



Title:	Deliverable D2.1 IPv6 network topologies and addressing types	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 presents an evaluation of network topologies in public administrations, covering for example aspects such as address types used inside the network.

Keywords:

IPv6, Governments, Network Topology, IPv6 Address.

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 and small corrections	Antonio Skarmeta (UMU)
v0.5	28/01/2013	Added Turkish Pilot contribution	Alvaro Vives (Consulintel)
v0.6	30/01/2012	Updates on Greek Pilot contribution	Anastasios Zafeiropoulos (GRNET) & Konstantinos Koumoutsos (CTI)
v0.7	05/02/2013	Added contribution	Martin Krengel (Citkomm)
v0.8	06/02/2013	Document revision	Alvaro Vives (Consulintel)
v0.9	06/02/2013	Document revision	Carsten Schmoll (FHG)
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 presents an evaluation of network topologies in public administrations, covering for example aspects such as address types used inside the network.

Some generic scenarios are shown to illustrate further discussions, which will cover the different routing and addressing options, depending on the types of addresses used in the network and the way of getting connectivity to the IPv6 Internet.

Some specific examples are included at the end of this document, showing real deployments in public administration networks.

Table of Contents

1.	<i>Introduction.....</i>	7
2.	<i>Options and Considered Scenarios.....</i>	8
2.1	<i>Options.....</i>	8
2.1.1	Follow the IPv4 network design?	8
2.1.2	Mix IPv4 and IPv6 on the Same Link?	8
2.1.3	Separation of IPv4 and IPv6	9
2.1.4	Use links with Only Link-Local Addresses?	9
2.1.5	Use Link-Local Next-Hop in a Static Route?	10
2.1.6	Separate or combined eBGP Sessions	11
2.1.7	eBGP Endpoints: Global or Link-Local Addresses?	12
2.2	<i>Scenario 1.....</i>	12
2.3	<i>Scenario 2.....</i>	13
3.	<i>Routing and Addressing options.....</i>	15
3.1	<i>Option 1: Dependant scenario.....</i>	15
3.2	<i>Option 2: Independent scenario</i>	16
3.3	<i>Option 3: Mixed scenarios.....</i>	16
3.4	<i>Option A: Use of ULA.....</i>	19
3.5	<i>Option B: Use of GUA</i>	20
3.6	<i>Option C: Use of both GUA and ULA</i>	21
3.7	<i>Routing Considerations</i>	21
4.	<i>Examples</i>	25
4.1	<i>German Example</i>	25
4.2	<i>Greek Example</i>	29
4.3	<i>Spanish Example.....</i>	35
4.3.1	SARA Network	36
4.4	<i>Turkish Example</i>	38
5.	<i>Conclusions.....</i>	40
6.	<i>References.....</i>	41

Figure Index

Figure 2-1: Scenario 1 scheme: small public organization	13
Figure 2-2: Scenario 2 scheme: big public organization	14
Figure 3-1: Option 1 scheme: Dependant scenario	15
Figure 3-2: Option 2 scheme: Independent scenario.....	16
Figure 3-3: Option 3 scheme: own prefix and dependant routing.....	17
Figure 3-4: Option 3 scheme: own routing and another's prefix.....	18
Figure 3-5: Option 3 scheme: dual connectivity service	18
Figure 3-6: Internal and external addresses scheme	19
Figure 3-7: Internal and external address scheme for huge service networks	22
Figure 3-8: Internal and external routing	24
Figure 4-1: German Example - 1.....	28
Figure 4-2: German Example - 2.....	28
Figure 4-3: German Example - 3.....	29
Figure 4-4: Greek Example - 1	31
Figure 4-5: Greek Example - 2	32
Figure 4-6: Greek Example - 3	32
Figure 4-7: GRNET Network Topology.....	33
Figure 4-8: GSN architecture	34
Figure 4-9: Spanish Example - 1	36
Figure 4-10: Spanish Example - SARA Network Architecture 1	37
Figure 4-11: Spanish Example - SARA Network Architecture 2	37
Figure 4-12: Spanish Example - SARA Network Architecture 3	38
Figure 4-13: Turkish Example Network Structure	39

1. INTRODUCTION

Different options exist when deploying a public administration network. Specifically, from the IPv6 point of view, there exist different design options for routing and for address types.

This document shows recommended options, their characteristics and pros and cons. When we talk about implementing IPv6, two options exist, dual-stack or IPv6-only. The first one is the commonest case, being IPv6-only used for some specific new service or part of the network although it is expected that IPv6-only networks will become more common sooner than later. In any case, considerations showed in this document apply for both types of implementation unless it is specifically stated.

This document ends with some real examples of deployments made in different countries by public administrations that already have implemented IPv6 on their network.

2. OPTIONS AND CONSIDERED SCENARIOS

Available options or design choices are shown and discussed to clarify the pros and cons, allowing for a good decision to the network designer. Generic scenarios to be used to illustrate further discussions are also described.

2.1 Options

2.1.1 Follow the IPv4 network design?

Three options:

1. **Follow the IPv4 network design for the IPv6 network design:** use the same topology, network devices, routing protocols, monitoring tools, etc.
2. **Create an independent design for IPv6:** that mostly follows its own topology, using different routing protocols, network devices, and monitoring tools.
3. **Mixed:** following the existent IPv4 network design in some parts and a different one for IPv6 in others.

