www.euro6ix.net

Title:	Document Version:
Deliverable D2.2 **Specification of the Backbone Network Architecture**	1.7

Project Number:	Project Acronym:	Project Title:
IST-2001-32161	Euro6IX	European IPv6 Internet Exchanges Backbone

Contractual Delivery Date:	Actual Delivery Date:	Deliverable Type* - Security**:
31/12/2002	25/05/2005	R – PU

* Type: P - Prototype, R - Report, D - Demonstrator, O - Other

** Security Class: PU- Public, PP – Restricted to other programme participants (including the Commission), RE – Restricted to a group defined by the consortium (including the Commission), CO – Confidential, only for members of the consortium (including the Commission)

Responsible and Editor/Author:	Organization:	Contributing WP:
Carlos Ralli	TID	WP2

Authors (organizations):

Peter Hovell (BT), Cesar Olvera (Consulintel), Jordi Palet (Consulintel), Alain Baudot (FTRD), Francisco Fontes (PTIN), Carlos Parada (PTIN), Sathya Rao (Telscom), Alberto Escolano (TID), Antonio Lucientes (TID), Cristina Peña (TID), Raffaele D'Albenzio (TILAB), Mario Morelli (TILAB), Olaf Bonness (T-Systems), Antonio Skarmeta (UMU), Tim Chown (UoS), Javier Sedano (UPM), Tomás de Miguel (UPM), Juan Carlos Valverde (Vodafone).

Abstract:

This deliverable describes different solutions in order to implement the Euro6IX IPv6 european backbone.

Includes de decisions taken by the project regarding the Euro6IX routing policy.

Keywords:

Backbone, Euro6IX, IPv6, Routing Policies.

Revision History

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

Revision	Date	Description	Author (Organization)
v0.1	20/09/2002	Document creation	Carlos Ralli Ucendo (TID)
v0.2	27/09/2002	Added comments discussed in Torino meeting	Carlos Ralli Ucendo (TID)
v0.3,4,5	09/12/2002	Internal Changes due to document structure re-thinking	Alberto Escolano (TID) Antonio Lucientes (TID)
v0.6	13/12/2002	Changes to the main index , work distribution, TID contrib added as example and starting thoughts compiled from WP2 mailing list.	Alberto Escolano (TID) Antonio Lucientes (TID) Carlos Ralli Ucendo (TID)
v0.7	13/12/2002	New contrib added to 3.1.2 point.	Alberto Escolano (TID) Antonio Lucientes (TID) Carlos Ralli Ucendo (TID)
v0.8	18/12/202	Internal Changes, added point 1.2, point 2.3.3 completed, point 3.1.2 completed.	Alberto Escolano (TID) Antonio Lucientes (TID)
v0.9	05/02/2003	Added TID and TILAB contributtions	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.0	05/02/2003	Added BT and TILAB contributtions	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.1	05/02/2003	Added Examples of Routing Policy	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.2	05/02/2003	Added contributions.	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.3	05/02/2003	Added suggestions for Routing Policy. Added TILAB proposal.	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.4	05/02/2003	Integrated Peter Hovell routing document. FTR&D contribution added. Telscom corrections reviewed. Changed the structure of point 3.	Alberto Escolano (TID) Antonio Lucientes (TID)
v1.5	25/03/2003	Final TID revision and integration .	Alberto Escolano (TID)
v1.6	30/03/2003	Updates from last version.	Alberto Escolano (TID)
v1.7	20/03/2003	Added information about 6Bone phase-out, WIDE, JGN and SINET. Info about GÈANT. Graphics updaded. Ortographic and gramatical corrections. Final review and updates.	Jordi Palet (Consulintel)

Executive Summary

This deliverable represents the final result of the work done in the context of the A2.2 activity inside the WP2 in first year of the Euro6IX project.

The A2.2 activity has the main objective fixed in the design of the Euro6IX backbone network. Once the IX model to be implemented was chosen in the six months of the Euro6IX project, A2.2 covers the way to interconnect the different Euro6IX Exchanges (IXs), and the routing policies of the whole Euro6IX backbone. Deliverable D2.2 aims to give the technical basis needed for the design and deployment of the routing policies needed for the Euro6IX backbone.

After a presentation of several IPv6 backbones, a comparison between IPv4 and IPv6 backbones, and the explanation of the objectives of the Euro6IX backbone, several sections complete this document.

First section (Chapter 2: Study of IPv6 IXs Models) covers the architecture of the different IXs models proposed since the beginning of the Euro6IX project.

Second section (Chapter 3: Backbone Euro6IX Network Architecture) contains the information related to the implementation of the routing policy applied to the Euro6IX backbone.

At the end, the third section (Chapter 4: Interconnection with other IPv6 Backbones) is focused in the different connections between Euro6IX partners and external networks, such as 6NET, 6Bone, etc.

Table of Contents

Table of Figures

1. INTRODUCTION

The Euro6IX network is composed of several Internet Exchanges (IXs) connected by IPv6 native links. All these links, and all the IXs conform the Euro6IX backbone.

Since the beginning of the Euro6IX project, there are several IX models that have been proposed and discussed between all the partners. This document describes all the IX models proposed and the routing policies applied to the Euro6IX backbone.

There are other IPv6 backbones in addition to Euro6IX. Some of them are studied or commented in this document. Moreover, the interconnection between some of those IPv6 backbones and Euro6IX is studied and contemplated in this deliverable.

1.1 State of the Art: Study of IPv6 Large Backbones

For these studies there are several features to focus in:

- Backbone Organization.
- Connectivity and network scheme (layer 2 and layer 3, IXs, POPs).
- IPv6 addressing schemes used or tested.
- IPv6 routing policies.
- Interaction with other IPv6 networks.
- Available services.

The backbones studied in this document are: 6Bone, 6NET, Abilene and WIDE. The possible conclusion to this study can be the interests in the IPv6 technology for the different continental organizations.

1.1.1 Study of a Worldwide Initiative: 6Bone

6Bone is an experimental IPv6 network mainly based on top of portions of the physical IPv4-based Internet to support routing of IPv6 packets within an environment where this function is not yet integrated into the production routers. The network is composed of islands that, in general, can directly support IPv6 packets, linked mainly by virtual point-to-point links called "tunnels".

The idea to set up an experimental IPv6 backbone over the Internet was the result of a spontaneous initiative of several research institutes involved in the experimentation of the first implementations of the IPv6 protocol. The network became a reality in March 1996 with the establishment of the first tunnels between the IPv6 laboratories of G6 (France), UNI-C (Denmark) and WIDE (Japan).

The 6Bone is the place where the most interesting geographical experimentation on the IPv6 protocol is currently taking place. The experimental activities carried out inside the 6Bone have been coordinated by the IETF in order to provide feedback to various IETF IPv6-related activities and to IPv6 product developers based on test-bed experience.

1.1.1.1 Topology and Addressing

The 6Bone is structured as a hierarchical network (Figure 1-1) with backbone nodes, transit nodes and leaf nodes. The 6Bone backbone is made of a mesh of IPv6 over IPv4 tunnels (and some direct links) interconnecting backbone nodes only. IPv6 routing inside the backbone is based mainly on the BGP4+ protocol, and most of backbone nodes provide a full transit service. Transit nodes connect to one or more backbone nodes and provide transit service for leaf sites. Routing outside the backbone is mainly static but the number of non-backbone sites making use of IPv6 routing protocols such as BGP4+ or an IGP routing protocol (RIPng, IS-ISng and OSPFv3) is rapidly growing.

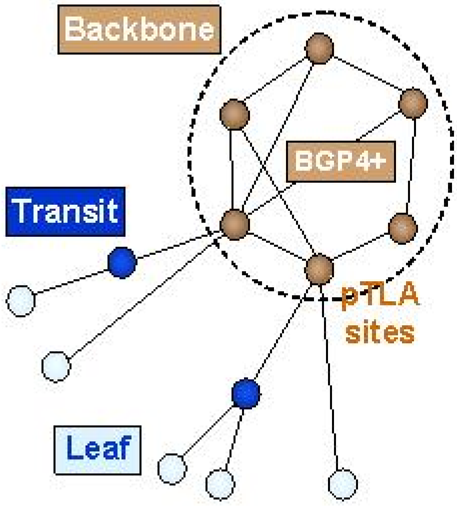

Figure 1-1: Logical Structure of the 6Bone Network

IPv6 addressing inside the 6Bone is based on the IPv6 Aggregatable Global Unicast Address Format and it matches the hierarchical topology described above. Backbone nodes play the role of experimental Top Level Aggregators (pTLAs, pseudo Top Level Aggregators) and they are responsible for assigning IPv6 addresses to non-backbone sites in such a way to establish an addressing hierarchy capable to enforce aggregation of routing information.

The whole 6Bone is identified by the IPv6 prefix 3ffe::/16 and every backbone node is assigned a 24, 28 or 32 bit long prefix (pTLA prefix) identifying an IPv6 addressing space which must be administered following all the rules defined for the TLAs. According to this model every pTLA plays the role of experimental top level ISP and has to assign chunks of its addressing space to directly connected transit and leaf sites without breaking aggregation inside the 6Bone backbone.

1.1.1.2 Network Growth

Since its creation in 1996 the 6Bone has been steadily growing in the number of connected sites (See Figure 1-2 and Figure 1-3). In July 1997 the network encompassed about 150 sites; in October 2002 more than 1200 sites distributed in 59 countries all over the world were officially registered in the 6Bone registry database. Over the same period of time the number of 6Bone backbone sites (i.e. assigned pTLAs) has increased from 36 to 133.

Date	Backbone Sites	All sites
7/2/1997	36	150
12/9/1997	43	203
3/31/1998	45	240
8/25/1998	47	302
12/12/1998	51	332
3/17/1999	55	361
7/17/1999	62	412
1/18/2000	67	505
3/21/2000	67	536
7/5/2000	68	590
12/6/2000	75	688
2/27/2001	82	759
12/10/2001	101	952
7/11/2002	126	1133
10/21/2002	133	1225

Figure 1-2: Growth of the 6Bone over the Last Five Years (Source: 6Bone Registry)

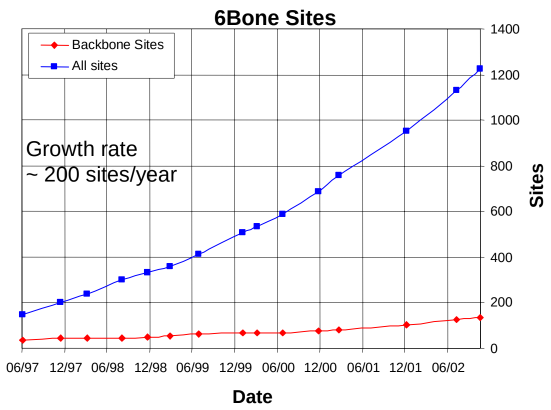

Figure 1-3: Graphical representation of 6Bone Growth (Source: 6Bone Registry)

1.1.1.3 6Bone Phase-out Plan

6Bone routing performance and stability is not always the best (e.g., see the Savola 6bone-mess I-D), and there is a clear need of a more robust core that can support test-beds.

In early 2002 discussions were started with the RIRs driven by two issues: Clarifying the role the 6Bone address registry has with respect to the RIRs IPv6 address registry and gaining access to the ip6.arpa reverse registry.

During the course of early discussions, the RIRs' management made it clear that they could not speak to the issue of how long the 6Bone allocation authority would last. Rather, it was an issue that the IETF and/or the IANA would have to deal with. To this end, a discussion was opened within the IETF on 6Bone phase-out planning.

Also there was the issue of the 6Bone operation. Even though the 6Bone came under the IETF ngtrans wg, it really had almost nothing to do with its operational policies. The 6bone community itself controls its policies and everyone expects that it will continue to do so.

Many comments came from the 6Bone community, but most relevant ones focus on having to pay for test-bed addressing when they haven't had to pay in the past. Note that many 6Bone participants (at all levels) do so to get experience and in the process convince their organizations that there is something worth paying for, i.e., the price is an issue, no matter how small it is.

There was concern about having to go through more complexity. It isn't clear if this is a real issue as we don't know what a pTLA-level request process might be with the RIRs. This may be a holdover of dislike of necessary procedures for scarce IPv4 address space.

Another concern is what is pay for service when the 6Bone is a volunteer effort… RIR services aren't needed. There is unwillingness to pay for service and then be expected to hand our free address services to downstream users.

Many comments have come from the RIR community as well: "…why should the 6Bone community get cheaper services than the dues paying members ?".

Also, the RIRs are supposed to recover costs for providing their services. Giving away any service would seem to go against this. A corollary to the above is, if the RIRs are just covering costs for a special service to the 6Bone, what are the RIRs doing for their regular customers. The feeling seemed to be, why should RIR members care about the 6Bone? Let 6Bone do their thing, and the RIRs theirs.

It isn't clear this proposal should proceed, given the opinions expressed on both sides, a soon to be in place 6Bone phase-out plan, a decline in the request rate for 6Bone prefixes, and a steady increase in allocated production prefixes.

Also, there is now the ability for the RIRs to temporarily allocate IPv6 addresses for Internet experiments.

As for e.f.f.3.ip6.arpa support, Bob Fink has proposed that the 6Bone operate the servers for this themselves, which would mean that the 6Bone community would sustain the cost of entering and maintaining the pTLA data in the e.f.f.3.ip6.arpa server, and that when phase-out is complete, the RIRs simply pull the eff3.ip6.arpa delegation and they have reclaimed it.

The RIRs have agreed that in light of the foregoing that there is no need to continue planning for a 6Bone RIR integration, and that the 6Bone would continue to manage its own allocations throughout the life of the phase-out plan.

The RIRs will delegate e.f.f.3.ip6.arpa to name servers that the 6Bone community provides.

Bob Hinden then presented the 6bone phase-out plan I-D, under the subtitle "be careful what you start". RFC2471 says 6bone addresses would be temporary addresses that would be reclaimed in the future (with implied renumbering for sites using 6bone addresses). The RIRs have been allocating production SubTLAs since 1999.

In 2002 more production allocations were made than 6Bone ones. The v6ops WG replaced ngtrans, which used to oversee the 6Bone, but the 6Bone is not in the v6ops WG Charter.

The current plan outlined in the draft-fink-6bone-phaseout-01.txt is to allocate 6Bone addresses until January 1, 2004, and for these to remain valid until June 6th 2006, after which no 6Bone prefixes should be carried on the Internet. The plan obsoletes RFC2471 that will become historic. It is up to IPv6 address holders to gain new address space, of whatever prefix length is appropriate (some pTLA holders may only require a site prefix, for example). In addition, IANA must not reallocate 3ffe:/16 for at least two years to avoid confusion with new allocations.

All these decisions were taken, finally, in the IETF-56, 6Bone BOF, chaired by Jordi Palet, from Euro6IX. Minutes available at http://www.ietf.org/proceedings/03mar/minutes/6bone.htm.

1.1.1.4 Conclusion

After years of work within the IETF, the standardization of IPv6 and the related components is coming to an end. Although at the same time the existing IPv4-based technology has been enhanced to partially cope with the problem of IP addressing space depletion as well as with the growing demand for new services like security, mobility and QoS, this has not removed but just delayed the need for a new network protocol as a long-term solution. In fact, it is quite clear that no low cost IPv4 patches will ever be able to guarantee the end-to-end network transparency and the huge amount of globally unique IP addresses required by the Internet evolution (e.g. the future xDSL or wireless data services based on the always-on paradigm). This is the real strong reason for deploying IPv6 within the Internet.

This is why most of major Internet ISPs are already looking with great attention at IPv6 and have been involved in the experimentation of the new protocol within the 6Bone for years now. Their contribution, together with the increasing effort coming from manufactures, universities and research centers from all over the world, is making the experimental IPv6 Internet growing fast. Nowadays, more than an environment to test IPv6 implementations, debug vendor equipment and make practice with it, the 6Bone and the ISP trials look like the down of the transition process.

Nevertheless, it is important to note that some other issues still need to be solved before a large-scale deployment of IPv6 within the Internet can take place. The 6Bone experience shows that multi-homing is still a problem, given that many of the un-aggregated IPv6 prefixes advertised within the BGP4+ cloud are due to lack of an alternative to the current inefficient IPv4 practice. Moreover, also the renumbering issue is worth further investigation and standardization effort.

