


IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

Project Number:	IST-2001-32603
Project Title:	6NET
CEC Deliverable Number:	32603/SURFnet/DS/3.1.1/A1
Contractual Date of Delivery to the CEC:	February 28th 2002
Actual Date of Delivery to the CEC:	March 2002
Title of Deliverable:	IPv6 Routing Plan for the 6NET Network
Work package contributing to Deliverable:	WP3
Type of Deliverable*:	R
Deliverable Security Class**:	PU
Editors:	Wim Biemolt and Ronald van der Pol
Contributors:	Bruno Ciscato, Havard Eidnes, Alexander Gall, Janos Mohacsi, Martin Potts, Pekka Savola, Gunter van de Velde and Wilfried Woeber

* 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:

This document describes the initial routing and naming setup of the 6NET core network. Chapter 1 describes the topology of the 6NET core network and the local loops to the NREN PoPs. Chapter 2 describes the number plan for the IPv6 prefix used in the 6NET project. The convention for the domain names of the equipment is also described. Chapter 3 describes the IGP routing setup and describes the setup and tuning parameters of IS-IS. Chapter 4 describes the EGP routing setup. The setup of both iBGP and eBGP is described. Finally, chapter 5 describes the initial network management setup regarding access, security and configuration templates.

Keywords:

IPv6, IS-IS, iBGP, eBGP, addressing

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

Executive Summary

This document describes the initial configuration of the 6NET core network. It describes the topology, which consists of a ring between the UK, FR, CH, IT, DE, SW, and back to the UK. NL is connected to both the UK and DE, AT is connected to both CH and DE. GR is connected initially to DE.

RIPE has allocated IPv6 address space to DANTE, and DANTE has made part of this address space available to the 6NET project for use in the core network. All participants have their own address space allocated to them by RIPE.

Each 6NET PoP has its own address range for use on links, LANs, etc. There is also a naming convention for the equipment used in the PoPs, which basically consists of the two letter country codes with *6net.org* appended to it.


The IS-IS routing protocol will be used for both IPv4 and IPv6 as the Interior Gateway Protocol for distributing the addresses of the core links. There will be some initial tuning of the IS-IS parameters to achieve a fast convergence when there is an outage. iBGP will be used to distribute the NREN IPv6 prefixes within the core network. eBGP will be used to exchange IPv6 prefixes between the 6NET core and NRENs and between the 6NET core and external organisations.

Finally, the initial setup for management is described. At each core PoP, the 6NET router will be connected to the GEANT router to provide inband IPv4 access to the 6NET router. The 6NET router will also be connected to the GEANT out of band network. An initial configuration template is also described for the routers.

All the associated configuration commands are given for the 6NET core Cisco routers.

Table of Contents

1.	Introduction.....	4
2.	Topology Overview.....	5
2.1	6NET Core PoP.....	5
2.2	6NET NREN PoP.....	6
2.3	Local Loop.....	6
3.	Addressing and Naming.....	6
3.1	Naming.....	6
3.1.1	PoP Naming.....	6
3.2	Addressing.....	7
3.2.1	6NET PoP Addressing.....	8
3.2.2	SLA addressing.....	10
3.3	Point-to-point links between 6NET PoPs.....	13
4.	IGP Routing.....	14
4.1	Choice of routing protocol.....	14
4.2	6NET IS-IS High Level View.....	14
4.2.1	Introduction.....	14
4.2.2	Implementation details.....	15
4.3	6NET CLNS addressing.....	21
4.3.1	Convention.....	21
5.	EGP Routing.....	23
5.1	Choice of routing protocol.....	23
5.2	6NET BGP High Level View.....	23
5.3	6NET General BGP Implementation Details.....	24
5.4	6NET iBGP Peering.....	28
6.	6NET Core Management.....	29
6.1	Initial Setup.....	30
6.2	Configuration Template.....	32
7.	Conclusion.....	34
8.	Glossary.....	35

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

1. Introduction

This Deliverable describes the IPv6 routing plan for the initial network configuration used to rollout 6NET. It focuses especially on the 6NET core network. It does not cover the design and implementation considerations required for connecting the local NREN to the NREN PoP.


This document is complemented by Deliverable D1.1: “Design Specification of 6NET”.

2. Topology Overview


2.1 6NET Core PoP

The 6NET core network spans several European countries. From each of these 6NET core routers, there is connectivity (local loops) towards the 6NET NREN PoP routers. These local loops and the NREN PoP routers are not visualised in the picture below.

Initially, the connectivity utilised to connect the 6NET core routers is based on POS STM1 interfaces. It is however intended to upgrade the line speed of some currently undefined links to higher capacity during the second year of the project. A detailed description of the 6NET core topology and links can be found in the WP1 design document (Deliverable D1.1), which will be produced at the end of March 2002.



The dotted line indicates Layer-2 VPN connectivity via the GEANT network, this is not visible to the 6NET core routers.

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

2.2 6NET NREN PoP

Each NREN PoP router is connected by means of a local loop to a 6NET core router of the Pan-European network.

The NREN PoP router aggregates the traffic of universities and/or NRENs towards the 6NET core network. The NREN PoP router falls under the NREN management and this is different from the management and co-ordination domain of the core network. The precise management arrangement will be defined in the framework of WP6.

2.3 Local Loop

The local loop is the connection between the 6NET core router and the NREN PoP router.

The local loop specification is part of the aforementioned WP1 design document (D1.1).

3. Addressing and Naming

3.1 Naming

The domain name *6net.org* was registered on behalf of 6NET. It is used to name all equipment that is part of the 6NET core network.

