


IST-2001-32603	Deliverable D 2.1.1	
----------------	---------------------	---

Project Number:	IST-2001-32603
Project Title:	6NET
CEC Deliverable Number:	32603/DANTE/DS/2.1.1/A1
Contractual Date of Delivery to the CEC:	30 th June 2002
Actual Date of Delivery to the CEC:	9 th August 2002
Title of Deliverable:	IPv4 to IPv6 migration scoping report for core networks
Work package contributing to Deliverable:	WP2
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Agnés Pouéle, Sabine Kühn
Contributors:	Tim Chown, Bruno Ciscato (Cisco), Pekka Savola (CSC/FUNET)

* 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)

Abstract: The goal of this deliverable is to provide an overview about existing transition mechanisms which are useful for a smooth integration of IPv6 into IPv4 core (backbone) networks. Transition scenarios are discussed on the example of the GÉANT network as a pan-European network, which is a potential candidate for an IPv4 to IPv6 transition in the near future.

Moreover, experiences and implementation details about the MPLS/CCC encapsulated IPv6 as one of the deployment strategies are provided, which is used to connect Greece, Poland and Hungary to the 6NET core network.

Key words: IPv4 to IPv6 transition mechanism, deployment strategies, core, backbone networks

Table of contents

1 INTRODUCTION3

2 REVIEW OF CORE NETWORK TRANSITION MECHANIS MS3

2.1 OVERVIEW4

2.2 DUAL STACK4

 2.2.1 *Addressing a dual stack network*.....4

 2.2.2 *Routing protocols*.....4

2.3 TUNNELING METHODS.....5

 2.3.1 *IPv4 tunnels*.....6

 2.3.2 *MPLS and Layer 2*6

2.4 TRANSLATION MECHANISMS.....7

3 SCENARIOS FOR CORE TRANSITION7

3.1 DEFINITION OF THE REQUIREMENTS.....8

3.2 CORE DESIGN8

 3.2.1 *Deploying IPv6 over MPLS Backbones*.....8

 3.2.2 *Deploying IPv6 Using Dual-Stack Backbones*..... 12

 3.2.3 *Different steps for the migration*..... 12

4 CONSIDERATIONS FOR MECHANISM SELECTION..... 14

5 CONCLUSIONS 15

6 REFERENCES 15

1 Introduction

The 6NET network is a native IPv6 network. Its core is built using native STM-1 links between DE, CH, AT, IT, SE and UK connected by Cisco 12404 series routers. However some parts of the network operate using a transition mechanism for carrying the IPv6 traffic. Providing 6NET connectivity to Greece, Hungary and Poland for example, Multi-Protocol Label Switching (MPLS) encapsulation over the GÉANT network is used.

Rather than complete a full upgrade to a native IPv6 network there exist several integration strategies for the deployment of IPv6. These strategies allow a controlled integration of IPv6 in IPv4 networks starting from the edges of the network and moving towards the core. While deliverable D2.3.1 [2] deals with the IPv4 to IPv6 migration at the edge of the 6NET network, the university networks, and D2.2.1 [1] with the transition in organisational networks, this deliverable addresses the transition mechanisms (section 2) and deployment strategies (section 3) suitable for core (backbone) networks.

Although this document covers some of the same topics as D2.2.1 it gives some special consideration to core networks here¹. In particular, this deliverable focuses on investigating deployment strategies for backbone networks with respect to their configuration effort, their complexity to network management, their impact to layer 2 and 3, and their performance implications. The next deliverable, D2.1.2, due in M12 of the project will be a combined deliverable with D2.2.2 (i.e. a combined NREN/ISP/core networks transition cookbook).

Fundamental to the successful market adoption of IPv6 is the ability to integrate it with the existing IPv4 infrastructure without significant disruption to services. Given the magnitude of the task involved in replacing today's IPv4 with the new IPv6, integration and coexistence need to be well defined and planned and this has been the focus of the IETF Next Generation Transition (NGtrans) Working Group for several years. There is however no reason to expect a full migration to the new IPv6 and for an indefinite period both IPv4 and IPv6 nodes will coexist. In tackling IPv6 integration into IPv4 networks, the approach taken has been to partition the tasks into the host/client portion and the network portion. Several techniques have been identified and are explained in the following sub-sections and it is reasonable to anticipate that significant deployments of IPv6 will employ a combination of these mechanisms.

