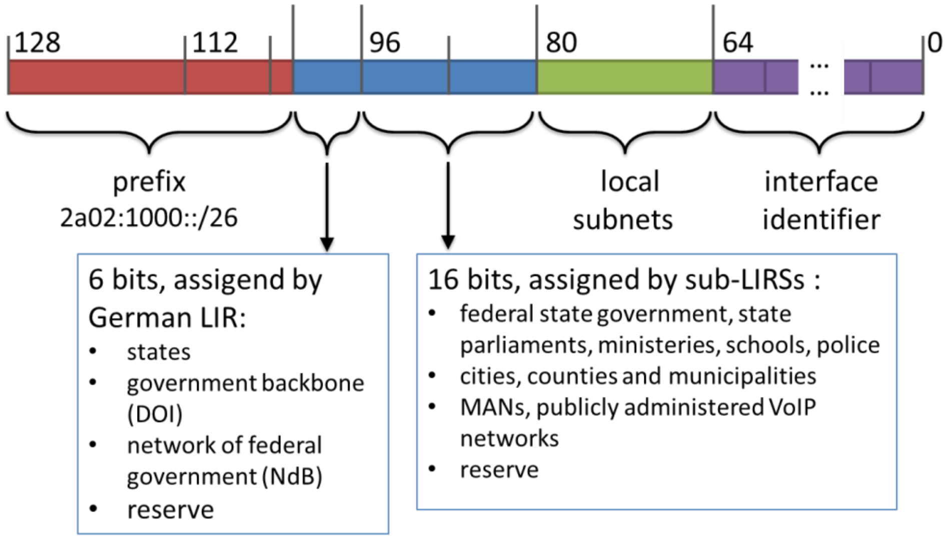
128
112
96
80
64
0
prefix
2a02:1000::/26
local subnets
interface identifier
6 bits, assigend by German LIR:
• states
• government backbone (DOI)
• network of federal government (NdB)
• reserve
16 bits, assigned by sub-LIRSs :
• federal state government, state parliaments, ministeries, schools, police
• cities, counties and municipalities
• MANs, publicly administered VoIP networks
• reserve

Figure 3-2: IPv6 address parts for a German public administration

This means that from the point of view of a public administration, the red and blue parts of their IPv6 addresses which make up the prefix of their site, are assigned by an external entity. The green part, they can choose themselves freely, and the purple part will be set individually for each interface of their systems' network devices, either manually or by some automated address configuration technique like DHCPv6 or SLAAC.

The following paragraphs take a look at the situation in public administrations of different size, from the small office to the data center. Due to different numbers of host systems and different IT requirements, the recommended ways of handling the IPv6 address scheme and the techniques used for address assignment differ. The basic way for structuring one's subnets with IPv6 addresses, however, remains.

The proposal for an IPv6 address framework, prepared by the IPv6 working group in February 2010, is set out below. The prefixes for the individual sub-allocations are shown in the table.