3.1.1 PoP Naming

Every PoP has its own subdomain within *6net.org*. The subdomain name corresponds to the two letter country code of the country where the PoP is located, i.e. <cc>.6net.org. The country codes are the following:

- at - Austria
- be - Belgium
- ch - Switzerland
- cz - Czech Republic

- de - Germany
- es - Spain
- fr - France
- gr - Greece
- hu - Hungary
- ie - Ireland
- it - Italy
- lu - Luxemburg
- nl - Netherlands
- pl - Poland
- pt - Portugal
- se - Sweden
- si - Slovenia
- sk - Slovakia
- uk - United Kingdom


3.2 Addressing

The Regional Internet Registries (RIRs) are currently allocating /35s and RIPE has allocated '2001:0798::/35' to DANTE. From this address space, DANTE has made '2001:0798::/40' available to the 6NET project.

The following hierarchy is used:

3 bits	26 bits	5 bits	8 bits	16 bits
FP	RIR	project	PoPs	SLA
		/35	/40	/48
				/64

For 6NET, the 5 bit <project> part consists of all zeroes. Other <project> allocations will be made by DANTE.

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network</p>	
----------------	---	---

To summarize, the assigned IPv6 address consists of the following parts:

Structure = <xyz>.<PoP>.<SLA>

Where:

<xyz> = <40 bit provider prefix> = 2001:0798::/40

<PoP> = assigned address range per PoP

<SLA> = segmentation of use within the PoP


SLA =Site Level Address

3.2.1 6NET PoP Addressing

The 6NET locations that have a 6NET core router get their own prefixes. The 8 bits for the PoP part are used as follows:

- 00. core links
- 01. reserved core links

- 10. AT
- 11. BE
- 12. CH
- 13. CZ
- 14. DE
- 15. ES
- 16. FR
- 17. GR
- 18. HU
- 19. IE
- 20. IT
- 21. LU
- 22. NL
- 23. PL
- 24. PT

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---


- 25. SE/NORDUNET
- 26. SI
- 27. SK
- 28. UK

Note: The start of the country PoP areas started at 10 for visual reasons when reading the IPv6 prefix addresses. The hexadecimal addresses A-F are not used.

This gives the following prefixes for the PoPs:

PoP Location	IPv6 PoP addressing: 2001:0798:<PoP>::/48
Core	2001:0798:00::/48
Sweden	2001:0798:25::/48
Netherlands	2001:0798:22::/48
Germany	2001:0798:14::/48
Austria	2001:0798:10::/48
Italy	2001:0798:20::/48
Switzerland	2001:0798:12::/48
France	2001:0798:16::/48
UK	2001:0798:28::/48
Greece	2001:0798:17::/48

The only 'special' address range, which is not really bound to a geographical location is the 'Core' address range. This address range is used for the connections between PoPs (see chapter 3.3).

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

3.2.2 SLA addressing

The SLA (Site Level Aggregate) is used for various prefixes, (mostly) within a PoP:

Range	use
0000 - 00FF	Loopback addresses
0100 - 01FF	Intra-PoP point-to-points
0200 - 02FF	connections to NREN PoPs
0300 - 03FF	external 6NET connectivity
0400 - 04FF	PoP LANs

Note: Using this convention there is room for 256 prefixes with /64 address space for each point-to-point link or broadcast media.


The prefix length selected for point-to-point connections will be '/64'. A '/64' was chosen because it seems to be the best current practice. It make the number plan easy because every interface gets a '/64'. There is no need for a '/126', '/127' or '/128' because address conservation is not a deciding factor with IPv6.

The potential future use of a '/127' address space (or other address space) which initially seems to use only minimal address range will not work in the long run. (More study material can be found in <http://www.ietf.org/internet-drafts/draft-savola-ipv6-127-prefixlen-01.txt>).

It is also possible to use '/128' on point-to-point links. The drawback here is that in Cisco IOS you need to add a static route to the remote side.

For Switzerland with <PoP>=12 this would give the following prefixes:

Loopback addresses:	2001:0798:0012:0000::/56
Intra-PoP Point-to-points:	2001:0798:0012:0100::/56
Connections to NREN PoP:	2001:0798:0012:0200::/56
External 6NET connection:	2001:0798:0012:0300::/56
PoP LANs:	2001:0798:0012:0400::/56

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

3.2.2.1 6NET Loopback addresses

In networking environments, it is seen as good practice to give each device an IPv6 loopback address. This is an IPv6 address that is not directly assigned to any physical interface and will typically be reachable when the networking appliance (in this case a router) is up and running.

This loopback address is also used for operational and management actions on the equipment, and for routing protocols like eBGP, which use these addresses for terminating the peering sessions.

Loopback addresses typically have a prefix mask of '/128'. This avoids unnecessary unused addresses although address conservation is not really an issue in IPv6.


In the initial 6NET core network, the 6NET PoP routers are the only pieces of equipment that need a loopback address. Below is a list which details the addresses used.

The <SLA> for the loopback addresses is '0000'.