2 Review of core network transition mechanisms

In this section an overview of existing transition mechanisms are given. Transition mechanisms define in general a set of techniques that IPv6 hosts and routers may implement in order to be compatible with IPv4 hosts and routers, for example tunneling mechanisms and dual stack mechanisms.

Integration strategies on the other hand, which are described in the following section 3) are guidelines where to implement services or to start providing them in the network and how to

¹ In M12 for the transition cookbook, a single deliverable will be released to cover both activities due to the overlap (as demonstrated in the similarities between D2.1.1 and D2.2.1).

combine such transition mechanisms to provide IPv6 services over specific IPv4 networks, for example using IPv6 tunnels on the customer edge routers or on the provider edge router.

2.1 Overview

The solutions standardised and implemented to migrate a core network from IPv4 to IPv6 have to be considered as a long-term transition mechanism rather than short term. This is due to the fact that it will take time for the Internet to converge homogeneously to IPv6 and main core backbones will have various IPv4 and IPv6 islands connected to them during this transition period.

Mainly, two families of core backbones can be foreseen in the first case. The ones that will stay IPv4 and deliver IPv6 services over tunneling methods and the ones that will run the two stacks of protocols and deliver native IPv4 and IPv6 services.

No general rule can be applied to the IPv4 to IPv6 transition process. In some cases, moving directly to IPv6 will be the answer. To be able to continue to provide IPv4 services over a native IPv6 core network, IPv4 over IPv6 tunneling mechanisms can be used at the edges or on the other hand, translation mechanisms like NAT-PT can be used at end sites.

2.2 Dual Stack

Dual stack mechanisms are already implemented on Internet hosts and are embedded on the different suites of Operating systems. The IP dual stack is now available for most router platforms.

Unfortunately, enabling networks to be dual stack is a bit more complex than just configuring a computer to be dual-stack. Having a dual stack network means that routers have IPv4 and IPv6 forwarding mechanisms embedded plus the suite of routing protocols adapted for both IPv4 and IPv6.

Therefore, a dual stack router will be more stressed than a pure IPv4 router or a pure IPv6 router in terms of forwarding performance, memory usage and routing protocols adaptability.

2.2.1 Addressing a dual stack network

The whole IPv4 addressing plan has to be ported to IPv6 and finally tuned according to the IPv6 services that will be available.

Each trunk, each access and logical interfaces needs at least two global IP addresses: one for IPv4 and one for IPv6.

The network addressing scheme can be even be more complex with the usage of private addresses for IPv4 and link-local and site-local addresses for IPv6.

2.2.2 Routing protocols

To discover the topology and route traffic across the network, routing protocols have to be implemented for both stacks of protocols.

Generally, a core backbone has two types of routing protocols, the Interior Gateway Protocols (IGPs) and the Exterior Gateway Protocols (EGPs).

The IGP is used to learn the topology of the network and forward traffic across the backbone through the best path. Currently IS-IS, OSPF and RIP are the classical routing protocols used by IPv4.

IS-IS and RIPng are now available for routing IPv6 while OSPF version 3, which is a complete rewrite of OSPFv2 for IPv4, is starting to be available in beta versions on some platforms.

Accordingly to which IGP is implemented in the core backbone, the efforts required to migrate the network from an IPv4 network to a dual stack network won't be the same.

A core backbone installed with IS-IS, after the dual stack upgrade of the routers, will have one IS-IS process building one database in which the links have different attributes according to which IP protocols (v4 or v6) are used for the forwarding. That means, that a link which has IPv6 and IPv4 addresses will have different TLVs in IS-IS describing it. Moreover, if a service provider runs IS-IS for IPv4 and wants to use it for IPv6 but doesn't intend to align both topologies, then a multi-topology IS-IS is needed.