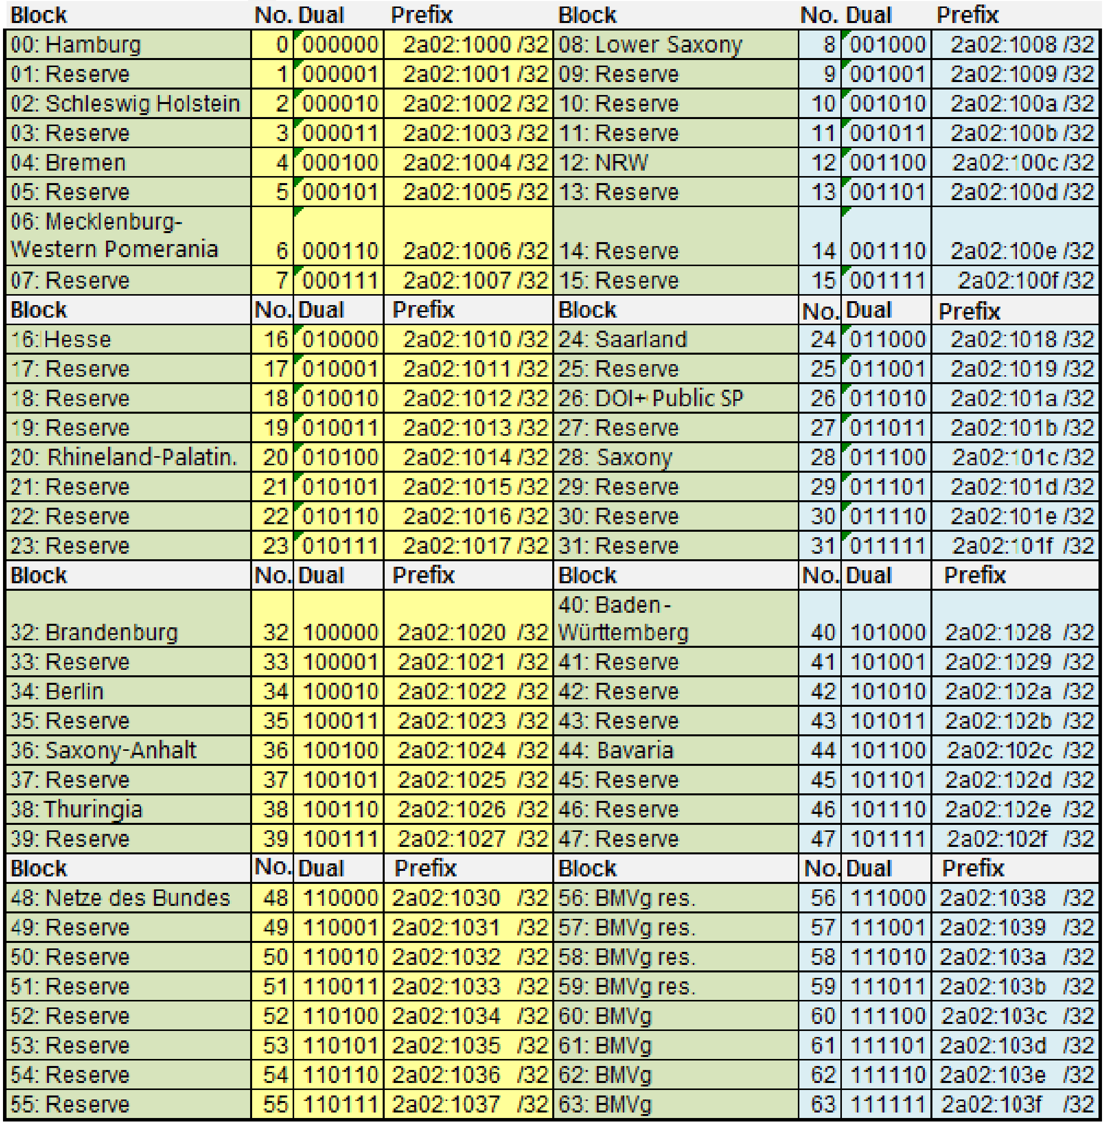

Block	No.	Dual	Prefix	Block	No.	Dual	Prefix
00: Hamburg	0	000000	2a02:1000 /32	08: Lower Saxony	8	001000	2a02:1008 /32
01: Reserve	1	000001	2a02:1001 /32	09: Reserve	9	001001	2a02:1009 /32
02: Schleswig Holstein	2	000010	2a02:1002 /32	10: Reserve	10	001010	2a02:100a /32
03: Reserve	3	000011	2a02:1003 /32	11: Reserve	11	001011	2a02:100b /32
04: Bremen	4	000100	2a02:1004 /32	12: NRW	12	001100	2a02:100c /32
05: Reserve	5	000101	2a02:1005 /32	13: Reserve	13	001101	2a02:100d /32
06: Mecklenburg-Western Pomerania	6	000110	2a02:1006 /32	14: Reserve	14	001110	2a02:100e /32
07: Reserve	7	000111	2a02:1007 /32	15: Reserve	15	001111	2a02:100f /32
Block	No.	Dual	Prefix	Block	No.	Dual	Prefix
16: Hesse	16	010000	2a02:1010 /32	24: Saarland	24	011000	2a02:1018 /32
17: Reserve	17	010001	2a02:1011 /32	25: Reserve	25	011001	2a02:1019 /32
18: Reserve	18	010010	2a02:1012 /32	26: DOI+ Public SP	26	011010	2a02:101a /32
19: Reserve	19	010011	2a02:1013 /32	27: Reserve	27	011011	2a02:101b /32
20: Rhineland-Palatin.	20	010100	2a02:1014 /32	28: Saxony	28	011100	2a02:101c /32
21: Reserve	21	010101	2a02:1015 /32	29: Reserve	29	011101	2a02:101d /32
22: Reserve	22	010110	2a02:1016 /32	30: Reserve	30	011110	2a02:101e /32
23: Reserve	23	010111	2a02:1017 /32	31: Reserve	31	011111	2a02:101f /32
Block	No.	Dual	Prefix	Block	No.	Dual	Prefix
32: Brandenburg	32	100000	2a02:1020 /32	40: Baden-Württemberg	40	101000	2a02:1028 /32
33: Reserve	33	100001	2a02:1021 /32	41: Reserve	41	101001	2a02:1029 /32
34: Berlin	34	100010	2a02:1022 /32	42: Reserve	42	101010	2a02:102a /32
35: Reserve	35	100011	2a02:1023 /32	43: Reserve	43	101011	2a02:102b /32
36: Saxony-Anhalt	36	100100	2a02:1024 /32	44: Bavaria	44	101100	2a02:102c /32
37: Reserve	37	100101	2a02:1025 /32	45: Reserve	45	101101	2a02:102d /32
38: Thuringia	38	100110	2a02:1026 /32	46: Reserve	46	101110	2a02:102e /32
39: Reserve	39	100111	2a02:1027 /32	47: Reserve	47	101111	2a02:102f /32
Block	No.	Dual	Prefix	Block	No.	Dual	Prefix
48: Netze des Bundes	48	110000	2a02:1030 /32	56: BMVg res.	56	111000	2a02:1038 /32
49: Reserve	49	110001	2a02:1031 /32	57: BMVg res.	57	111001	2a02:1039 /32
50: Reserve	50	110010	2a02:1032 /32	58: BMVg res.	58	111010	2a02:103a /32
51: Reserve	51	110011	2a02:1033 /32	59: BMVg res.	59	111011	2a02:103b /32
52: Reserve	52	110100	2a02:1034 /32	60: BMVg	60	111100	2a02:103c /32
53: Reserve	53	110101	2a02:1035 /32	61: BMVg	61	111101	2a02:103d /32
54: Reserve	54	110110	2a02:1036 /32	62: BMVg	62	111110	2a02:103e /32
55: Reserve	55	110111	2a02:1037 /32	63: BMVg	63	111111	2a02:103f /32

Table 3-1: German Example - IPv6 address sub-allocations

3.1.2 Small public administration

A small public administration in Germany will usually receive a /56 IPv6 prefix. A „small administration“ can be for example:

(a) A residents’ registration office in a municipality

(b) A geographically remote part (site) of a larger administration , e.g. one specific school in a district

(c) A geographically separate subsidiary of a larger administration, for example an office in a small village being part of a larger municipality

“Small” in this context refers to the number of distinct devices that potentially will get an IPv6 address assigned.

The following figure shows the parts of an IPv6 address in the case of a /56 prefix.

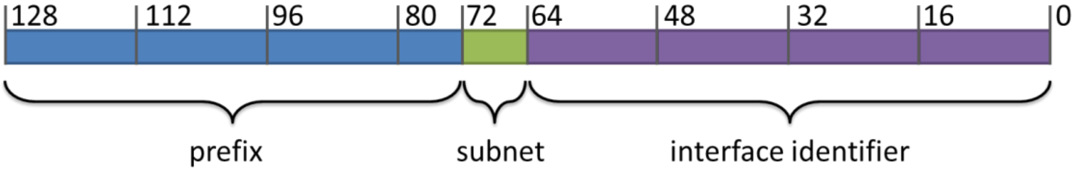

Figure 3-3: Parts of an IPv6 address when using a /56 prefix

Analogously to the previous figure, there are now 8 bits which can be used to structure and enumerate the IPv6 subnets of the administration. This means that “only” 256 subnets can be defined in total, but each still with space for 2^64 IPv6-enabled network devices (based on the 64 bit interface identifier).

If different network types with different use cases and access permissions exist, then it is advantageous to use the 4 upper (leftmost) bits from the “green part” of the address for structuring one’s address space into network types and use the 4 lower bits for subnet numbering. This allows for up to 16 network types (Intranet, DMZ …) with up to 16 subnets per type. If there are no different subnet types, for example if only workstation-type networks exist, then all 8 bits can be used for numbering them.

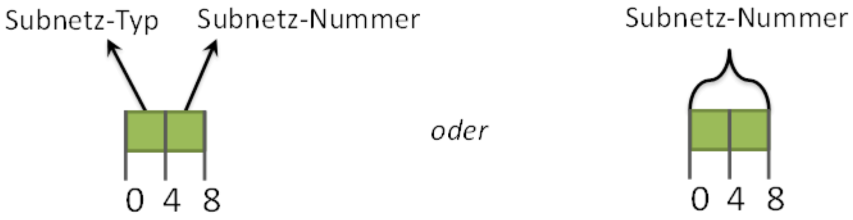

Figure 3-4: Options for splitting an 8 bit subnet mask into 4 bit parts

A split into 2 bits (for 4 different types) and 6 bits for numbering of (up to 64) networks is technically possible. However, it is not recommended because it complicates reading of literal IPv6 addresses and troubleshooting related to IPv6 addresses very much for a human administrator.