6NET PoP location	6NET PoP address space 2001:0798:<PoP>:<SLA= 0>::/64	IPv6 loopback address
Sweden	2001:0798:25::/64	2001:0798:25::1/128
Netherlands	2001:0798:22::/64	2001:0798:22::1/128
Germany	2001:0798:14::/64	2001:0798:14::1/128
Austria	2001:0798:10::/64	2001:0798:10::1/128
Italy	2001:0798:20::/64	2001:0798:20::1/128
Switzerland	2001:0798:12::/64	2001:0798:12::1/128
France	2001:0798:16::/64	2001:0798:16::1/128
United Kingdom	2001:0798:28::/64	2001:0798:28::1/128
Greece	2001:0798:17::/64	2001:0798:17::1/128

3.2.2.2 Intra-PoP point-to-points

Intra-PoP point-to-point links are numbered from the PoP prefix with a <SLA> = '01xx', where 'xx' is a sequence number. This allows for 256 point-to-point links per PoP with a '/64' prefix.

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

The table below gives an example of intra-PoP point-to-point prefixes which have a <SLA>=01xx for the Germany PoP with <PoP>=14:

2001:0798:0014:0100::/64
2001:0798:0014:0101::/64
2001:0798:0014:0102::/64
2001:0798:0014:0103::/64

3.2.2.3 6NET PoP to NREN PoP point-to-points


The 6NET core routers are connected to the NREN routers as seen previously in this document. A special address range is selected for this type of point-to-point connectivity:

2001:0798:<PoP>:<SLA=02xx>::/56, where 'xx' is a sequence number.

The addresses utilised on the point-to-point links between the 6NET core and the NREN PoP have a prefix-length '/64'. Initially there is only one NREN PoP router attached to a 6NET PoP router.

The host part of the address is '::1' for the 6NET core PoP side, and '::2' for the NREN PoP side.

6NET PoP location	6NET PoP address space	1st Point-to-Point prefix 2001:0798:<PoP>:<SLA=0200>::/64
Sweden	2001:0798:0025:0200::/56	2001:0798:0025:0200::/64
Netherlands	2001:0798:0022:0200::/56	2001:0798:0022:0200::/64
Germany	2001:0798:0014:0200::/56	2001:0798:0014:0200::/64
Austria	2001:0798:0010:0200::/56	2001:0798:0010:0200::/64
Italy	2001:0798:0020:0200::/56	2001:0798:0020:0200::/64
Switzerland	2001:0798:0012:0200::/56	2001:0798:0012:0200::/64
France	2001:0798:0016:0200::/56	2001:0798:0016:0200::/64
United Kingdom	2001:0798:0028:0200::/56	2001:0798:0028:0200::/64
Greece	2001:0798:0017:0200::/56	2001:0798:0017:0200::/64

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network</p>	
----------------	---	---

3.2.2.4 6NET PoP to external networks

Currently, interconnectivity is planned to the 6BONE, 6TAP and the WIDE network.


The connectivity will be provided by the 6NET participants, and will provide transit to the whole 6NET project. Interconnection to 6BONE is also planned in London with IPv6 over IPv4 tunnels. For the initial setup this does not have to be taken into account.

3.3 Point-to-point links between 6NET PoPs

Inter-PoP point-to-point links are numbered from the Core prefix (2001:0798:00::/48) with a <SLA> = '00xx' , where 'xx' is a sequence number

For the initial rollout of 6NET, this results in following table:

Connectivity Between	IPv6 prefix	IPv6 address side (1)	IPv6 address side (2)
UK (1) - FR (2)	2001:0798:0:1::/64	2001:0798:0:1::1/64	2001:0798:0:1::2/64
FR (1) - CH (2)	2001:0798:0:2::/64	2001:0798:0:2::1/64	2001:0798:0:2::2/64
CH (1) - IT (2)	2001:0798:0:3::/64	2001:0798:0:3::1/64	2001:0798:0:4::2/64
IT (1) - DE (2)	2001:0798:0:4::/64	2001:0798:0:4::1/64	2001:0798:0:4::2/64
DE (1) - NL (2)	2001:0798:0:5::/64	2001:0798:0:5::1/64	2001:0798:0:5::2/64
NL (1) - UK (2)	2001:0798:0:6::/64	2001:0798:0:6::1/64	2001:0798:0:6::2/64
UK (1) - SE (2)	2001:0798:0:7::/64	2001:0798:0:7::1/64	2001:0798:0:7::2/64
SE (1) - DE (2)	2001:0798:0:8::/64	2001:0798:0:8::1/64	2001:0798:0:8::2/64
DE (1) - AT (2)	2001:0798:0:9::/64	2001:0798:0:9::1/64	2001:0798:0:9::2/64
AT (1) - CH (2)	2001:0798:0:a::/64	2001:0798:0:a::1/64	2001:0798:0:a::2/64
DE (1) - GR (2)	2001:0798:0:b::/64	2001:0798:0:b::1/64	2001:0798:0:b::2/64

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

4. IGP Routing

4.1 Choice of routing protocol

The options available for an IGP routing protocol within 6NET are: static routing, RIPv6 or IS-IS. The option 'static routing' would be very difficult to manage in practice, and would not be very scalable. The dynamic routing options available are 'RIPv6' and 'IS-IS'.

The first routing protocol is a distance vector routing protocol while 'IS-IS' is a link-state protocol. Although a distance vector routing protocol is easier to troubleshoot and the operation simpler to understand, it is preferred to utilise a link-state protocol due to its advantages in convergence, tuning and additional features (like opaque information, enhanced TLV (Type/Length, Value) information for Traffic Engineering, etc.).

To benefit from the advanced functionality that a link-state protocol gives, the choice for IPv6 IGP routing protocol within 6NET has fallen on 'IS-IS'. This decision was taken by WP1, WP3 and WP6 participants cooperatively.

