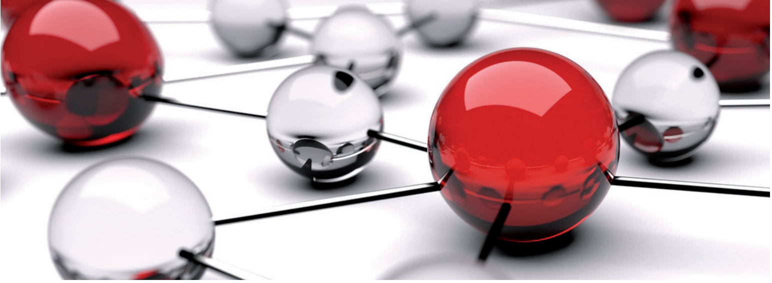

GOVERNMENTS ENABLED WITH IPv6

GEN6

# ADDRESSING AND TRANSITION

## FROM IPv4 TO IPv6 IN GOVERNMENT NETWORKS

This project has received
funding from the
European Union´s

European
Commission

# MANAGEMENT SUMMARY

Addresses are one of the most useful inventions of the past. They are used to locate somebody or something spatially or geographically, to identify a fixed or mobile device or to send messages to electronic mailboxes.

Most of these addresses have a textual representation being easy to remember or they may consist of numbers easy to forget – therefore the phonebook was invented. In case of electronic communication, no user needs to care much about the decimal or hexadecimal representation of addresses being used by technical systems to connect to each other or to send messages around. We are used to writing names such as www.ec.europa.eu to address a server instead. Yet, basically all communication between all stakeholders takes place based on IP addresses.

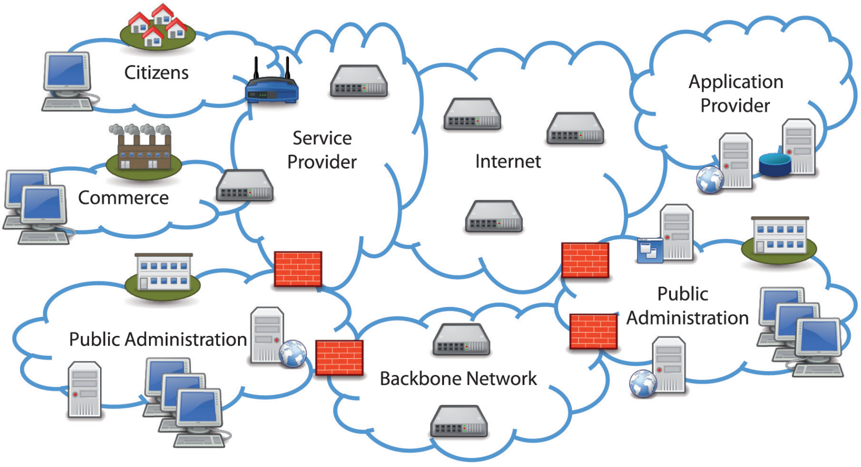

This booklet will give a short introduction to IP addressing and the structure of IPv6 addresses, on how you can get IPv6 address space and the options you have to structure and tailor the address space to your needs. The German address plan will be used as an example.

Furthermore, this document gives an overview of useful options for changing an IPv4-only network to a dual-stack network; for a public administration as well as for a data centre supporting a number of public administrations. For further detailed descriptions, have a look at the documentation published by the GEN6 project on the web page www.gen6-project.eu

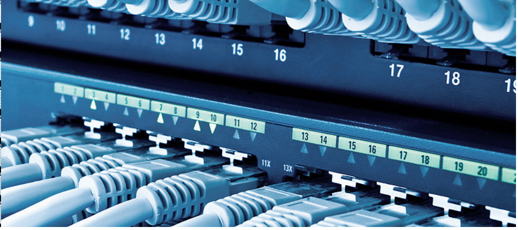

# ADDRESS SPACE AND ADDRESSING

The number of stakeholders involved in addressing is enormous, where communication is based on the Internet protocol (IP). The original idea behind the Internet protocol is to give every network node a globally unique IP address to allow end-to-end connectivity between any two nodes. Nowadays, this principle does not hold true anymore due to the global IPv4 address shortage and due to policy constraints that limit connectivity between nodes (mainly for security reasons).

Public servers and services must be reachable for all their clients. Therefore, they need to have a globally valid IP address. As the number of such Internet-connected servers (and connected devices in general) increases, the number of unused IPv4 addresses diminishes. This scarcity of IPv4 addresses is the main driving factor behind the standardization and rollout of the newer IPv6 protocol.

Each of the entities in the figure on page 2 uses a part of the overall network address space – for themselves or for their customers in the case of an Internet Service Provider. Depending on the number of IP addresses, an address block of different size is assigned to an organization. This organization can be an ISP or a commercial company, as well as the government of a country. In the latter case, a government can receive a rather large set of so-called provider-independent IP addresses for IPv6. IP address ranges that are subsections of this large set can then be distributed among the country's governmental institutions. This independent approach eases the management of governmental address space, because addresses never have to change, compared to the alternate approach. There, each governmental organization obtains their own small set of IP addresses from their local ISP. In the latter case, changing the ISP requires a renumbering of the whole network.

# IP ADDRESS STRUCTURE / FORMAT

While the current IP version 4 (IPv4) addresses are always 32 bits long, the newer IP version 6 (IPv6) addresses are always 128 bits in size. This allows for a vastly larger number of addressable end systems and paves the way to a much more structured addressing scheme.

IPv6 addresses are always made up out of a prefix part (for routing) and a host part (in IPv6, called "interface identifier"). For IPv4, the prefix part is usually 8, 16, or 24 bits long. For IPv6 the prefix can be at most 64 bits long, because the IPv6 interface identifier part is fixed at 64 bits. For IPv6 end systems, the prefix is always exactly 64 bits in size. This is shown in the following two figures:

The first figure shows the IPv6 address format used on end systems (e.g. a web server, a PC, a mobile phone, or a networked sensor) within an IPv6 subnet.

The second figure shows a /48 prefix, which is common for one networked site, e.g. one data centre:

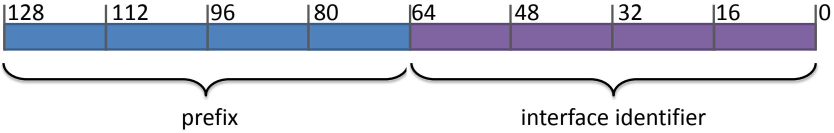

In this case, the /48 prefix would be announced via routing protocols to the outside world, to make the data centre reachable for traffic coming from the outside, e.g. from the Internet.

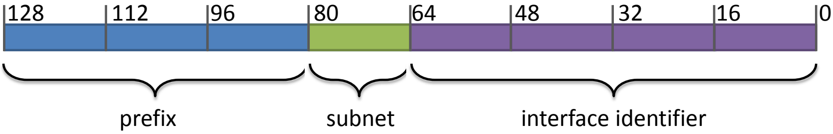

This leaves 16 bits for numbering IPv6 subnets inside the data centre. A later section in this booklet explains how an administrator could use these 16 bits to handle network management.

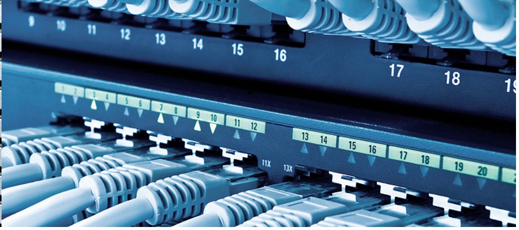

# HOW TO OBTAIN A BLOCK OF IPv6 ADDRESSES

The Internet Assigned Numbers Authority (IANA) manages the global IP address space on the top level. On request, it allocates IP (version 4 and version 6) address blocks in large chunks to the five major Regional Internet Registries (RIRs, see the following figure). The RIRs further subdivide these address blocks and distribute the smaller address ranges to Local Internet Registries (LIRs).

The example shown in the following figure has three LIRs per RIR, but many more do exist. Finally, these LIRs can be contacted by domestic entities in need of public IP addresses for their own, local private or public use. To obtain address space in Europe - please contact the RIPE NCC.