The recommended and most used is option 1, because it makes easier and cheaper the implementation and management of the IPv6 network. Option 2 is not usually considered, but for some parts or services it could make sense if the IPv4 support is not needed or is going to disappear.

The mixed scenario could be necessary because lack of IPv6 support in some network devices.

2.1.2 Mix IPv4 and IPv6 on the Same Link?

Two options:

1. **Mix IPv4 and IPv6 traffic on the same layer 2 connections:** only one layer three interface is needed with both IPv4 and IPv6 addresses.
2. **Separate IPv4 and IPv6 by using separate physical or logical links:** two layer 3 interfaces are needed, one for IPv4 addresses and one with IPv6 addresses.

There is a quite strong consensus in the operator community that option 1 is the preferred way to go because it has several advantages:

- Requires only half as many layer 3 interfaces as option 2, thus providing better scaling.
- May require fewer physical ports, thus saving money.
- Can make the QoS implementation much easier (for example, rate limiting the combined IPv4 and IPv6 traffic to or from a customer).

- Provides better support for the expected future of increasing IPv6 traffic and decreasing IPv4 traffic.
- Is generally conceptually simpler.

However, there can be situations where option 1 is the pragmatic choice, for example, to work around limitations in network equipment. One big example is the generally poor level of support today for individual statistics on IPv4 traffic vs. IPv6 traffic when option 1 is used. Other, device-specific, limitations exist as well. It is expected that these limitations will go away as support for IPv6 matures; making option 2 less and less attractive until the day that IPv4 is finally turned off.

2.1.3 Separation of IPv4 and IPv6

There is a general consensus around that IPv4 and IPv6 traffic should generally be mixed together. This recommendation is driven by the operational simplicity of mixing the traffic, plus the general observation that the service being offered to the end user is Internet connectivity and most users do not know or care about the differences between IPv4 and IPv6. Thus it is very desirable to mix IPv4 and IPv6 on the same link to the end user. On other links, separation is possible but more operationally complex, though it does occasionally allow the operator to work around limitations on network devices. The situation here is roughly comparable to IP and MPLS traffic: many networks mix the two traffic types on the same links without issues.

By contrast, there is more of an argument for carrying IPv6 routing information over IPv6 transport, while leaving IPv4 routing information on IPv4 transport. By doing this, one gets fate-sharing between the control and data plane for each IP protocol version.

2.1.4 Use links with Only Link-Local Addresses?

Two options:

1. **Use only link-local addresses** ("unnumbered")
2. **Have global or unique-local addresses** assigned in addition to link-locals

There are two advantages of unnumbered links:

- **Ease of configuration:** In a network with a large number of unnumbered links, the operator can just enable an IGP on each router, without going through the tedious process of assigning and tracking the addresses for each link.
- **Security:** Since link-local addresses are unroutable, the associated interfaces cannot be attacked from an off-link device. This implies less effort around maintaining security ACLs.

There are various disadvantages to unnumbered links in IPv6:

- Troubleshooting is more difficult: It is not possible to ping an interface that has only a link-local address from a device that is not directly attached to the link. Thus, to troubleshoot, one must typically log into a device that is directly attached to the device in question, and execute the ping from there.
- A traceroute passing over the unnumbered link will return the loopback or system address of the router, rather than the address of the interface itself.
- On some devices, by default the link-layer address of the interface is derived from the MAC address assigned to interface. When this is done, swapping out the interface hardware (e.g. interface card) will cause the link-layer address to change. In some cases (peering config, ACLs, etc) this may require additional changes. However, many devices allow the link-layer address of an interface to be explicitly configured, which avoids this issue.
- The practice of naming router interfaces using DNS names is difficult-to-impossible when using LLAs only.
- It is not possible to identify the interface or link (in a database, email, etc.) by just giving its address.

Today, most operators use numbered links (option 2) using global unicast addresses.

2.1.5 Use Link-Local Next-Hop in a Static Route?

Two options:

1. **Use the far-end's link-local address** as the next-hop address.
2. **Use the far-end's GUA/ULA address** as the next-hop address.

Recall that the IPv6 specs for OSPF [RFC5340] and ISIS [RFC5308] dictate that they always use link-locals for next-hop addresses. For static routes, [RFC4861] section 8 says:

A router MUST be able to determine the link-local address for each of its neighbouring routers in order to ensure that the target address in a Redirect message identifies the neighbour router by its link-local address. For static routing, this requirement implies that the next-hop router's address should be specified using the link-local address of the router.

This implies that using a GUA or ULA as the next hop will prevent a router from sending Redirect messages for packets that "hit" this static route. All this argues for using a link-local as the next-hop address in a static route.

However, there are two cases where using a link-local address as the next-hop clearly does not

work. One is when the static route is an indirect (or multi-hop) static route. The second is when the static route is redistributed into another routing protocol. In these cases, the above text from RFC 4861 notwithstanding, either a GUA or ULA must be used.

Furthermore, many network operators are concerned about the dependency of the default link-local address on an underlying MAC address, as described in the previous section.

There is also an specific issues related with using a link-local address as next-hop for a static route, the outgoing interface should also be specified. This happens because the same link-local prefix is used in all IP devices' interfaces (fe80::/64).