4.2 6NET IS-IS High Level View

4.2.1 Introduction

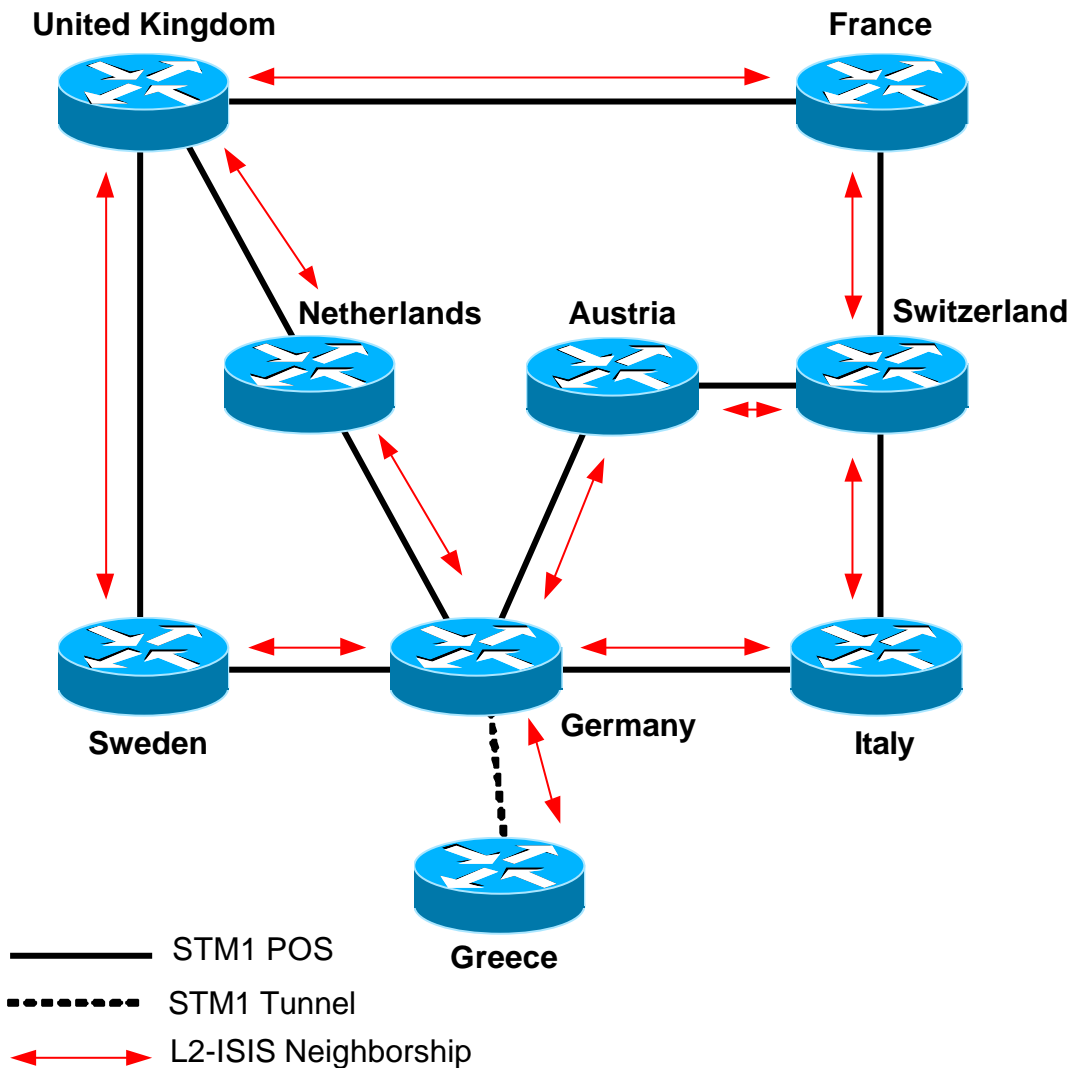
The operation of IS-IS is described in ISO10598. This standard specifies an OSI routing protocol for CLNS traffic. 'IS-IS' is a link state protocol with a two level hierarchical structure. With the use of TLVs, information is exchanged between two Intermediate Systems (= Routers).

IETF RFC1195 extended the ISO CLNS with support for IPv4. This type of 'IS-IS' is also known as Integrated or Dual IS-IS (abbreviated as 'I/IS-IS'). To make use of 'I/IS-IS' as the routing protocol to route IPv4 packets, it is required to configure CLNP parameters and the IPv4 parameters on the Intermediate Systems (= routers). This is required because 'I/IS-IS' runs directly on top of the CLNS datalink layer, which is different from protocols like RIP, OSPF and EIGRP (which use IP).

The addition of IPv6 as a protocol supported by 'I/IS-IS' is not fully standardised yet. The IETF document describing the current implementation is 'draft-ietf-isis-ipv6-02.txt'.


4.2.2 Implementation details

A 1000 mile Overview



One Single Level-2 Domain

The hierarchical structure of IS-IS enables us to utilise two different levels of routing, allowing the routing protocol to scale through summarisation. The 6NET core network initially exists of a small amount of routers and prefixes. This small amount justifies the use of a single level of routing. Typically a Level-2 backbone domain is configured since it has good scaling potential towards the future. Therefore, within 6NET a pure Level-2 routing domain is implemented.

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

To configure this on the Cisco equipment the following command lines are required:

```
interface POS2/0
  isis circuit-type level-2-only
!
router isis
  is-type level-2-only
!
```

Authentication

Within the IS-IS routing protocol there are three basic password authentication mechanisms possible. The only authentication used is a password to authenticate the IIH packets.