A core backbone with OSPFv2 implemented can not build a database for IPv6 and needs, in parallel, a second IGP which can be IS-IS or OSPFv3.

In some cases, if it is not possible to run an "IPv6-only" IGP database - to avoid having IPv4 routes in the second IGP either - two strategies have to be studied:

- The core backbone can run in parallel two IGPs.
- The core backbone has to change its IGP before running in dual stack mode.

2.2.2.1 Running in parallel two IGPs

The routers might be able to run two IGPs in parallel with good performance. However, having two IGPs in parallel will make it difficult to maintain consistency for both databases of the IPv4 traffic.

While IPv6 will be handled only by the second IGP, IPv4 will fill the databases of both IGPs and the administrator will have to configure preferences among the two link state protocols.

This might be difficult to manage and may create inconsistencies.

2.2.2.2 Migration of the core backbone from one IGP to another one

This is possibly the best choice for the reasons explained above. Technical solutions and methods are known for migrating the IGP. Unfortunately it adds complexity for migrating an IPv4 network to a dual stack mode.

BGPv4 is the external protocol to route IP traffic between domains. It has been enhanced with IPv6 NRLI and is available in BGP4+ after upgrade of the image of the router.

Nothing special has to be done, IPv4 and IPv6 BGP peerings can coexist. However vendor specific requirements and interoperability have to be taken into account (see Section 3 for more information about routing issues related to the core transition scenarios).

2.3 Tunneling methods

The second transition technique relies on tunneling. Considering the performance of its equipment or some other criteria a service provider often decide that part of its network won't be dual stack and instead it opts to deploy tunneling. Tunneling enables the interconnection of IP clouds. For

instance, separate IPv6 networks can be interconnected through a native IPv4 service by means of a tunnel. IPv6 packets are encapsulated by a border router, before their transportation across an IPv4 network, and de-capsulated at the border of the receiving IPv6 network. Tunnels can be statically or dynamically configured, or implicit and automatic (6to4). In later stages of transition, tunnels will also be used to interconnect any remaining IPv4 clouds through the IPv6 infrastructure.

2.3.1 IPv4 tunnels

Deploying IPv6 over IPv4 tunnels by encapsulating IPv6 packets in IPv4 packets is primarily used for carrying IPv6 traffic between isolated IPv6 sites or as a connection to remote IPv6 networks over an IPv4 backbone. The techniques comprise manually configured tunnels, generic routing encapsulation (GRE) and fully automatic tunnel set-up mechanisms such as 6to4. Such tunneling techniques can be used for example between the edge routers of a core network. The techniques used can be various and depend on the available features on the routers. However, the service provider will have to consider the forwarding performance of its edge equipment, as tunnel encapsulation/decapsulation is a heavier task than only forwarding IPv6 packets, and the mechanisms or routing protocols to be implemented. Applying eg. the 6PE method (which is explained in Section 3.2.1.3, and also in Deliverable D2.2.1 [1]) while using the IPv4 forwarding mechanisms results in a set of fully meshed GRE tunnels between the edge routers.

2.3.2 MPLS and Layer 2

IPv6 traffic can either be tunnelled over MPLS using deployment strategies like CE (using IPv6 tunnels on the customer edge - see Section 3.2.1.1) and 6PE (using IPv6 on the provider edge – see Section 3.2.1.3) or by the use of layer 2 VPNs. ATM core backbones or layer 2 framing carried over MPLS can allow IPv6 traffic to cross the core backbone in a completely transparent way. Point-to-point tunnels (ATM PVC or LSP cross connection) can be set up between IPv6 customer routers and deliver a primary IPv6 service until the backbone achieves its transition.

Currently GÉANT provides one point-to-point connection of that type to connect Greece to the 6NET backbone network.

The two 6NET routers gr6.gr and de6.de are interconnected via a MPLS point-to-point virtual link across GÉANT.