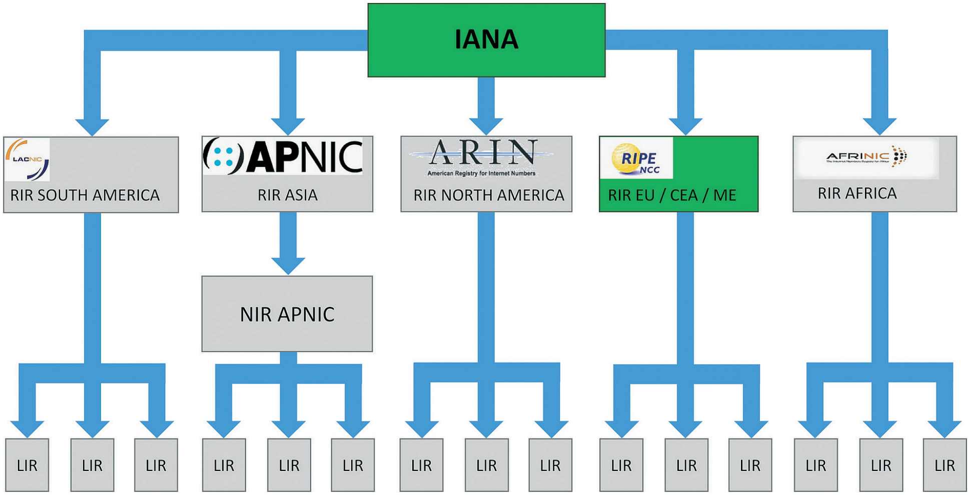

# IPv6 SUBNET ADDRESSING

So if you have an address space from your RIR or LIR – what can you do with those bits not being used for the prefix? Commonly a networked site running IPv6 will be assigned one /48 IPv6 prefix. This networked site can be, for example, a data centre, a small business office, or a government building. As the interface identifier within an IPv6 address occupies 64 bits, this leaves 16 bits for the site's subnet structuring – enough for numbering 65536 different IPv6 subnets! These 16 bits can be used for more things if they not only present a plain number (e.g. numbering subnets simply with consecutive numbers) but also convey some semantics, e.g. all workplace systems have „similar" numbers. Such a grouping of subnets of the same type allows for a clearer and freer specification of i.e. firewall filter rules and access control lists (ACLs) – given that the size of the groups is a power of two. For example, you can imagine the following split of the 16-bit subnet number (each green section represents 4 bits):

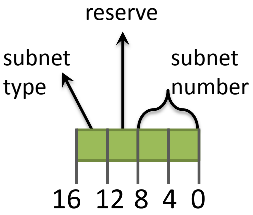

This leaves enough room for 16 different subnet types (0…9, a…f), for example 0 = internal workplace systems, 4 = internal servers network, 8 = DMZ, c = guest network subnets for externally visible servers.

Still, for each of these types, $2^8$ = 256 subnets could coexist at the site. Using this scheme, it is very easy to specify filter rules or ACLs, since one can aggregate all subnets of a given type into one wildcard expression (e.g. <48-bit-prefix>9/52 matches end systems in all subnets starting with '9'). In the above scheme, 4 bits are still unused, leaving more room for either more subnet types (256) or more subnets per type (4096), as shown as the two options in the following figure:

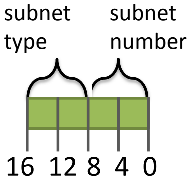

*or*

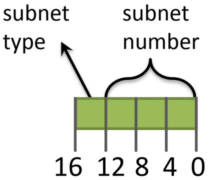

Even in the case that one's local network site would only receive a smaller/56 prefix, eight bits can be used for local subnet structuring. This allows e.g.16 network types with 16 subnets of each type or 256 subnets in total. The /56 prefix can often be found in the IPv6 access of residential customers.

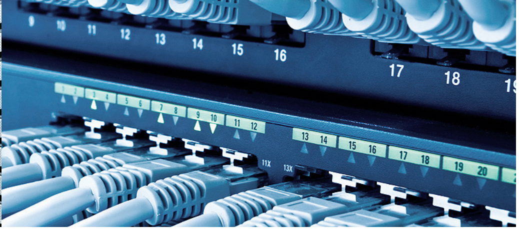

Each of the16 subnets corresponds to one digit of the hexadecimal notation that is widely used in IPv6 addressing.

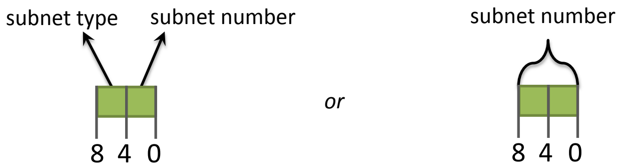

As mentioned previously above an organization may want to use the following subnet numbers as specified below:

| SUBNET | TYPE | REMARKS | |
|---|---|---|---|
| 0 | Workplace systems | 1,2,3 | reserved |
| 4 | Internal servers | 5,6,7 | reserved |
| 8 | DMZ subnets | 9,a,b | reserved |
| c | Guest networks | d,e,f | reserved |

This scheme is beneficial over the consecutive numbering of subnet types (0,1,2,3,…), since it allows an easier grouping and aggregating of types in case new subnet types have to be added later on (e.g. type 5 for a new set of internal server subnets). The figure on the right illustrates this:

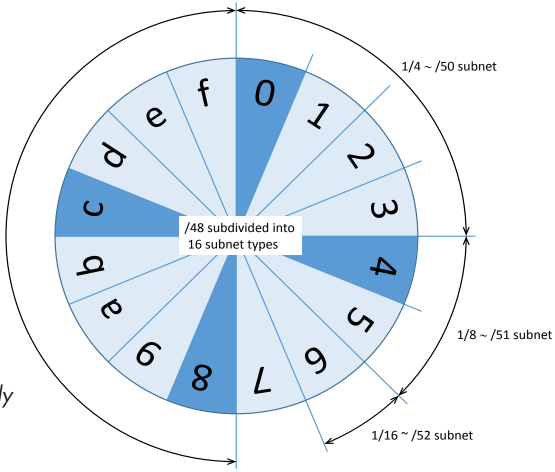

*This figure shows how consecutive subnet types can be aggregated easily in groups of 2, 4, or 8.*

The Dutch Research network Surfnet describes additional schemes for structuring addresses in detail in this document: http://www.ripe.net/lir-services/training/material/basic-ipv6-training-course/Basic-IPv6-Addressing-Plan-Howto.pdf

Always keep in mind that this IPv6 addressing scheme does not affect the number of potential end systems per subnet. In each of the above cases, each subnet may still contain up to $2^{64}$ hosts due to the separate 64-bit interface identifier.

# GOVERNMENTAL IP ADDRESS PLANNING

As mentioned previously, Local Internet Registries (LIRs) manage the domestic allocation of IP address ranges to Sub-LIRs and/or commercial users of these addresses. However, a country may want to use a separate address space for governmental use only. This allows for an overall simplified IP address structure across governmental institutions such as ministries, municipal services, schools, police etc. In effect, this approach allows for simpler, easier verifiable access permissions and routing structures. It facilitates the interoperability among the different levels of government and better supports initiatives aimed at infrastructure consolidation and shared services provisioning. This has important benefits, compared to the usual current situation where each governmental entity organises an IP address space for their local use, in a way dependant from the domestic Internet providers and independent from other governmental entities.

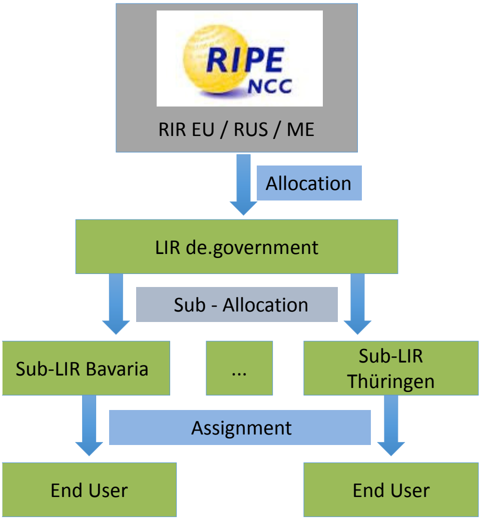

For a central management of domestically used governmental IP addresses, an institution of a country needs to become a LIR, registered with the RIPE NCC (see figure below). This way has been chosen in Germany, and it is currently under discussion in Spain. Germany set up its central LIR called „de.government" in 2009.

The RIPE NCC, the European RIR, will grant an IPv6 prefix up to a /29 without the need for an extensive justification for every initialized LIR. Only for a larger address space, i.e. a shorter prefix, an extended justification is mandatory.