Finally, to make IPv6 attractive for the users, and particularly for the new users who may foster the worldwide adoption of IPv6 by choosing the IPv6 only transition scenario, a suitable range of IPv6 applications must be available, starting from the basic services typical of the present Internet and Intranet environments. The present lack of such application services clearly indicates that the application developers and manufacturers will have a key role in breeding the transition process and making the new IPv6 Internet really happen.

1.1.2 Study of Asia Pacific Initiatives

There are several test-beds and production networks in Asia Pacific, that support native IPv6 connectivity, including dual stack approaches, and tunnels to 6Bone.

The older of these initiatives is the one started by the WIDE project, in Japan, started in August 1999. WIDE was connected to 6Bone even before, in June 1996.

It has connectivity to 6TAP, NSPIXP6 (Tokyo) and 6IIX/LA (Los Angeles, US). Under consideration is the connectivity to 6IIX/NY and S-IX (San Jose).

WIDE network uses AS2500 and prefers native IPv6 peering, whenever possible.

The 2nd Japanese network with offers IPv6 support is Japan Gigabit Network (JGN), developed by TAO (Telecommunications Advancement Organization), an authorized corporate body of the Ministry of Public Management, Home Affairs, Posts and Telecommunications.

JGN was started to conduct research and development of high-speed networking and high-performance application technologies such as the Next Generation Internet technology. JGN is easily accessible to both public and private entities for 5 fiscal years from 1999 to 2003, providing a wide range of research opportunities.

The last initiative, somehow equivalent to the European Research Networks, is SINET (Science Information Network), an information communication network dedicated to academic research. This network connects nationwide connection points (nodes) through high speed communication lines; at each node, facilities such as ATM switches and IP routers are installed.

SINET mutually connects with the Inter-Ministry Research Information Network (IMnet) and commercial Internet service providers to promote the international exchange of information as well as exchange of research data between the industrial, governmental, and academic sectors.

SINET started recently to support IPv6, initially with tunnels over IPv4.

1.2 Comparison between IPv6 and IPv4 Backbones

The IPv4 backbone example for this comparison is **GÉANT**. It is a four-year project, set up by a Consortium of 27 European national research and education networks (NRENs), with **DANTE** as its coordinating partner. Co-funded by the **European Commission** within the Fifth Framework Programme (FP5), its goal was the improvement of the previous generation of pan-European research network, TEN-155, by creating a backbone at Gigabit speeds.

Built using the latest DWDM technology operating at 10 Gbps, it will introduce new services covering a range of requirements, including guaranteed Quality of Service, Virtual Private Network, Security, Multicast, IPv6, etc.

An important element of GÉANT is the development of **connectivity with equivalent Research Networks in other world regions**. Connectivity is being consolidated with the existing equivalents of GÉANT in North America (Abilene, CA*net) and in Asia-Pacific (SINET, KOREN, SingAREN) and developed further between Europe and the Asia-Pacific, North American, South American and Mediterranean regions. Indeed, a total of three circuits of 2.5 Gbps, each one dedicated to research and education purposes, is currently being implemented in co-operation with North American research networks.

These developments will exploit much of the co-operative work done in the research networking community trough the next generation network task force (TF-NGN) and the works of development projects, such as SEQUIN, 6NET, etc.

GÉANT is based on a shared core network of nine 10 Gbps wavelengths. This is complemented by further twelve 2,5 Gbps wavelengths and additional SDH connectivity to those locations where wavelengths are not yet cost-effectively available. The core nodes where wavelengths are deployed, use Juniper M160 routers for provide traffic engine to IPv4 network. The planning team is responsible for the engineering of all network-related and for introduction of new services undergo a pilot phase. This team also leads the GÉANT Test Programme (TF-NGN) and its members have proven expertise in various fields such as Multicast and Multi-Protocol Label Switching.

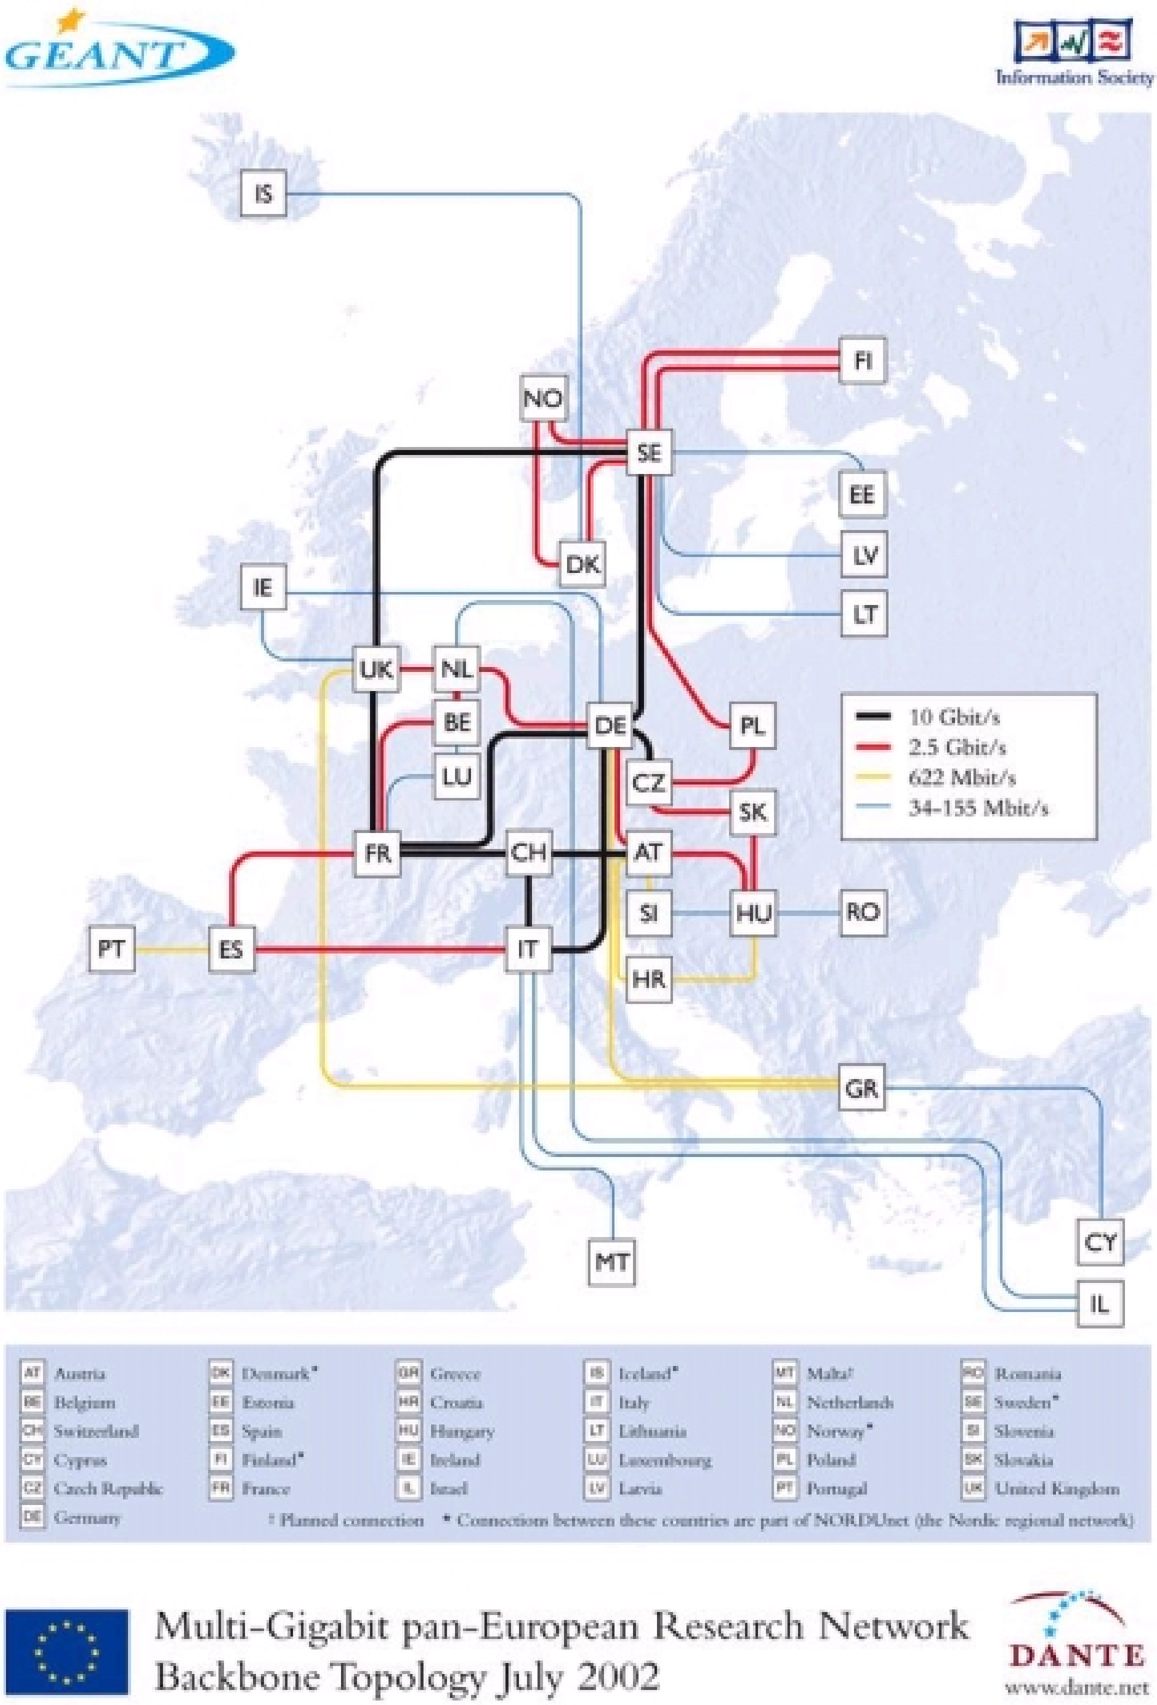

Figure 1-4: GEANT Backbone Topology

6NET is a three-year European project aiming to demonstrate that continued growth of the Internet can be met using new IPv6 technology. It also aims to help European research and industry play a leading role in defining and developing the next generation of networking technologies.

The 6NET objectives are install and operate an international IPv6 pilot network with both static and mobile components in order to gain a better understanding of IPv6 deployment issues. This network will primarily use native IPv6 links (initially running at 155 Mbps and increasing to 2.5 Gbps in the second year), although encapsulation over IPv4 infrastructure may be necessary in some cases. Some activities of the project will be to deploy the migration strategies for integrating IPv6 networks with existing IPv4 infrastructure, to introduce and test new IPv6 services and applications, to study legacy services and applications on IPv6 infrastructure, to evaluate address allocation, routing and DNS operation for IPv6 networks, to collaborate with other IPv6 activities and standardization bodies and to promote IPv6 technology.

Is important to note that GÉANT started moving to IPv6 in the third quarter of 2003, in part as a consequence of the results of the 6NET project, that prove the stability of the equipment, protocols and basic services needed to support IPv6 native (dual stack) production networks.

Indeed, several European NREN’s, including RedIRIS, Renater and FCCN, already support native IPv6 connectivity with the GÉANT backbone.

1.3 Objectives of Euro6IX Backbone

At the very highest level there are two types of geographically large networks:

- Homogeneous networks owned and operated by a single entity; Worldcoms global network being an example.
- Separate independently owned and operated networks that cooperate; the Internet being an example.

The Euro6IX network falls in the second category with a number of Internet Exchanges being owned and operated by separate organizations but which have agreed to cooperate to form a geographically larger network ie the Euro6IX network.

As such the Euro6IX project does not have a routed backbone network but rather a number of links that interconnect the co-operating Internet Exchanges. The Euro6IX network is therefore a consortium of co-operating Internet Exchanges interconnected by a number of links. Like any consortium there needs to be a set of guidelines that the partners adhere too. In the case of Euro6IX these guidelines needs to cover:

- Acceptable usage policy – to define the type of traffic to be exchanged.
- Peering policies – to control the route traffic takes.
- Interface guidelines – QoS, multicast, etc.
- Experiments – provider independent addressing etc.

2. STUDY OF IPV6 IXS MODELS

2.1 Introduction

One of the goals of the Euro6IX project is to analyze the different Internet Exchange models that can be adopted inside the Euro6IX network to better understand their main features and the reciprocal advantages.

In order to make a more detailed analysis, three different models (from now on called Model A, Model B and Model C) have been considered and analyzed. We started just from the classical model of the IX, described in detail in the following section, consisting of a simple switched architecture where the ISPs routers are connected. Then we considered newer models, where the IX itself can assign the addresses and the IX user can change the ISP without any change in the addressing plan, described in Model B and Model C sections.

However, inside the Euro6IX network, the model A won't be implemented because it does not introduce any new feature. It has been taken into account in this document because it could be interesting to compare it with the other models that will be implemented. Particularly, it could be interesting to compare the different routing frameworks and the adopted addressing policies to better understand the improvements and the advantages of the new models with reference to this well-known model.

2.2 Model A

The IX Model A, here described, can be considered as the traditional model of an Internet Exchange and by now it is the model more widely adopted in the Internet IPv4 community. The Figure 2-1 shows its internal architecture.

In this case, the IX is an "interconnection point" where the ISPs come to in order to exchange traffic between each other according to some defined routing policies.

Normally, inside the IX building, there is a layer 2 section (the dotted area that represents the neutral part of this IX model) consisting of a set of switches where the routers (belonging to the ISPs) are connected.

The routing policy is normally quite simple: Each ISP inside the IX belongs to a particular Autonomous System. The routers exchange the routing information using the eBGP protocol also deciding which kind of routing information to filter implementing the routing rules to guarantee (or not) the reachability of its own network. With this model, the IX does not have the possibility to assign the addresses and normally each ISP accessing the IX has its own internal addressing plan.

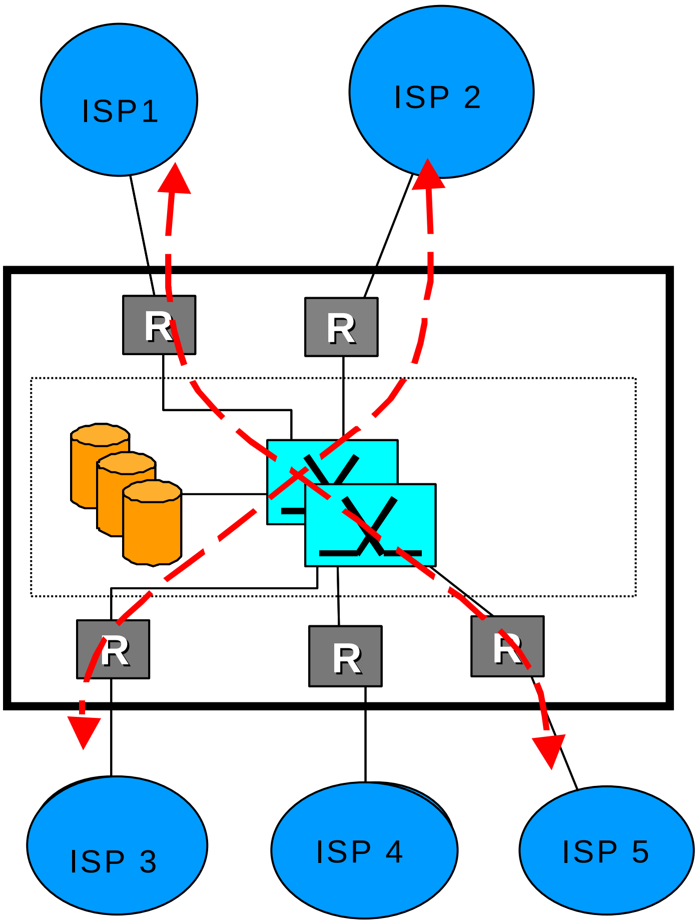

Figure 2-1: Architecture of IX Model A

2.3 Model B

2.3.1 IX Internal Network Architecture

The Internal Architecture for the IX Model B is shown in the Figure 2-2.

The Internal Architecture for the IX Model B in the figure above shows the infrastructure necessary to develop full redundancy inside the IX.