3.1.3 Medium size public administration

A medium size public administration will usually receive a /48 IPv6 prefix. Thus 16 bits remain for structuring and numbering all its IPv6-capable local networks, as shown in the following figure. With this setup, up to 65536 local IPv6 networks can be enumerated, each with room for 2^64 networked IPv6-capable interfaces.

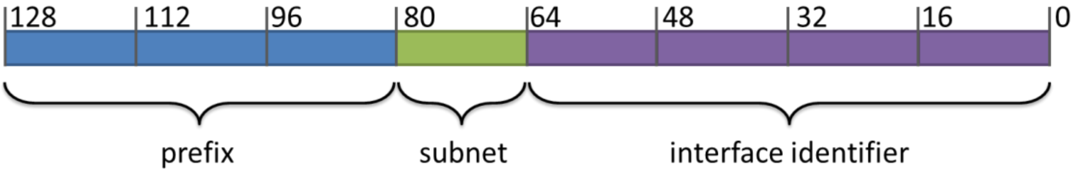

Figure 3-5: Parts of an IPv6 address when using a /48 prefix

Again, it is recommended to use the higher order (leftmost) bits of the "green part" of the IPv6 address as an identifier for the network type and to use the lower bits for enumerating all subnets of a given type. With 16 bits, a useful split is: 4 bits type (to distinguish up to 16 network types), 4 bits reserve (not yet used) and 16 bits for enumerating the subnets per network type.

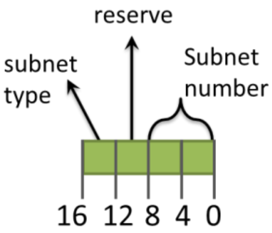

Figure 3-6: Recommended split of a 16 bit subnet mask into 4 bit parts

This address scheme is advantageous as it allows to specify network policies and filter rules per subnet type based on a simple 48 bits plus 4 bits = /52 network prefix (e.g. in routers' access control lists (ACL)).

If less than 16 network types are needed, then it is useful to leave „gaps“ in the type numbering. For example, in the case of only 4 network types, one should not use 0, 1, 2, 3 but 0, 4, 8, c instead. That way, new network types can later be added "next to" a similar type. This allows for prefix aggregation in the case of common ACLs, keeping router and firewall configurations neat.

If more than 16 network types are needed later on, then one can use the four reserve bits to allow a much higher number of network types (up to 256). Alternatively, the reserve bits can be used to enumerate 2^12 = 4096 subnets per type, instead of only 256 with 8 bits for the subnet number. Both alternatives are compared in the following figure.

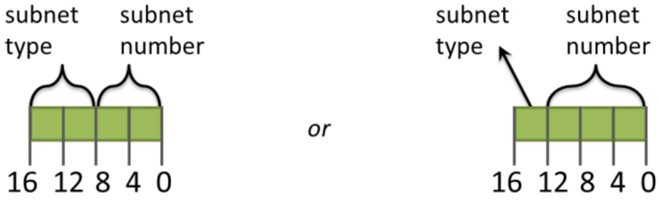

Figure 3-7: Options for splitting 16 bits for subnet identification

In the case of an IPv4/IPv6 dual-stack network, the bits of the IPv6 address containing the subnet number (the lower bits of the green part) can also be used to store the number of the existing **IPv4 subnet** of the end system. Using the same "well-known" network number in both, the literal IPv4 and the literal IPv6 addresses in use in one subnet, makes network maintenance easier for administrators (especially in the case of network troubleshooting).

Consider for example the case where a local IPv4 subnet 192.168.98.0/24 exists. Administrators will recognize that this is a subnet from the private IPv4 address range. Such a network could be remembered by them as the "number 98" private subnet. Combining the new IPv6 addresses with the subnet prefix <prefix>:0098::/64 in this subnet, one can easily detect the "98" again due to its prominent position as the last two nibbles of the first half (64 bits) of each IPv6 address.

Note that an IPv6 network with a three digit subnet number >99 like <prefix>:0124::/64 (with three digit number, >99, inside) would need to make use of the right scheme from the last figure, with 12 bits in use for the subnet number. Each literal digit in an IPv6 address represents one 4-bit nibble. Therefore, literal numbers larger than "99" in an IPv6 address ("99" here is hexadecimal, i.e. 0x99, to be precise) need three nibbles, i.e. 12 bits, to represent them (in contrast to 8 bits for the literal numbers 0 to 255 of an IPv4 address).

It is useful to represent well-known network prefix parts of existing IPv4 addresses also inside the new locally used IP6 addresses to support easy recognition and maintenance of these parts.

3.1.4 Very small office and home office

A different use case for an IPv6 address scheme is represented by (i) work places with a very small number of networked systems and (ii) the home office setup with IPv6. Very small offices exist for example in the form of a local citizens' advice office or the local (maybe even mobile) branch of a public library. Home offices are used by employees of the public administration working at home. In both cases, the number of networked systems if often smaller than five, and the employees need access to documents and services of their own administration plus sometimes to services of a remote public administration in order to fulfill their jobs.

Both cases have some similarities:

- Only a few local networked systems require an IP address.
- There are no services provided by these local devices – they all work stand-alone only or are used as clients for other, remote services.

- The criticality of the connectivity to the governing public administration is rather low, i.e. only the local part (e.g. home office) is negatively affected by an outage.
- The overall bandwidth requirements of an individual very small office or home office are comparatively low.

In such a scenario, it is recommended to access the Internet by using a publicly available Internet service in the same way as a private user would do. The connectivity can be provided by wired techniques such as DSL or cable modem or by wireless such as UMTS or LTE. The wired techniques should, however, be preferred due to better stability.

Ideally, the Internet service provides native IPv4 plus IPv6 access, i.e. dual-stack connectivity, without the use of tunneling techniques. In that case, the user's home gateway (customer premises equipment, CPE) must support (among others) ICMPv6 prefix delegation (ICMPv6-PD) to assign a IPv6 prefix to the local network and its devices.

Where IPv6 Internet access is not available natively, it can be added to the local environment by using a tunneling technique. Two ways are conceivable in this case: tunnel broker or VPN.

To use a tunnel broker means that the local CPE (e.g. home router) sets up an IP-in-IP tunnel, for example a 6in4 tunnel, with a trusted tunnel endpoint (broker). The tunnel broker used should be hosted and maintained directly by the governing local administration, or by a trusted data center run for the public administration. It is recommended NOT to rely on a public tunnel broker, no matter what technique it uses, because the performance may be poor and data integrity may not be protected sufficiently. An IPv6-capable tunnel broker allows all devices behind a connected CPE to make use of IPv4 and IPv6 connections independently from IPv6 support by the local Internet provider.