Authenticating the IS-IS Link State Packets either for Level-1 or for Level-2 routers will not be utilised within 6NET. The reason for this decision is that the operational aspects of password changes are not practical.

The selected authentication can be defined per interface. The initial selected password is '6NET_IIH_Password'. This initial password is used only in staging and initial deployment. Afterwards, it will be changed.


Interface Metric

In order for a routing protocol to decide the best path from a one network to another it is important to assign a cost or metric to each interface. Typically, the metric is a direct correlation between value and the speed of the link. Other correlations can be delay, reliability, etc. Initially, the link speed is used for the 6NET core.

To support the higher metric values placed on the interfaces, it is required on Cisco equipment to change the metric style from 'narrow' to 'wide'. This wide style metric is supported by all major router vendors that support I/IS-IS.

Metrics are assigned to interfaces and are configured as follows:

```
interface POS6/0
  isis metric 100 <level-1/level-2>
```


IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network</p>	
----------------	---	---

The metric style is configured at the general "router isis" command as follows:

```
router isis
 metric-style wide
```

The following metric convention will be used in 6NET:

Link Speed (Mbps)	IS-IS Metric
1	10000
10	1000
100	500
155	400
1000	300
2500	200
10000	100

Passive Interfaces


When a router interface is placed in the routing process it will send out IIH packets in order to form an IS-IS neighbourhood with peering IS-IS routers over that interface.

There are instances where it is desirable to include an interface in the IS-IS routing process to distribute the IPv6 prefix among the Level-2 area, but to never form an adjacency over that interface. To disable the sending of IIH packets the interface can be made passive with the command <passive-interface> under the IS-IS routing process. It is recommended to do this on LAN stub network interfaces.

In 6NET, all loopback interfaces and interfaces connecting to the NREN PoP routers and LAN stub interfaces will be made passive.

Example:

```
router IS-IS
 passive-interface loopback 3/0
 passive-interface pos 3/0
```

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

IS-IS Tuning

IS-IS can be run with default routing process values. However, it is advisable to tune certain parameters at the initial deployment of the network to achieve more optimal convergence and bandwidth utilisation.

A few preliminary remarks:

- Cisco IOS time granularity is 4 msec.
- This means that even if a command allows you to configure less than 4 msec, it is absolutely not guaranteed at all that such interval will be honoured.
- When playing with fast convergence you want to react very fast to "bad news" (link going down) and you can afford to wait a bit longer for "good news" (link coming up). However, the current backoff algorithm doesn't take into account the type of the change.

It has to be noted that not all described tuning parameters are available in all Cisco IOS versions. An overview of initial recommendation is put in next paragraphs.

Router Parameters


These values need to be identical across all routers in the IS-IS domain. Some of the tuning commands are new to the Cisco implementation. It needs to be verified if the IOS running on the routers supports changing the defaults.

`<lsp-refresh-interval>` - specifies the time in seconds the router will wait before refreshing and transmitting its own LSPs. The default is 900 seconds (15 minutes). This is to make sure that LSPs are refreshed before the `max-lsp-lifetime` expires. The recommendation is to increase this timer to 65000 seconds (18 hours). Make sure that the `lsp-refresh-interval` is lower than the `max-lsp-lifetime` in order for the LSP never to age-out.

The initial value for 6NET will be 65000 seconds and will be configured as follows:

```
router isis
  lsp-refresh-interval 65000
```

`<max-lsp-lifetime>` - specifies the maximum lifetime in seconds specified in the LSP header. Routers use this timer to age-out and purge old LSPs. The default is 1200 seconds (20 minutes). For

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

6NET the recommendation is to increase this timer to the maximum of 65535 seconds (~18.7 hours). This will decrease the number of unnecessary LSP re-flooding.

The initial value for 6NET can be configured as follows:

```
router isis
  max-lsp-lifetime 65535
```

The parameters below use 3 backoff algorithm timers. The syntax is

```
spf-interval <a> [<b> <c>]
  <a> seconds between SPF runs (seconds)
  <b> milliseconds between first trigger and SPF
  <c> milliseconds between first and second SPF
```

Same syntax for: `prc-interval` and `lsp-gen-interval`


The first timer `<a>` is the maximum interval, which is the maximum amount of time the router will wait between consecutive executions. The second timer `` is the initial delay, which is the time the router will wait before starting execution. The third timer `<c>` is the incremental interval, which is the time the router will wait between consecutive executions. This timer is variable and will increase until it reaches `maximum-interval`.

`<spf-interval>` - specifies the minimum time in seconds between two consecutive SPF computations (Dijkstra). The default is 5 seconds. A SPF computation is necessary when the topology changes. Increasing this timer in a large network where topology changes are frequent can increase stability but will decrease the network convergence rate. Decreasing the timer can increase the network convergence rate. However, the SPF computations are CPU intensive and a small `spf-interval` in a network with flapping links (ie. links that go down, come up, go down, come up, etc. or that otherwise change their status all the time) would cause havoc on the router CPU.

Due to the relatively small size of 6NET, the timers can be configured quite aggressive:

```
spf-interval 1 40 80
```

`<prc-interval>` - specifies the minimum time in seconds between two consecutive Partial Route Calculations (PRC). The default is 5 seconds. PRC is similar to SPF and therefore also CPU intensive.