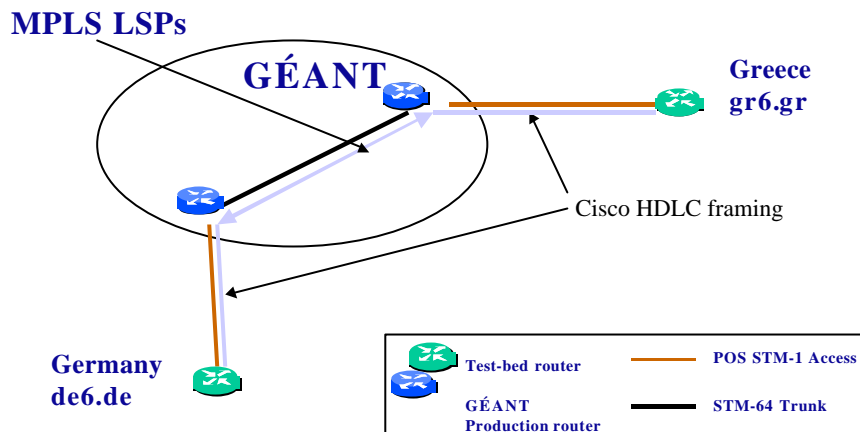


Figure 1: MPLS/CCC connection set-up via GÉANT between Greece and 6NET

The framing layer 2, “Cisco-HDLC”, is carried transparently over MPLS across GÉANT. The edge and core routers are only IPv4/MPLS capable.

The main advantage of MPLS layer 2 tunneling is the ability to manage the routing of the tunnel across the backbone with Traffic Engineering.

Additional services like bandwidth guarantees and backup tunnels can be provided.

2.4 Translation mechanisms

The last technique uses a translation mechanism. Translation is necessary when an IPv6-only host has to communicate with an IPv4 host. At least, the IP header has to be translated but the translation will be more complex if the application processes IP addresses; in fact such translation inherits most of the problems of IPv4 Network Address Translators. ALGs (Application-Level Gateways) are required to translate embedded IP addresses, recompute checksums, etc. SIIT (Stateless IP/ICMP Translation) and NAT-PT (Network Address Translation - Protocol Translation) are the associated translation techniques. A blend of translation and the dual stack model, known as DSTM (Dual Stack Transition Mechanism), has been defined to allow for the case where insufficient IPv4 addresses are available. Like tunneling techniques, translation can be implemented in border routers. Such techniques are covered in Deliverable D2.3.1 [2], and are most usually applied at the site (university) networks, rather than in the NREN or core networks.

3 Scenarios for core transition

In this Section we consider a non-exhaustive list of steps or actions that a service provider might have to go through in order to provide an IPv6 service.

3.1 Definition of the requirements

Providing an IPv6 service at the customer level, it is necessary to find out which areas and customers are most likely to want IPv6 services, and then identify the access routers that can be upgraded to be dual-stack (a technique for running both IPv4 and IPv6 protocols in the same router) so as to provide both an IPv4 and IPv6 service to these customer sites.

Alternatively separate dual-stack access routers could be specified and installed to provide solely an IPv6 service, thus minimising the impact on the existing IPv4 services even further. Other activities consist of setting up a Domain Name Server (DNS) that supports both the existing IPv4 A records and the new IPv6 AAAA records, and, if there is a need for intercommunication between IPv6-only and IPv4-only hosts, operating one of the protocol translation mechanisms such as NAT-PT in the router or a TCP-UDP Relay. It is quite likely that in the context of an academic network such as GÉANT, the access routers are the national NREN PoPs, and that the translation methods would be run in the site (university) networks who connect to the national PoPs via tunnelled or native links.

Initially, these access routers should be interconnected over the existing IPv4 core routers or infrastructure using one of the available deployment strategies to carry IPv6 over IPv4: carrying IPv6 packets inside IPv4 packets (tunneling), running IPv6 over a dedicated Layer 2 technology (such as ATM), or forwarding IPv6 packets over Multiprotocol Label Switching (MPLS) backbones. The choice of deployment strategy will determine the choice of an IPv4 or IPv6 routing protocol.