Beneath the LIR de.government a set of Sub-LIRs was founded to organize the IPv6 address deployment in Germany. Upon extensive requests, the RIPE NCC allocated one /26 prefix for use by de.government. Based on this prefix, the LIR takes care of the (top level) management of the IPv6 addresses for the public administrations in Germany. A domestic address plan determines the use of the next six bits, after the /26 prefix. This way, one or more /32 prefixes are allocated to sub-LIRs as the basis for /48 site prefixes they hand out on request to their customers.

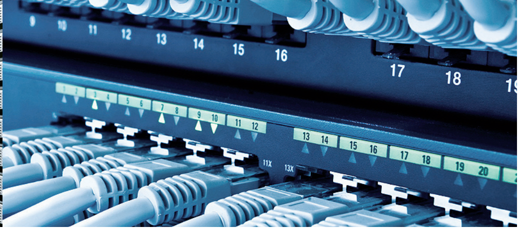

This split is shown in the next figure:

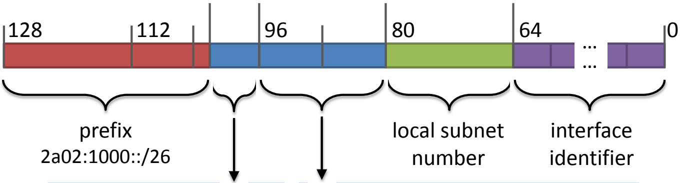

6 bits assigned by German LIR to the Sub-LIRs:

- Federal states
- Public services‘ networks
- Governmental networks
- Defence networks
- Reserve

16 bits, assigned by the Sub-LIRs to government agencies and public administrations:

- German federal government, state chancellery, state parliament, audit court, ministeries, schools, police
- Cities and municipalities
- WAN networks, gov. VoIP networks
- Reserve

With such an IPv6 address assignment, 6 + 16 = 22 bits are determined by the LIR and Sub-LIR (blue parts in the figure above), and each customer can assign another 16 bits (the green part) to structure their local IPv6 address space into local IPv6 subnets, as we described before. The following picture shows partially the assign-ment of addresses from the LIR de.government to the Sub-LIRs of some states in Germany. As you will recognize there is one /32 as reserve between each state, so another block can be assigned. This ensures a contiguous address space, which can be aggregated for routing and keeps rules for packet filters small and simple.

| Präfix: 2a02:1000 /26 | | | Regional Bits: | | 6 | | |
|---|---|---|---|---|---|---|---|
| **Block** | **Nr.** | **Dual** | **Präfix** | **Block** | **Nr.** | **Dual** | **Präfix** |
| 00: Hamburg | 0 | 0 | 2a02:1000 /32 | 08: Niedersachsen | 8 | 1000 | 2a02:1008 /32 |
| 01: Reserve | 1 | 1 | 2a02:1001/32 | 09: Reserve | 9 | 1001 | 2a02:1009 /32 |
| 02: Schleswig-Holstein | 2 | 10 | 2a02:1002 /32 | 10: Reserve | 10 | 1010 | 2a02:100a /32 |
| 03: Reserve | 3 | 11 | 2a02:1003 /32 | 11: Reserve | 11 | 1011 | 2a02:100b /32 |
| 04: Bremen | 4 | 100 | 2a02:1004 /32 | 12: NRW Land | 12 | 1100 | 2a02:100c /32 |
| 05: Reserve | 5 | 101 | 2a02:1005 /32 | 13: Reserve | 13 | 1101 | 2a02:100d /32 |
| 06: Mecklenburg-Vorpommern | 6 | 110 | 2a02:1006 /32 | 14: NRW Kommunen | 14 | 1110 | 2a02:100e /32 |
| 07: Reserve | 7 | 111 | 2a02:1007 /32 | 15: Reserve | 15 | 1111 | 2a02:100f /32 |

## RECOMMENDATION

Organizing IPv6 address in the public sector is mainly influenced by two criteria:

1. the number of address you need and
2. the degree of freedom you want to have

Let us see why. If you only need a small set of IP subnets and addresses and neither have to or want to manage a LIR on your own then you can ask your preferred provider and you are fine. He will give you addresses and do the announcements for routing. This may have some drawbacks for you, among them, that you need to renumber your IP address space upon changing your provider) – so have a look at option two. If you need a large address space (like the German government or a large institution), you can ask in Europe the RIPE NCC for addresses but you have to set up a LIR and possibly Sub-LIRs. Therefore, you get ‚provider-independent' addresses. What does this mean? You have to look for a provider to announce the address range for you. The advantages you get:

- A continuous IP address range you can be organized, as you desire.
- You can change your provider without losing addresses or to pay for them if you want to keep them and do not want to renumber everything.
- Routing will be much simpler if you do not have to cope with non-contiguous ranges and different providers.
- Dedicated networks can be set up with less effort.

Define your requirements and have an address plan in mind before you start your IPv6 transition!

## NEXT STEPS

Having addresses and an idea on how to structure subnets and devices, there is one question left: How do I deploy these addresses to an existing IPv4 network? Moreover, how can I do so with minimal hassle and operational interruptions? To find an answer to these questions you should know where you want to start. Making a service only visible from the outside can be handled different from changing the complete infrastructure of a data centre. A number of techniques exists for a transition process. Their principles and the advantages / disadvantages are presented in the remainder of this booklet.

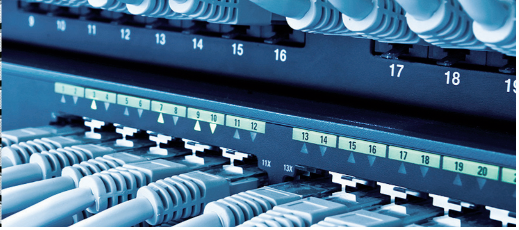

# TRANSITION

## IT-Situation Of The Public Authorities

The pressure increases to make governmental services available for the citizens via public networks such as the Internet, as more and more of these services become available in digital form. As part of this process, the public authorities (PA) governing the service are becoming IT service providers themselves. In the case of local government, PAs often implement IT services in collaboration with a municipal data centre (which aggregates services for several municipalities) for reasons of improved management and efficiency. Reliable, trustworthy networks are mandatory for such cooperative work. In many cases, multiple PAs must work together as shown in the following figure:

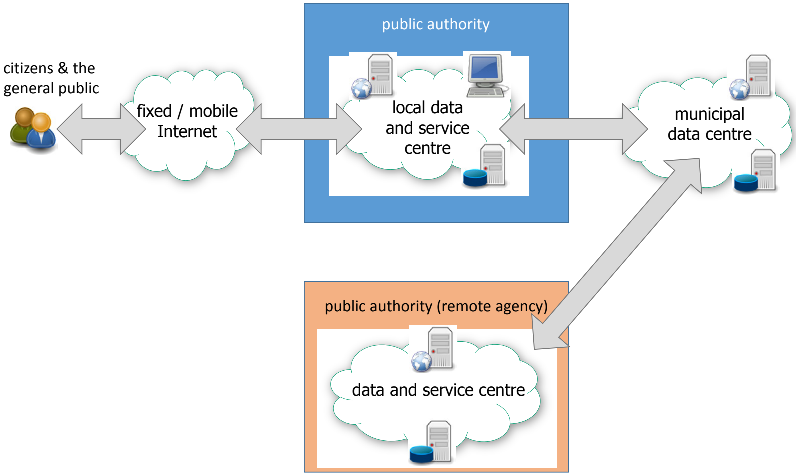
public authority
citizens & the general public
fixed / mobile Internet
local data and service centre
municipal data centre
public authority (remote agency)
data and service centre

As more and more citizens obtain IPv6-enabled Internet access, the pressure on the PAs increases to provide their e-government services via IPv6, too. Sometimes a citizen may even need to access e-government services from within an IPv6-only network. This challenge can of course be mitigated by presenting public services to citizens via IPv6, in addition to IPv4.Simple as that? Not really: How can the public authorities' access solution, internal networks, servers, and applications support IPv6 in practice? Well-tested transition techniques, such as using a reverse proxy in front of a web portal, can help to ease the needed transition steps by allowing a partial transition to IPv6 and with limited effort.

The next figure shows a setup in which the PA is connected with IPv4 and IPv6 in parallel to the Internet and features a local proxy server. In the form of a reverse proxy, it can reside upstream of the PA's web server and handle IPv4 plus IPv6 connections, while the web server "behind it" can stay unchanged.