The second option is to make use of a virtual private network (VPN). In this case, usually an end system such as a client PC sets up the VPN tunnel across the local CPE to a VPN gateway that is hosted and maintained by its governing public administration. The disadvantage of this solution is that each end system in the home office or in the small, remote office needs to open a separate VPN tunnel towards its administration. The advantage is that a VPN tunnel can simultaneously provide IPv6 support, traffic encryption support and access rights management based on individual user authentication.

The local VPN endpoint can also reside on the local CPE, a setup which circumvents the disadvantage of the "VPN on the client PC" solution. However, this way, the end-to-end security is lower compared to an endpoint on the client PC, because every system connected to the CPEs local network can potentially send traffic across the tunnel. This setup is also called transparent LAN-to-LAN coupling using a VPN.

3.1.5 Large administration or data center

For a large data center, the assignment of a /48 IPv6 prefix will be the norm. If it consists of multiple sites, then, depending on size and demand, even more than one /48 prefix can be assigned. However, due to the vast size of the IPv6 address space, with room for 2^64 hosts in each single IPv6 subnet, even big data centers can practically be covered using only one /48 prefix.

Compared to the medium size public administration, the large administration and the data center will have a much higher demand on the number of IPv6 subnets. This is especially true when virtualized infrastructures, hosts, and servers, plus multi-tenancy-capable solutions are in use in the data center.

For the IPv6 address scheme of such an administration, however, the same “rules” as documented in the subchapter about the medium size public administration apply. But in this environment, one will usually have a need for a more detailed differentiation of network types, for example if there are many tenants and their ID should be incorporated in the IPv6 subnet number to distinguish them more easily. Therefore, it is of high importance, especially for the data center, to assess the local network structure and, based on this, to decide on the addressing scheme (and to decide, for example, how to use the “green reserve bits” shown in previous figures in the most economical way).

3.1.6 Practical example for a medium size public administration

This subchapter shows an exemplary address scheme for a public administration intending to add IPv6 to their existing IPv4 networks (only to some at the start), and having been assigned a /48 prefix for their local networks.

At first, an exhaustive documentation of existing IPv4 networks and their connectivity is required. This does not necessarily need to include all hosts and devices as well, but it must include subnet numbers, prefixes, some information about routing and network types (client, server, DMZ …) plus connectivity to external networks. It is also very valuable to record the responsibilities for each network, especially in the case where they are maintained by different parties.

Further to these topics, a public administration should clarify the following points:

- Which prefix length is appropriate for the public administration? (commonly, /56 and /48 will be available)
- How and when will the Internet service provider of the public administration make IPv6 for external communication to the public administration available?

- In the case of multiple locations: How will IPv6 connectivity be reached between all the locations of the public administration?
- *In the case of a provider-dependent prefix*: By which process can the administration obtain an IPv6 prefix from its Internet service provider?
- *In the case of a provider -independent (PI) prefix*: From whom (from which LIR or Sub-LIR) can the public administration receive its IPv6 prefix? Which process must be initiated there?
- *In the case of a provider-independent (PI) prefix*: Which process must be started with the ISP, so that it will route the global PI prefix through its network?

In this example, we assume that the networks listed in table XXX exist in the public administration. It is further assumed that a /48 bit prefix (provider-dependent or PI) has been obtained by the public administration. This means that 16 bits will be used to identify the IPv6 subnets. We also use the above-mentioned principle of “visually reusing” certain network numbers of the existing IPv4 networks in order to improve recognizability of the IPv6 network addresses and to simplify troubleshooting for the network administrators.

	IPv4 subnet address / prefix length	Subnet type	Remarks
1	10.0.0.0 / 8	private, internal IPv4 /8 subnet	Class A subnet, for work stations
2	192.168.22.0 / 24	private, internal IPv4 /24 subnet	Class C subnet, for internal servers
3	193.193.96.0 / 23	public IPv4 address range of the administration, used for IPv4 source NAT	NAT address space to hide outgoing connections from work stations behind
4	193.193.98.0 / 24	public IPv4 address range, used for hosting services	Demilitarized zone (DMZ) hosting the public services (e.g. publicly accessible web servers) of the administration
5	193.193.99.8 / 28	public IPv4 address range	DMZ hosting servers only visible for other public administrations

Table 3-2: Exemplary capture of existing IPv4 subnets

Based on this assessment, it is quite easy to classify the existing subnets into subnet types like

- Internal networks without Internet access
- Internal networks with Internet access
- DMZ networks with services accessible from the public Internet

- DMZ networks with services only accessible from other administrations

and so on.

Here, a subnet type will usually denote a class of subnets which should be treated equally by the administration's routers, gateways and firewalls / security devices.

Generally, there will be no IPv6 support for the N:1 source NAT one is used to from the current IPv4 world. The immense number of available addresses makes it obsolete.

IPv6 security for the hosts that are behind the existing NAT (with regards to IPv4) can be provided by filter rules that are similar to the existing ones, for example to prohibit incoming TCP connections. Application Layer gateways and proxies can, in principle, be used in the same way as with IPv4, i.e. these middle boxes may terminate TCP connections, thereby effectively protecting the identity and privacy of internal hosts and their addresses.

When using a /48 prefix, it is recommended to structure one's 16 bit subnet numbers into the following parts:

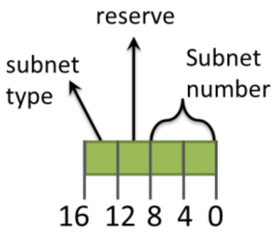

Figure 3-8: Partitioning of address bits for IPv6 subnets

By partitioning the 16 bits in the shown way, one can distinguish 2^4 = 16 different subnet types, with up to 256 subnets per type. If either of those numbers is not large enough then 1 to 4 bits from the 4 bit "reserve" part can be assigned to either the type or the subnet numbering. Whenever possible, splitting the "reserve" part between subnet type and subnet number should be avoided for reasons of literal IPv6 address interpretability.