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

PRC is used when just some leaf routes (i.e.: IP prefix) have changed on one or more LSPs. Then we just flag that particular LSP for recalculation and instead of re-running SPF we just recompute all leaf routes of that particular LSP. So, in general this involves a few LSPs and doesn't take that much time. We can afford here to have much more aggressive timers than in SPF but still we can't predict how many LSPs we will have to recalculate:

```
prc-interval 1 40 20
```

`<lsp-gen-interval>` - specifies the minimum interval in seconds at which LSPs are generated. The default is 5 seconds. When a topology change occurs in a router, the router will generate a LSP and flood it. A chain reaction occurs forcing all routers in the area to perform a SPF computation. A flapping link can increase the number of LSPs sent and, as a result, SPF computations.

In 6NET, it is desirable for each router to generate a new LSP as soon as something changes.

```
lsp-gen 1 20 1000
```

The following commands are recommended configuration across the whole IS-IS domain:


`<ignore-lsp-errors>` - LSP packets received with a bad checksum are purged by default and the initiator must regenerate the LSP. This command allows the receiving router to ignore those LSPs. A link corrupting packets could cause a large number of purging and re-generating of LSP in a network.

`<log-adjacency-changes>` - causes IS-IS to generate a log message when an adjacency changes state (up or down).

`<no-hello-padding>` - allows hello packets to be sent without padding the rest of the packet to MTU size. This substantially reduces the amount of bandwidth used by IIS.

The commands listed above are configured as follows:

```
router isis
  ignore-lsp-errors
  log-adjacency-changes
  no-hello-padding
```

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

Interface Parameters

The IS-IS interface parameters will remain default during the initial rollout.

There is one POS interface parameter that has impact on the convergence of the routing protocol but is not really an IS-IS tuning parameter.

`<carrier-delay>` - A value we should look at is the carrier-delay on the POS interfaces. The default is 2 seconds. So if we wait 2 seconds that an interface is down before informing IS-IS, the tuning won't do much good. Advice here would be to set a short carrier-delay to detect link down events fast. Value of about 100 msec would be recommended ("`carrier-delay msec 100`"). That way there is a fast reaction to interface down situations.

If it is expected that the interfaces are to be physically unstable for periods shorter than it takes for the routing protocol to converge, you may want to put this value higher, and that will effectively "filter out" these interface transitions as far as the routing protocol is concerned.

4.3 6NET CLNS addressing

4.3.1 Convention


Within a routing domain with IS-IS as its IGP, it is a prerequisite that each Intermediate System has a CLNS address, that can be used as identifier for the LSP packets.

The CLNS address has the following format:

3 byte	6 byte	1 byte
<domain>	<System ID>	<NSEL>

<domain>

The domain will consist of a 1 byte AFI (Address Format Identifier) of '49' (= CLNS AFI for private addressing) and a 2 byte postfix of '0001'. So the domain is '49.0001'.

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

<System ID>

Each router in a routing area needs a unique 6 byte system ID in order for proper routing operation. The system ID chosen consists of a 2 byte PoP ID followed by a 4 byte sequence number.

Using this convention the system ID for the first router in the PoP Germany would be:

<System ID> = <00014.0000.0001>

<NSEL>

In 6NET, this is always '00'.

6NET CLNS Addresses

The 6NET core network consists of 9 routers that form the pan-European backbone. The following CLNS addresses are placed on the 9 routers:

Country	<domain>	<System ID>	<NSEL>
Sweden	49.0001	0025.0000.0001	00
Netherlands	49.0001	0022.0000.0001	00
Germany	49.0001	0014.0000.0001	00
Austria	49.0001	0010.0000.0001	00
Italy	49.0001	0020.0000.0001	00
Switzerland	49.0001	0012.0000.0001	00
France	49.0001	0016.0000.0001	00
United Kingdom	49.0001	0028.0000.0001	00
Greece	49.0001	0017.0000.0001	00

5. EGP Routing

5.1 Choice of routing protocol

In 99.9% of the interconnections between IPv4 Autonomous Systems, BGP (Border Gateway Protocol) is used. For the interconnection of Autonomous Systems in IPv6 the same protocol is used with some updated exchange and database structures to support the IPv6 address format.


A clear protocol overview is described in the following RFCs:


- RFC 2858: Multiprotocol Extensions for BGP-4
- RFC 2545: Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing

Currently BGP+ is the only EGP supported for IPv6 when using Cisco IOS code. New EGP protocols may be supported as the standards emerge.

5.2 6NET BGP High Level View

The Autonomous System number of 6NET is '6680'. Each connecting NREN will use (one of) its own Autonomous System identification number. An overview of these can found in the picture.



IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

Within the 6NET Autonomous System, all 6NET PoP nodes form fully meshed iBGP sessions. More implementation details can be found in a later chapter. Between the 6NET Autonomous System and the NREN, an eBGP session runs to exchange IPv6 routes between the Autonomous Systems.

Autonomous System Overview table:

NREN PoP	Autonomous System Number
United Kingdom	786
France	2200
Switzerland	559
Italy	137
Germany	680
Netherlands	1103
Nordunet	2603
Austria	1853
Greece	5408


5.3 6NET General BGP Implementation Details

This chapter contains BGP implementation guidelines that have common aspects between the iBGP and eBGP routing strategies.

Router-ID

When starting up a BGP process in a IPv6 only environment, it is required to define a router-id on the Cisco router. The router-id has the same format as an IPv4 IP address.