For high-level service providers, a registration for their own IPv6 address prefix is useful/necessary using the relevant International Regional Internet Registry (RIR) Process. The intermediate and mid-level service providers should contact their high-level service provider. Alternatively, organisations wanting to connect only to the IPv6 6bone for testing before formal registration should apply for a prefix from the 6bone community.

In view of service continuity the IPv4 to IPv6 transition is not only an addressing or a routing issue. Available and emerging enhanced IPv4 services such as IP QoS, IP security, telephony over IP or multicast have to be continuously provided whatever the IP infrastructure might be.

3.2 Core Design

At first several core transition scenarios will be explained in detail followed by the different steps and requirements to proceed with the migration.

3.2.1 Deploying IPv6 over MPLS Backbones

IPv6 over MPLS backbones enables isolated IPv6 domains to communicate with each other over an MPLS IPv4 core network. This implementation requires far fewer backbone infrastructure upgrades and less reconfiguration of core routers because forwarding is based on labels rather than the IP header itself, providing a very cost-effective strategy for the deployment of IPv6.

Additionally, the inherent Virtual Private Network (VPN) and traffic engineering services available within an MPLS environment allow IPv6 networks to be combined into VPNs or extranets over an infrastructure supporting IPv4 VPNs and MPLS-TE. This of course assumes that MPLS is already available on the migrating network; deploying MPLS purely for IPv6 migration is not a recommended solution.

A variety of deployment strategies are available or under development, as follows:

- IPv6 using tunnels on the customer edge (CE) routers
- IPv6 over a circuit transport over MPLS
- IPv6 on the provider edge (PE) routers (known as 6PE)

The first of these strategies has no impact on and requires no changes to the MPLS provider (P) or PE routers because the strategy uses IPv4 tunnels to encapsulate the IPv6 traffic, thus appearing as IPv4 traffic within the network. The second of these strategies also requires no change to the core routing mechanisms. The last strategy requires changes to the PE routers to support a dual-stack implementation, but all the core functions remain IPv4. Table 1 summarizes the primary use, benefits, and limitations for each MPLS mechanism.

MPLS Mechanism	Primary Use	Benefits	Limitations	Requirements
IPv6 Using Tunnels on CE Routers	Enterprise customers wanting to use IPv6 over existing MPLS services.	No impact on MPLS infrastructure.	Routers use IPv4-compatible or 6to4 addresses.	Dual-stack CE routers.
IPv6 over a Circuit Transport over MPLS	Service providers with ATM or Ethernet links to CE routers.	Fully transparent IPv6 communication.	No mix of IPv4 and IPv6 traffic.	Cisco 12000 or 7600 Internet routers in the core.
IPv6 on PE Routers	Internet and mobile service providers wanting to offer an IPv6 service.	Low cost and low risk upgrade to the PE routers. No impact on MPLS core.	No VPN or VRF support currently planned.	IPv6 software upgrade for PE routers.

Table 1: MPLS Mechanisms: Primary Uses, Benefits and Limitations

IPv6 for Cisco IOS software currently supports the first two of these strategies, and Cisco has plans to support IPv6 provider edge routers in Phase II of its IPv6 for Cisco IOS software strategy.

A final strategy would be to run a native IPv6 MPLS core, but this strategy would require a full network upgrade to all P and PE routers, with dual control planes for IPv4 and IPv6.

The following sections describe each mechanism in more detail.

3.2.1.1 IPv6 Using Tunnels on the Customer Edge Routers

Using tunnels on the CE routers is the simplest way of deploying IPv6 over MPLS networks, having no impact on the operation or infrastructure of MPLS, and requiring no changes to either the P routers in the core or the PE routers connected to the customers.

Communication between the remote IPv6 domains uses standard tunneling mechanisms, running IPv6 over IPv4 tunnels in a similar way that MPLS VPNs support native IPv4 tunnels. The CE routers need to be upgraded to be dual-stack, and configured for IPv4-compatible or 6to4 tunnels, but communication with the PE routers is IPv4, and the traffic appears to the MPLS domain to be IPv4. The dual-stack routers use the IPv4-compatible or 6to4 address, rather than an IPv6 address supplied by the service provider. Figure 2 shows the configuration using tunnels on the CE routers.