The different elements composing the IX Model B are the following:

- Layer 2 Infrastructure (L2) composed by SW1 and SW2, which provides fully redundant physical connectivity inside the IX. There are also several routers (Ra, Rb, Rc and Rd), necessary to establish peerings among them the different Telcos and/or ISPs. In this model, this is the neutral area of the Exchange, so its not operated by a neutral organization.
- Layer 3 Infrastructure (L3), which provides new functionalities to the IX compared with the traditional IX (Model A) and is composed by R1, R2, R3, R4 and optionally the IPv6 Route Server. L3 provides IPv6 connectivity among the equipment. This is the no-neutral area of the IX.
- Application Services: Available because of the Layer 3 capabilities of the IX. These services are, for example:
    - Basic Internet Services: NTP, DNS.
    - Content Delivery Services: HTTP, FTP.

- Network Access Services: RADIUS.
- Other services: POP3, SMTP, IRC.

- Monitoring Applications and Statistics Systems. The following one is a list of possible applications:
    - Routing Monitoring: (AS-Path tree).
    - Reachability Monitoring (Ping view).
    - Management Systems (Magalia).
    - Traffic Monitoring (Cricket, MRTG).
- Customers of the IX that can be divided in three main groups:
    - Other national Telcos connected to the Layer 2 IX Infrastructure.
    - Large Customers of the incumbent Telco connected to the Layer 3 IX Infrastructure.
    - Standard Customers connected to the Layer 3 IX Infrastructure through the national backbone of the incumbent Telco.
- Transit Providers, such as connections to the other Euro6IX IXs and connections to external IPv6 Networks (6Bone, 6NET, …).

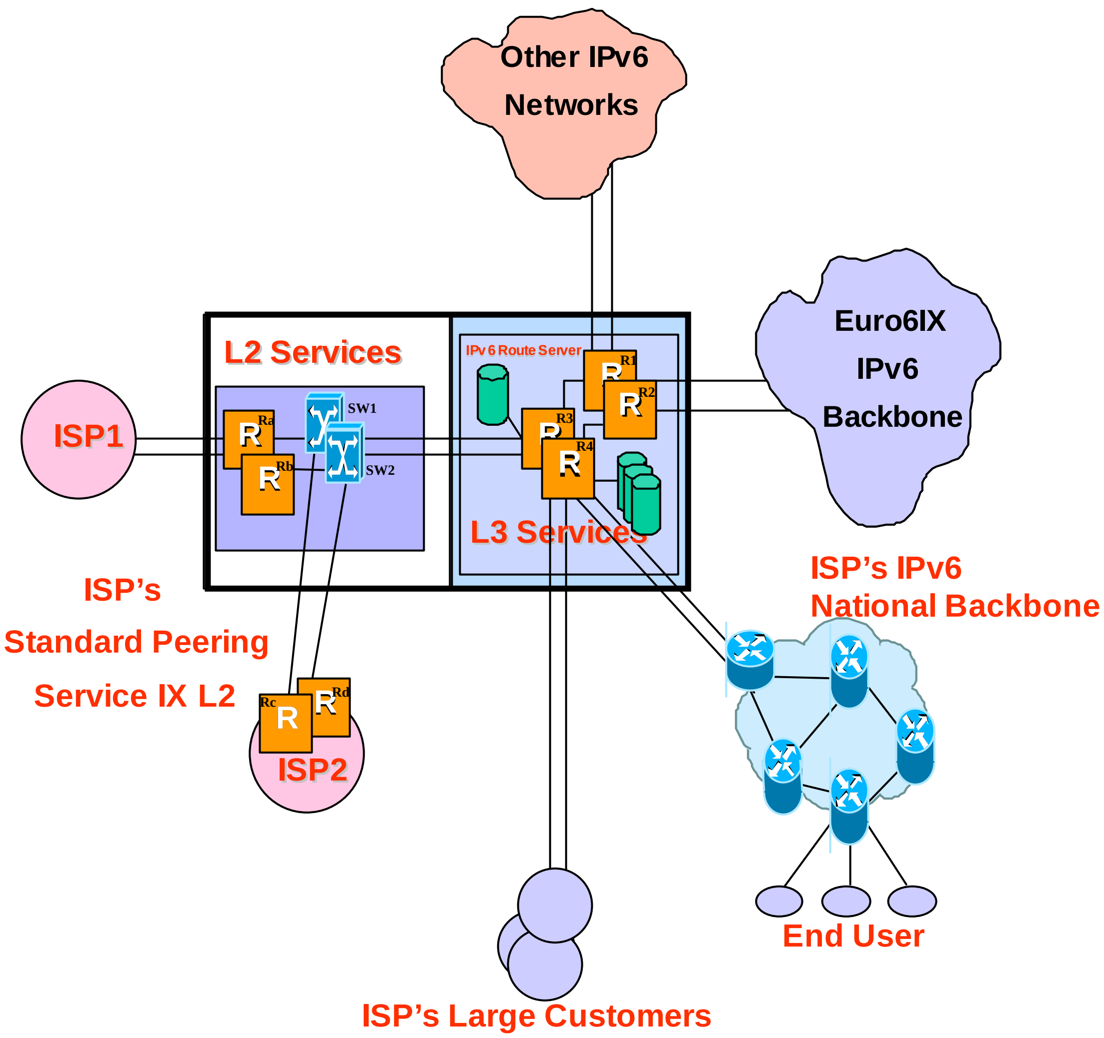

Figure 2-2: Architecture of IX Model B

2.3.2 IX Functionality

There are two differentiated areas according to the functionality in the IX Model B.

2.3.2.1 Layer 2 Infrastructure

The layer 2 infrastructure is based on the concept of the traditional IXs (Model A), where several ISPs are present in the same geographical area, and interested in peering agreements, will be connected. This infrastructure is the core of the IX and must be high performance, switched and fully redundant infrastructure, connecting all the equipment in the IX. To provide redundancy, each IX node should consist of two switches supporting local connectivity inside the IX.

Routers are used to establish peering between these ISPs.

The interest of a Telco in peering with another one is based on the amount of traffic exchanged between them. If they do not peer in the national IX, all the traffic exchanged between them will flow through their own international links. On the other hand, if they peer in the national IX, traffic will be exchanged there and thus transmission costs could be reduced.

This Layer 2 side must offer neutral services for those customers connected, and must be managed by a neutral organization with clear Statements and Operation Rules. All ISPs interested in being connected to the Layer 2 side must agree and follow these rules that preserve the correct operation of the neutral L2 node.

Finally, the development of the architecture of the IX consists in several stages regarding to the complexity of the infrastructure used:

- The first of these stages is developing the IX with one switch for the layer 2 infrastructure and one link among the equipment.
- The second stage will consist in improving the equipment of the IX by implementing redundancy, so that two switches should provide connectivity to the various blocks inside the IX.

2.3.2.2 Layer 3 Infrastructure

Layer 3 infrastructure (L3) provides new added functionalities to the traditional IX. With this topology, the IX works like an IPv6 Network Access Point (NAP) or like an IPv6 Point of Presence (PoP) of the owner. Layer 3 infrastructure consists, basically, of Router Equipment and optional elements, such as an IPv6 Route Server.

The Router Equipment of the IX owner provides IPv6 connectivity within the network inside the IX, as well as establishing peerings with other IXs and other IPv6 networks. Besides, it can provide IPv6 services to new L3 customers.

IPv6 Route Server is the equipment that centralizes the interchange of eBGP routes between the different elements in the IX. All the peers communicate only with the Route Server, instead of peering among each other, so it discharges Router Equipment to make that function, saving router resources.

2.3.3 Addressing Plan

The addressing plan for the Exchanges and for the whole Euro6IX network is in full conformance with the RFC2374 “An IPv6 Aggregatable Global Unicast Address Format”.

In the commercial phase, is expected that every IX Model B will have a commercial prefix delegated from 2001::/16. The aggregatable address format is designed to support long-haul providers, Exchanges, multiple levels of providers and subscribers, as shown in the following figure.

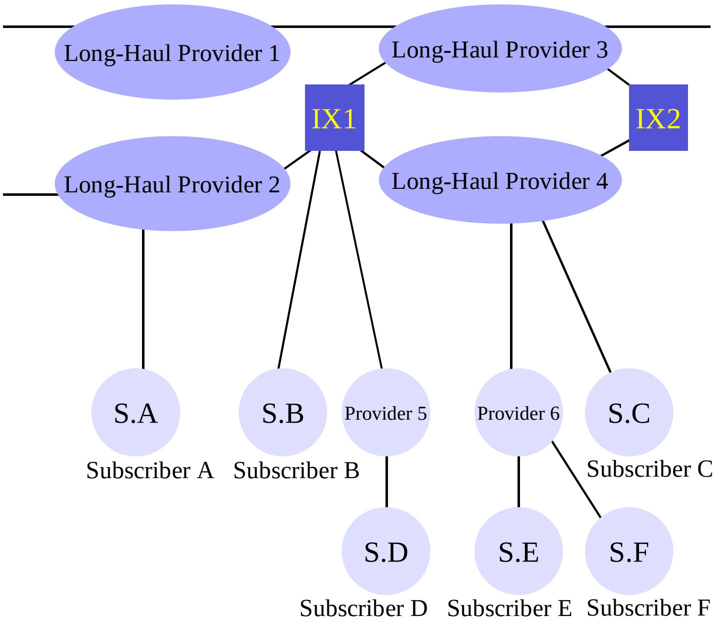

Figure 2-3: Public Topology Hierarchy

Exchanges will allocate IPv6 addresses. In IX Model B, this condition is fulfilled since the L3 IX is owned by one of the providers. This ISP assigns addresses to its customers delegated from its own 2001:xxxx::/32 prefix(es).

Based on RFC2374 statements a new concept to study and research on is that, considering that the delegated addresses belong to the ISP in charge of the IX, organizations that connect to these IXs will achieve addressing independence from long-haul providers. Then, they will be able to change long-haul providers without renumbering their organization. They can also be multihomed via the IX to more than one long-haul provider.

The next figure shows the Aggregatable Global Unicast Address Structure.

3	13	8	24	16	64 bits
FP	TLA ID	RES	NLA ID	SLA ID	Interface ID
Public Topology				Site Topology	Interface Identifier

Figure 2-4: Aggregatable Global Unicast Structure

The different fields of the previous figure are:

- FP: Format Prefix (001).
- TLA ID: Top-Level Aggregation Identifier.
- RES: Reserved for future use.
- NLA ID: Next-Level Aggregation Identifier.
- SLA ID: Site-Level Aggregation Identifier.
- INTERFACE ID: Interface Identifier.

The following sections specify each part of the IPv6 Aggregatable Global Unicast address format.

- Top-Level Aggregation ID

  Top-Level Aggregation Identifiers (TLA ID) are the top level in the routing hierarchy. The routing topology at all levels must be designed to minimize the number of routes into the routing tables.

- Reserved

  The Reserved field is reserved for future use and must be set to zero.

- Next-Level Aggregation Identifier

  Next-Level Aggregation Identifiers are used by organizations assigned a TLA ID to create an addressing hierarchy and to identify sites. The organization can assign the top part of the NLA ID in a manner to create an addressing hierarchy appropriate to its network.

  Each organization assigned a TLA ID receives 24 bits of NLA ID space. This space can be delegated to approximately as many organizations as the current IPv4 Internet.

Organizations assigned a TLA ID can provide service to organizations providing public transit service and to organizations that do not provide public transit service. The organizations receiving an NLA ID may also choose to delegate their space to another NLA ID's. This is shown in the following picture.

n	24-n bits	16	64 bits
NLA1	Site ID	SLA ID	Interface ID

m	24-n-m bits	16	64 bits
NLA2	Site ID	SLA ID	Interface ID

o	24-n-m-o bits	16	64 bits
NLA3	Site ID	SLA ID	Interface ID

Figure 2-5: Next-Level Aggregation Identifier

- Site-Level Aggregation Identifier

  The SLA ID field is used by an individual organization to create its own local addressing hierarchy and to identify subnets. It is a 16 bit field, so it supports 65.535 subnets.

  The approach chosen for structuring an SLA ID field is the responsibility of the individual organization. This is show in the next figure.

n	16-n bits	64 bits
SLA1	Subnet	Interface ID

m	16-n-m bits	64 bits
SLA2	Subnet	Interface ID

Figure 2-6: Site-Level Aggregation Identifier

- Interface ID

  Interface identifiers are used to identify interfaces on a link. They are required to be unique on that link.

2.3.3.1 Example of Addressing Plan (MAD6IX)

The Addressing Plan for MAD6IX has been made following the above guidelines. The addresses used have been delegated by Telefónica Data who owns a 2001:800::/32 prefix. The whole prefix for all Euro6IX networks connecting to MAD6IX is 2001:800:40:2000::/52, also including point-to-point links and MAD6IX networks.

2.3.3.2 Point-to-Point Links

The prefix 2001:800:2fXY::Z/126 has been used for the point-to-point links. The format for these addresses has been elected to correctly identify them, with the criteria explained below.

MAD6IX and every organization connecting to MAD6IX have been identified with an integer number. Every point-to-point link has the following structure:

- 2001:800:40:2fXY::Z/126.

The fields XY are designed according to the number assigned to the organizations, and in numerical order.

For example, MAD6IX has been assigned the number 0 and TID has been assigned the number 1. So the point-to-point link between MAD6IX and TID will be 2001:800:40:2f01::/126.

The criteria elected for the addressing of each side link (field Z) has been to use ::1 for MAD6IX side and ::2 for the other side.

For the point-to-point connections to other IXs, the prefix used has been 2001:800:40:2eXY::Z/126. Every IX connecting to MAD6IX and MAD6IX has been identified with an integer number. The fields XY are designed according to the number assigned to the IXs, and in numerical order. MAD6IX has been assigned number 0, LON6IX has been assigned number 2 and LIS6IX has been assigned number 3. The criteria elected for the addressing of each side link (field Z) has been to use ::1 for MAD6IX side and ::2 for the other side.

2.3.3.3 MAD6IX Attached Networks

For those networks connected to MAD6IX a ::/56 prefix has been delegated from the ::/52 available. So, the prefix delegation is as follows:

- MAD6IX Services Network: 2001:800:40:2000::/56
- MAD6IX Management Network: 2001:800:40:2100::/56
- TID Services Network: 2001:800:40:2200::/56
- TID Management Network: 2001:800:40:2300::/56
- TID Customers Network: 2001:800:40:2400::/56
- Consulintel Network: 2001:800:40:2a00::/56
- UPM Network: 2001:800:40:2b00::/56
- UMU Network: 2001:800:40:2c00::/56
- Vodafone Network: 2001:800:40:2d00::/56

All the above considerations are only suggestions and every IX administrator should choose their own address delegation architecture.

2.3.4 Services and Monitoring

The IX Model B provides large variety of services to offer to the customers of the Exchanges, due to the added layer 3 functionalities. Those application services are:

- Basic Internet Services: Synchronization Time Services (NTP) and Domain Name Service (DNS). These services were traditionally offered by the carrier's national backbones and ISPs, and now they can be delivered directly from the IX as an innovative issue.
- Content Delivery Services: Web Servers (HTTP), File Transfer Protocol Servers (FTP) and Video Streaming Servers. These services were traditionally offered by ISPs. The high demand by customers of these mentioned services gives the Exchanges more presence, and increases the performance of the carrier's network.
- Network Access Services: Remote Authentication Dial In User Service (RADIUS). This service, offered in Internet by carriers and ISPs, can be provided directly in the IX, providing specially to the ISPs more reliability to connect directly to the Exchanges instead of connecting through a national carrier. Moreover, it increases the possibilities of potential customers accessing the IX.
- Other services: Mail Servers (POP3, SMTP) and Internet Relay Chat Servers (IRC). These are value added services that complete the variety of new possibilities of the Layer 3 Exchanges.

Layer 3 capabilities of the IX provides the use of monitoring applications inside de Exchanges such as:

- Routing Monitoring Systems:
  - AS-Path tree (developed by TILAB) that permits controlling the AS-path of BGP4+ routes.
  - Looking Glass that permits executing predefined commands, such as ping or traceroute, over network elements.
- Reachability Monitoring Systems (Ping view) that provides information about alive and dead machines in the network.
- Management Systems (Magalia developed by TID) that is a distributed environment for managing and monitoring links and machines of the whole network.
- Traffic Monitoring Systems (Cricket, MRTG ported to IPv6 by nGn) that provides information and statistics about the traffic flow of the network.

2.3.5 Customers

One of the differences between traditional L2 IXs (Model A) and new L3 IXs (ModelB) is the possibility of offering a wide range of new services to new customers. The traditional IXs, only hosted L2 equipment are offered and the ISP should bring their own router equipment for peering with another ISPs. So, the only customers connecting to L2 IXs were ISPs.