Today most operators use GUAs as next-hop addresses.

2.1.6 Separate or combined eBGP Sessions

For a dual-stack peering connection where eBGP is used as the routing protocol, there are two options:

1. Use **one BGP session to carry both IPv4 and IPv6 routes**.
2. Use **two BGP sessions, a session over IPv4 carrying IPv4 routes and a session over IPv6 carrying IPv6 routes**.

The main advantage of 1 is a reduction in the number of BGP sessions compared with 2.

However, there are three main concerns with option 1:

- On most existing implementations, adding or removing an address family to an established BGP session will cause the router to tear down and re-establish the session. Thus adding the IPv6 family to an existing session carrying just IPv4 routes will disrupt the session, and the eventual removal of IPv4 from the dual IPv4/IPv6 session will also disrupt the session.
- There is the question of which protocol to use to carry the dual IPv4/IPv6 session: over IPv4 or over IPv6? Carrying it over IPv4 makes sense initially from a stability and troubleshooting perspective, but will eventually seem out-of-date.
- Carrying (for example) IPv6 routes over IPv4 means that route information is transported over a different transport plane than the data packets themselves. If the IPv6 data plane was to fail, then IPv6 routes would still be exchanged, but any IPv6 traffic resulting from these routes would be dropped.

Given these disadvantages, option 2 is the better choice in most situations, and this is the choice selected in most networks today.

2.1.7 eBGP Endpoints: Global or Link-Local Addresses?

When running eBGP over IPv6, there are two options for the addresses to use at each end of the eBGP session:

1. Use link-local addresses for the eBGP session.
2. Use global addresses for the eBGP session.

Note that the choice here is the addresses to use for the eBGP sessions, and not whether the link itself has global (or unique-local) addresses. In particular, it is quite possible for the eBGP session to use link-local addresses even when the link has global addresses.

The big attraction for option 1 is security: an eBGP session using link-local addresses is impossible to attack from a device that is off-link. This provides very strong protection against TCP RST and similar attacks. Although there are other ways to get an equivalent level of security (e.g. GTSM [RFC5082], MD5 [RFC5925], or ACLs), these other ways require additional configuration which can be forgotten or potentially miss-configured.

However, there are a number of small disadvantages to using link-local addresses:

- Using link-local addresses only works for single-hop eBGP sessions; it does not work for multi-hop sessions.
- One must use "next-hop self" at both endpoints, otherwise redistributing routes learned via eBGP into iBGP will not work.
- Operators and their tools are used to referring to eBGP sessions by address only, something that is not possible with link-local addresses.
- If one is configuring parallel eBGP sessions for IPv4 and IPv6 routes, then using link-local addresses for the IPv6 session introduces an extra difference between the two sessions, which could otherwise be avoided.
- On some products, an eBGP session using a link-local address is more complex to configure than a session that use a global address.
- A strict interpretation of RFC 2545 can be seen as forbidding running eBGP between link-local addresses, as RFC 2545 requires the BGP next-hop field to contain at least a global address.

For these reasons, most operators today choose to have their eBGP sessions use global addresses.

2.2 Scenario 1

The first generic scenario we will describe is the smallest one, where the public administration

network is not very big nor expands over a big geographical area. This kind of networks is usually served by another bigger public organization, that in some cases are dedicated to provide the connectivity service.

Example of this scenario could be a University that has its own campus network, but connectivity is obtained through commercial ISPs or a NREN (National Research Network).

The following figure shows this scenario:




Figure 2-1: Scenario 1 scheme: small public organization

The network has its users that are supposed to connect to Internet, to other public organizations inside its country or in some cases in other countries, and to services published by them. The services published by the public organization could be divided in two types, for internal use only and also for public access.

2.3 Scenario 2

The second generic scenario is a network of a big public organization that expands over a big geographical area that could cover a whole country. This network could be used for the organization own needs or could be used to provide connectivity to other, usually smaller, organizations.

Example of this scenario could be a NREN (National Research Network) used to connect educational and research institutions all over a country, or a government network used to give connectivity to local institutions all over a country.

The following figure shows this scenario:




Figure 2-2: Scenario 2 scheme: big public organization

The network has a core network that interconnects all the parts of the network: the Internet connection infrastructure, some datacenter for internal and external use, and different networks where users and some services are located.

3. ROUTING AND ADDRESSING OPTIONS

We will consider the different choices regarding the addressing and routing, depending on who provides the connectivity to our public organization network, and related with that which IPv6 addresses are used within the network.

Regarding the connectivity to the IPv6 Internet there will be different options:

3.1 Option 1: Dependant scenario

The public administration depends on other network, usually only one, for their addressing and routing. The service could be provided by another public administration, a private company, or a mix.

This scenario is most common in the case of small public organization networks, that do not have strong requirements on addressing and routing.




Figure 3-1: Option 1 scheme: Dependant scenario

As could be seen in the figure, the ISP has one big prefix (P1) that is announced to the IPv6 Internet using BGP. The ISP assigns a sub-prefix (SP1) to the public organization network, that announces it to the ISP, or static routing is used.

If there is a change on the ISP, then the network of the public organization needs to be renumbered to the new sub-prefix of the new ISP. By other side, only one route (P1) is announced to the IPv6 Internet.