Let's assume the following definition of the subnet types '0' to 'f' for our administration:

Subnet type (hexadecimal)	Meaning	Remarks
0	workstations network	
4	internal servers network	
8	public DMZ	e.g. for publicly accessible web servers
C	non-public DMZ	for services only available for other public administrations
Others		not yet assigned

Table 3-3: Exemplary definition of subnet types (4 bits of IPv6 address)

For the example networks defined in table 3-2, this means that the following IPv6 subnet IDs would be assigned:

	IPv4 subnet / prefix	IPv6 subnet / prefix	Remarks
1	**10**.0.0.0 / 8	<prefix>:**0010**:: / 64	the “10“ network
2	192.168.**22**.0 / 24	<prefix>:**4022**:: / 64	the “22“ network
3	193.193.96.0 / 23	---	Source NAT is not available with IPv6. Therefor use the work stations IPv6 addresses “as is” for outgoing connections or use an application level gateway for proxying them.
4	193.193.**98**.0 / 24	<prefix>:**8098**:: / 64	the “98“ DMZ network
5	193.193.**99**.8 / 28	<prefix>:**C099**:: / 64	the “99“ DMZ network

Table 3-4: Exemplary assignment of IPv6 subnet addresses for existing subnets

Additional note:

When mapping an IPv4 network number (here: 10, 22, 98, 99) to the IPv6 subnet prefixes for easier recognition, one has to remember that literal IPv4 addresses use a decimal notation, while IPv6 uses hexadecimal notation. This means that e.g. a “10” in an IPv4 address has different bits set than a “10” which is part of an IPv6 address:

IPv4: (dec) 10 = (binary) **0000 1010**; IPv6 (hex) 10 = (binary) **0001 0000**

IPv4: (dec) 99 = (binary) **0110 0011**; IPv6 (hex) 99 = (binary) **1001 1001**

In literal representation, which is also often the output of network analyzer tools, these numbers look the same, and the difference between the bits does not play a role. However, if IPv4 and IPv6 addresses are stored and analyzed in binary form then this must be considered.

This difference also raised the question how to represent numbers ≥ 100 within an IPv6 address (e.g. if the original IPv4 network is 192.168.**175**.0/24). In that case, one can either use 3 hex digits,(12 bits) to represent the “175” as part of the literal IPv6 address, or stay with two digits and use the numbers > (hex)99, i.e. the values from “a0” to “ff”.

While planning and using the presented IPv6 address scheme, the following additional advices should be considered:

The minimum size of an IPv6 subnet containing end systems is a /64 subnet, which has space for 2^{64} interface identifiers, square the number of end system addresses of the current Internet. Therefore, sizing a subnet just for the few systems inside, like e.g. using a small /28 IPv4 subnet, to save a few valuable addresses is not an issue anymore with IPv6.

IPv4 subnets which are semantically equivalent should preferentially be merged before a migration towards IPv4/IPv6 dual-stack subnets. This approach simplifies routing tables and access filter lists and is easier than merging subnets when they are already dual-stack enabled.

Without a NAT solution in place for IPv6, the globally unique addresses (GUA) of end systems are in general globally routable and, without an appropriate security setup, also globally reachable. Therefore, access and filter rules must be configured and activated so that a security level at least as high as with IPv4-only operation is realized.

The proposed IPv6 address scheme which uses 4 bits to distinguish an internal network type allows a very concise filter expression. To block for example the types 0, 1, 2, and 3 one would only need to specify one filter <prefix>:0/50 (aggregated from :0/52…:3/52); to filter only type 8 the expression would be <prefix>:8/52 (48 bits prefix plus 4 bits subnet type).

For IPv6, there are no real “private IP address“ ranges as for IPv4. For purely internal communication, one can use the “unique local” IPv6 addresses (XXX Reference). They are usually not routed towards and within the Internet, but filters should still be installed to prevent leaking of these internal IPv6 addresses into 3rd party networks or into the public.

In general, when creating a concrete IPv6 address scheme, it should be a goal to find a solution which closely couples the new IPv6 to the existing IPv4 address scheme, while still maintaining the network functionality and following existing directives.

3.1.7 IPv6 address management systems

The extended number of IPv6 addresses compared to IPv4 addresses and the longer literal address format is much harder to read and to remember for the human user. Although many users will use DNS in their daily work and thus not get in touch with IP address management, this makes the work of system administrators more error-prone. This is intensified by the fact that almost all hosts will have at least two IPv6 addresses, a link local address and a global unicast address (or, alternatively, a unique local address).

To ensure a safe operation of networks and hosts under these circumstances (including documentation and troubleshooting tasks), it is strongly recommended to use of a good IP address management (IPAM) tool for operating all but the small and very small public administrations. "Good" here means that certain requirements on stability, usability and integration with other network entities, such as DNS and DHCPv6 servers must be fulfilled to make such a tool really useful.

Using an IPAM system is usually accompanied by the use of a central DHCPv6 server (instead of manual, static configuration of addresses on hosts and even servers). This central DHCPv6 server can also be complemented by additional downstream DHCPv6 servers, depending on the size of the administration. Neither the client hosts' IP configuration nor the configuration of the DHCPv6 servers needs to be done manually, but the central IPAM system will hand this information automatically to the systems. Similarly, an administration's DNS server has to be configured by the IPAM system, not by manual editing of the server's A and AAAA records. The following figure shows an exemplary setup of an IPAM-enabled IT infrastructure.

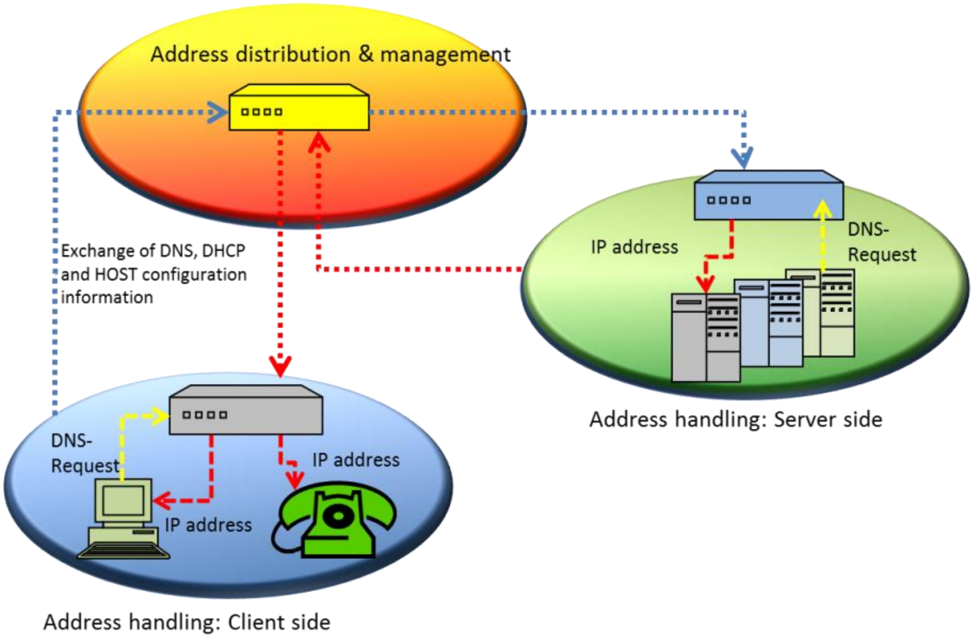
Address distribution & management
Exchange of DNS, DHCP and HOST configuration information
IP address
DNS-Request
Address handling: Server side
DNS-Request
IP address
IP address
Address handling: Client side

Figure 3-9: Generic structure of an IPAM solution

IPAM technology is sometimes also called "DDI management", standing for "DNS, DHCP, and IP address management". Before introducing IPAM technology to an administration, it must be ensured that the intended system can (a) interact with existing (or slightly adapted) DNS and DHCP servers, and (b) support IPv4 and IPv6 networks and their hosts. It should also be checked that the DHCPv6 server supports both, the distribution of dynamic addresses out of an address pool and the distribution of static addresses. In the latter case, a preconfigured IP address is commonly bound to an interface's MAC address. If needed in the concrete case, additional access techniques, for example 802.1X with certificate based authentication, have to be supported, too.