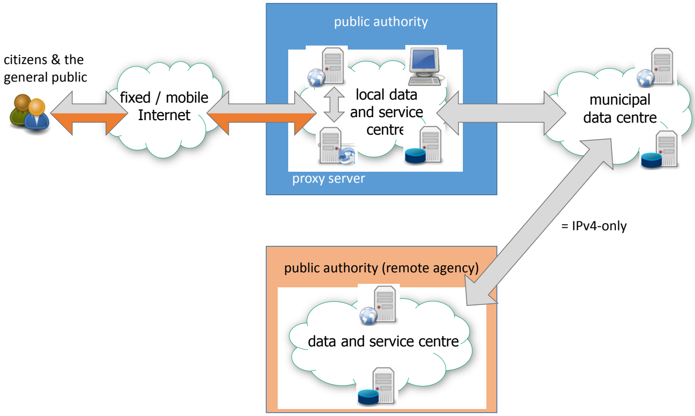

*IPv6 traffic* 🟧 *IPv4 traffic* ⬜

In this case, the Internet access, the reverse proxy server and the network(s) in between need to support IPv4 and IPv6, but the web server itself does not. Still, its contents will be reachable via IPv4 and IPv6 from the outside, with the reverse proxy being in charge of "translating" IPv6 connections into IPv4 ones. Similarly, a separate local proxy server might be used to access IPv6-enabled sites in the Internet from IPv4-only clients from inside the PA (not separately depicted).

The use of tunnelling solutions is another intermediate solution to couple IPv6-enabled services. Such tunnels can be used e.g. across IPv4-only infrastructures, thereby implementing a connectivity that is equivalent to native IPv6 support from the point of view of local clients. This is achieved by encapsulating IPv6 traffic into IPv4 packets, so that the underlying IPv4-only infrastructure can manage them as usual.

In the following scenario (figure below), this would allow direct IPv6-connectivity between systems in the municipal data centre and the remote public authority.

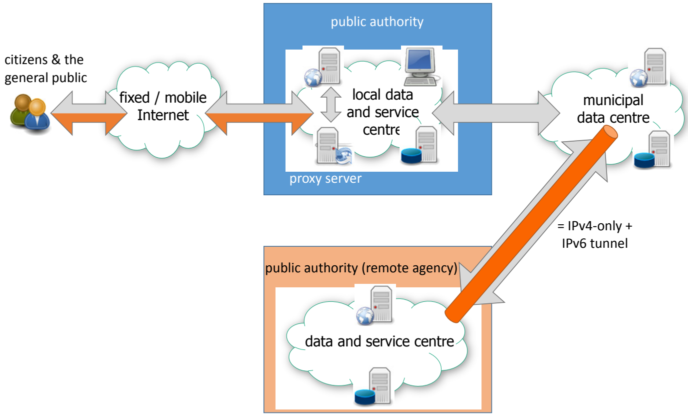
public authority
citizens & the general public
fixed / mobile Internet
local data and service centre
municipal data centre
proxy server
= IPv4-only + IPv6 tunnel
public authority (remote agency)
data and service centre

This on-going development finally leads to a scenario where IPv6 will be available everywhere, in public networks (the Internet) as well as in local networks (Intranets). Then, users can connect to all services they need whether they connect from an IPv4 or an IPv6 endpoint. A proxy for translating connections between IPv4 and IPv6 is not required anymore. This can be achieved when in all affected nodes and network equipment both protocols (v4 and v6) are implemented This is called dual-stack operation.

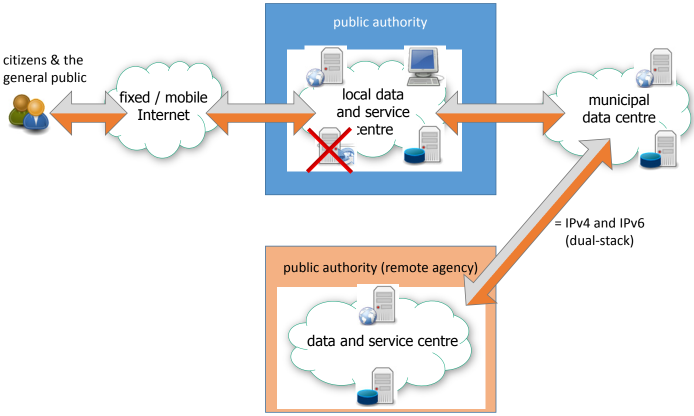
public authority
citizens & the general public
fixed / mobile Internet
local data and service centre
municipal data centre
= IPv4 and IPv6 (dual-stack)
public authority (remote agency)
data and service centre

These techniques will be shown in more detail later on. Other techniques are explained in the transition section where we also give recommendations about usage and risks.

# IPv6 TRANSITION APPROACH – HOW TO GET FROM IPv4 TO IPv6

We expect IPv4 and IPv6 to co-exist for quite a long time in the Internet and in local networks. Coming from an IPv4-only network, there are several options to change to a network where you can access IPv4 and IPv6 services. Each different technique has advantages and disadvantages, and there is no best overall technique. Depending on the use case, e.g., whether you are a carrier, an Internet Service Provider (ISP), a content provider / web-hoster, running a company network, or using the Internet from your home, a different technique might be best for you.

In practice, there also exist mixed forms, such as Internet content providers that also consume Internet services and so act as Internet users. Data centres are providing Internet connectivity to their customers but are also hosting services so they are (technical) content providers.

From a top-level point of view, the transition process should look like this (inculding procurement of new hardware, in cases where existing hardware cannot be upgraded or updated to support IPv6):

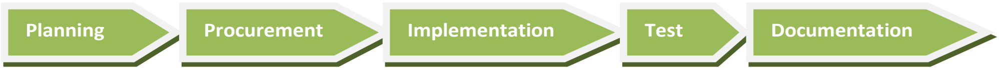

During this process, the following questions have to be answered:

- What are the transition objectives (e.g. making WWW servers available via IPv6)?
  - Where do you need IPv6 for the mission of the administration?
  - Are you going for a full or partial transition?
  - What are your time frame and planned milestones?

- What is my inventory in terms of hardware as well as used applications?
  - What is the current network sub-division and (re-) design?
  - What services are offered and their life cycles?
  - What is the state of the network access points (e.g. WAN access) and the capabilities of the external provider concerning IPv6?

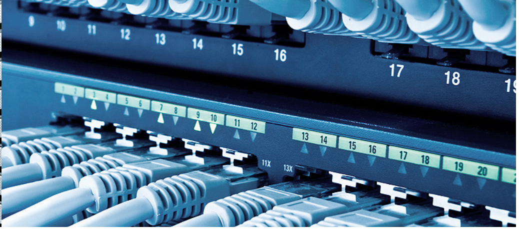

- What contracts (e.g. maintenance and services) are affected?
  o Including contact information, terms of use, depreciation and service-level agreements (SLAs)

- What is the impact on human resources?
  o Do my technicians have the needed IPv6 skills?
  o Did we plan resources for managing both protocols during the coexistence period?

Because the transition can affect several stakeholders involved in service provisioning, it requires sufficient time for preparation and must be addressed early by all parties.

Make sure that an overall plan and timetable is available and known to all involved entities, so that (a) the goals of the transition project are clearly visible and (b) the targeted timeline is well known.

In the next sections, we list in more detail the most common scenarios for using IPv6 in public administration and approaches for enabling IPv6 in an IP network, or parts thereof. In some cases this means adding IPv6 to the existing IPv4-capabilities to a gateway, sometimes directly on the end systems as well. We identify the advantages and disadvantages with three different fields of application in mind:

- Small office home office (SoHo),
- Internet Service Provider (ISP), and
- Public Administration (PA).

The most common scenarios / use cases for adding IPv6 to an existing IPv4-only infrastructure are the following three:

1 A service that is reachable externally, e.g. public web presence or governmental service should be reachable via IPv6, in addition to IPv4.

2 Internal clients need to access an external web site or service via IPv6.

3 Internal clients need to access an internal service via IPv6.

The third scenario is probably the least urgent, as it is only a must-have case if the client or the service is IPv6-only (or both are) – therefore we will concentrate on the first and second cases.

In both scenarios, the partial and full transitions exist.