3.2 Option 2: Independent scenario

The public administration has its own ASN, prefix and routing towards the IPv6 Internet, usually through transit providers. This scenario is most common in case of big public administration network or networks that service other public administrations, like the NRENs. This scenario requires the public administration to become a LIR to get all the resources from its RIR (Regional Internet Registry).




Figure 3-2: Option 2 scheme: Independent scenario

As could be seen in the figure, the public organization has its own prefix (P1) that is announced using BGP to the IPv6 Internet. Inside its network different sub-prefixes (SPi) belonging to P1 are assigned to different parts of the network or to the serviced public organizations. An IGP (Internal Gateway Protocol) is used to announce internally all the subprefixes used, in order to be able to reach all parts of the network.

Only one route is announced to the IPv6 Internet, and there is no problem related with network renumbering. By other side, this scenario requires bigger and expensive equipment, more configurations and management and an internal know-how with dedicated IT staff.

3.3 Option 3: Mixed scenarios

There are other scenarios that have a mix of the characteristics of the two seen before. We will

show some of them to illustrate possible scenarios a public organization could find.

The public organization could have its own prefix but be dependant for the routing and announcement




Figure 3-3: Option 3 scheme: own prefix and dependant routing

The figure shows this scenario, where the public organization has its own prefix (P1) that announces to two ISPs (ISP1 and ISP2) that will announce the prefix to the IPv6 Internet. There are two cases where this scenario could happen:

- **Independent Scenario:** Similar to the one already seen before, where the public organization is a LIR with its own prefix and ASN, but uses two ISPs to propagate the announcement of its prefix to the IPv6 Internet.
- **Multihoming scenario:** One option that exists for a multihomed network, in our example using two ISPs, is to obtain a PI (Provider Independent) prefix that the ISPs should announce to the IPv6 Internet.

Another possible scenario occurs when the public organization has its own routing and announcing capabilities but use another's IPv6 prefix.




Figure 3-4: Option 3 scheme: own routing and another's prefix

In the figure we have a public organization network that obtains a sub-prefix (SP2) from one ISP (ISP2) that will announce the main prefix (P2) to the IPv6 Internet. The public organization has routing capabilities and even its own ASN (Autonomous System Number) and can announce its sub-prefix to the IPv6 Internet or for example, announce it in an Internet Exchange (IX) to optimize the routing.

The last option we will consider is the scenario where the public organization has commercial connectivity provided by an ISP for IPv6 Internet traffic, and internal traffic to other public organizations inside its country or even other countries' public organizations is provided by another public organization. An example could be educational institutions in Europe that have their commercial traffic through an ISP and connect themselves using the NREN in each country and using DANTE/GEANT through all Europe.




Figure 3-5: Option 3 scheme: dual connectivity service

In the figure the service network refers to a public organization network that offers connectivity to other public organizations. A sub-prefix is allocated from the service network (SP2). The public organization network, in the figure, has routing capabilities, its own ASN and prefix (P1), that is announced to the IPv6 Internet.

Another option could be to have the service network and an ISP, from which another sub-prefix is received (for example SP1).

In both cases, care should be taken on which addresses are used for which services and clients.

3.4 Option A: Use of ULA

Regarding internal connectivity, a kind of e-government IPv6 Intranet within one country, there could be also different options, because of the possibility of using ULA.

The following figure illustrates the scenario used for this discussion.




Figure 3-6: Internal and external addresses scheme

In the figure, the public organization network is connected to the IPv6 Internet using its own global unicast addresses (GUA) (P1) and can connect to internal services from other public organization networks using internal prefixes (P2 and P3).

Unique Local Addresses (ULAs) are defined in RFC 4193 [RFC4193] as provider-independent prefixes that can be used on isolated networks, internal networks, and VPNs. Although ULAs may be treated like global scope by applications, normally they should not be used on the publicly routable Internet.

The uniqueness is provided by using a random part of 40 bits of length to create a /48 ULA

prefix, what is considered enough to avoid collisions when merging networks using ULA addresses. This was one of the problems with the IPv4 private range [RFC1918], and ULA was designed to overcome this deficiency.

Being the ULA prefix (FC00::/7) well known, it is easy to be identified and easy to be filtered.

The ULA prefixes defined locally will have the eighth bit set to one, resulting in a local ULA prefix FD00::/8, followed by random 40 bits will give a local ULA prefix of 48 bits (/48).

So, in case of using ULAs:

- The biggest prefix that could be defined is a /48. If more address space is needed, for example for a big network with different PoPs, then several prefixes should be defined. If the pseudo-random requirement is followed, these ULA prefixes will not be contiguous and aggregatable, what makes address management a little bit more complicated.
- Devices with only one ULA address (and a link-local that is always present) will never be able to connect to the IPv6 Internet, or being accessed from other networks because ULA prefixes are not routed and probably will be filtered in the border of sites.

There are two possible ways to provide connectivity to a ULA addressed network, one is using a kind of NAT for IPv6 called Network Prefix Translation (NPTv6) [RFC6296] that provides a one-to-one translation. The other is the use of application-layer proxies. Both are not recommended because introduce more problems than solutions, problems that could be avoided using global addresses.

For our example scenario, we are considering to use GUA for public services and for users that want to connect to the IPv6 Internet and use ULA for services only used internally. This means that you need a ULA to connect to/from networks connected through the service network. In our figure, this means making P2 and P3, both a ULA prefix.