The mentioned approach that includes a central IPAM system, a central DHCP server and several downstream DHCP servers combines the advantages of a central management user interface with load balancing across multiple management servers. Address distribution via DHCP can additionally be used to only allow authorized mobile devices to connect to the administrations' network. The advantage of using DHCPv6 over stateless IPv6 autoconfiguration (SLAAC) is that the system always has control over valid IPv6 address allocation and that it can log which host (e.g. identified by its MAC address) had which IPv6 address assigned during which time interval. This recording is sometimes mandated in order to have reliable information in the case of misuse or misbehavior of systems.

If an IPAM system logs the used IP address over time per associated MAC address then the IP address can be resolved to a technical system. When it is needed to log IP address use per person, then a different solution e.g. based on 802.1X and user certificates for network access can be deployed.

3.1.8 Special considerations for network operators

The basic points for addressing above are mainly focussed for government organizations operating networks as a necessary infrastructure. But there exist also a number of governmental organizations that are positioned as information technology provider and – most time also –network provider for their government customer. For all these organizations operating wide networks special consideration on the network addressing must be taken.

In all wide area networks or other kind of structured routed networks a number of transfer networks exist. In IPv4 transfer networks always were addressed using a very small Subnet with only a few hosts. In IPv6 it is basically asked to use the maximum subnet length of /64, even if the subnet afterwards will only need two hosts addressed. For network management reasons the use of ULA for this purpose most time will be discarded (see discussion on former chapter).

Looking at the number of network gateways and the number of transfer networks still existing this yields to allocate a significant part of the usable address space for transfer networks only.

The IT provider Citkomm is operating a data centre and a wide area network with more than thousand satellite locations. Even in the backbone area of the data centre several transfer networks exist to create the necessary security infrastructure with several differentiated security zones. Citkomm therefore decided to reserve half of the allocated IPv6 addresses of its /48 only for network, which means in fact for transfer networks.

Other government IT provider in Germany with some more satellite locations even reported that they need to claim a /40 only for network addressing. So the addressing of the network components and the transfer networks must be given special attendance during address planning in IPv6, because – other than in IPv4 – they will allocate significant address space.

3.2 Greek Example

In this section, it is provided information regarding the addressing plans of the Greek public networks.

Regarding the governmental network (SYZEYXIS network) connecting all the public agencies, the following information and diagrams are proposals. SYZEXIS has not officially published its future IPv6 addressing plan.

Two approaches are under study:

- Aggregation using regions
  - Allocate a specific address space per region
  - Minimum number of entries into the routing table
- Aggregation using agencies
  - Allocate a specific address space per agency
  - Easier to apply security policies
  - Inject more entries into the routing tables

Following, two figures are presented based on the above-mentioned approaches.