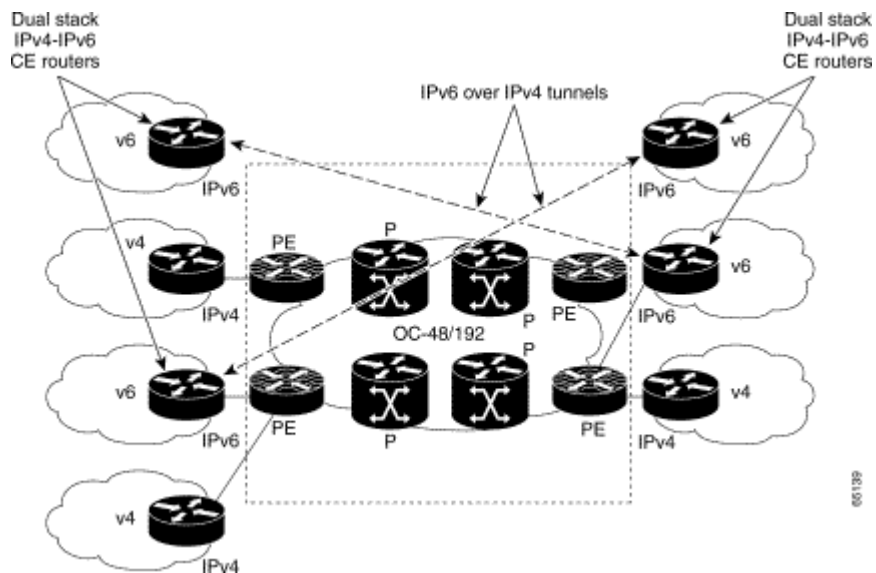


Figure 2: IPv6 Using Tunnels on the CE Routers

It is quite probable that the IETF will phase out or deprecate the use of IPv4-compatible addresses in the near future.

3.2.1.2 IPv6 over a Circuit Transport over MPLS

Using any circuit transport for deploying IPv6 over MPLS networks has no impact on the operation or infrastructure of MPLS. It requires no changes to either the P routers in the core or the PE routers connected to the customers.

Communication between the remote IPv6 domains runs native IPv6 protocols over a dedicated link, where the underlying mechanisms are fully transparent to IPv6. The IPv6 traffic is tunneled using Any Transport over MPLS (MPLS/AToM) or Ethernet over MPLS (EoMPLS), with the IPv6 routers connected through an ATM OC-3 or Ethernet interface, respectively. Figure 3 shows the configuration for IPv6 over any circuit transport over MPLS.




Figure 3: IPv6 over a Circuit Transport over MPLS

3.2.1.3 IPv6 on the Provider Edge Routers

A further deployment strategy is to configure IPv6 on the MPLS PE routers. This strategy has the advantage for service providers in that there is no need to upgrade either the hardware or software of the core network (P router), and it thus minimises the impact on the operation of or the revenue generated from the existing IPv4 traffic. The strategy maintains the benefits of the current MPLS features (for example, MPLS or VPN services for IPv4) while appearing to provide a native IPv6 service for enterprise customers (using ISP-supplied IPv6 prefixes). Figure 4 shows the configuration for IPv6 on the PE routers.




Figure 4: IPv6 on the Provider Edge Routers

The IPv6 forwarding is done by label switching, eliminating the need for either IPv6 over IPv4 tunnels or for an additional Layer 2 encapsulation, and allowing the appearance of a native IPv6 service to be offered across the network. The core network continues to run MPLS and any of the Cisco IOS software-supported IPv4 interior routing protocols, eliminating the requirement for upgrades to the hardware for native IPv6 forwarding and allowing the network to continue with current proven releases of Cisco IOS software.