3.5 Option B: Use of GUA

The use of IPv6 Global Unicast Addresses (GUAs) allows all the network devices to access and be accessed from the IPv6 Internet. It must be clear that it allows but do not obligate, so filtering should be applied to protect networks / prefixes following a good security policy.

Following with the scenario seen in the figure showed in previous section, one option is to use GUA for the internal prefixes as well (P2 and P3). For example, the public organization in the figure could make P2 a sub-prefix of its own GUA prefix (P1), and announce it to the service network.

3.6 Option C: Use of both GUA and ULA

Another option is to use both ULAs and GUAs. Two options exist:

1. **IP devices with both GUA and ULA addresses in the same interface:** This could be an interesting scenario and will be described in more detail below.
2. **IP devices with only GUA addresses or with only ULA addresses on their interfaces:** This option has the same considerations seen above for the use of only ULAs and only GUAs, but for different parts of the network. The only difference is that internally to that network using both types of addresses, routing between them is allowed, resulting in the ULAs being reachable by GUAs with the appropriated routing configuration.

As described in [RFC4864], in practice, applications may treat ULAs like global-scope addresses, but address selection algorithms may need to distinguish between ULAs and ordinary global-scope unicast addresses to ensure bidirectional communications.

In our example scenario, we will have a GUA prefix (P1) and a ULA prefix (P2) announced to the service network. The users within our network will use both a GUA and an ULA address.

If the source address selection algorithm works properly, when a user within the public organization tries to connect to a GUA somewhere in the IPv6 Internet will use its GUA as source address. When tries to connect to some internal service reachable using a ULA, then will use its ULA address.

3.7 Routing Considerations

Using a network design as shown in figure 3-7 in large scale networks occurs some challenges for routing. If there are several organizations connected to the Service Network, as shown in the figure, it will be necessary to hold routing information for each connected partner. This means, that all Prefixes P 2 ... Px have to held in the routing table of each edge router of a public organization network.




Figure 3-7: Internal and external address scheme for huge service networks

In fact the public governments use separated secured networks for some kind of internal communication. Examples are DOI in Germany, SARA in Spain or on European level the sTESTA network. The IPv4 addressing in national networks most time bases on private addresses. sTESTA still uses a IPv4 public address space that is not announced to the internet. This way each connected site has to support an exception in routing to the global mask for the default route to the internet, showing the special route to sTESTA sites using the secured networks.

Nevertheless security shall be transferred to the applications itself then being offered (only) by using secured networks those secured networks are still a existing basic security feature for government communication now an also for several years ahead. On the other hand IPv6 is designed as and to end communication. As mentioned above features similar to NAT in IPv4 occur for IPv6 also, but are still in strong discussion and in no kind common use until know. So the impact of the decision ULA vs. GUA should be considered regarding that end to end communication in IPv6 is a necessary fact for a functional network.

The ULA address space covers a subnet of /7. The non-collision idea in ULA bases on random choice of the (small) networks for the users. For the addressing of the whole national government organizations of a nation random choice for ULA will not be a good opportunity, so the address space has to be assigned planned. Considering further that Germany received a /26 for their national governments and Spain is working on an subnet of /24 for the Spain governments it is getting clear, that such a concept will not work in European dimension, when end to end communication over secured service networks and within unique addresses also

must be enabled. In that case, for each state a subnet of the ULA address space has to be assigned and obligatory to be used to make communication possible over the secured networks. This is expected as being nearly impossible.

To get the national networks routable it is still necessary, that the used addresses can be proper aggregated at the edge of the network. It will be easy possible to route one (short) prefix to one national network. If the random system of ULA or a free addressing with GUA by each government - using its local provider or own PI address space per governmental organization - will be used the routing information has to cover every little network connected to the service network system. This will force every router at network edge points, which has to make the routing decision between internet connection and service network, to hold a huge number of routes and therefore scales to expensive carrier grade systems. Especially networks with a source near decision for the differentiation between routing to the internet vs. using closed service networks will not be able to cover these routing tables, because there are small and medium sized SOHO routers in use at presence. Therefore the whole network architecture has to be redesigned fundamental, disproportional equipment has to be installed - or routes simply have to be dropped. Since the last alternative will be the easiest and cheapest choice it can be assumed, that the idea of the seamless secured communication end to end over secured service networks in European government will get heavy damage in practical implementation as far as there is no addressing directive.

Therefore, Germany as well as Spain still decided to request a consolidated address space of GUA by RIPE NCC. This will make routing transparent to all users and especially the inter-nation communication much more easy and transparent by using highly aggregated prefixes for international routing.




Figure 3-8: Internal and external routing

At this moment the national claimed GUA are planned to be used for internal communication over (secured) service networks as well as for common internet communication. This brings up another mandatory aspect of routing design. Because using the same addresses over the service network as over the Internet the addresses of common internet services will be routed over the service network. So each service offered to the Internet must be made available - using the same address - over the service network as well. This affects the security design of some connected governments, because until now the security area providing services only to other governments over service networks is strictly separated from the security area for Internet services.

As an alternative solution servers only used for offering services to the Internet can use addresses out of the scope of the harmonized national addressing scheme, so that in any case routing is using the Internet connection. In those cases, the routing to the service networks over the aggregated subnets will not be affected.