With L3 IXs (Model B) adding to traditional customers of L2 IXs (ISPs), appears a new set of potential customers. These new customers can be:

- Large Customers (Corporate Networks of large/world-wide companies). These customers have the suitable infrastructure to connect themselves directly to the Exchanges (avoiding the usage of a national carrier). Layer 3 IX (Model B) brings them the possibility of connecting directly to the IX node.
- Standard Customers and ISPs connected to the L3 IX through the carrier's IPv6 national backbone. This is the traditional service offered by L2 IX (Model A), and it is also offered with the new L3 IX (Model B).
- ISPs offer L3 dependent services. They need L3 infrastructure to connect to. With L2 IX (Model A) they need the carrier's national backbone to connect to the Exchanges. With

L3 IX (Model B), L3 capabilities are fulfilled inside the IX, so ISPs can connect directly without carrier's backbone dependence.

2.4 Model C

The third model that will be investigated inside the Euro6IX network is related to an advanced concept of Internet Exchange (IX) that, assuming an IX as a neutral point of traffic exchange among IP networks (in general Internet Service Providers), takes in consideration the importance of prefix aggregation inside the backbone of the next generation IPv6 Internet. The proposal starts from a suggestion inside RFC2374 "IPv6 Global Unicast Address Format", where (textually copied) "… the IPv6 aggregatable address format is designed to support long-haul providers (shown as P1, P2, P3, and P4), exchanges (shown as IX), multiple levels of providers (shown at P5 and P6), and Exchanges (unlike current NAPs, FIXes, etc.) will allocate IPv6 addresses. Organizations who connect to these exchanges will also subscribe (directly, indirectly via the exchange, etc.) for long-haul service from one or more long-haul providers. Doing so, they will achieve addressing independence from long-haul transit providers. They will be able to change long-haul providers without having to renumber their organization. They can also be multihomed via the exchange to more than one long-haul provider without having to have address prefixes from each long-haul provider."

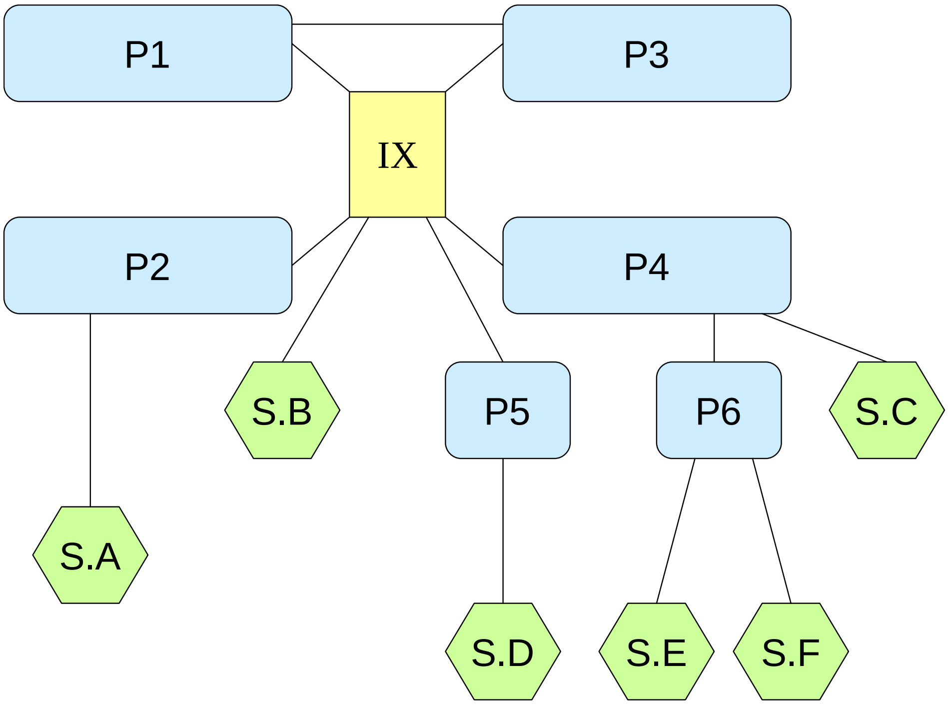

Figure 2-7: Network Architecture as Shown in RFC2374

It is easy to understand that the current models of Layer 2 based Internet Exchange cannot be considered an optimal solution to reproduce this architecture. In fact, even if it should be possible to:

- Configure routing between Exchange Subscribers and Long Haul providers announcing disaggregated prefixes.
- The Long Haul Providers take care to announce to the Big Internetv6 the aggregated IX prefix.

In this way the Internet Exchanges looses the possibility to control the routing of its own prefixes.

For this reason, the Euro6IX project is going to study a new functionality inside the architecture of an Internet Exchange: This functionality is called "Layer 3 mediation function".

2.4.1 Internet Exchanges Architecture

The following figure depicts the proposed model of the IX interconnecting Long Haul Providers and Exchange Subscribers (indicated in the figure as "Next Generation Customers").

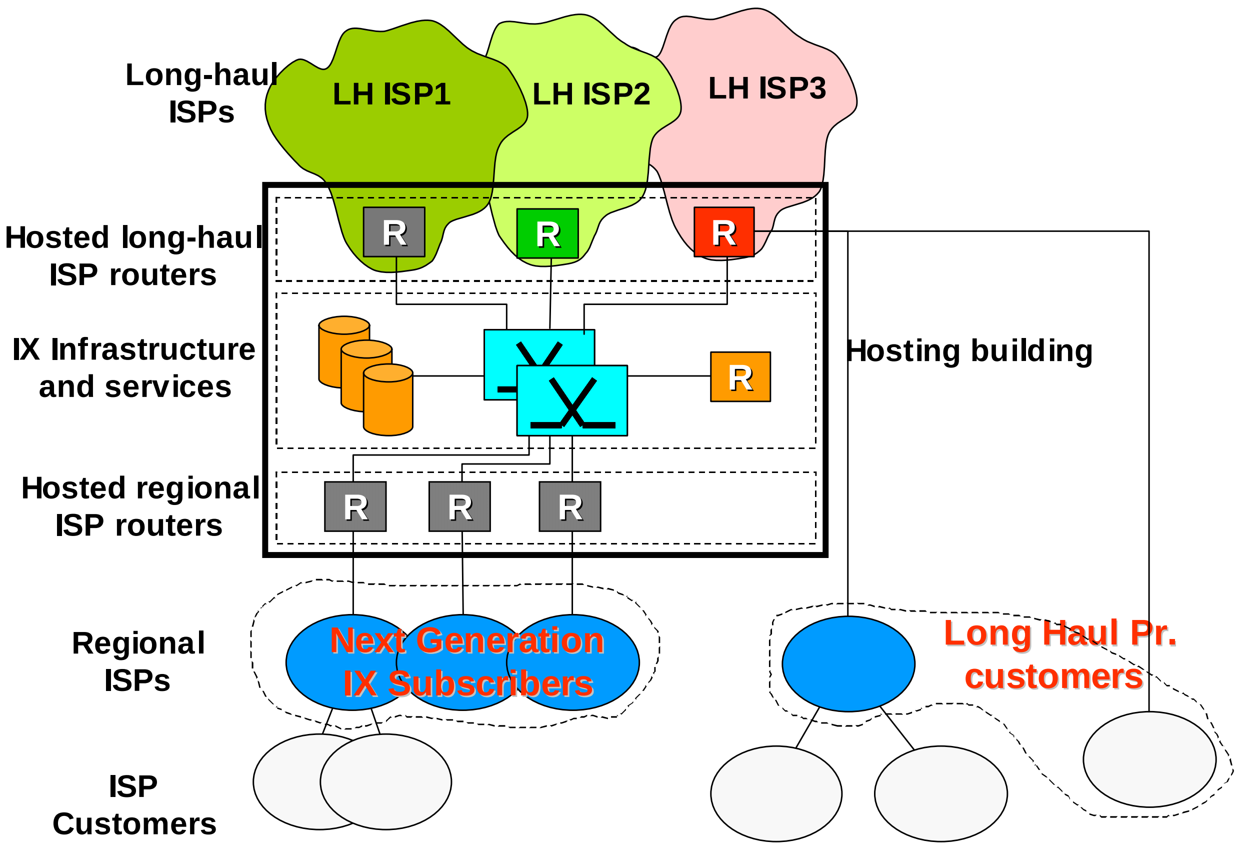

Figure 2-8: Internet Exchanges Structure

The Layer 3 device on the right side of the Figure performs the "Layer 3 Mediation Function". It can be just a router (but in this case the IX forwarding performances are influenced by the performance of this router) or can perform only a control function leaving the forwarding to the high speed Layer 2 infrastructure of Internet Exchange (in this case how to define the "Layer 3 Mediation Function element" will be investigated during the research activities inside the Euro6IX project).

As shown in Figure 2-8 the Internet Exchange is installed inside a building that in general contains:

- The IX equipment (Switches, Layer 3 device performing "L3 Mediation Function", and management devices such as Route Server, Monitoring Workstation etc.) shown inside the central dotted area.
- Routers belonging to the Long Haul Providers (linked to the switched infrastructure).
- Routers belonging to the Internet Exchanges Customers (linked to the layer 2 switch infrastructure).

2.4.2 Goals of the Model

As already indicated, this model has been proposed in order to experiment the new IPv6 IX model defined in RFC2374. This model considers a new conception of IX, which is no more simply a point where ISPs meet each other and exchange their traffic but it can be considered like an entity that assigns IPv6 prefixes that are not dependent of the Long Haul Provider used by the user for the long haul connection. A Long Haul Provider places its own router inside the IX building (outside the dotted area) and uses the high-speed layer 2 connections to connect with its users. This scenario make easier the renumbering process since if an user wants to change its provider it has only to modify its routing policies while the addresses are always the same because they are assigned by the IX and not by the Long Haul Provider.

2.4.3 Routing and Addressing

The model here described basically relies on the idea to consider the Euro6IX backbone or part of it (at least the router of the Euro6IX network collocated with an IX) like an autonomous entity that can be seen, by the IXs, as one of the providers that connect to.

In this way, the part of Euro6IX backbone emulating the long haul provider will have an AS number whereas each IX will use its own AS number.

Even if the Euro6IX project doesn't define a common policy for the routing inside the Internet Exchange is useful to define some guidelines about the routing suggested for the model of IXs with new Layer 3 Mediation Function. In this model the routing guidelines are shown in Figure and are:

- eBGP4+ between layer three mediation function router and backbone routers.
- eBGP4+ between layer three mediation function router and Next Generation IX Customers routers.
- eBGP4+ between routers belonging to the backbone and Standard IX Customers (accessing IX without using layer 3 mediation service).
- eBGP4+ between routers belonging to the backbone and Standard IX Customers (Standard Long Haul Provider Customers).
- IGP and iBGP4+ inside the part of Euro6IX backbone emulating the Long Haul provider. IGP needs to be used to guarantee the loop-back interfaces reachability inside the backbone.

The routing inside the overall network depends on which part of the Euro6IX backbone will be used to emulate the long haul provider. In order to guarantee the maximum independency on which model every partner will reproduce inside the Euro6IX and in order to guarantee the stability of the routing, the majority of partners hosting the Internet Exchanges agreed that the Long Haul provider will be emulated only by the router collocated inside the Internet Exchanges. This Long Haul provider will have one or more direct peering with other providers or other Internet Exchanges.

The routing framework to be adopted is shown in the following picture:

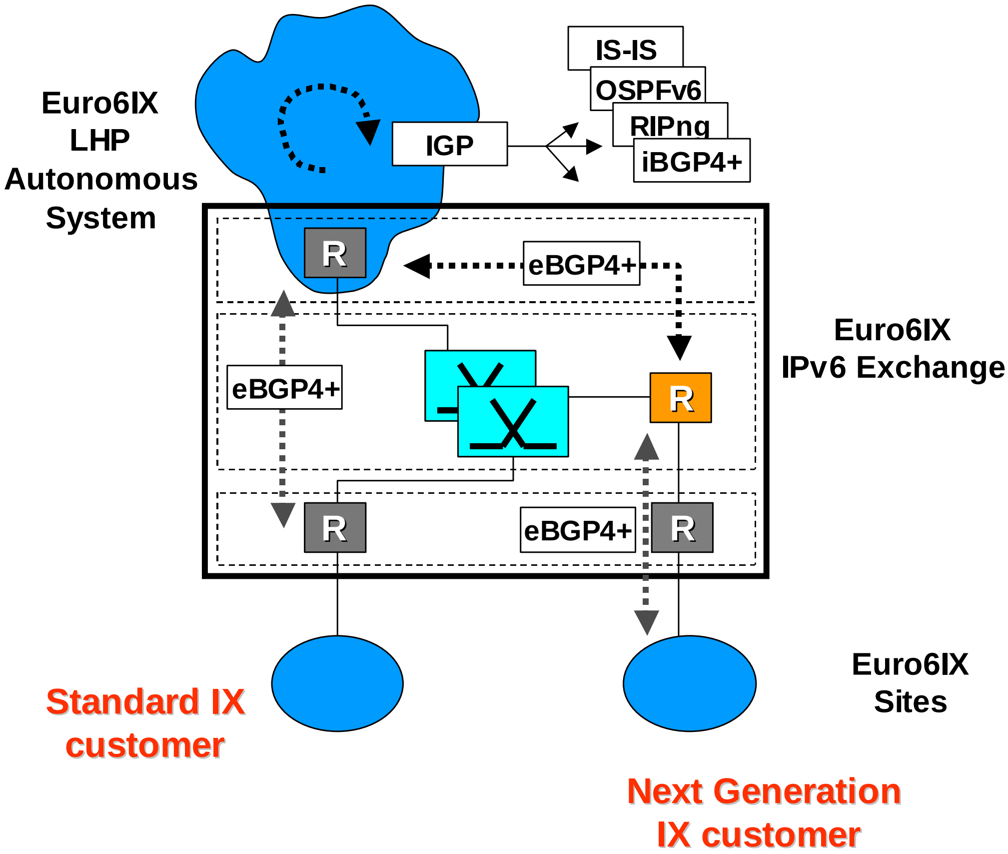

Figure 2-9: Routing Architecture Inside an IX

From the addressing point of view, every long haul provider and every Internet Exchange will have its own prefix. For example, a prefix assigned to one of the Long Haul Providers could be the Euro6IX 6Bone prefix (3ffe:4011::/32) and prefixes already owned by the partners hosting Internet Exchanges could be used as IX owned prefixes.

Every Long Haul Provider will use its own prefix to number all the links and the routers belonging to its backbone and to assign prefixes to their direct connected customers.

In particular the part of Euro6IX emulating a Long Haul provider will number with the same prefixes:

- The routers belonging to the Long Haul Provider collocated in the IX building outside the dotted area.
- The links connecting the Long Haul Provider routers between each other's.
- Providing the IPv6 addresses to the Standard Long Haul Provider Customers.

The s/pTLA belonging to each IX will be used for:

- Assigning the addresses to the Next Generation Customers.
- Numbering the layer 3 part of the Internet Exchange (including the Layer 3 mediation Function Router).