**Let us check scenario 1 first:**

Using a partial transition can be done using a reverse proxy for incoming requests. In this case, the Internet / WAN access needs to support IPv6 as well the local DMZ hosting the proxy server (and the proxy of course) – the web server or service "behind" the proxy however can stay as-is.
A full transition does not need such a reverse proxy at all; however, the lo
cal web server that hosts governmental services and its network then need to support IPv6, too:

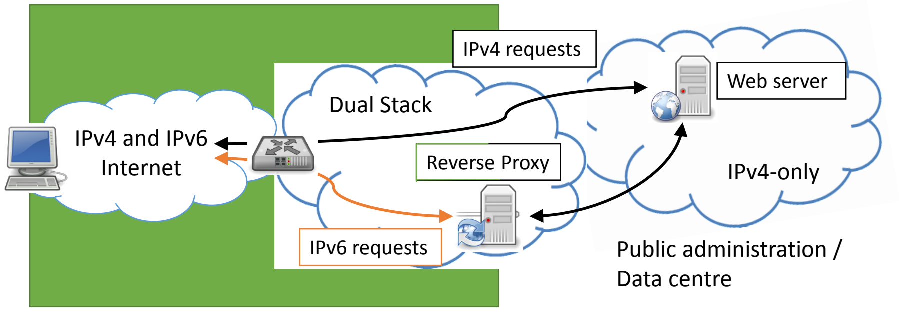

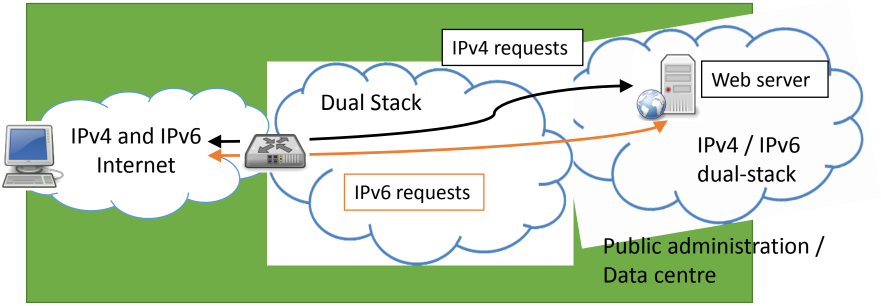

We give more details on running a reverse proxy solution later on in this document.

Which option will cost you less effort in the short term and in the long term depends on your local network and server(s) setup. Choose wisely!

**Let us look at scenario 2 now:**

A setup similar to the above consists of an added outgoing proxy that converts between IPv4 and IPv6 – but only for ‚proxy-compatible' protocols! In this case, the local Intranet can stay as-is.

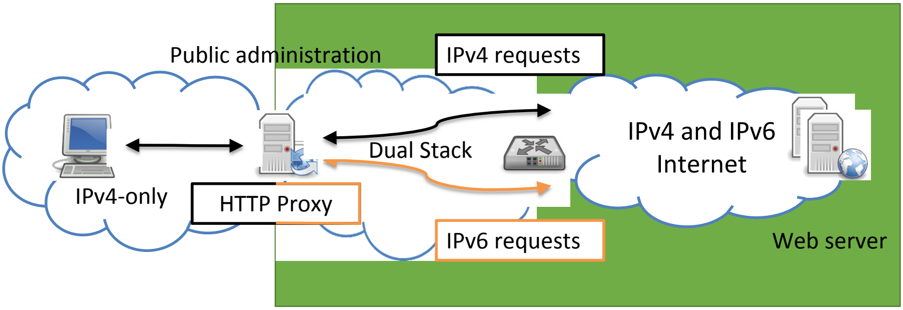

A full transition to IPv4 / IPv6 dual-stack support does not need the proxy in between the Intranet and external networks such as the Internet.

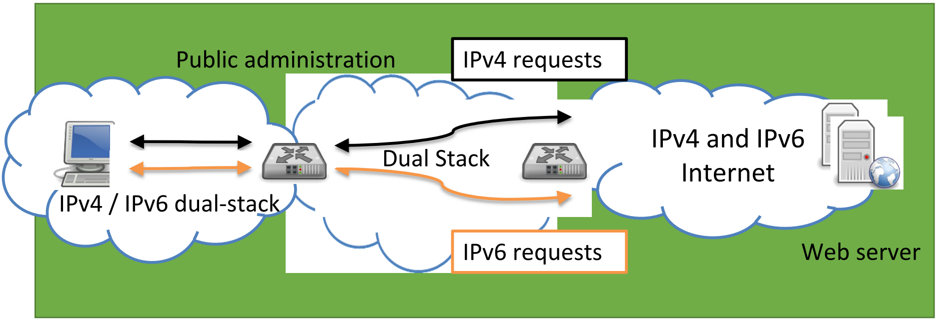

On the other hand, full dual-stack support including the end systems means that all affected systems (switches, routers, security devices, end systems …) must support both IP protocol versions. The future will look like this for quite some time – until IPv4 is not needed anymore.

However, setting up a proxy first can save valuable time during which clients and client networks can be migrated, while they already can access IPv6 servers and services „on the outside". Choose your way on the base of your environment and requirements.

## GO AHEAD – GO YOUR OWN WAY

This small booklet gives an overview on IPv6 addresses, address planning and transition planning based on the experiences GEN6 made from different pilots. Introducing IPv6 has nothing to do with miracles nor is it rocket science. Starting now with an inventory of your network, equipment and service will enable you to move gradually to dual-stack. This will bring you to the safe side whatever protocol version your customer or you will use. Yes, you are right - dual-stack means additional management duties and tasks but you get the chance to change to IPv6-only as soon as possible and to reduce this additional work in the future. Do it now with enough time and you will have fewer problems during the transition process. Don't do it in a hurry!

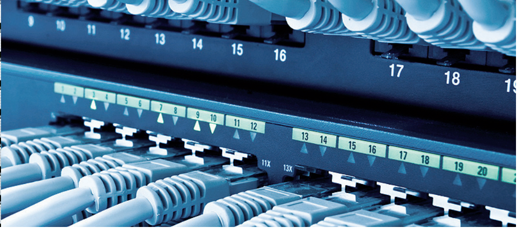

# TRANSITION TECHNIQUES – AN OVERVIEW

On the top level, the following four classes of techniques can be distinguished for the transition to IPv6:

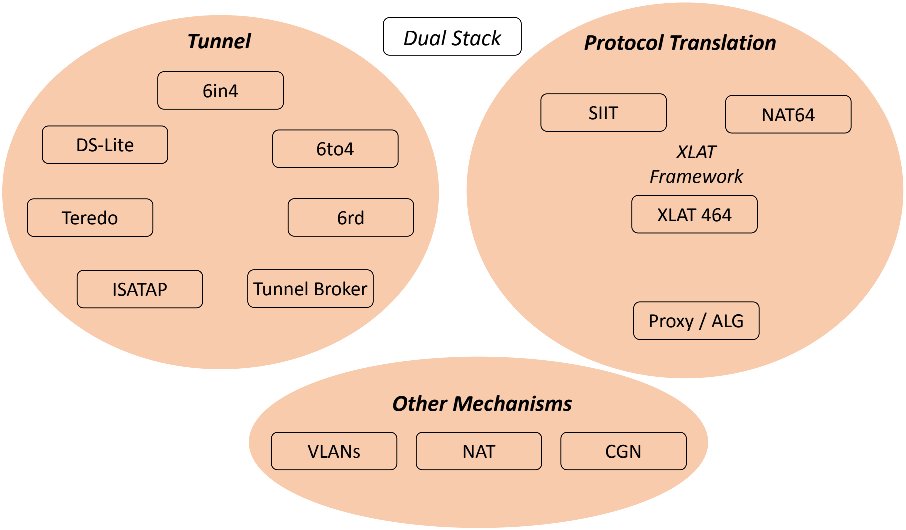

Dual-stack means using IPv4 and IPv6 in parallel, i.e. making both available in native form on IT systems, albeit at the cost of more maintenance, but free of any tunnelling or translation hassles. Additionally, using dual-stack, i.e. native IPv4 plus native IPv6, for access networks, keeps the possibility for end-to-end connectivity intact.

Tunnelling is mostly use to bridge IPv6 islands across networks that are IPv4-only (or vice versa). They can be considered a valid bridge technology, usable, before native IPv6 comes into widespread use.

Protocol translation tries to bring the two incompatible versions of the Internet protocol together by making IPv6 packets out of IPv4 packets, and vice versa (in the opposite traffic direction). It avoids the dual maintenance of Dual-stack, but the cost of running a complicated translator gateway. Such translation is inherently complicated due to inherent protocol version differences.