For the further design of IPv6 implementation in national government networks in European member states and for the connection using the backbone network sTESTA itself these considerations have to be taken in account, as far as they operate with closed service networks.

4. EXAMPLES

4.1 German Example

The German government network structure is composed of different parts.

Deutschland Online Infrastruktur – DOI (German Online Infrastructure)

DOI is the network of the German national authorities. Like sTESTA connecting the EU administrations, it is a secure way to connect the different administrations levels (federal government, federal states (Länder), and municipalities). The legal basis for the collaboration between the Federation and the federal states in matters of information technology was established by the “Act on Connecting the IT Networks of the Federation and the federal states” (implementing Article 91c of the Basic Law) [IT-NetzG] and the “State Treaty on IT” which took effect on 1 April 2010.

DOI is still IPv6 ready since end of 2012. The roll out for all access points is running. The DOI is under control of the federal government. The network uses MPLS technology, and it is operated by T-Systems. An additional IPsec based cryptographic layer is under operation of the Federal Administration Office. User fees for every connected access point fund DOI. The fees scale with the used bandwidth.

Federal state networks

Some federal states operate networks to connect their sites but also for secured access to municipalities in their states. There are several models of operation (self-operated or outsourced), also of financing (central vs. based on access points).

North Rhine Westphalia does not provide its own federal state network. The federal state government provides all services that must be accessed by municipalities over DOI. Therefore, in North Rhine Westphalia each municipality has either a direct or an indirect connection to DOI. For this reason, Citkomm is connected to most other governments using its DOI access point.

Data centre networks

In Germany, there are several data centres that are part of the government. Many of them do not stick to one single government. There are in ownership by groups of municipalities. Usually for the group of the owners, but also for other governmental institutions, these data centres operate networks of their own. Most networks are based on rented infrastructure (lines, networks, parts of the equipment), operated and maintained by different providers. In rare cases own networks (based on own fibre or radio relay systems) are operated.

Citkomm itself operates a network based on MPLS and internet VPN that connects about 250 locations.

Citkomm networks as example for a municipality data centre

Within the Citkomm responsibility, several networks have to be considered when talking about IPv6 transition.

- Backbone
 - Core backbone network
 - Several networks connect application servers to the customers
 - Several infrastructure services (DNS, Proxy, AD, Management, ...)
- WAN
 - MPLS or Internet based VPNs
 - leased lines
 - Linux based router appliances (self-developed, system line called “iWAN”)
 - Cisco routers
- DMZ
 - Several DMZ networks, representing different levels of security and access providers
- Citkomm LAN
 - Flat layer 2 network
 - VPN users
- Customer LAN
 - Typically flat layer 2 network
 - Satellite locations possible
 - Mobile VPN users possible
- All networks mentioned above are under control of Citkomm, with the exception of the customer’s LAN. We intend to cooperate with a selected customer for the GEN6 project whose network, esp. servers and business applications, is closely managed by Citkomm.
- We intend to use an address space from the de.government assignment. Currently we can expect a /48 for Citkomm itself and another /48 for the involved customer
 - Use of de.government will give perspectives of
 - easy routing to other governments over the network aggregate

de.government

- using end to end connectivity features of IPv6 to other governments
- Use of multiple IPv6 addresses per client should be avoided to minimize possible occurring problems of correct address usage at client
- IPv6 Addresses from the de.government assignment will be used for public services (DMZ) to give positive proof of concept for other users, too.
- GUA will be used for end-to-end communication. This also reduces the need to introduce several generic proxy systems.
- The Federal Office for Information Security (BSI) recommends ULA for purely internal network communication.
- During the further roll out each customer will get a /48 from de.government
- Beneath the address concept for the subnets for each governmental location, there are no further addressing conventions inside an authorities network at this moment. Attention should be given to different projects (IPv6 profile (Federal Ministry of Interior/Fraunhofer FOKUS), ISI-LANv6 (Federal Office for Information Security), reference handbook of IPv6 working group (organized by Federal Ministry of Interior)), but all these projects are still under work and results not published yet. Citkomm can use the results in that sense that the drafts of all projects are available for the GEN6 project work.
- Therefore, the design and approval of a detailed addressing schema within a data centre and a customer network will be part and outcome of the GEN6 pilot.

The following figure shows German network structure for IPv4.



Figure 4-2: German Example - 2

Page 28 of 41

might be possible, that a network operator far away from Europe will only route aggregated routes into the RIPE region. Because there is no specific policy, network operators have to stick to in this case small networks / long prefixes may get connection problems.

For the German addressing scheme it is discussed to implement a backup mechanism for routing for long subnet prefixes in a way, that one provider will announce the whole de.government Subnet to the internet. To minimize network traffic to this provider all other owners of a /32 will be asked to publish their /32 also. Due to the nearly geographical structure of the /32 deployment this will result in a 'near to end point' routing on upper network levels, eg. the routing from regions far away from Germany.




Figure 4-3: German Example - 3

This concept secures routing to German network providers. To ensure further routing also for long prefixes to national governments access points it is mandatory, that the access provider for this government and all transfer providers to the announcement point of the aggregated prefix /32 support routing of long prefixes, as far as they regard to the national government address space maintained by de.government. The IPv6 working group in Germany just started a discussion with the national network operators to achieve a general commitment to support this routing concept. If this can be performed successful, the operators can show their commitment directly when applying as network access provider for government organizations.