Each PE router that must support IPv6 connectivity needs to be upgraded to be dual-stack (becoming a 6PE router) and configured to run MPLS on the interfaces connected to the core. Depending on the site requirements, each router can be configured to forward IPv6 or IPv6 and

IPv4 traffic on the interfaces to the CE routers, thus providing the ability to offer only native IPv6 or both IPv6 and native IPv4 services. The 6PE router exchanges either IPv4 or IPv6 routing information through any of the supported routing protocols, depending on the connection, and switches IPv4 and IPv6 traffic using the respective fast switching path (either Cisco Express Forwarding (CEF) or distributed CEF (dCEF) for IPv4 or CEF or dCEF for IPv6) over the native IPv4 and IPv6 interfaces not running MPLS.

The 6PE router exchanges reachability information with the other 6PE routers in the MPLS domain using multiprotocol BGP, and shares a common IPv4 routing protocol (such as Open Shortest Path First (OSPF) or i/IS-IS) with the other P and PE devices in the domain.

The 6PE routers encapsulate IPv6 traffic using two levels of MPLS labels. The top label is distributed by the label distribution protocol (LDP) used by the devices in the core to carry the packet to the destination 6PE using IPv4 routing information. The second or bottom label is associated with the IPv6 prefix of the destination through multiprotocol BGP-4.

The 6PE architecture allows support for IPv6 VPNs. Refer to the Internet-Draft *draft-ietf-ngtrans-bgp-tunnel-02.txt* for further information on 6PE routers.

3.2.2 Deploying IPv6 Using Dual-Stack Backbones

Using dual-stack backbones is a basic strategy for routing both IPv4 and IPv6. All routers in the network need to be upgraded to be dual-stack. IPv4 communication uses the IPv4 protocol stack (with forwarding of IPv4 packets based on routes learned through running IPv4-specific routing protocols), and IPv6 communication uses the IPv6 stack with routes learned through the IPv6-specific routing protocols.

The key requirements are that each site has an IPv6 unicast global prefix and appropriate entries in a DNS that map between host names and IP addresses for both IPv4 and IPv6. Applications choose between using IPv4 or IPv6 based on the response from the DNS resolver library, with the application selecting the correct address based on the type of IP traffic and particular requirements of the communication.

Today, dual-stack routing is a valid deployment strategy for specific network infrastructures with a mixture of IPv4 and IPv6 applications (such as on a campus or an aggregation point of presence), requiring both protocols to be configured. However, apart from the obvious need to upgrade all routers in the network, limitations to this approach are that the routers require a dual addressing scheme to be defined, require dual management of the IPv4 and IPv6 routing protocols, and must be configured with enough memory for both the IPv4 and IPv6 routing tables. One should also consider requirements on additional services such as Multicast and QoS.

3.2.3 Different steps for the migration

3.2.3.1 Router performance

A first task consists of evaluating the performance and interoperability of various router platforms.

The parameters to consider include:

- Forwarding performance in dual stack mode at line rate.

The idea is to measure the level of IPv6 traffic a router can forward (and with which kind of CPU consumption), and compare that to estimated IPv6 traffic load (multiplied by some factor for safety). This evaluation can be done in a laboratory with a simple IPv4/IPv6 setup to get a first idea. Then, a more complex test, which takes into account the other types of forwarding mechanisms, which are currently in production, can be achieved.

The results expected are that there is no impact on the router's memory and that forwarding performance is close to the estimated traffic load.

- IGP tests

According to which IGP the production network is running, tests can be done to measure the performance of the routers in the case where several IGPs are running in parallel for both stacks of protocols and compared with one IGP handling the both stacks.

The results expected are that there is stability and no impact on the router's memory.

- Tunneling methods

Tunnels are part of the technical solutions for transitioning the network. One strategy could consist in having only the edge routers enabled dual stack.

From edge to edge a tunnel IPv6/IPv4 can be established, core routers stay pure IPv4.

Tunneling and en/de-capsulation performance have to be evaluated, the parameters to look at are again forwarding performance and memory usage.

- Interoperability

Interoperability tests have to be done for a backbone based on multi-vendor platforms. Tunneling techniques and routing protocols have to interoperate.

3.2.3.2 Field trial and test-bed

Once the product evaluation is achieved a laboratory can be set up to evaluate the best design for the production network.

Based on the current architecture, several designs can be evaluated for provisioning the IPv6 service and transitioning the network.