# DUAL-STACK

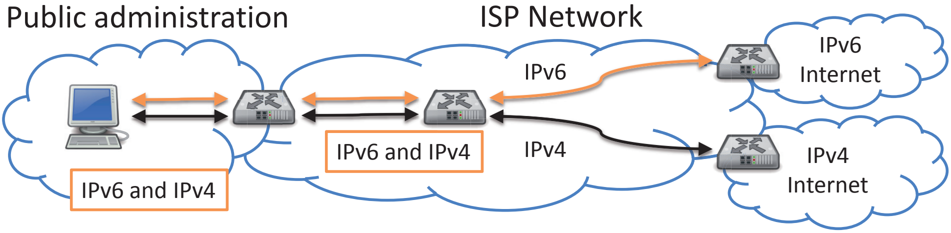

- see (among other RFCs) RFC4213 Section 2 and following
- not a transition technology per se, as it still keeps IPv4 up and running as-is
- parallel operation of IPv4 and IPv6 to „bridge“ the non-compatibility between IPv4 and IPv6
- universally applicable …
  - in the Intranet (for workstations, servers, applications)
  - for the Internet access / for Internet-based services
  - in the backbone networks (IPv6 is already broadly used there today)
- especially in the case of Internet browsers the selected IP version for communication between dual-stack systems is not easy to predict → this can make errors analysis tricky

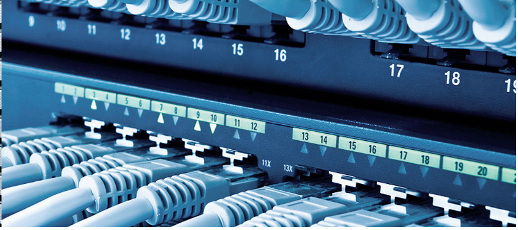

- dual-stack *can be deployed in multiple, consecutive steps (also as partial transition):*
    - e.g. outside-in: WAN at first; then DMZ; then till proxy server(s); then till end systems
    - additionally use bottom-up approach: dual-stack first to routers and security devices, then end systems, then applications
    - deployment needs to be oriented on existing network zone boundaries, to avoid any security hassles
- dual-stack approach very usable for ISPs, public administration, and SoHo alike

| | |
|---|---|
| 🏠 | Dual-stack networking is a very applicable approach, as it allows the seamless, parallel operation of older IPv4-only, newer dual-stack, and bleeding-edge IPv6-only components inside the home network, including Internet access for them. |
| 🏢 | Dual-stack networking is similarly useful in the public administration use case, as it allows the use of IPv4-only, dual-stack, and IPv6-only components. Depending on the actual setup different IP subnets in the PA may be running as IPv4-only (legacy only systems), IPv6-only (e.g. new VoIP infrastructure), or with both IP protocols supported in parallel. |
| ISP | For ISPs, the use of dual-stack networks is one valid potential solution to support customers with IPv4 plus customers with IPv4 and IPv6 (IPv6-only customers being extremely rare yet). However, due to the high number of network entities and endpoints the added management overhead for running both protocols in parallel everywhere also makes other (tunnelling) solutions attractive for ISPs. |

# 6RD

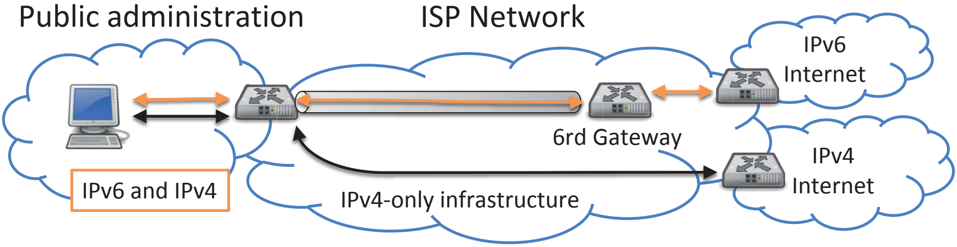

- RFC 5969 – IPv6 Rapid Deployment on IPv4 Infrastructures (6rd) – Protocol Specification
- 6rd is a provider technology; provides IPv6 to ISP's customers via an IPv4-only infrastructure
  - can be seen as an advancement of 6to4, with the gateway now inside the ISP's network
  - initially developed by the French provider „Free"
  - customers keep IPv4 and get (slightly limited) IPv6 support in addition to it
    - IPv6 is tunnelled to the customer via IPv4 till ISP's gateway routers
  - IPv6 addresses for customers are taken from the IPS's IPv6 address range
    - better reachability and stability of IPv6 nodes, compared to 6to4

- not directly applicable for SoHo or public administration; often combined with local dual-stack

- allows the rollout of IPv6 to customers before making the ISP network completely IPv6 capable

| | |
|---|---|
| 🏠 | 6rd not applicable locally, if 6rd is used by an ISP, native IPv6-support is still needed in the home |
| 🏢 | 6rd not applicable, (a) because tunnelling IPv6 over IPv4 may impose worse transmission quality than native IPv6, and (b) because no provider-independent IPv6 addresses could be used by the PA |
| ISP | Useful as “bridge technology” for a quick rollout of IPv6 to the ISP’s customers, even before native IPv6 support is available in all infrastructure components of the ISP |

# 6TO4

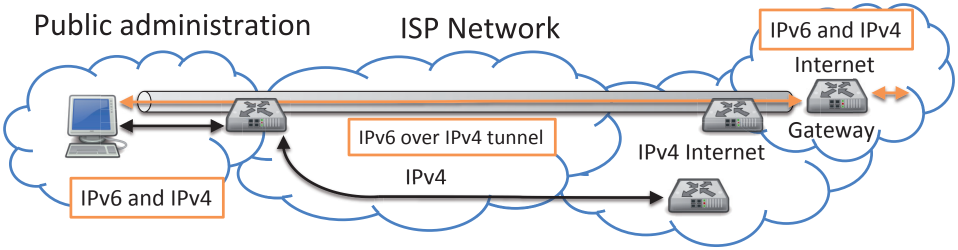

- 6to4 standard: RFC 3056 – Connection of IPv6 Domains via IPv4 Clouds
- allows connection of dual-stack clients to external IPv4 and IPv6 Internet
  - IPv4 natively, IPv6 via tunnel and gateway
- IPv6 packet tunnelled from client via tunnel to 6to4 relay; no IPv6 support needed from ISP
- End system of local router creates tunnel towards the public 6to4 gateway
- Security of 6to4 can be considered critical, since the gateway to the Internet (and thereby its location) is chosen arbitrarily, i.e. completely out of control of the end system
- 6to4 is prone to attacks since incoming 6to4 traffic must be accepted from any 6to4 gateway in the Internet

| | |
|---|---|
| 🏠 | Possible „compromise solution“, if no better solution (native IPv6, 6rd, DS-Lite) allows to access the IPv6-Internet otherwise |
| 🏢 | Security risks of 6to4 are much too high – do not use |
| ISP | Can usually pass 6to4 traffic, but commonly does not provide 6to4 gateways themselves |

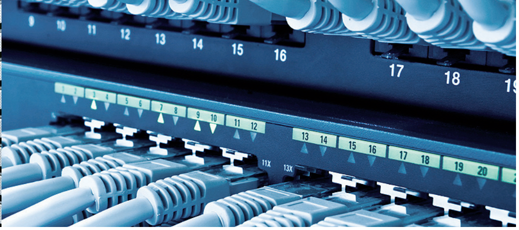

# ISATAP

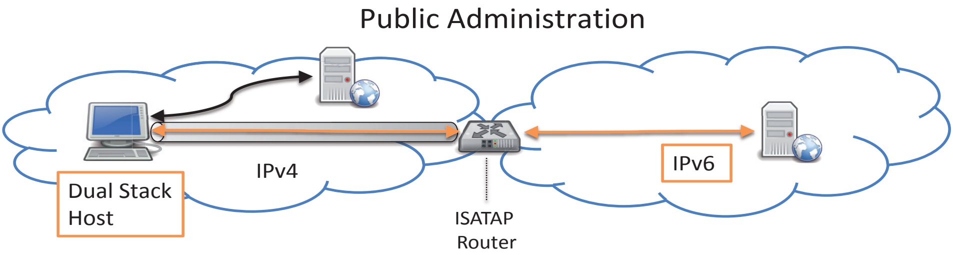

- RFC 5214 – Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)
- Usable inside an organization, where networks are IPv4-only but some hosts need IPv6
- Uses IPv4 as link layer for IPv6 traffic; generated IPv6 address, based on local IPv4 address
- ISATAP can be used to give a set of Dual-stack Hosts IPv6 inside an IPv4 environment

| | |
|---|---|
| 🏠 | Dual-stack support is common enough in hosts and home routers so that complexity of using ISATAP can be avoided |
| 🏢 | ISATAP can be used in exceptional cases for hosts in IPv4 networks in need of IPv6 |
| ISP | Not applicable / not usedprovide 6to4 gateways themselves |