4.2 Greek Example

SYZEFXIS-I is the Greek Public Administration Network and is operational since late 2005. It

offers the following main services to 4.500 public agencies:

- Broadband and secure IP connections so that they can offer their e-government services and have access to the Greek State Intranet and the web.
- Low cost voice services (telephony) (for free within the network)
- Video Services
- sTESTA pan-European connectivity

SYZEFXIS-II Network is designed as the successor of SYZEFXIS-I. The new project aims to:

- Cover all Greek Agencies (at least 34.000 actors)
- Provide broadband access
- Use of MANs (Metropolitan Area Networks) fibre infrastructures
- Provide upgraded/updated services
- Provide new value added services – emphasis on security - video – collaboration – mobility of users
- Maximization of “aggregation of demand” in telecom services for the Greek public sector

All 34.000 SYZEFXIS-II actors cover almost the entire spectrum of the Public Sector (except from Greek Army’s classified network and Ministry of Foreign Affairs’ NETVIS Network) and are categorized, depending on their supervising authority and the services offered to them by SYZEFXIS II.

The current public network is split into six regions, as shown in the following figure. Similar architecture will be followed in the future upgrades (in terms of footprint and services) for the public network infrastructure SYZEFXIS-II but the regions will be increased to nine.




Figure 4-4: Greek Example - 1

The services provided are:

- Internet connectivity
- VPNs
- Voice (SIP-based, mobile, etc.)
- Collaboration (videoconference, instant messaging, etc.)
- Security
- Data centre (hosting, IaaS, etc.)
- Wireless access

MPLS-based ISPs will provide interconnection services. All the regions will be connected in a L3 “exchange point”, in which the Internet upstream provided will also be connected.




Figure 4-5: Greek Example - 2

The logical network of each region is provided in the following diagram. The ISP should provide access based on xDSL, public MAN fibre infrastructures, or his own network infrastructure.




Figure 4-6: Greek Example - 3

A single network will interconnect all the public agencies. Several ISPs have planned to provide interconnection (as well as security, videoconferencing, voice, etc.) services based on a contract / SLAs.

Regarding the public academic and research institutions in Greece, they are interconnected through a completely separated network, called GRNET, while the public schools in Greece are interconnected through the Greek School Network (GSN).

GRNET constitutes an “intelligent” network that provides advanced internet services to the Greek Research & Education Community. Based on leading-edge technologies, the infrastructure is constantly upgraded to meet the growing demands of its users in terms of transmission, traffic and network capacity.

The GRNET network is a new generation optical fiber network based on Wavelength Division Multiplexing – WDM technology at high speeds (1-10 Gbps). The core network is formed by IP routers that are interconnected with PoS 2.5 Gbps circuits over 10Gbps wavelengths that are implemented via owned DWDM equipment. Since 2008, GRNET dark fiber network is extended all over Greece, with total length of dark fiber more than 9000km and optical equipment that may support speeds up to 21x10 Gbps per link.

The GRNET IP network topology including the established Layer 2 Ethernet links for the interconnection of GRNET clients is shown in the following Figure. The GRNET network can be divided into core and access network parts. The access network consists of dark fiber pairs between the point of presence (PoP) of GRNET in each major city in Greece and the PoP of the connected university or research institute. Around 100 clients are connected to the GRNET network. Thus, the GRNET network topology can be considered as a flat network topology without large aggregation points. Alternative backup paths are available for the majority of the network nodes while more than one alternative paths exist for the central network nodes.



Figure 4-7: GRNET Network Topology

The Greek School Network (GSN) is the educational intranet of the Ministry of Education, Life Long Learning and Religious Affairs that interconnects all schools and a large number of

educational administrative units and organizations and provides to them high quality electronic services. It is the biggest public network in the country, having the largest number of users, and has been recognized internationally as a remarkable educational network that promotes the introduction and exploitation of Information and Communication Technologies (ICT) in the Greek educational system.

The logical architecture of the GSN, operated by CTI, is shown in the following Figure.




Figure 4-8: GSN architecture

The figure depicts the six different technologies that are used in order to interconnect schools into the GSN and, thus, to the Internet. So, every school is connected to the Internet using one of the following technologies:

- ADSL links with access bandwidth at 2-24Mbps,
- Ethernet with access bandwidth at 1Gbps, through Metropolitan Area Networks of the public sector, available to numerous municipalities across Greece,
- Wireless link with access bandwidth at 11-54Mbps,
- Leased Lines with access bandwidth at 0,5-2Mbps,
- VDSL with access bandwidth at 4-12Mbps,
- ISDN/Dialup access with access bandwidth at 64-128Kbps.

The design model and the operational specifications of the GSN are based on the TCP/IP

protocol. The network's topology has a hierarchical structure and consists of four levels: the Backbone network, the Distribution network, the Access network and the Local area networks within the school laboratories.

GSN uses the GRNET network, with forty interconnection points and link capacities up to 10 Gbps, as its Backbone Network (BN). The choice of GRNET as the GSN's provider was a strategic choice of the Ministry of Education, absolutely compatible with the international practice.