The addressing framework to be adopted is shown in the following picture:

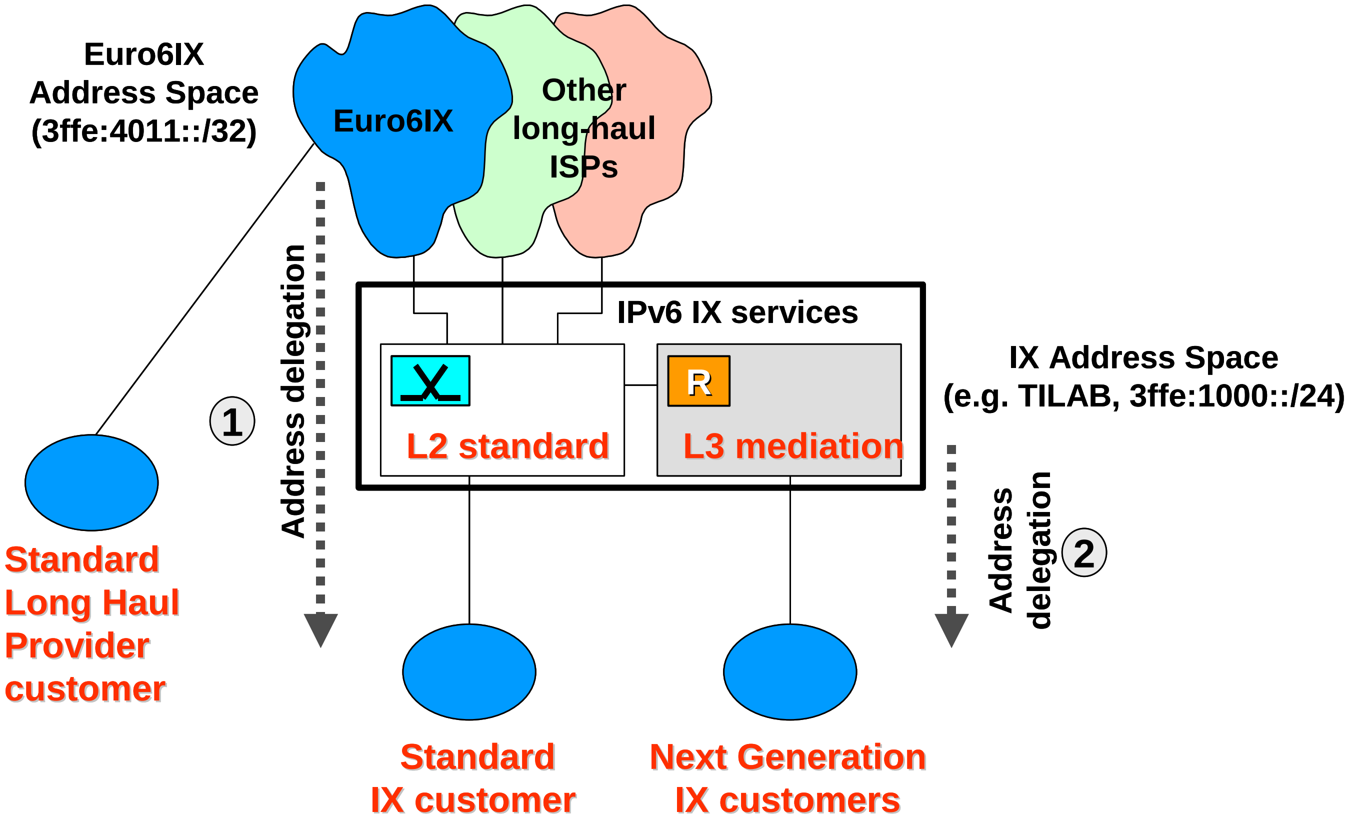

Figure 2-10: Addressing Framework

3. BACKBONE EURO6IX NETWORK ARCHITECTURE

3.1 Definition and Requirements of a Backbone Model

Euro6IX project deploys a Pan-European IPv6 native network to research on advanced network services and applications. One of the objectives of the project is to research and develop a European pre-commercial Internet made of IPv6 Exchanges (IX).

In order to control the routing inside the Euro6IX network and between Euro6IX network and external networks, which Euro6IX is connected to routing, policies have to be defined. Every IX administrator must fulfill all the routing policies designed for Euro6IX backbone.

Some premises must be considered before designing the routing policy. One of them, and maybe the most important, is network topology and the interconnection between Euro6IX network and other networks.

3.2 Network Description

Network architecture of Euro6IX is basically made by Internet IPv6 Exchanges (6IX or IX from now on). Several European Telcos of Euro6IX consortium provides the necessary infrastructure to deploy and to place the IXs.

The Euro6IX network is composed by:

- IPv6 Exchanges (IXs): Hosted by the Telcos.
- Links among IXs: Sponsored by the Telcos and forming a ring topology.
- Links between the Euro6IX Network and other Networks (6NET, 6Bone, etc.).

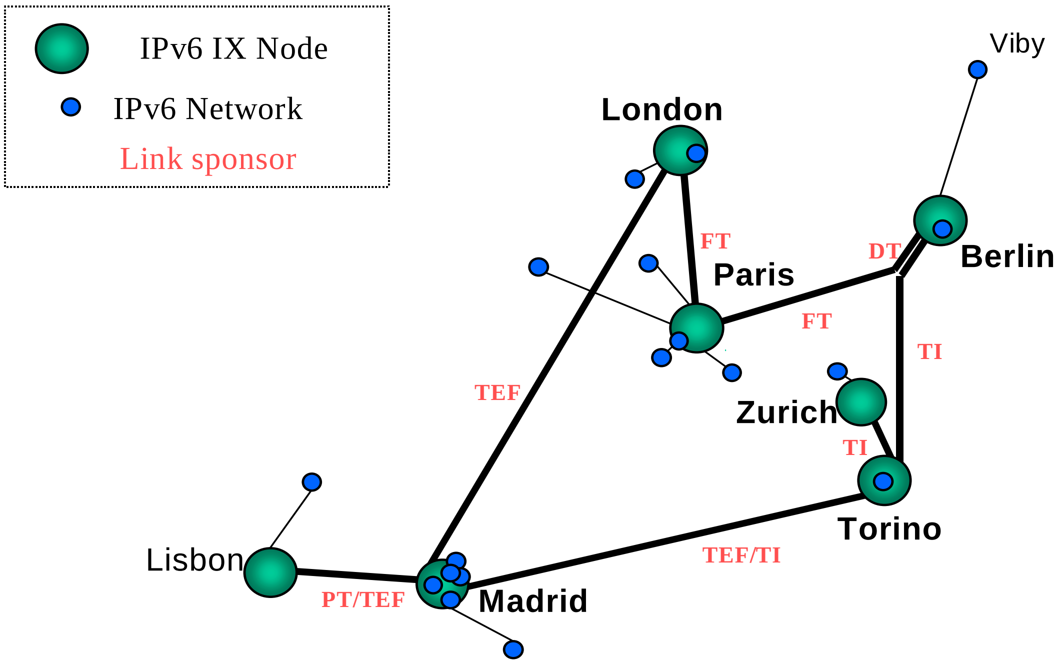

Figure 3-1: Euro6IX Network

The network is interconnected with other IPv6 networks, in general, from different points as depicted in Figure 3-2. Each external network connecting to Euro6IX, i.e. 6NET, must have a BGP4+ peering with one or more Euro6IX IXs to establish that connectivity. There will be also national networks connected to any IX. Euro6IX connectivity to these networks can be achieved establishing dynamic routing (BGP4+ peering) or just with static routing.

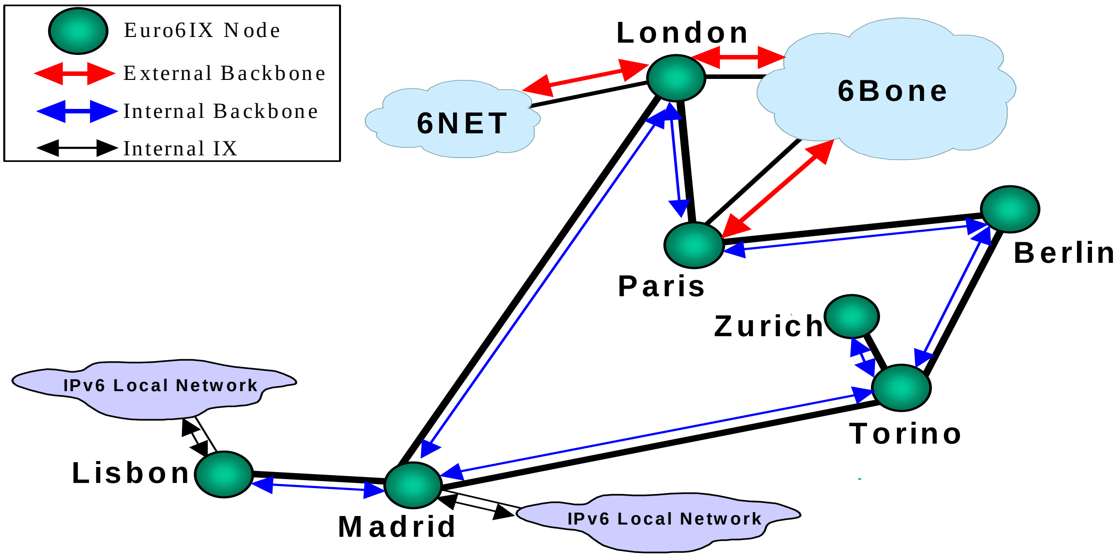

Figure 3-2: Example of Network Architecture

3.3 Routing Policy Description

The elements that compose the network and the links among them identify the traffic flows to be exchanged. All these flows must be identified and controlled, so that a Network Routing Policy definition is needed.

Network Routing Policy is mandatory to identify, classify and manage network traffic flows. This work must consider the following Routing Levels and policies associated with those levels:

- **Internal Euro6IX Backbone routing**: Describes the routing policy in the context of traffic exchange among IXs. The policy defined for this type of routing is:
    - Every IX have to reach another IX using as much as possible the Euro6IX network infrastructure. For this reason Euro6IX prefixes received from Euro6IX links have to be always preferred respect to the same prefixes received from another link (e.g. 6Bone).
    - Routes will be aggregated as much as possible and non-IANA allocated prefixes will be filtered out.
- **External Euro6IX Backbone routing**: Describes the routing policy needed to exchange traffic between IXs and other networks like 6Bone, 6NET, Asia Pacific Networks, etc. refer from now on with the term “external networks”. Respect to this networks the routing policies defined are:
    - In general IX doesn’t provide transit between two external networks (prefixes coming from an external network aren’t announced to another external network), except for some special networks.

- Every IX have to reach external networks exiting from the interconnection point between Euro6IX network and the External network. If more than one interconnection point is present between Euro6IX network and external network the routers will make a decision based on the nearest link (ASPath based or IGP based).
- Special procedures will be specified for special routes to/from particular experimental networks like 6Bone, 6NET:
    - **6Bone:** Many of the Euro6IX IXs peer with parts of the 6Bone. IXs will inject 6Bone routes onto the Euro6IX backbone. Euro6IX IX will prefer 6Bone routes received from 6Bone peers. IX will not provide transit to 6Bone network. Euro6IX prefixes will be announced to 6Bone peers so that the 6Bone PoPs will choose (in the most cases) 6Bone path instead of Euro6IX path. In the future, additional policies (e.g. preferring 6Bone native links respect to tunneled) may be specified.
    - **6NET:** Some of the Euro6IX IXs peer either directly with 6NET or via a transit provider (both cases will be treated the same). 6NET routes received from an IX will be announced Euro6IX backbone only to Euro6IX partners: Routes will be preferred using normal metric (e.g. BGP AS path length). Euro6IX prefixes will be announced to 6NET using no metric change. This is a requirement of 6NET that Euro6IX only advertise their routes to companies in the Euro6IX consortium.

- **Internal IX Routing**: Describes the routing policy needed to exchange traffic between one IX and other national networks, such as NRENs or other National Telcos, etc.
    - This policy is freely implemented by Internet Exchanges Administrators.
    - Since some IX can have already defined their own routing policies, this point is open to the election of the administrators of each IX.

3.4 Routing Policies Implementations

3.4.1 Available Mechanisms to Perform Routing Policies

Controlling the routing levels described above requires an implementation of some routing protocol mechanism. The routing protocol used in Euro6IX Backbone to exchange network routes is external BGP4+, which include multiprotocol extensions for IPv6.

There are some mechanisms to control the traffic exchange, such as:

- Local Preference: Mechanism used to degree the preference given to a route to compare it with other routes for the same destination. A higher Local Preference value is an indicator that the route is more preferred. This attribute is local to an Autonomous System.
- AS Prepend: BGP4+ mechanism used to modify AS-Path length. It is a sequence of Autonomous System Numbers a route has traversed to reach a destination. Prepending is the act of adding the AS Number to the beginning of the list.
- Multi Exit Discriminator (MED): BGP4+ attribute used to perform a best path selection when the same route is received from multiples Autonomous System with the same AS path length.
- Communities: BGP4+ mechanism used to tag the routes in order to classify and group traffic flows.
- Traffic filtering: Mechanism to filter a set of router advertisements (prefix list).

3.4.2 Routing Control Mechanism used in Euro6IX Backbone

Using external BGP4+ the selected mechanisms for the Euro6IX network to control the traffic exchange are:

- Extended Communities.
- Local Preference.
- AS-Prepend.

All of them are available in most BGP4+ implementations and the main mechanism used in this model to control traffic flows is the Community tagging.

Local Preference (LP) permits every IX administrator to choose the best link or way to reach a destination. It is fundamental for setting the higher preference to the Euro6IX links and controlling traffic flow among them. A higher LP value for Euro6IX routes will assure that Euro6IX traffic will flow among Euro6IX links instead of other links, such as 6Bone connections.

AS-Prepend provides an added value to the Backbone. It permits to install a specific route to a concrete network using a desired path instead of the preferred one for the BGP4+ routing protocol. The selection of this path depends, in first case, on the Internal Euro6IX Backbone Routing Policy definition, and in second case, on the administrator criteria if possible.

3.4.3 Euro6IX Backbone Routing Policy: First Proposal

The first proposal for the implementation of Euro6IX Routing Policy, uses the BGP4+ routing control mechanisms to control the traffic flow inside Euro6IX network.

Given the necessity of controlling the traffic flows, they will be grouped by using “community tagging” mechanism. The “tag” will be chosen taking into account the kind of traffic, its source and the required handling.

The community tagging will be done when network announcements enter Euro6IX network and the same pattern will work throughout the whole network.

All routes exchanged among IXs must use a community and this community must not be changed in all the Euro6IX Backbone. Peers will discard all the routes announced by IXs with no community tag.

Every AS has to be known and described within the Internal Euro6IX Backbone Routing Policy. When any community inherits from another network the IX that is injecting those routes with that community must convert the remote community to a known community defined in this document.

Extended communities format will be used to tag routes. They will be used as an identifier of the network, so that customized information for that route could be generated.

Extended communities will be composed by two numerical parts separated by two points (<ASNumber>:<Value>), representing:

- ASNumber → Autonomous System Number of the announcer of the network. This field will be used to indentify the AS Number of the IX that is injecting a concrete route into the Euro6IX backbone.
- Value → Numerical value XXYZ.

Value field can be divided into:

- XX → numerical value between 00 and 99, which defines the nature of a route.
- Y → numerical value between 0 and 9, reserved for future use. It can be used for defining the kind of traffic associated to a route (experimental, commercial, testing, etc.).
- Z → numerical value between 0 and 9, which represents the action to be taken for that route.

As stated before, these communities will remain the same through two Euro6IX's IXs.

3.4.3.1 AS Number Values

The following table shows the possible values that "AS Number" can take.

IX Name:	AS Number
LIS6IX	3243
MAD6IX	3352
LON6IX	1752
PAR6IX	5511
BER6IX	3320
TOR6IX	5069
ZUR6IX	3303

Figure 3-3: AS Number Values

3.4.3.2 XX Field Values

XX will have the value of the peer from which it receives the route. The information of the peering such as:

- Peer with an Internal backbone element (IX).
- Peer with an External network (6Bone, 6NET).
- Peer with a national IX network (TID-IX, Vodafone-IX).

The following table shows the possible values that XX field can take. They define the scope of each route within each level.

XX Range	Detailed XX Value	Nature
00 —09 Not to be used	00 - 09	Not to be used
10 —20 Internal Euro6IX Routing	10	TLA of Euro6IX IX
	11 - 20	Reserved
21 —60 External Euro6IX Routing	21 —30	Reserved
	31	6bone Native Connection
	32	6NET Native Connection
	33	Abilene Native Connection
	34 —40	Reserved
	41	6Bone Tunneled Connection
	42	6NET Tunneled Connection
	43	Abilene Tunneled Connection
	44 —58	Reserved
	59	Other Native Connection
	60	Other Tunneled Connection
61 —90 Internal IX routing	61 - 90	Internal to IX domain
91 —99 Reserved	91 - 99	Reserved

Figure 3-4: XX Field Values

3.4.3.3 Y Field Values

The following table shows the possible values that Y field can have.

Y Value:	Action
0	Non commercial traffic
1	Commercial traffic
2	Testing traffic
3-9	Reserved

Figure 3-5: Y Field Values

As an example of future use, Y field could be used to define: different qualities of service (QoS), traffic test, commercial or non-commercial traffic, etc.

3.4.3.4 Z Field Values

The following table shows the possible values that Z field can take.