The best design will be the one that achieves the smoothest transition and provides the best stability.

3.2.3.3 Monitoring and troubleshooting

The introduction of IPv6 on the backbone implies having monitoring and troubleshooting tools in place.

The basic tool set would be a DNS server for the resolution of IPv6 addresses and router names.

Telnet (ssh), TFTP/FTP, Ping and Traceroute have to be available for IPv6 on the routers.

In the perspective of having a pure IPv6 backbone, monitoring tool platforms like HP-Openview, Infovista etc. have to be ported and evaluated in dual stack mode and pure IPv6. Other useful applications like netflow monitoring and SNMP polling (Cricket, MRTG) have to be evaluated.

These issues are covered in 6NET WP6, and reported on by the D6 deliverable set [3,4].

3.2.3.4 Migration

After the field trials another set of tests can be organized to evaluate the steps for the migration from the current architecture to the selected design.

Based on these tests a road map is established for the deployment.

This road map can be built in two major phases.

- The migration of the core
- The connection of the customers

3.2.3.5 Connection of the users

As soon as the core is ready for delivering IPv6 services the connection of the users can start.

This second phase can be planned on a short or long period, it depends on how many people (NRENs) there are to be connected.

Each client has to be compliant to the agreed service specification and some procedures have to be defined to control the compliance of the customers during their connections.

At this stage, the backbone is enabled to deliver IPv4/IPv6 service. It can be fully dual stack or partially dual stack, in some cases (an ATM backbone for example) it can even remain IPv4 only.

4 Considerations for mechanism selection

The mechanism selections are closely related to the features available on the routers and the services already in production on the network.

Therefore dual stack routers and tunneling mechanisms can be deployed if performances are acceptable.

Mechanisms	Benefits	Limitation	Migration
Dual stack backbone	Native IPv4/IPv6 service	Router performance Routing protocols adaptability	Heavy effort for the migration
Partial dual stack backbone (Only edge routers) and tunneling at the edge	Core routers will be less stressed	Router performance for forwarding performance and tunneling	Flexible migration in different phases
VPN layer 3	Core routers will be less stressed and optimised with MPLS forwarding	Interesting if a VPN layer 3 service is already deployed in the network	Heavy effort for the migration if layer 3 VPNs don't exist

VPN layer 2	<p>Completely transparent to the network</p> <p>Easy to set up if the functionality is available (backbone already ATM or router's features available)</p>	Postpone the migration of the core network to IPv6	Transparent
-------------	--	--	-------------

5 Conclusions

This deliverable gave an overview about IPv4 to IPv6 transition mechanisms, which can be deployed in backbone networks to allow a smooth integration of IPv6. Which transition mechanism will be chosen depends on the customer requirements and current design. But on the other hand, if there is a choice to deploy one or the other mechanism in a backbone network, field trials are useful and necessary to evaluate which of them is scalable enough, is compliant with other services required, etc.

In view of the introduction of IPv6 in GÉANT, it will depend on the trial field test for which migration strategy and transition mechanism the core backbone will go for a migration from a pure IPv4 network to a core backbone providing IPv6 service.

According to our experiences, research networks backbones are upgraded every three or four years. That means that the next generation networks might have to consider the problem on the other way around and design its native IPv6 network to support IPv4 services.

That why we recommend that one study area we can pursue is to perform tests for the proposed scenarios to be prepared for the migration of GÉANT from IPv4 to integrated IPv4 and IPv6. However, we will also focus on combined studies of transition for core and NREN (ISP) networks in the upcoming joint deliverables in M12 (D2.1.2 and D2.2.2), which will be the first version of a transition cookbook for NREN and core networks.

6 References

- [1] 6NET Deliverable 2.2.1: "IPv4 to IPv6 migration scoping report for end site networks/universities"
- [2] 6NET Deliverable 2.3.1: "IPv4 to IPv6 migration scoping report for organisational (NREN) networks"
- [3] 6NET Deliverable 6.1.1: "6NET Network Management Architecture"
- [4] 6NET Deliverable 6.2.1: "Management and monitoring tools requirements and specifications"