The Distribution Network (DN) is the part of GSN that interconnects its points of presence with the backbone network. The Access Network (AN) is the part of GSN that interconnects educational and administrative units to their nearest node. In the case of wired connections, the common broadband technology used is ADSL, which supports a maximum capacity of 24/1Mbps (downstream/upstream), using the existing local loop infrastructure. VDSL is also lately becoming an attractive choice, especially for connecting large school units. Nowadays, there is a significant growth in the number of schools connected to GSN by using Ethernet technologies. This happened because GSN currently pursues to utilize as much as possible the optical Metropolitan Area Networks that were deployed by several municipalities across Greece in the previous years. Moreover, the usage of the optical fibers gives GSN the flexibility to select the desirable data technology transmission, thanks to the absence of a certain telecommunication provider interconnecting its nodes. Finally, wireless access technologies have also been used, mainly for the local access in the schools.

4.3 Spanish Example

The Spanish Government Network Structure is organized according to the territorial organization of the country, which is based on three levels of Public Administration:

- National: 13 Ministries, more than 130 Agencies.
- Regional: 17 Autonomous Communities plus 2 Autonomous Cities.
- Local: More than 8,000 municipalities.

This implies the existence of networks at different tiers:

- Municipalities' networks, which connect municipalities in the same geographic area (typically in the same province). These networks are managed by local entities that group municipalities in order to provide them with share services (e.g. consortiums)
- Regional networks, which connect municipalities' networks of the same region, and the entities of the regional Public Administration (e.g. NEREA network in Andalusia, Consorci AOC in Catalonia). These networks are managed by regional entities controlled by the Autonomous Communities.
- National network, Red SARA, which connects regional networks, the entities of the

national Public Administration (Ministries, Agencies, etc.), constitutional bodies (e.g. Parliament, Crown) and sTESTA network. SARA network is managed by MINHAP (Ministry of Treasury and Public Administration)

The previous is a simplified view, since there can be some exceptions, such as municipalities which connect directly to a regional network.




Figure 4-9: Spanish Example - 1

4.3.1 SARA Network


- Set of communications infrastructure and basic services (not just a telecommunications network)
 - E-signature validation, verification of identity and residence data, e-notification, etc.
- Allows the interconnection among the 3 levels of Spanish Public Administrations, facilitating the exchange of information and services
- Reliable, secure, capable and flexible
- Key tool to keep moving towards more ambitious e-government goals
 - All Regional Governments connected
 - 3.707 Local Councils
 - 90% population coverage
- Legally supported by 11/2007 Law



SARA Network - Architecture



- Overall view



1




Figure 4-10: Spanish Example - SARA Network Architecture 1



SARA Network - Architecture



- Backbone



8




Figure 4-11: Spanish Example - SARA Network Architecture 2



SARA Network - Architecture



• Backbone



9



Figure 4-12: Spanish Example - SARA Network Architecture 3

4.4 Turkish Example

Turkish pilot topology consists of TURKSAT network and participating governmental institutions (ULAKBİM, SGK and PTT). TURKSAT network provides a frontend for the governmental services. Once a citizen connects and starts a service, on the backend TURKSAT communicates with the related institution and brings the information citizen requests. Hence, TURKSAT topology and connection between TURKSAT and the participating institutions are in the scope of the pilot.

eGovernment Gateway (EGG) Web portal, which has been run by TURKSAT, has over 14 million registered users and includes over 550 services in connection with 70 governmental institutions.

TURKSAT has four main units regarding network operations namely:

- eGovernment Gateway (EGG)
- Satellite Operations (VSAT, TV and radio streaming, etc.)
- TURKSAT Local Network Operations
- Cable TV and Internet

The TURKSAT network structure is presented in detail as follows. The service provider Turk Telekom is abbreviated as TT in the figure. In addition, the figure shows the EGG frontend and

backend structures.




Figure 4-13: Turkish Example Network Structure

5. CONCLUSIONS

Different options are available for the design of an IPv6 network for a public administration. To date there is little or no practical experience with the introduction and operation of IPv6 in Europe. In fact, we are entering uncharted waters here with the public administration sector taking on a pioneering role in the deployment of IPv6 in Europe.

6. REFERENCES

[RFC1918]	Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear , "Address Allocation for Private Internets", RFC 1918 / BCP 0005, February 1996
[RFC4193]	R. Hinden, B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC4193, October 2005
[RFC4861]	T. Narten, E. Nordmark, W. Simpson, H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", RFC 4861, September 2007
[RFC4864]	G. Van de Velde, T. Hain, R. Droms, B. Carpenter, E. Klein, "Local Network Protection for IPv6", RFC4864, May 2007
[RFC5082]	V. Gill, J. Heasley, D. Meyer, P. Savola, Ed., C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC5082, October 2007
[RFC5308]	C. Hopps. "Routing IPv6 with IS-IS", RFC5308, October 2008
[RFC5340]	R. Coltun, D. Ferguson, J. Moy, A. Lindem, "OSPF for IPv6", RFC5340, July 2008
[RFC5925]	J. Touch, A. Mankin, R. Bonica, "The TCP Authentication Option", RFC5925, June 2010
[RFC6296]	M. Wasserman, F. Baker, "IPv6-to-IPv6 Network Prefix Translation", RFC6296, June 2011