For the 6NET environment, the BGP router-id is identical to the initial IPv4 address placed on the Ethernet port for the initial management of the equipment. An overview of the initial router-ids is found in the following table:

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

6NET PoP router location	BGP router-id
Austria	62.40.98.14
Switzerland	62.40.98.78
Germany	62.40.98.142
France	62.40.98.206
Greece	62.40.98.238
Italy	62.40.99.78
Netherlands	62.40.99.142
Sweden	62.40.99.238
United Kingdom	62.40.100.78

The IOS command to set the <router-id> are under the BGP router process:

```
router bgp 786
  bgp router-id 62.40.100.78
```

Authentication

To secure the neighbourhood, it can be useful to use password authentication from the initial start of 6NET. This is done on the iBGP sessions and on the eBGP sessions. The goal is not to secure the data send over the BGP tcp sessions, but to avoid the formation of malicious peerings by accident.

The following initial passwords have been elected:


- For eBGP sessions: External_BGP
- For iBGP sessions: Internal_BGP

These initial passwords will be changed after initial deployment.

NLRI Filtering

During the initial rollout of 6NET there is no NLRI filtering.

This results in a not very secure initial deployment. However it is the responsibility of the NCC (WP1, WP3, WP6) to define secure policies that can be referred to as workable and offering

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

sufficient security to achieve production grade security as soon as possible. Installing initial policies without understanding project expectations on security could potentially break some research work and hence no NLRI filtering will happen except for the filtering required to summarise the 6NET prefix.

Peer group usage

Often, in a BGP speaker, many neighbours are configured with the same update policies (that is, the same outbound route maps, distribute lists, filter lists, update source, and so on). Neighbours with the same update policies can be grouped into peer groups to simplify configuration and, more importantly, to make updating more efficient.

The usage of peer group functionality has two main advantages:

- Simplify the configuration and enable easier reading of the configuration file
- Optimise the performance on Cisco routers

With this motivation the concept of peer group is used in the 6NET environment.


Two initial peer groups are defined initially:

- Peer group for peerings between 6NET core PoP and NREN PoP:
 - 6NET_EXTERNAL_NREN_PEER
- Peer group for peerings between 6NET core PoPs:
 - 6NET_INTERNAL_6NET_PEER

6NET Prefix aggregation

The complete 6NET IPv6 address range can be aggregated to '2001:0798::/40'. However, due to the aggressive filtering policies on the IPv6 Internet (see RFC 2772), initially the address range is aggregated to '2001:0798::/35'. To do this type of aggregation in a controllable way, some NLRI filtering is required. To enable the aggregation, 3 basic steps are required:

- Configure static IPv6 route for '2001:0798::/35' to the null0 interface
- Redistribute the static summary route to the IPv6 BGP table (this will require NLRI filtering in order to redistribute only the summary prefix)
- Use BGP NLRI filtering to the iBGP peers to exclude the summary prefix from being exchanged

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

From the explanation above, we see that two initial filters definitions are required:

- **STATIC_TO_BGP_FILTER_ROUTE_MAP:** This filter is used during the redistribution phase when redistributing from static IPv6 routes to the BGP process
- **IBGP_PEER_FILTER_ROUTE_MAP:** This filter is placed on the iBGP neighbourship to exclude the summary route of being redistributed through iBGP sessions. This filter prevents multiple summary IPv6 prefixes in the BGP routing table

To aid in the initial filtering there is also a need for initial prefix-lists:

- **ANY_PREFIX:** This prefix-list is used when there is requirement to match any prefix
- **STATIC_TO_BGP_FILTER:** This prefix-list is used to explicit match those prefixes that require to be blocked between internal BGP peers

The template to add the NLRI filter to the neighbourship is:


```
neighbour <iBGP neighbour/Peer-group Name> route-map IBGP_PEER_FILTER_ROUTE_MAP
out
```

The template to add the filtering to the redistribution is:

```
redistribute static route-map STATIC_TO_BGP_FILTER_ROUTE_MAP
```

The template for the initial filter route-maps and prefix-lists is:

```
route-map IBGP_PEER_FILTER_ROUTE_MAP deny 10
  match ipv6 address prefix-list STATIC_TO_BGP_FILTER
!
route-map IBGP_PEER_FILTER_ROUTE_MAP permit 20
  match ipv6 address prefix-list ANY_PREFIX
!
route-map STATIC_TO_BGP_FILTER_ROUTE_MAP permit 10
  match ipv6 address prefix-list STATIC_TO_BGP_FILTER
  set origin igp
!
route-map STATIC_TO_BGP_FILTER_ROUTE_MAP permit 20
  match ipv6 address prefix-list ANY_PREFIX
  set origin igp
!
!
ipv6 route 2001:0798::/35 Null0
!
!
ipv6 prefix-list ANY_PREFIX seq 5 permit ::/0 ge 1
!
ipv6 prefix-list STATIC_TO_BGP_FILTER seq 5 permit 2001:0798::/35
!
```

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

BGP Synchronisation

If your Autonomous System is passing traffic through from another Autonomous System to a third Autonomous System, it is very important that your Autonomous System is consistent about the routes that it advertises. For example, if your BGP were to advertise a route before all routers in your network had learned about the route through your IGP, your Autonomous System could receive traffic that some routers cannot route yet. To prevent this from happening, BGP must wait until the IGP has propagated routing information across your Autonomous System. This causes BGP to be synchronised with the IGP. Synchronisation is enabled by default on a Cisco router.

In some cases, you do not need synchronisation. If you will not be passing traffic from a different Autonomous System through your Autonomous System, or if all routers in your Autonomous System will be running BGP, you can disable synchronisation.