Z Value:	Action
0	Announce everywhere (FULL TRANSIT)
1	Announce everywhere with a NATIVE link
2	Announce everywhere with a TUNNELED link
3-4	Reserved
5	Announce only to all Euro6IX partners
6	Reserved
7	Don't announce the route to any NATIVE link
8	Don't announce the route to any TUNNELED link
9	Don't announce the route

Figure 3-6: Z Field Values

Note: Value 5 (announce only to Euro6IX partners) means announce to sites where the controlling organization is part of the Euro6IX consortium. This stops routes going to organizations that the IX has allocated addresses which are outside the Euro6IX consortium (this is a requirement of 6NET).

This field value must not change throughout the Euro6IX network. In this case, only the IX that announces the network sets the community value for each route.

3.4.3.5 Internal Euro6IX Prefixes

The value to use on "XX" field for Internal Euro6IX Backbone is defined on the next table:

Peering Backbone level	XX
Internal Euro6IX Backbone	10

Figure 3-7: XX Field Value

The whole set of Euro6IX routes exchanged among IXs must be tagged with "10" value into the XX field. So, it can be affirmed that a route with XX=10 is an Euro6IX route. With this classification all the routes belonging to any Euro6IX partner/IX are grouped and tagged. As defined before the IX announce IX prefixes outside Euro6IX network (Z=0) and the Euro6IX traffic will be considered as no commercial (Y=0).

The "ASN" field defines the information about what IX announces any route into the Euro6IX backbone.

In summary:

- Euro6IX prefix – ASN:1000

3.4.3.6 Injecting Routes into the Euro6IX Backbone

As specified in the previous section, only aggregated 6Bone and IANA assigned prefixes should circulate inside Euro6IX backbone routing table. For this reason, before any route is injected into the Euro6IX backbone, it must be filtered considering the following rules:

- Permit the prefixes belonging to 3FFE::/18 with length equal to /24.
- Permit the prefixes belonging to 3FFE:4000::/18 with length equal to /32.
- Permit the prefixes belonging to 3FFE:8000::/22 with length equal to /28.
- Permit the prefixes belonging to 2001::/16 with length between /29 and /35.
- Permit 2002::/16 only.

3.4.3.7 Use of Local Preference to Control Routing

In general the Euro6IX backbone network will use standard BGP4+ AS length to decide the preferred route. This however would lead to traffic taking a not optimum route (in terms of defined routing policies) in certain circumstances and hence certain specific actions will be taken in each Euro6IX IX to ensure:

- Optimum routing wherever possible applying policies specified in §3.3.
- Fulfilling condition imposed by external network peerings (6NET, etc.).
- Allowing new routing paradigms to be trialed.

In particular Local Preference will be used inside each IX to “favor” certain routes.

The Local Preference BGP attribute for various sets of routes will be set internally by each IX and this will cause traffic to flow in accordance with the routing policy (ie Euro6IX network preferred over external networks, locally connected 6Bone preferred over 6Bone routes received via Euro6IX etc.).

The following LP values are recommended:

LP Value	Route Set
140	Internal Euro6IX route ie TLAs of IXs identified by XX=01
100	All external Euro6IX routes besides 6Bone
60	6Bone route connected directly to IX
40	6Bone route received from Euro6IX peer

Figure 3-8: LP Values

Note: Euro6IX prefixes announced/received via 6Bone (direct tunnels i.e. Italy-Spain or Italy-Germany, or indirect tunnels) will not be tagged by any community value. In this case these prefixes should be considered as 6Bone route using suggested local preference value 60.

3.4.3.8 Routing Policy Applications: 6Bone

According to the specified routing policies and to the specified community tagging scheme 6Bone routes will be injected onto the Euro6IX and backbone with the following identifications:

- <ASNumber>:XXYZ where XX = 31 or 41 (6Bone native connection or 6Bone tunneled connection), Y = 0 (non commercial traffic) and Z = 5 (announce only to Euro6IX partners).

For establishing the 6Bone peering in each IX, the Local Preference (LP) of the routes received from 6Bone should be set to 60. The default LP for Euro6IX routes will be set to 140. So, the Euro6IX routes are preferred over 6Bone ones and traffic will flow through Euro6IX links instead of external ones among Euro6IX partners. The LP for the 6Bone routes received from another Euro6IX peer must be set to 40. So, traffic flow through 6Bone must go via the nearest IX, and if the nearest IX link to 6Bone fails, then 6Bone connectivity is reached via another IX.

3.4.3.9 Routing Policy Applications: 6NET

For 6NET prefixes the following identifications will be used:

- <ASNumber>:XXYZ where XX = 32 (6NET native connection), Y = 0 (non commercial traffic) and Z = 5 (announce only to Euro6IX partners).

When the Euro6IX IX identifies these routes it will use normal BGP AS path length decision to identify its preferred routes, it will advertise these routes only to Euro6IX consortium partners. This is a requirement of 6NET that we only advertise their routes to companies in the Euro6IX consortium; in many cases, this will mean not all the entities connected at a particular IX.

3.4.3.10 Routing Policy Applications: General External Euro6IX Routes

Many of the IXs will peer with numerous networks that do not come into the categories of either other Euro6IX IXs or special external networks. Routes from these peers will be injected onto the Euro6IX backbone with (in general) the following identifications:

- <ASNumber>:XXYZ where XX = 59 or 60 (other native connection or other tunneled connection, Y = 0 (non commercial traffic) and Z = 0 (announce everywhere).

When the Euro6IX IX identifies these routes it will use normal BGP AS path length decision to identify its preferred routes, it will advertise these routes to all the other networks it peers with i.e. Euro6IX provides transit, unless the IX has agreed additional local preferences with its other networks. In the future additional policies on whether the link is native or tunneled may be specified and recommendations made i.e. Euro6IX could provide transit to just natively connected other networks.

This community tagging must be only local to each IX, and must be changed when announcing these routes into the Euro6IX backbone.

3.4.3.11 Suggested Internal IX Community Tagging Scheme

The local administrator of each IX must define the internal IX Routing Policy. The way of defining and making the peering connections with other National networks or big customers is only responsibility of IX administrator.

In any case, this section suggests an Internal IX Routing community scheme. It is only a recommendation and could be used like a configuration guide.

The following table shows the possible values that XX field can take.

Peering Id.	XX
Internal IX routing	61 —90

Figure 3-9: Internal IX Suggested XX Field Value

A specific community is assigned to each internal peer connection defined locally to IX. For example:

Peer Id.	Community
VODAFONE	61
CONSULINTEL	62
UPM	63
TID	64
…	…

Figure 3-10: Internal IX Suggested Community Value

This community tagging must be only local to each IX, and must be changed when announcing these routes into the Euro6IX backbone.

3.4.3.12 Case Examples

The following pictures show several case studies to illustrate the routing policy defined previously.

Case 1: Abilene

This case studies the peering between an Euro6IX IX (LON6IX) and an external network (Abilene).

This example shows how to define the communities for providing connectivity only to Euro6IX partners with Abilene network. So, the routes belonging to Abilene network, must not be announced outside Euro6IX.

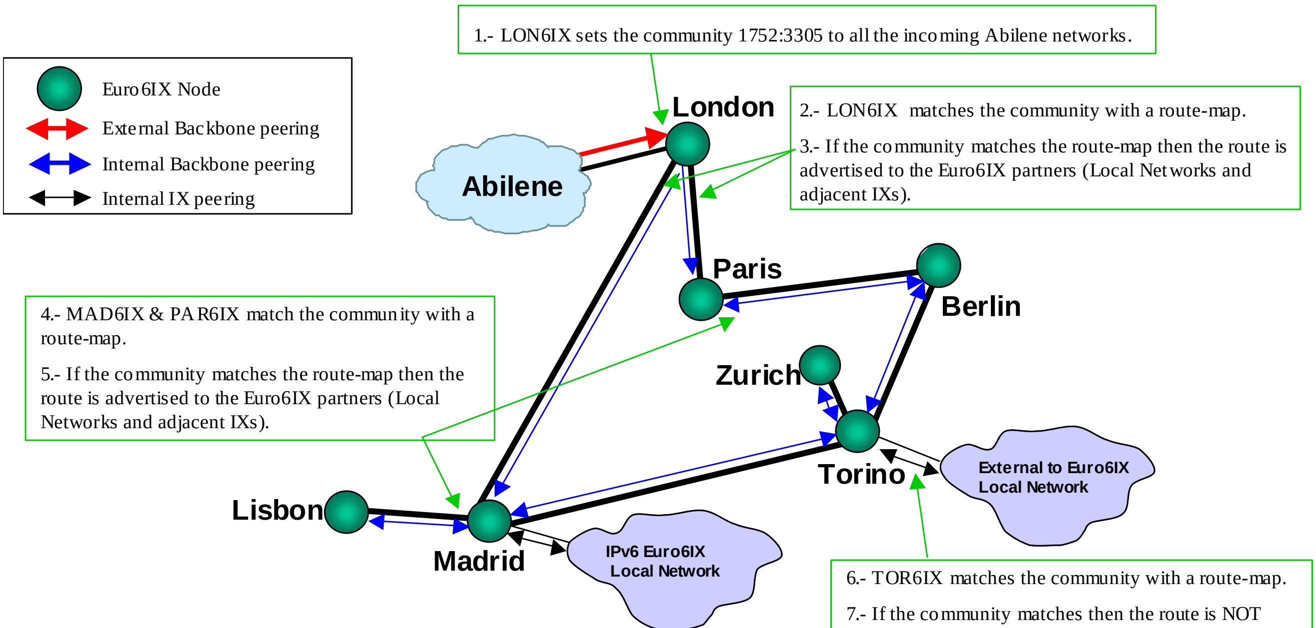

Why is the community value defined to 1752:3305 ?

Community Format: ASN:XXYZ

In this case:

ASN=1752 → LON6IX Autonomous System Number
XX=33 → External Peer (values from 31 to 60). Abilene is identified in “D2.2” as 33.
Y=0 → This value indicates that the traffic is non-commercial.
Z=5 → This value indicates the action for the route. This action refers to announce the route to all Euro6IX partners.

Figure 3-11: Example Case 1: Peering with Abilene

In this example, LON6IX (ASN=1752) is peering Abilene (XX=33) network, and the action to do with the routes received is to announce them only to Euro6IX partners (Z=5). LON6IX sets community value for Abilene routes to 1752:3305 and injects them into Euro6IX BackBone.

All the IXs receiving these routes only must match the action field "Z" and announce the routes only to Euro6IX partners. In the following text box the configuration for Cisco and Juniper equipment for this example is shown.

Cisco

```
!
ip community-list expanded Abilene permit (*:...5)+
!
route-map Abilene permit 10
match community Abilene
!
```

Juniper

```
//under policy-statement sub-tree
community abilene members *:....5;
```

NOTE: this route map could be applied not only to Abilene but all routes received with the community field value "Z=5" (announce only to all Euro6IX partners). With only one route-map we could define an action to take with all routes of the same kind.

So, the "Z" field could used like the last four fields in TILABs proposal defining the action to do with the routes. TILAB includes in these four fields the kind of the route and TID includes this information in the "XX" fields. These fields are only matched if the IX administrator wants to use them.

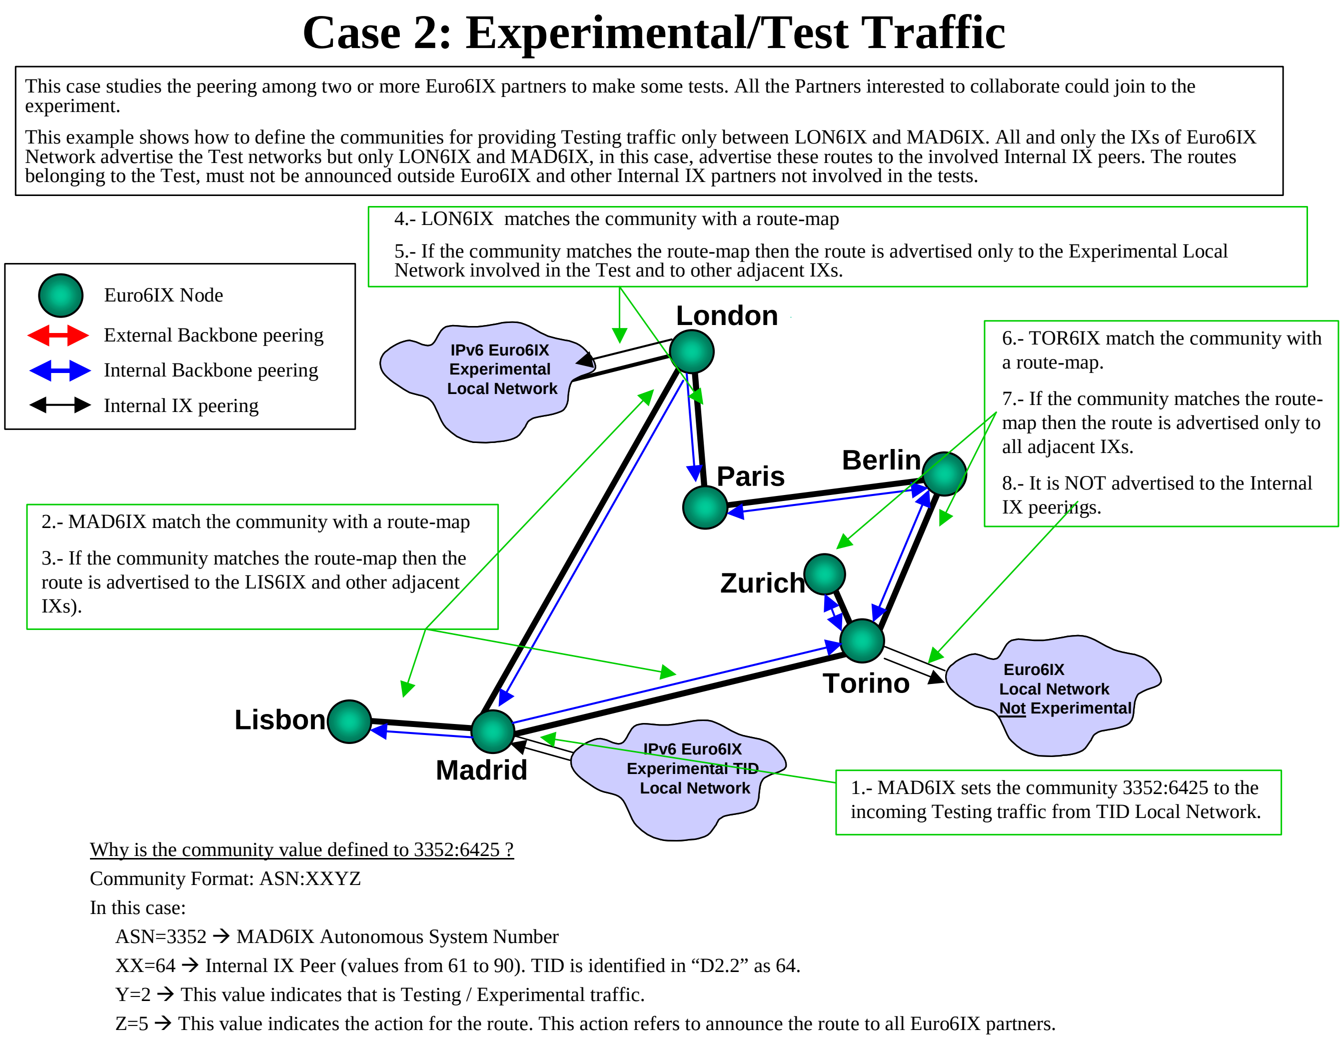

Figure 3-12: Example Case 2: Experimental/Test Traffic

With this example, the functionality of the “Y” field is shown.

The following example shows how to deny all commercial (Y=1) routes from any partner.

Cisco
```
!
config:ip community-list expanded ALL_COMMERCIAL permit (*:..1.)+
!
route-map FILTER_COMMERCIAL deny 10
match community ALL_COMMERCIAL
!
```

Juniper
```
//under policy-statement sub-tree
community ALL_COMMERCIAL members *:..1.;
```

The following example shows how to permit some experimental/testing routes only from 6NET and announce them only to Euro6IX partners. Each IX receives these routes and announce to the adjacent IXs. The IXs announce the routes only to the partners involved in the experiment.

Cisco
```
!
config:ip community-list expanded 6NET_TEST permit (*:322.)+
!
route-map FILTER_TESTING deny 10
match community 6NET_TEST
!
```

Juniper
```
//under policy-statement sub-tree
community 6NET_TEST members *:322.;
```

The following example shows how to accept all 6NET (XX=32) routes from BT (ASN=1752).