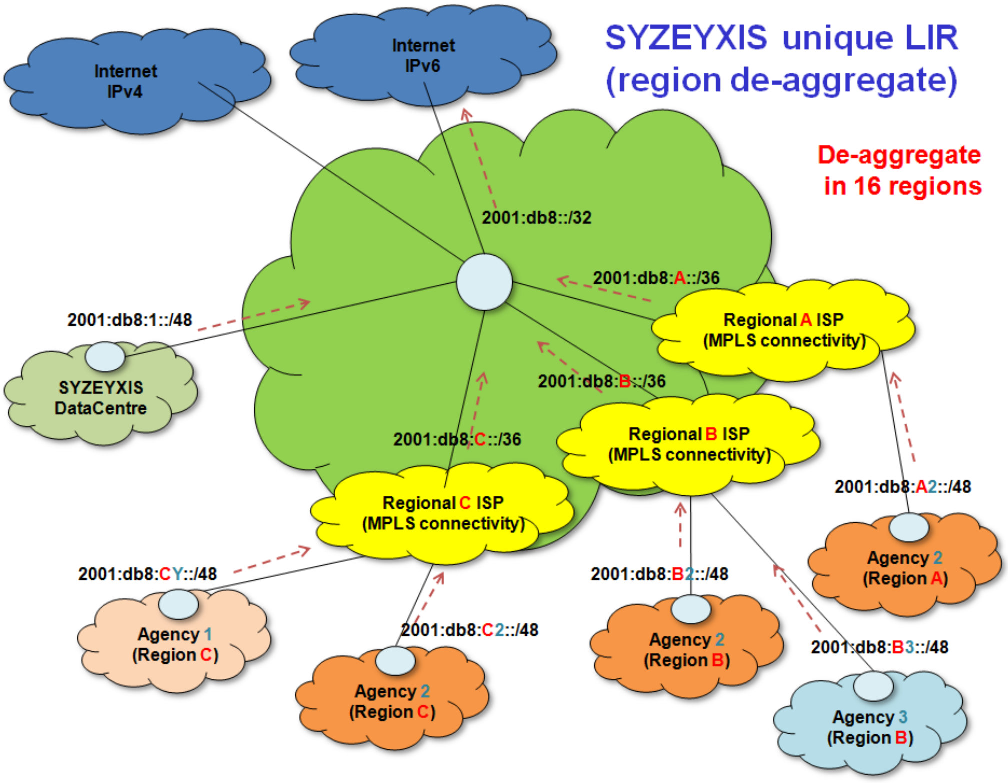

Figure 3-10: Greek Example - Region de-aggregate

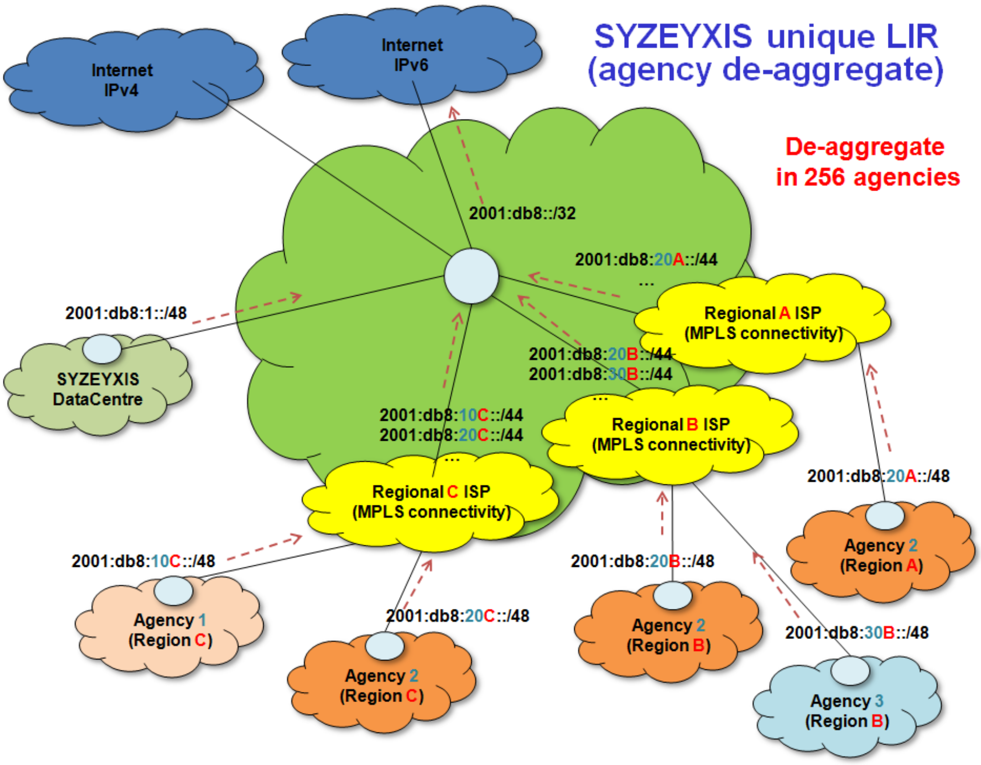

Figure 3-11: Greek Example - Agency de-aggregate

GRNET -the LIR for the research networking infrastructures in the country- is using 3x/48 address spaces for providing IPv6 addresses to the core network, the LANs and the access-links of the GRNET network respectively. For the core network and the LANs in GRNET a /64 address space is provided for each link. Regarding the GSN, since 2003 it has been delegated a /47 and one /48 IPv6 address spaces from GRNET. However, the current address space in GSN is not adequate to fulfill the future needs (e.g. Personal Area Networks into schools) of the school networks. So, since January 2013, GSN has been delegated a new /40 address space.

The previous IPv6 addressing schema has been provided by the GSN's NOC, as follows:

- 2001:db8::1300/47 assigned to access network
- 2001:db8::1302/48 assigned to backbone network

After the new delegation of /40 IPv6 address space from GRNET, the current IPv6 addressing schema will have been implemented as follows:

- 2001:db8:3400::/40 assigned to access network
- 2001:db8::1302/48 assigned to backbone network

The /48 ipv6 address space that has been assigned to backbone remains the same with previous allocation. The new /40 will be assigned to access network. The new length of new 40 will be leaded to bellow scenarios:

- The fact that each school provides for the use of a / 56 (ie 8 bits up to 64).
- The fact that the number of units in any case not expected to exceed 3.
- The table below are the subdivisions of space 2001: db8:3400 :: / 40 prefixes in 16/44, with intermediate lengths prefixes appear in the intermediate table cells. In that way, each prefix / 44 split other longer prefixes up to 16 x / 48.

2001:db8:3400::/40	2001:db8:3400::/41	2001:db8:3400::/42	2001:db8:3400::/43	2001:db8:3400::/44
				2001:db8:3410::/44
			2001:db8:3420::/43	2001:db8:3420::/44
				2001:db8:3430::/44
		2001:db8:3440::/42	2001:db8:3440::/43	2001:db8:3440::/44
				2001:db8:3450::/44
			2001:db8:3460::/43	2001:db8:3460::/44
				2001:db8:3470::/44
	2001:db8:3480::/41	2001:db8:3480::/42	2001:db8:3480::/43	2001:db8:3480::/44
				2001:db8:3490::/44
			2001:db8:34a0::/43	2001:db8:34a0::/44
				2001:db8:34b0::/44
		2001:db8:34c0::/42	2001:db8:34c0::/43	2001:db8:34c0::/44
				2001:db8:34d0::/44
			2001:db8:34e0::/43	2001:db8:34e0::/44
				2001:db8:34f0::/44

Table 3-5: Greek pilot - *Subdivisions of space 2001:db8:3400::/40*

According to the previous table, different address blocks are allocated to different school categories, e.g. primary schools, secondary schools, etc. Every school connected to the GSN is provided a /56 subnet to be used for its own internal LAN. As mentioned previously, the ipv6 space of GSN divided into pools prefixes. The smallest possible group of similar prefixes will be / 48. Each prefix / suffixes containing 28 48/56, ie 256 can group similar units.

Note that address block 2001:db8:1301::/48 remains unassigned for future use.

3.3 Spanish Example

3.3.1 IPv4 Addressing

Since each network tier is managed by entities of different organizational level, each network can have its own IPv4 addressing policy.

However, to assure the interconnection through SARA network, a Public Administration

Interconnection and Addressing Plan has been developed.

This plan defines a common private addressing space for Public Administration entities, allowing each entity to set up independently its own addressing plan, based on its network infrastructure or its internal organization, but at the same time maintaining a coordinated action to prevent the use of duplicated addresses.

According to the plan, the private addressing range 10.0.0.0 a 10.255.255.255 is divided into blocks (typically /16), and these blocks are assigned to the different entities that connect to SARA network, depending on their size and needs. Hence, the Ministries, the constitutional bodies and the Autonomous Communities have their own blocks, as well as other entities which, due to the nature of their work, need a separate access to the SARA network (such as the Tax Collection Agency or the Social Security).

3.3.2 IPv6 Addressing

Regarding IPv6 implementations, we have the experience of MINETUR (providing IPv6 access to some of the Ministry web portals) and MINHAP (providing IPv6 access to 060 Portal):

- In the case of MINETUR, it uses global addressing provided by Red-IRIS, the Spanish scientific and research network (2001:0720:0438::/64).
- In the case of MINHAP, it uses global addressing provided by its ISP.

However, the overall strategy for addressing is still under discussion.

As far as ULA / GUA addresses are concerned, though the final decision has not been made, it seems that GUA addresses will be preferred, since the use of ULA poses some problems when interconnecting to external networks (and SARA network is connected to Internet and sTESTA).

In the case of whether using PI or PA addresses, the option of PI addresses seems to have more support, since that option minimizes the impact of changing the ISP (something quite usual when the telecommunications contract is about to finish).

The main issue then is to determine the governance model of the network with GUA addresses, with three main alternatives on the table:

- A common addressing scheme for all the Public Administrations. In this scenario SARA network registers as a LIR and gets a /32 prefix from RIPE, splitting it up into /48 prefixes to be assigned to the different entities. This allows full aggregation in the network, but constrains the autonomy of the entities.
- A consolidated addressing scheme, in which each entity gets its own /48 prefix from its ISP, and report it to the rest of the network. It gives more autonomy to the entities, but,

being address management completely distributed, it makes it also more complicated.

- A mixed scenario, in which some entities get their prefixes from SARA network, and other entities get them directly from RIPE. Since some entities have already registered as LIR in RIPE in IPv4, obtaining an IPv6 prefix for them would be straight forward.