Disabling this feature will allow you to carry fewer routes in your IGP and allow BGP to converge more quickly. In 6NET it is recommended to disable synchronisation. Use the following command in router configuration mode:

```
!  
router bgp 786  
  no synchronization  
!
```


5.4 6NET iBGP Peering

The Internal BGP peering happens via full mesh neighbourships between the initial nine 6NET PoP routers. The peering sessions are built 'to and from' the IPv6 loopback interfaces configured on each router.

The only filtering initially performed is to prevent the aggregated IPv6 prefix to be exchanged between iBGP peers.

Route Reflector

Typically in ISP environments BGP Route Reflector functionality is utilised to enhance future scalability and to ease the logical topology. For the initial 6NET rollout this functionality is not used. It will be investigated further later in the project.

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---


Next-hop-self

The Next-hop-self functionality will not be used in the initial 6NET rollout. The initial 'next-hop' attribute checking of the NLRI is based on the IPv6 address from the NREN to which 6NET is peering.

6NET eBGP Peering

The eBGP peering is set up directly between the 6NET PoP router and the NREN PoP router. To establish the eBGP peering, the global aggregate addresses of the point-to-point link are used. This implementation does not require the 'eBGP multihop' feature because BGP session termination points are directly connected.

A general peering session is visualised in following picture:



Within 6NET, each NREN PoP has its own Autonomous System number. This was described in the 'BGP high Level View' Chapter.


Route Dampening

During the initial rollout no route flap dampening is implemented. It will be investigated further later in the project.

6. 6NET Core Management

During the initial deployment of the 6NET core, a marginal Network Management configuration is applied. One of the tasks of the 6NET project is to work on the NMS tools for IPv6 as part of WP6. The initial management applied is direct console access to the 6NET PoP routers through a terminal server and IPv4 telnet access to the management Ethernet port of the route processors on the 6NET core routers.


The detailed design of the 6NET core PoPs can be found in the WP1 design document. A representation of the initial management is shown in the picture below:



6.1 Initial Setup

IPv4 addressing

For management and monitoring purposes, the 6NET core routers are reachable via GEANT via IPv4. For that purpose all Cisco 1240x routers have an IPv4 address on their management Ethernet interface. The IPv4 address is assigned by DANTE from their current IPv4 backbone allocation.

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

Management interface IPv4 addresses (note: This IPv4 address is also used for the BGP router-id):


et0.at6.at.6net.org	62.40.98.14/29
et0.ch6.ch.6net.org	62.40.98.78/29
et0.de6.de.6net.org	62.40.98.142/29
et0.fr6.fr.6net.org	62.40.98.206/29
et0.gr6.gr.6net.org	62.40.98.238/29
et0.it6.it.6net.org	62.40.99.78/29
et0.nl6.nl.6net.org	62.40.99.142/29
et0.se6.se.6net.org	62.40.99.238/29
et0.uk6.uk.6net.org	62.40.100.78/29

To be able to communicate via the management Ethernet port on the route processor, it is required to place a default route pointing towards the OSS (Operational Support Systems) segment. This IPv4 default route destination is always the first IPv4 address in the respective /29 network.

The IPv4 DNS server addresses also need to be configured.

The initial IPv4 addresses are:

193.63.211.19
130.59.10.30
130.59.1.30

IST-2000-32603	<p style="text-align: center;">Deliverable D 3.1.1</p> <p style="text-align: center;">IPv6 Routing Plan for the 6NET Network</p>	
----------------	--	---

6.2 Configuration Template

The following standard generic configuration template is applied during initial configuration. This configuration may change once 6NET is up and running.

```

version 12.0
service password-encryption
!
hostname <name_of_router>
!
boot system flash slot0:
enable secret 5 $1$<secret>
!
username manager password <removed>
username operator privilege 2 password <removed>
!
ip subnet-zero
ip cef distributed
ip domain-list 6net.org
ip name-server 193.63.211.19
ip name-server 130.59.10.30
ip name-server 130.59.1.30
no ip source-route
no service tcp-small-servers
no service udp-small-servers
no service finger
no cdp running
no ip bootp server
!
ip route 0.0.0.0 0.0.0.0 <next hop IP address>
!
ipv6 unicast-routing
!
banner motd ^C

```

This is cc6.cc.6net.org, a 6NET Router in City, cc

Warning : Unauthorized access to this equipment is strictly forbidden
and will lead to prosecution


```

^C
privilege exec level 2 show configuration
privilege exec level 1 show
!
line con 0
  exec-timeout 15 0
  login local
  transport input none
  escape-character 27
line aux 0
  exec-timeout 15 0

```



```
login local
escape-character 27
line vty 0 4
exec-timeout 15 0
login local
escape-character 27
!
end
```

IST-2000-32603	Deliverable D 3.1.1 IPv6 Routing Plan for the 6NET Network	
----------------	---	---

7. Conclusion

This document has described the 6NET core topology, the configuration and tuning parameters of IS-IS and iBGP and eBGP configurations. It also defined a DNS naming convention for the 6NET PoPs. Finally, it showed the initial network management setup.

8. Glossary

DNS:	Domain Name System
eBGP:	external Border Gateway Protocol
iBGP:	internal Border Gateway Protocol
IIH:	IS-IS Hello packet
IS-IS:	Intermediate System - Intermediate System
LSP:	Link State Packet
NLRI:	Network Layer Reachability Information
NREN:	National Research and Education Network
PoP:	Point of Presence
POS:	Packets Over SDH/SONET
PRC:	Partial Route Calculations
RIPE:	Reseaux IP Europeens
RIR:	Regional Internet Registry
SLA:	Site Level Aggregate
SPF:	Shortest Path First
TLV:	Type/Length Value