Cisco
```
!
config:ip community-list expanded 6NET_BT permit (1752:32..)+
!
route-map BT_IN permit 10
match community 6NET_BT
!
```

Juniper
```
//under policy-statement sub-tree
community 6NET_BT members 1732:32..;
```

NOTE: In this case, the fields "YZ" do not have to be used, and they are not used in the match sentences.

The following example shows how to accept all 6NET (XX=32) routes from any partner.

Cisco

```
!
config: ip community-list expanded 6NET_BT permit (*:32..)+
!
route-map BT_IN permit 10
match community 6NET_BT
!
```

Juniper

```
//under policy-statement sub-tree
community 6NET_BT members *:32..;
```

Case 3: 6NET

This case studies the peering between Euro6IX 6NET. This example assumes that there are two IXs

This example shows how to define the communities for providing connectivity only to Euro6IX partners with Abilene network. So, the routes belonging to Abilene network, must not be announced outside Euro6IX.

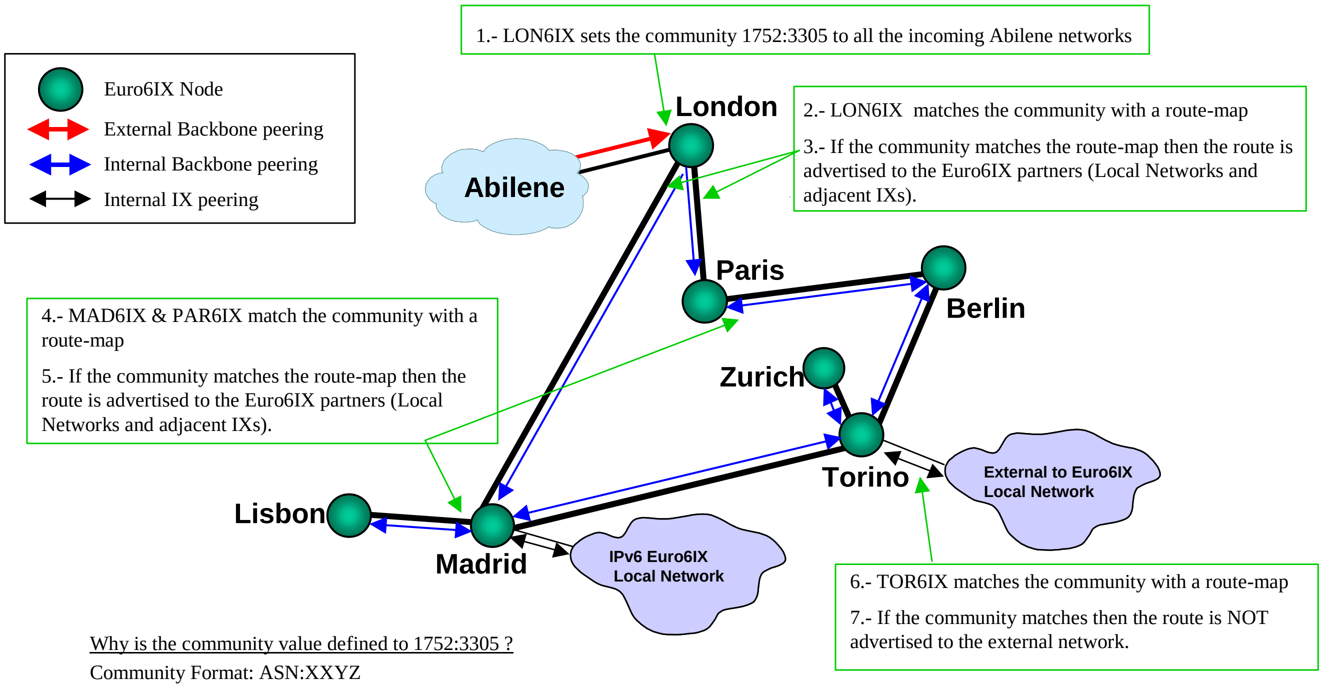

<u>Why is the community value defined to 1752:3305 ?</u>

Community Format: ASN:XXYZ

In this case:

ASN=1752 → UK6X Autonomous System Number

XX=33 → External Peer (values from 31 to 60). Abilene is identified in "D2.2" as 33.

Y=0 → This value indicates that the traffic is non-commercial.

Z=5 → This value indicates the action for the route. This action refers to announce the route to all Euro6IX partners.

Figure 3-13: Example Case 3: Peering with 6NET

The following example shows how to establish the peering (ie: in LON6IX) with 6bone. The routes injected from 6bone must be marked with the following community: 1752:3109. The Local Preference for those routes has to be set to 90.

Cisco

```
!
route-map 6BONE_IN permit 10
  set local-preference 90
  set community 1752:3109
!
```

Juniper

```
//under policy-statement sub-tree
   community 6bone-community members 1752:3109
   policy-statement 6bone-tag {
    then {
        local-preference 90;
        community add 6bone-community;
        accept;
    }
}
```

Case 4: 6Bone

This case studies the peering between an Euro6IX IXs and 6Bone.

This example shows how to define the communities for providing connectivity to the national partners and announcing the 6Bone routes to other IXs. The example assumes that all the IXs have their own peering with 6bone.

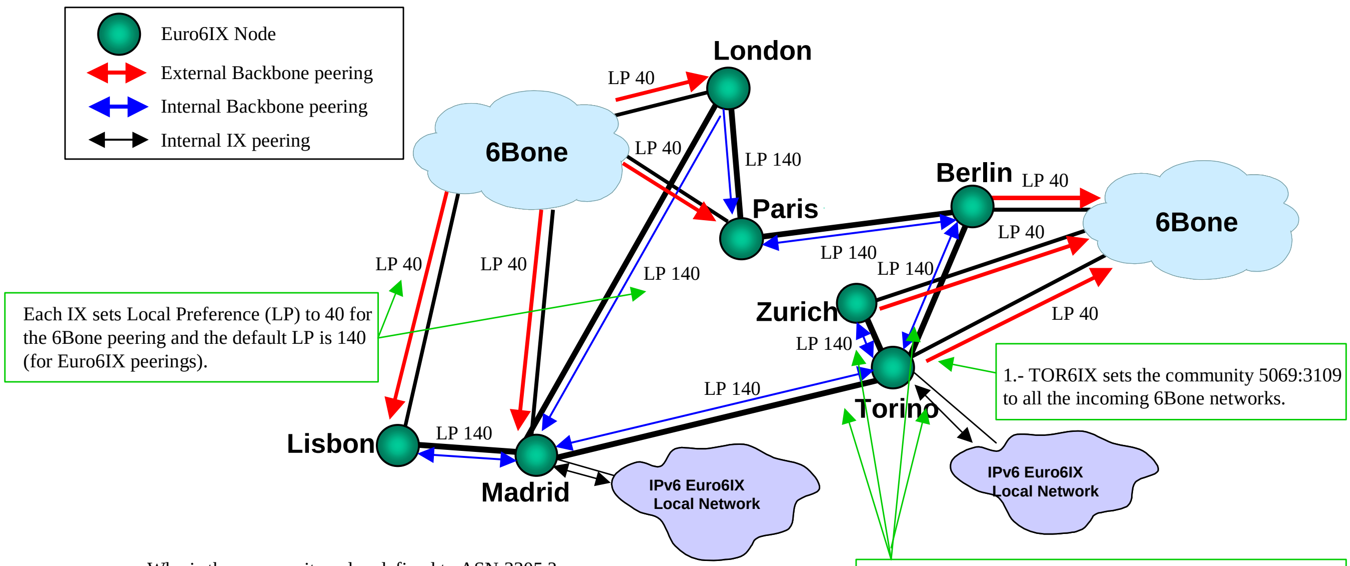

Why is the community value defined to ASN:3305 ?

Community Format: ASN:XXYZ

In this case:

- ASN=The Autonomous System Number of each IX. In the example: TOR6IX=5069
- XX=31 → External Peer (values from 31 to 60). 6Bone is identified in "D2.2" as 31.
- Y=0 → This value indicates that the traffic is non-commercial.
- Z=5 → This value indicates the action for the route. This action refers to announce the route to all Euro6IX partners.

All the IXs must do the same as shown in the figure for TOR6IX.

Figure 3-14: Example Case 4: Peering with 6Bone

General Case: Study of MAD6IX

This case studies the concrete case of MAD6IX. MAD6IX is an Euro6IX exchanger who have INTERNAL BACKBONE, EXTERNAL BACKBONE and INTERNAL IX peers.

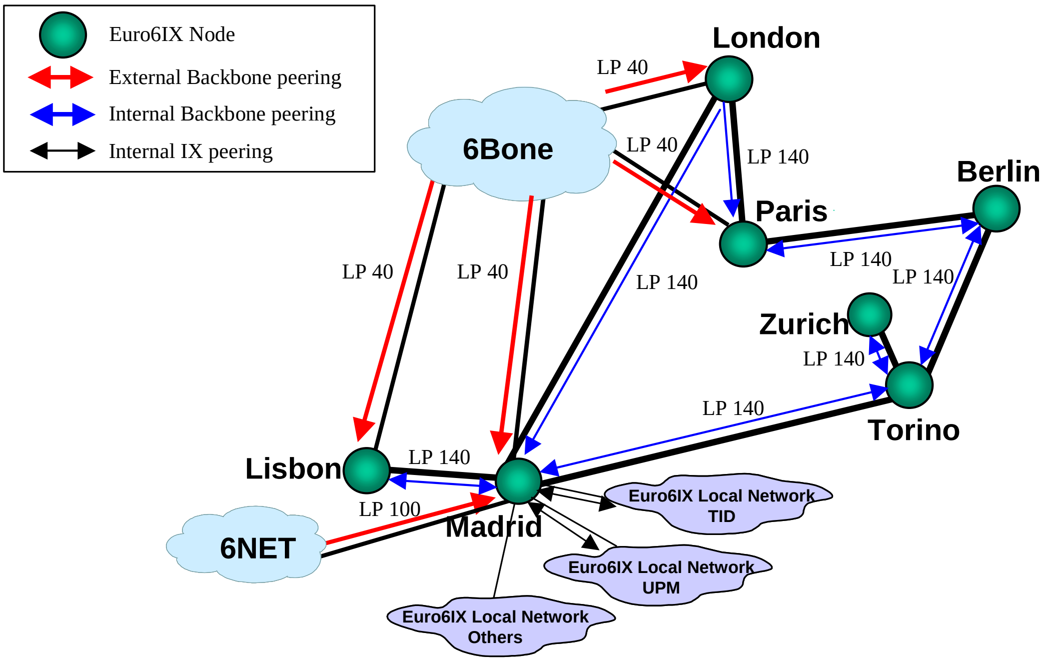

Figure 3-15: General Case: Study of MAD6IX

3.4.3.12.1 Example: MAD6IX Router Configuration

The following configuration example shows the general configuration for the routers of the IXs. It is focused in MAD6IX, but could be applied to each IX. This configuration example is valid for Cisco routers, and it includes several cases studied previously.

```
!
! #### Global configuration Parameters ####
ip bgp-community new-format
!
!
router bgp 3352
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
 neighbor 2001:800:40:2E02::2 remote-as 1752 ← LON6IX
 neighbor 2001:800:40:2E03::2 remote-as 3243 ← LIS6IX
 neighbor 2001:800:40:2F01::2 remote-as 64574 ← TID
 neighbor 2001:800:40:2F03::2 remote-as 65455 ← UPM
 neighbor 2001:800:40:2FFF::2 remote as 65515 ← 6BONE (Not established)
 neighbor 2001:800:40:2FFE::2 remote as 65514 ← 6NET (Not established)
!
 address-family ipv6
! ####
! ####  INTERNAL BACKBONE PEERINGS (between IXs)####
! ####
! ####  Activate the BGP session with _LON6IX_ ####
 neighbor 2001:800:40:2E02::2 activate
! ####  Suppress the private AS Number in the AS path ####
 neighbor 2001:800:40:2E02::2 remove-private-AS
```

```
! ####  Advertise the community attribute  ####
 neighbor 2001:800:40:2E02::2 send-community
! ####  Route-maps applied to LON6IX peering ####
 neighbor 2001:800:40:2E02::2 route-map INTBB_OUT out
 neighbor 2001:800:40:2E02::2 route-map INTBB_IN in
!
! #### The same configuration for LISBON (LIS6IX) ####
!
! ####  Activate the BGP session ####
 neighbor 2001:800:40:2E03::2 activate
! ####  Suppress the private AS Number in the AS path ####
 neighbor 2001:800:40:2E03::2 remove-private-AS
! ####  Advertise the community attribute  ####
 neighbor 2001:800:40:2E03::2 send-community
! ####  Route-maps applied to LIS6IX peering ####
 neighbor 2001:800:40:2E03::2 route-map INTBB_OUT out
 neighbor 2001:800:40:2E03::2 route-map INTBB_IN in
!
!
! ####
! ####  EXTERNAL BACKBONE PEERINGS (with other Networks)####
! ####
!
! ####
! #### 6bone Peering
! ####
! ####  Activate the _6BONE_ BGP session ####
 neighbor 2001:800:40:2FFF::2 activate
! ####  Suppress the private AS Numbers in the AS path ####
 neighbor 2001:800:40:2FFF::2 remove-private-AS
! ####  Advertise the community attribute
! #### neighbor 2001:800:40:2FFF::2 send-community
! ####  Route-maps applied to 6BONE peering ####
 neighbor 2001:800:40:2FFF::2 route-map 6BONE_OUT out
 neighbor 2001:800:40:2FFF::2 route-map 6BONE_IN in
!
!
! ####
! #### 6NET Peering
! ####
! ####  Activate the _6NET_ BGP session ####
 neighbor 2001:800:40:2FFE::2 activate
! ####  Suppress the private AS Numbers in the AS path ####
 neighbor 2001:800:40:2FFE::2 remove-private-AS
! ####  Advertise the community attribute  ####
 neighbor 2001:800:40:2FFE::2 send-community
! ####  Route-maps applied to 6NET peering ####
neighbor 2001:800:40:2FFE::2 route-map EXTBB_OUT out
neighbor 2001:800:40:2FFE::2 route-map 6NET_IN in
!
!
!
! ####
! ####  INTERNAL IX PEERINGS ####
! ####
! ####  Activate the session with TID ####
 neighbor 2001:800:40:2F01::2 activate
! ####  Advertise the community attribute  ####
 neighbor 2001:800:40:2F01::2 send-community
! ####  Suppress the private AS Numbers in the AS path ####
 neighbor 2001:800:40:2F01::2 remove-private-AS
! ####  Route-maps applied to _TID_  peering ####
 neighbor 2001:800:40:2F01::2 route-map INTPEERS_NO6BONE_OUT out
 neighbor 2001:800:40:2F01::2 route-map TID_IN in
!
! ####  Activate the session with UPM ####
 neighbor 2001:800:40:2F03::2 activate
! ####  Advertise the community attribute  ####
 neighbor 2001:800:40:2F03::2 send-community
! ####  Suppress the private AS Numbers in the AS path ####
 neighbor 2001:800:40:2F03::2 remove-private-AS
! ####  Route-maps applied to _UPM_  peering ####
```

```
 neighbor 2001:800:40:2F03::2 route-map INTPEERS _OUT out
 neighbor 2001:800:40:2F03::2 route-map UPM_IN in
!
exit-address-family
!
```

```
!
ip community-list expanded BB_EURO6IX permit
   (^(3243|3352|1752|5511|3320|5069|3303)+(:10..)+$)+
!
ip community-list expanded 6BONE permit
   (^(3243|3352|1752|5511|3320|5069|3303)+(:[3-4]1..)+$)+
!
ip community-list expanded EURO6IX_EXT permit
   (^(3243|3352|1752|5511|3320|5069|3303)+(:[31-60]..)+$)+
!
```

```
!####
!####  INTERNAL BACKBONE ROUTEMAPS ####
!####
!
!####  Advertise to other IXs the internal routes with XX=10 ####
route-map INTBB_OUT permit 10
  match ipv6 address prefix-list TLF_AGGREGATE
  set community 3352:1000

!####  Advertise to other IXs the rest of the routes ####
route-map INTBB_OUT permit 20
  match ipv6 address prefix-list ALL_IPv6
!

!####  Sets to 140 the Local Preference for Euro6IX routes ####
route-map INTBB_IN permit 10
  match community BB_EURO6IX
  set local-preference 140

!####  Sets to 40 the Local Preference for 6Bone routes from other IXs ####
route-map INTBB_IN permit 20
  match community 6BONE
  set local-preference 40

####  Permits all the other routes ####
route-map INTBB_IN permit 30
  match ipv6 address prefix-list ALL_IPV6
!


!####
!####  EXTERNAL BACKBONE ROUTEMAPS ####
!####
!#### _6bone_ PEERING ####
!####
```

```
!#### Sets the BGP attributes for 6bone routes ####
route-map 6BONE_IN permit 10
 set local-preference 90
 set community 3352:3109
!
!#### Advertise _ONLY_ MAD6IX prefix ####
route-map 6BONE_OUT permit 10
 match ipv6 address prefix-list TLF_AGGREGATE
!#### Do _NOT_ advertise to 6BONE other prefixes ####
route-map 6BONE_OUT deny 20
 match ipv6 address prefix-list ALL_IPv6
!
!####
!#### Example of 6NET Peering ← Not established, it is just an example
!####
!####  This example could be used for other External Peerings
!####  (Abilene, …) ####
!####
!
!#### Sets the BGP attributes for 6NET routes ####
```

```
route-map 6NET_IN permit 10
 set community 3352:3205
!
!#### Do _NOT_ advertise Euro6IX External routes ####
route-map EXTBB_OUT deny 10
 match community EURO6IX_EXT
!#### Advertise to 6NET all other prefixes (Euro6IX prefixes) ####
route-map EXTBB_OUT permit 20
 match ipv6 address prefix-list ALL_IPv6
!
!####
!####  INTERNAL IXs ROUTEMAPS ####
!####
!
!#### UPM ####
!#### Sets the community attribute to _UPM_ routes ####
route-map UPM_IN permit 10
 set community 3352:6305
!
!#### Advertise _ALL_ the routes to _UPM_ ####
route-map INTPEERS_OUT permit 10
 match ipv6 address prefix-list ALL_IPv6
!
!#### TID ####
!#### Sets the community attribute to _TID_ routes ####
route-map TID_IN permit 10
 set community 3352:6405
!
!#### Do _NOT_ advertise 6BONE routes (optional) to _TID_ ####
route-map INTPEERS_NO6BONE_OUT deny 10
 match community 6BONE
!#### Advertise the rest of the routes ####
route-map INTPEERS_NO6BONE_OUT permit 20
 match ipv6 address prefix-list ALL_IPv6
!
```

```
!
ipv6 prefix-list ALL_IPv6 description *** ALL IPv6 ROUTES ***
ipv6 prefix-list ALL_IPv6 seq 100 permit ::/0 le 128
!
! #### Select only Telefonica´s aggregated route ####
ipv6 prefix-list TLF_AGGREGATE description *** TDATA AGGREGATE***
ipv6 prefix-list TLF_AGGREGATE seq 100 permit 2001:800::/32

!
```