In the first alternative, different criteria for address assignment are being considered:

- Organizational criteria, similar to the current addressing plan in IPv4, so that each independent entity gets an independent IPv6 prefix.
- Functional criteria (voice, data, multimedia, etc.), so that network traffic of the same type can be easily managed.
- Territorial criteria, taking into account in which province the entity is located.

3.4 Turkish Example

As a result of “Design of National IPv6 Infrastructure and Transition to IPv6 Protocol” project, a circular which stating IPv6 transition plan for Turkey has been published on December 2010. This circular aims to lead to the governmental organizations while enabling IPv6 in their networks. However this circular does not include recommendations about the addressing, in other words there is no commonly defined IPv6 address management scheme in Turkey. In general, governmental institutions manage their IPv6 address blocks in parallel to their IPv4 address blocks. Also in this phase, the institutions may consult to the experienced institutions such as ULAKBİM.

For the Turkish pilot case, TURKSAT has allocated 2a00:1d58::/32 address block from RIPE NCC in order to use for both eGovernment Gateway (EGG) services and other services that TURKSAT provides. 2a00:1d58::/36 block has been reserved for EGG network and is being announced with AS47524 to the Internet. For the time being, network global unicast addresses (GUA) are being used. Although unique local addresses (ULA) are not required for the time being, throughout the pilot ULA may also be deployed where necessary.

TURKSAT IPv6 network, including EGG network, has been structured using by-service sub-netting method as follows:

- /36 for eGovernment Gateway (EGG) Datacentre (2a00:1d58:0::/36)
- /36 for VSAT (2a00:1d58:2000::/36)
- /36 for TURKSAT Local Services (2a00:1d58:1000::/36)
- /36 for Cable TV and Internet (2a00:1d58:8000::/36)

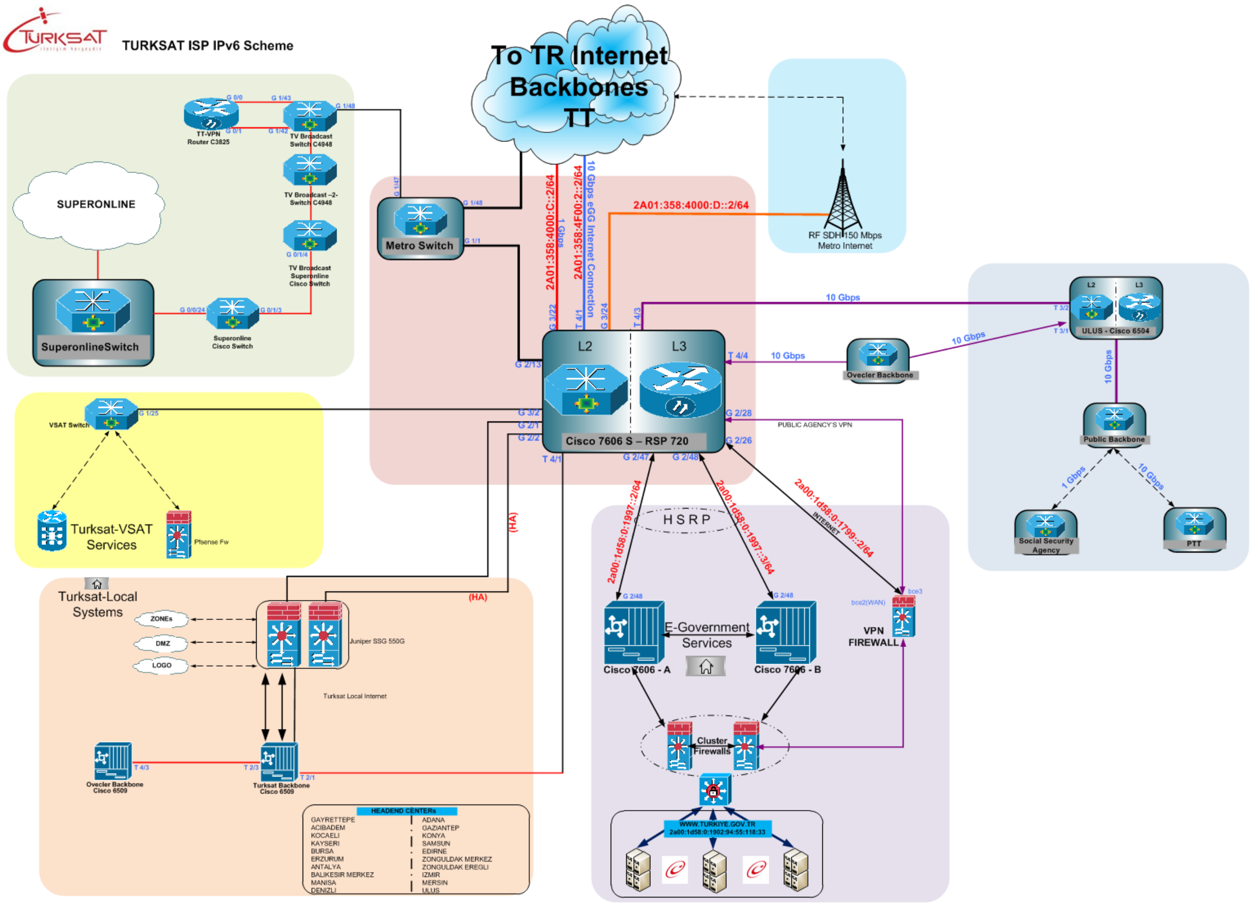

Figure 3-12: Turkish Example Network Structure

4. CONCLUSIONS

One of the first things any IPv6 deployment will have to deal with is the addressing plan for the considered network. In our case, in the scope of public administration networks, there is no exception, an addressing plan is needed.

There is no template to be used for all the IPv6 addressing plans, so each network will need some previous study of its topology, services, IPv4 addressing plan, internal procedures and organization, and foreseen changes to create a full custom IPv6 addressing plan.

There are different options to be taken into account when creating an addressing plan for a public administration, like in any other network, and a subset of them should be used.

5. REFERENCES

[RFC4007]	S. Deering, B. Haberman, T. Jinmei, E. Nordmark, B. Zill, "IPv6 Scoped Address Architecture", March 2005
[RFC4038]	M-K. Shin, Ed., Y-G. Hong, J. Hagino, P. Savola, E. M. Castro, "Application Aspects of IPv6 Transition", March 2005
[RFC4193]	R. Hinden, B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC4193, October 2005
[RFC4291]	R. Hinden, S. Deering, "IP Version 6 Addressing Architecture", February 2006
[RFC5952]	S. Kawamura, M. Kawashima, "A Recommendation for IPv6 Address Text Representation", August 2010
[IPv6RefManGermany]	IPv6 reference manual, Ministry of the Interior, Germany, January 2011