# DS-LITE

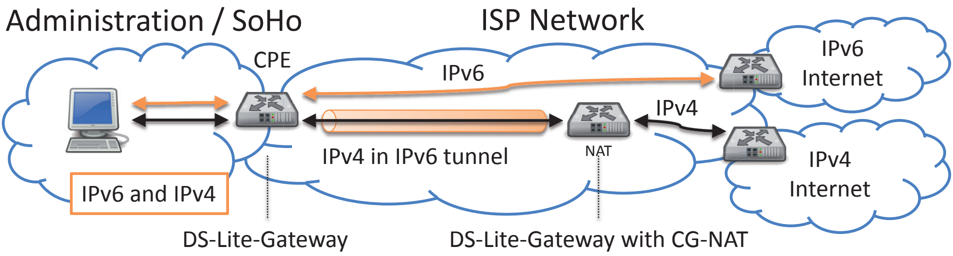

- Standard: RFC 6333 – Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion

- Provider technology: Customers get native IPv6 plus IPv4 tunnelled over IPv6, with Carrier Grade NAT

    - Customer premises equipment (CPE) uses only native IPv6 on the WAN-side, dual-stack towards the internal, customer network, simplifying management.

    - Most desirable technology for providers with limited amount of IPv4 addresses

    - Use of Carrier Grade NAT (CGN) means extra efforts for provider, and functional disadvantages for customer (e.g. reachability; more complex port forwarding)

    - CGN also means that customers may share an external visible IPv4 address, which can negatively affect some applications

- Overall, DS-Lite is a provider technology. Main goals: Reduce/save number of needed public IPv4 addresses and at the same time transition provider networks towards full IPv6 support (or even build new provider networks as IPv6-only, thus reducing the management efforts needed for operating IPv4).

| | |
|---|---|
| 🏠 | Customer can get IPv4 plus IPv6 from a provider with DS-Lite, uses dual-stack inside the home network. Gets native IPv6 plus IPv4 with some new limitations. |
| 🏢 | Recommended to demand native IPv4 plus native IPv6 from ISP, since the IPv4-related drawbacks of DS-Lite + CGN may introduce new problems with existing governmental applications. |
| ISP | ISP may use DS-Lite where IPv6-capable infrastructure exists; needs new CGN gateways too, but saves IPv4 addresses and reduces some management efforts for its own networks. |

# NAT64

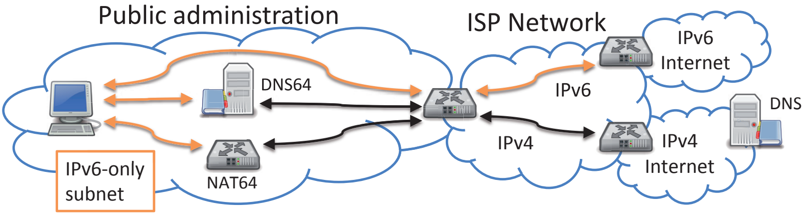

- Standard: RFC 6146 – Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers

- Allows end systems in an IPv6-only environment to access the IPv4 Internet, in addition to the IPv6 Internet

  - Needs a DNS64 server plus a NAT64 gateway inside the local Intranet

  - Only works with services which can be accessed via a hostname (e.g. URI/URL containing a hostname, no literal IPv4)

- DNS64 server returns IPv6 address with embedded IPv4 address for each requested external host that has only an IPv4 address

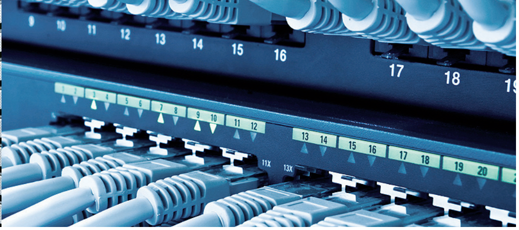

| | |
|---|---|
| 🏠 | Too complex for the home use; may inhibit connection to some IPv4 services; dual-stack recommended instead |
| 🏢 | Useful technology for application in newly built, IPv6-only workplace subnets<br>• reduces management effort in new IPv6-only networks (compared to dual-stack)<br>• useful for making legacy IPv4-only systems still accessible to IPv6-only end systems<br>Pitfalls: IPv4-only client software with embedded IPv4-addresses will fail to work (e.g. active FTP, Skype, and many computer games) |
| 🌐 | NAT64 can be deployed by providers as part of a XLAT464 setup |

## APPLICATION LEVEL GATEWAY (ALG) VIA REVERSE HTTP PROXY

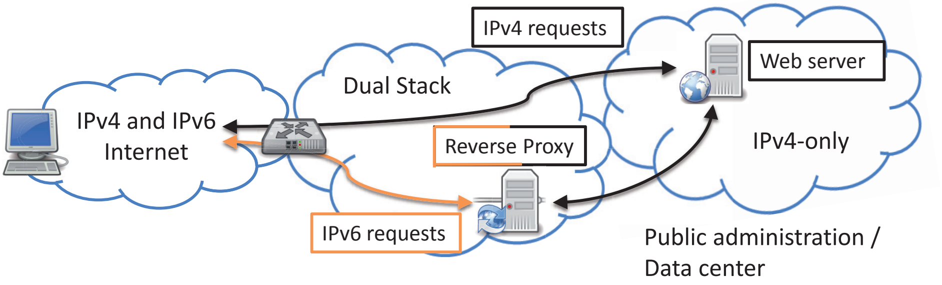

- Integrates an additional reverse proxy server into the local network in order to make content hosted on a local IPv4-only webserver available for IPv6 clients too
- Usually the configuration of the existing IPv4-only webserver can stay untouched

- Quickly deployable bridge-technology, usable until webserver supports  IPv4 and IPv6 natively
- Precondition: IPv4 plus IPv6 available on the WAN access (preferably both in native form)
- Applicable for HTTP and for other "proxy-able" services and protocols, e.g. FTP, SMTP and POP3
- Partial migration with reverse proxy can also be applied e.g. in load balancer environment
- Needs monitoring of amount of IPv6 requests vs. IPv4 requests to dimension and scale reverse proxy web server accordingly

| | |
|---|---|
| 🏠 | Uncommon, since home users rarely provide services to the external Internet, however if services are provided then reverse proxy might be used |
| 🏢 | Useful „bridge-technology" to quickly setup IPv6-availability of web servers/services |
| ISP | Useful as „bridge technology" for own hosted services; could also be provided as a service for customer servers/services hosted by an ISP |

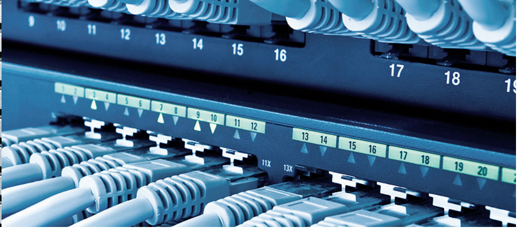

# OUTGOING HTTP PROXY

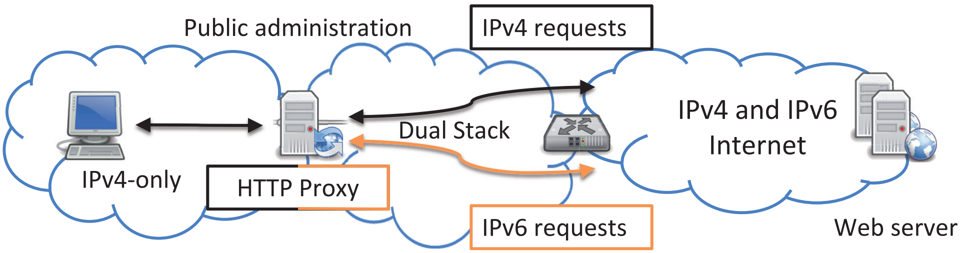

- Bridge solution to access external IPv4 and IPv6 servers and services from an IPv4-only Intranet environment
- Quickly deployable, useful before the local Intranet gets upgraded to support IPv6
- Existing external servers/services need no configuration adjustments
- Reasonable effort for setting up a proxy server; highly mature solutions and software exist
- Related case: outgoing proxy can also allow clients in IPv6-only Intranet clients to access IPv4- and IPv6-based external servers and services

| | |
|---|---|
| 🏠 | Use of outgoing proxy not common; still, use at home is possible. |
| 🏢 | Useful bridge technology for allowing existing IPv4-only clients in the Intranet to access IPv4 and IPv6 services (via proxy-compatible protocols)<br>• (possibly existing) proxy provides useful border for partial transition to dual-stack<br>• can also be used for accessing internal IPv4-based servers that cannot be migrated to support IPv6 themselves from Intranet IPv6 clients |
| 🌐 ISP | Proxies can be used for internal services and could also be provided as a service<br>• However: Confidentiality of transmitted (proxies) information must be guaranteed |

# SUMMARY–VALIDATION OF TRANSITION TECHNIQUES

The following table lists currently proposed transition techniques for making use of IPv6 by either a small or home office (SoHo) environment, a Public Administration (use in Intranet and DMZs), and by Internet Service Providers. An empty field denotes that this combination does typically not apply – and is therefore not evaluated. A "++" means that this combination is in broad use and is very recommended for practical use, "+" a bit less so, until "–" means, "Not recommended, avoid!" for either security reasons or because of high effort (for little gain).