3.4.3.12.2 Example: TOR6IX Router Configuration

This example of TOr6IX router configuration, it is a Juniper router so the configuration syntax is different from the previous example:

```
under "protocols bgp" sub tree
group internal-peers {
    type internal;
    local-address 2001:6b8:kkkk;
    export nhs;
    peer-as 5609;
    neighbor 2001:6b8:0:jjjj;
    neighbor 2001:6b8:0:llll;
}
group Euro6IX {
    type external;
    export Euro6IX-policy-out;
```

```
    family inet6 {
        unicast;
    }
    remove-private;
    neighbor 2001:7A0:100:FF00::2B { "Ber6IX"
        peer-as 3320;
    }
    neighbor aaaa:bbb:cccc:dddd { "Zur6IX" not running at the moment
peer-as aabb;
    }
}
group 6Net {
    type external;
    export 6Net-policy-out;
    family inet6 {
        unicast;
    }
    remove-private;
    neighbor aaabbbcc { "6Net peering" not running at the moment
        peer-as 3320;
}


under "policy-options" sub tree

policy-statement Euro6IX-policy-out {
    term 6bone-prefixes {
        from {
    family inet6;
            community 6bone-prefixes;
        }
        then reject;
    }
then accept;
}

policy-statement 6Net-policy-out {
    term Euro6IX-prefixes {
        from {
    family inet6;
            community Euro6IX-prefixes;
        }
        then accept;
    }
then reject;
}
community Euro6IX-prefixes members "(5609|3320):[10-20].."
community 6bone-prefixes member "5609:3105"
```

3.4.4 Routing Policy Implementation: Second Proposal

This second approach tries to define community scheme based on a KISS approach (“Keep It Simple Stupid”) using a simple and easy-to-implement community tagging scheme to mark the routes coming from the external of Euro6IX network with the aim to avoid any overlapping with communities that every partners could have defined inside the their internal network (Autonomous System). It is based on the same approach of the first proposal, identifying the prefixes using community tags, filtering out disaggregated routes and use of Local preference attribute.

Basically, this proposal foresees that the community tagging will be applied for network announcements entering Euro6IX network. Moreover these announcements will be propagated, with no changes, through the overall Euro6IX routers.

The extended communities structure proposed in this scheme will be composed of two numerical parts that are, in the standard notation, divided by a ":" symbol. The proposed structure is <Euro6IX_ID:commid> where:

- Euro6IX_ID is a number that identifies the Euro6IX network and cannot be overlapped with any of the identifiers of other IPv6 public networks; in particular it cannot be overlapped with the identifiers of networks of partners hosting an Internet Exchange. In fact, if a partner hosting the Internet Exchange has defined some communities having the structure <ASNumber:value> and if Euro6IX_ID was the AS number of the partner, the probability that this community value is overlapped with one of Euro6IX communities is not zero. For this reason the value proposed for Euro6IX_ID is 65535 (the IANA AS reserved value).
- comm_id is a value used to identify what kind of prefixes are announced from one Euro6IX internet exchange to another Euro6IX internet exchange.

The values assumed by comm_id depend on routing policies defined in this proposal. The proposed routing policy is to individuate the following prefixes:

- Euro6IX Prefix community prefixes belonging to Euro6IX network (e.g 65535:0001).
- External Networks transit provided within Euro6IX (e.g. 65535:1000-1999) can be in the future divided in tunneled connections versus native connections.
- External Networks transit within Euro6IX not provided (e.g. 65535:2000-2999). Also in this case in the future it can be divided in tunneled connections versus native connections. Some examples are:
  - 6Bone - 65535:2001
  - 6NET - 65535:2002
  - Abilene – 65535:2003
- Experimental Networks transit within internal IXs not provided (e.g. 65535:3000-3999).

The current community-tagging proposal leaves other community values that could be defined in the next network revision documents because it could be useful to identify some different type of network types of network connections.

As far as we can see, the IX routers processing the incoming route have only to implement a minimal set of operations because no modification is required to the incoming routes.

Moreover this proposal seems to be well scalable because the requested configuration over the router is not dependent on the number of the IXs inside the network (since there is only a single Euro6IX identifier, i.e. the first part of community).

3.4.4.1 Use of Local Preference to Control Routing

The Local Preference BGP attribute is used in the same way as the previous mechanism. In particular the following LP values are recommended:

LP Value	Route Set
140	Internal Euro6IX route
100	All external Euro6IX routes besides 6Bone
60	6Bone route connected directly to IX
40	6Bone route received from Euro6IX peer

Figure 3-16: LP Value to Control Routing

Note: Euro6IX prefixes announced/received via 6Bone (direct i.e. Italy-Spain or Italy-Germany, or indirect) tunnels will not be tagged by any community value. In this case these prefixes should be considered as 6Bone route using suggested local preference value 60.

The application of this community tagging scheme to routing policies defined in §3.3 are:

3.4.4.2 Prefixes Received by other IXs:

`match community 65535:0001 (Internal) ->`

set local preference to 140

announce to other IXs (leaving the community) and to other networks (providing transit among IXs)

`match community 65535:1000-1999 (External-Transit among Euro6IX) ->`

announce to other IXs (leaving the community)

`match community 65535:2000-2999 (External No Transit among Euro6IX) ->`

announce to other IXs (leaving the community)

`match community 65535:2000 (6bone prefix) ->`

set local preference to 40

announce to other IXs (leaving the community)

3.4.4.3 Prefixes Received by other Networks:

external network

`add community 65535:1000-2999 (depending on the network)`

if 6bone set local preference to 60

3.4.4.4 Prefixes Advertised to all Networks (Euro6IX and others):

Euro6IX IXs ->

```
      announcing prefixes matching community 65535:1 and 65535:1000-1999 and
65535:
```

network (transit provided to all IXs) ->

```
      announcing prefixes matching community 65535:1 and 65535:1000-1999
```

Other custom actions (that are in general parts of Internet Exchange Administrator) can be done for particular networks can match particular value of communities.

3.4.5 Election of the Routing Policy Proposal

The two proposal defined in the previous points have been discussed since December 2002. After several months of discussion and with the contribution of all the Euro6IX partners, the first proposal was chosen.

All the partners agreed including both proposals into this document, so both of them could be compared, even if first of them is the Routing Proposal to be implemented in the Euro6IX network.

4. INTERCONNECTION WITH OTHER IPV6 BACKBONES

4.1 Concrete Interconnections Design

4.1.1 6Bone

Commercial (Desired) Situation: All IX having a native connection and peering with the 6Bone.

Intermediate solution: Some IX peer with the 6Bone natively, others peer through a tunnel and others use temporarily another IX connection and peering.

4.1.1.1 BER6IX

BER6IX is interconnected with the 6Bone using IPv6-in-IPv4 tunnel. BER6IX announces the prefix 3ffe:80a0::/28 which is assigned to Berkom and RIR (2001:7a0:/32) space assigned to Deutsche Telekom.

Most full peerings are within European region to avoid long RTT between 6Bone sites. Offers 6Bone NLA space to many leaf sites and run an IPv6 Showcase for German leaf sites using the RIR space. Aggregation and ingress-, egress-filtering of invalid IPv6 address space is made on all external peering routers.

4.1.1.2 LIS6IX

LIS6IX is announcing the sTLA 3ffe:4000::/32, and the commercial pTLA 2001:08A0::/32 (RIR) prefixes, both assigned to Telepac (the PT Service Provider). The peering with 6Bone is made using BGP4+ via an IPv6/IPv4 tunnel through the Internet. This connection is not actually performed on the router co-located on the LIS6IX, but in another one on Telepac backbone.

4.1.1.3 LON6IX

The 6Bone is a global IPv6 Test-bed that is predominately built by tunneling IPv6 over the existing global IPv4 Internet. Its purpose is to test standards and implementations and increasing transition and operational procedures.

The LON6IX (layer 3 part) has a 6Bone address allocation (pTLA) of 3ffe:2c00::/24 and as such peers with other sites in the 6Bone default zone. The LON6IX also allocates 6Bone address space to its customers and downstream providers and peers privately with a number of other 6Bone sites. The resulting 6Bone routing table that is filtered to only contain 6Bone backbone sites (/24, /28 and /32) is combined with the filtered commercial routing table to form the “standard” routing table. There are a number of other routing tables generated at the LON6IX for specialist purposes.

4.1.1.4 MAD6IX

MAD6IX is receiving the full routing table of 6Bone routes from TID. TID is receiving it from Telefónica Data (Internet Service Provider). The peering between TID and Telefónica Data is made using BGP4+ via a native IPv6 connection. The peering between Telefónica Data and 6Bone is made via an IPv6/IPv4 tunnel over the Internet. Telefónica de España is planning to peer with 6Bone directly in the MAD6IX premises.

4.1.1.5 PAR6IX

PAR6IX is receiving full routing tables from various peering of OpenTransit network, to which PAR6IX belongs to, with different 6Bone sites. PAR6IX does not announce any pTLA, and FTR&D hasn't got any 6Bone prefix allocated.

4.1.1.6 TOR6IX

The TOR6IX is connected to 6Bone network and it is owner of a pTLA as a 6Bone backbone site. It offers connection to some labs and universities that are configured as leaf sites of TILAB 6Bone pTLA. For this reason TILAB announces over 6Bone, using BGP4+ only its aggregated pTLA address and doesn't break the aggregation inside the 6Bone network.

4.1.1.7 ZUR6IX

Basel is currently connected to 6Bone and 6NET, through the 100 Mbps link to SWITCH, using the following address/prefix: 2001:620:204::/48. Within the middle of March the ZUR6IX connectivity will be made available by Swisscom (Main Switzerland telecom provider) and the routing exchange information will be set up in conformity with the other Euro6IX nodes. At that time the Basel node will be configured in order to be interconnected to the ZUR6IX node as well.

4.1.2 6NET

4.1.2.1 BER6IX

A 6NET connection using the native peering over 155 Mbits POS with the DFN in Germany is still under negotiation. If the decision is made to use this connection we must define a routing policy inside the Euro6IX network to keep the routing in proper state and avoid transit traffic inside the Euro6IX network and inside every IX.

4.1.2.2 LON6IX

Interconnection between Euro6IX and 6NET is achieved in LON6IX.

Interconnection between the LON6IX (layer 3 part) and the 6NET network is achieved by one of the LON6IX customers (UKERNA) being a project participant of the 6NET project and providing transit between the LON6IX and 6NET network. UKERNA has a native IPv6 fast Ethernet connection to the LON6IX and hence the transit connection is of adequate speed. Routing between the LON6IX (and its peers) and 6NET is carefully controlled via a combination of community tagging and filtering. In the 6NET to LON6IX direction, UKERNA sends the LON6IX a routing table that has the 6NET routes marked with a community tag, the LON6IX uses this information to only forward the tagged routes (i.e. 6NET routes) to the Euro6IX partners that the LON6IX natively peers with, i.e. France Telecom and Telefónica. In the LON6IX to 6NET direction, a list of Euro6IX AS numbers has been given to UKERNA and they filter the full routing table received from the LON6IX before forwarding to the 6NET network. The filtering on AS numbers may be changed to community tagging as it provides a more flexible approach where the LON6IX can control the injection of Euro6IX routes into the 6NET network, i.e. when a new Euro6IX AS number has native connectivity to the LON6IX transit to the 6NET network can be activated without reconfiguration of URERNA's network.

4.1.2.3 PAR6IX

PAR6IX is getting routes to 6NET via Renater, through out a peering connection between OpenTransit and Renater 3 network, over a tunneled link. The upgrade of the tunneled link to an IPv6 native link is under study.

4.1.2.4 TOR6IX

TOR6IX planned to be interconnected to 6NET network using a gigabit Ethernet interconnection between Telecom Italia Lab and Politecnico di Torino, one of the sites of the Italian part of 6NET network.

At the moment there is no additional routing policy defined to exchange routing information between TOR6IX (or in general Euro6IX) and 6NET.

4.1.3 Other Networks

The LON6IX peers with an ever increasing range of networks, transit to these networks is provided via the full routing table being made available to the Euro6IX partners that the LON6IX natively peers with (France Telecom and Telefónica). Of course it depends on the routing policies of these partners if they forward these routes to other Euro6IX partners.

For the publicly available list of LON6IX members please look at: http://www.uk6x.com under Operational Info, Public Members List.

The LIS6IX has another important peering with FCCN, the Portuguese NREN and the owner of the actual Portuguese IXv4 (GigaPIX). They have already started a pilot of the public PIXv6 a few months ago.

5. SUMMARY AND CONCLUSIONS

The technical analysis realized in the first year inside the A2.2 activity has been summarized in this D2.2 deliverable.

This analysis has been focused in the different IX models proposed, and the routing policies to apply in the Euro6IX backbone. After several months of discussion in the WP2 mailing list, the majority of the Euro6IX partners selected one of the proposals. That proposal is the one to be implemented in all the Euro6IX IXs.

From this discussion, Euro6IX network must grow with the IXs models proposed in the first six months of the project, and the routing policies detailed in this D2.2 deliverable.

6. REFERENCES

[1] 6Bone web. See: http://www.6bone.net

[2] 6NET web. See: http://www.6net.org/overview.html

[3] An IPv6 Aggregatable Global Unicast Address Format. RFC 2374. July 1998.

[4] GÉANT web. See: http://www.dante.net/geant/about-geant.html

[5] Abilene Network Operations Center web. See: http://www.abilene.iu.edu/ndoc.html

[6] "Specification of the Internal Network Architecture of each IX point". Euro6IX D2.1 Deliverable. August 2002. See: http://www.euro6ix.org