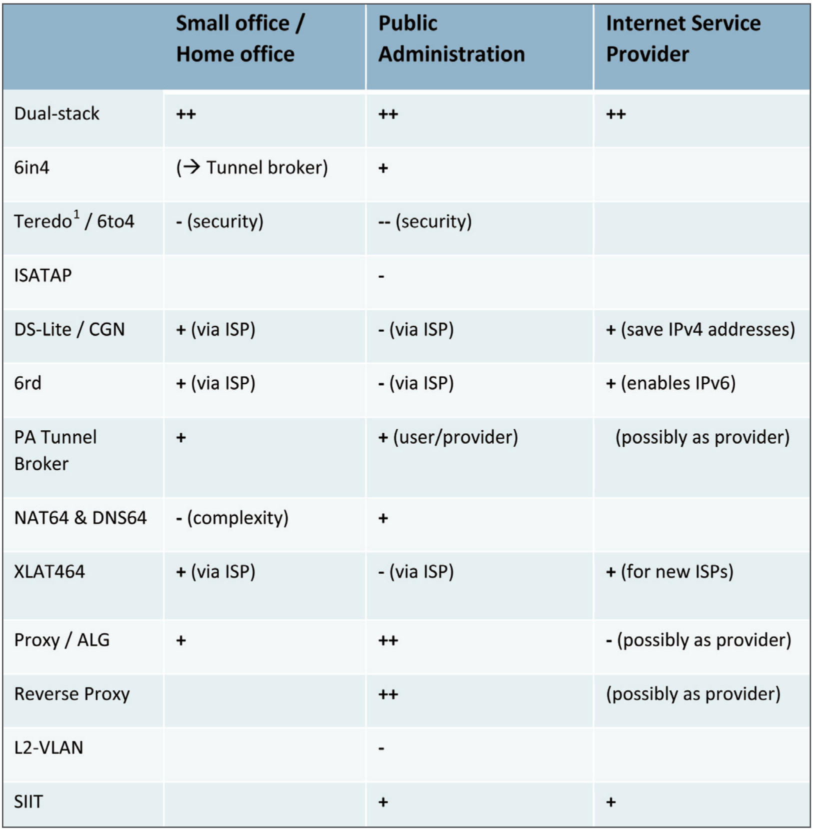

| | Small office / Home office | Public Administration | Internet Service Provider |
|---|---|---|---|
| Dual-stack | ++ | ++ | ++ |
| 6in4 | (→ Tunnel broker) | + | |
| Teredo[1] / 6to4 | - (security) | -- (security) | |
| ISATAP | | - | |
| DS-Lite / CGN | + (via ISP) | - (via ISP) | + (save IPv4 addresses) |
| 6rd | + (via ISP) | - (via ISP) | + (enables IPv6) |
| PA Tunnel Broker | + | + (user/provider) | (possibly as provider) |
| NAT64 & DNS64 | - (complexity) | + | |
| XLAT464 | + (via ISP) | - (via ISP) | + (for new ISPs) |
| Proxy / ALG | + | ++ | - (possibly as provider) |
| Reverse Proxy | | ++ | (possibly as provider) |
| L2-VLAN | | - | |
| SIIT | | + | + |

# RECOMMENDATIONS

You can see from these technical descriptions, the recommendations and the previous table that the Internet research community has devised a huge number of transition techniques – but, depending on the use case, only a small set can be recommended per case.

For a user's internal network at home, dual-stack support is the best choice. We recommend dual-stack also for the Internet access at home. However, some tunnelling solutions such as 6rd and DS-Lite will give you practically the same IP (v4+v6) Internet access quality at home, as long as the ISP dedicated enough resources to its tunnelling infra-structure. However, this strongly depends on one's Internet service provider.

For public administration networks different requirements apply and therefore different solutions are also applicable: Dual-stack is the preferred way to go! The tunnelling solutions that are okay for the home user are not advisable for PAs. Instead a PA should either rely on full dual-stack (access + internal) or on dual-stack access plus local proxy solutions for internal IPv4 nodes.

Tunnelling techniques have to be handled carefully! Access via a trusted tunnel broker can be used to enable IPv6 access of a PA. For security reasons, ISATAP, 6to4 and Teredo interfaces on networked nodes should be manually disabled, unless they are explicitly needed.

We did not add a separate section on the Teredo technique to the previous section due to the major security drawbacks of its firewall-punching nature. The use of Teredo is not recommended. Tunnelling techniques must be handled carefully in any case! Access via the manual setup of tunnels in controlled environments for specific cases of international connections between PA networks can be necessary. In GEN6, this was done in a cross-border pilot between Germany and Spain where tunnelling was used because the sTESTA network does not support native IPv6.

Newly created IP networks inside the PA may also operate with NAT46 plus DNS64, which reduces network management overheads compared to dual-stack when used at a larger scale.

All transition pilots from within GEN6, the German data centre transition, the ERCS system from Slovenia, the Turkish government portal services, the Spanish government backbone transition and the Greek schools' energy efficiency network are following three basic rules and are currently favouring rule number 2:

(1) If you can run it on IPv6 only then go for IPv6.

(2) If you have to support a mix of IPv4 and IPv6 then go for dual-stack and fade out IPv4 as quickly as possible.

(3) If you have special requirements or rule (1) and (2) doesn't work for you then choose a solution which allows a smooth transition to either rule (1) or (2) in the future.

Detailed information about the German addressing scheme has been copied from:

DEUTSCHLAND-ONLINE INFRASTRUKTUR
IPv6 Referenzhandbuch January, 2011

http://www.bva.bund.de/DE/Organisation/Abteilungen/Abteilung_BIT/Leistungen/IT_Beratungsleistungen/IPv6/best_practice/ipv6migrationsleitfaden/download/IPv6_Referenzhandbuch_2011_Version_KW45.pdf?__blob=publicationFile&v=2

# 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-ND 3.0) licence. Consequently, you are 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-project.eu"), and use the document for non-commercial purposes only, and without any alteration, transformation or building derivatives upon this work.

The information herein does not necessarily express the opinion of the EC. The EC and the document authors are not responsible for any use that might be made of data appearing herein and effects that result from doing so. 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 any direct nor indirect loss or damage suffered by any person using this information.

This booklet was created by

Fraunhofer
FOKUS

GOBIERNO
DE ESPAÑA

MINISTERIO
DE HACIENDA
Y ADMINISTRACIONES PUBLICA

This work was in part supported by the European Commission as part of the project "Governments Enabled with IPv6" (GEN6). GEN6 is about stimulating EU-wide deployment of IPv6 by means of best practices and guidelines in existing eGovernment infrastructures.

**Authors**
Carsten Schmoll (FhG FOKUS, Germany)
Uwe Holzmann-Kaiser (FhG FOKUS, Germany)
Carlos Gómez Muñoz (MINHAP, Spain)
Martin Krengel (Citkomm, Germany)

**Contact**
To get in contact with the GEN6 project or
the partners please contact us in:
**info@gen6-project.eu**
**www.gen6-project.eu**

GOVERNMENTS ENABLED WITH IPv6
GEN6

*Layout by Citkomm, all photographs © 2014 by fotolia.com*
*or by Steffen Konegen (Fraunhofer FOKUS)*