

Title:	Deliverable D2.4 Result on Research on New Address Delegation Strategies for IPv6 IXs	Document Version: 1.0
---------------	--	-------------------------------------

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

Contractual Delivery Date: 15/01/2004	Actual Delivery Date: 13/06/2004	Deliverable Type* - Security**: R – PU
---	--	--

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

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

Responsible and Editor/Author: Jonathan Stevens	Organization: BT	Contributing WP: WP2
---	----------------------------	--------------------------------

Authors (organizations): Jordi Palet Martinez (Consulintel), Mario Morelli (TILAB), David Fernández (UPM), Jose Luis Rubio (UPM), Jose Antonio García (UPM).
--

Abstract: IPv6 IXs can perform address delegation. This deliverable discusses what IX address delegation means, how it can be implemented in theory in practice.

Keywords: IPv6 address delegation, IX, L3MF, Route Server.
--

Revision History

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

Revision	Date	Description	Author (Organization)
v0.1	13/11/2003	Document structure created	Jonathan Stevens (BT)
v0.2	14/11/2003	Contributors identified	Jonathan Stevens (BT)
v0.3	19/01/2004	TILAB and contribution added	Jonathan Stevens (BT)
v0.4	30/01/2004	BT contribution added	Jonathan Stevens (BT)
v0.5	27/02/2004	UPM contribution added	David Fernández (UPM)
v0.6	04/03/2004	Reorganized and tidied up	Jonathan Stevens (BT)
v0.7	23/03/2004	Conclusions added	Jonathan Stevens (BT)
v0.8	06/05/2004	Extra Diagram Added	Jonathan Stevens (BT)
v0.9	28/05/2004	Merge of two parallel versions. Summary	Jordi Palet (Consulintel)
v1.0	13/06/2004	Final Review	Jordi Palet (Consulintel)

Executive Summary

This document describes the requirements for the address delegation process of the IPv6 IX and the Layer 3 Mediation Function (L3MF) as a means of performing IPv6 IX based address delegation.

It goes on to describe possible implementations of the L3MF, and some currently implemented solutions that are variations of L3MF.

Table of Contents

1.	<i>Introduction</i>	6
2.	<i>Requirements</i>	8
2.1	Requirements of IX Address Delegation Process	8
3.	<i>IX Business Models</i>	9
3.1	Business models	9
4.	<i>A Solution: Layer 3 Mediation Function</i>	10
4.1	Introduction to Layer 3 Mediation Function	10
4.2	IX Route Server based Architecture	12
5.	<i>Implementations</i>	15
5.1	TILAB TOR6IX Implementation	15
5.2	UPM's Route Server based IX Address Delegation Architecture	17
5.3	BT LON6IX Implementation	19
5.3.1	LON6IX summary	19
5.3.2	Summary of variations to L3MF	20
5.3.3	Summary of issues with variations to L3MF	20
6.	<i>Address Delegation Tools</i>	21
7.	<i>Summary and Conclusions</i>	22
8.	<i>References</i>	23

Table of Figures

Figure 1-1:	IPv6 Routing Model following RFC2374.....	6
Figure 4-1:	D2.2 IX Model	11
Figure 4-2:	Route Server based IX Model.....	12
Figure 5-1:	TOR6IX Test-bed Network	15
Figure 5-2:	BGP Peering Scheme.....	18
Figure 5-3:	LON6IX Architecture	20

1. INTRODUCTION

This document examines the area of address delegation by an IPv6 IX. It describes a particular mechanism that has been developed, in detail, and then describes some variations to this mechanism.

IPv6 proposes a strictly hierarchical routing and addressing model that essentially follows the principles stated in CIDR: Hierarchical assignment of addresses and routing based on aggregation. Figure 1-1 shows the overall structure of this routing model.

Roughly, IPv6 addresses assigned to an organization are divided in three parts:

- A global routing prefix, which defines precisely where the organization is connected to.
- A subnet identifier, assigned following the internal topology of the organization.
- An interface identifier, which identifies a system inside a link.

The novelty comes from the fact that the hierarchical assignment of global routing prefixes is compulsory. This means that the prefix which an organization receives depends on the provider it is connected to. For example, in Figure 1-1 prefix assigned to Site E must be a sub-prefix of provider P6, and the prefix assigned to P6 must be a sub-prefix of P4. In this way, all addresses assigned to providers and customers below provider P4 can be aggregated in a few prefixes, reducing considerably the routing table size of “big” routers in the DFZ zone.

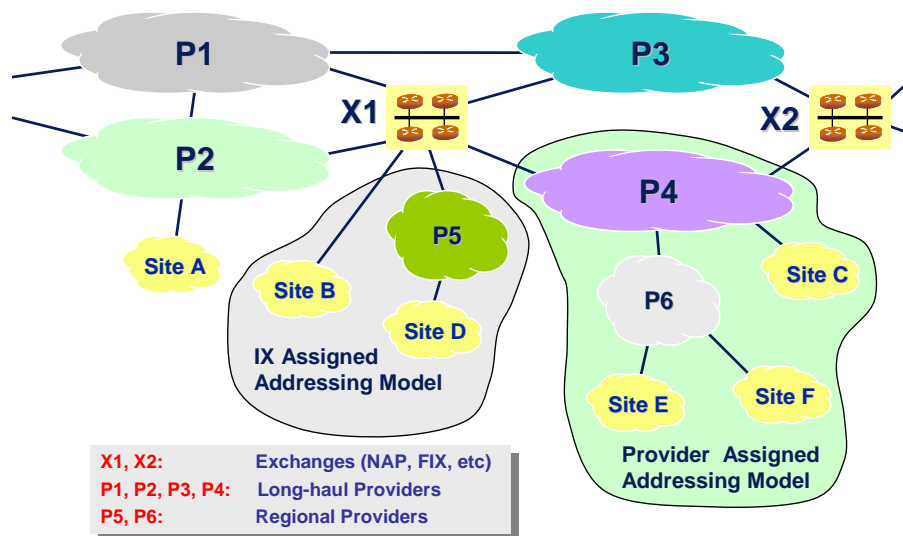


Figure 1-1: IPv6 Routing Model following RFC2374

Provider independent (PI) addresses (that is, addresses assigned directly to organizations independently of its location), which are nowadays the base for IPv4 multi-homed sites, do not exist in IPv6 anymore. As a consequence, if a site changes provider its global prefix must be changed according to its new location in the global topology.

Changing the prefix means renumbering the whole site network, which is normally a difficult and costly task. Fortunately, IPv6 has been developed keeping renumbering in mind, and several autoconfiguration methods have been defined to make network renumbering easier.

Another consequence of IPv6 routing model is that IPv4 multi-homing techniques are not allowed. Multi-homing in IPv4 is mostly based on the use of PI address prefixes that are announced through all organization's provider connections. New techniques to allow IPv6 multi-homing in a scalable way are being investigated, as multi-homing is a much-demanded feature to achieve reliability or load balancing. Although no clear solution has been defined yet, multi-homing is a key feature that clearly will have to be worked out for the success of IPv6.

In order to avoid renumbering when changing provider, as well as to allow a geographically restricted way of multi-homing, IPv6 routing model defines a new way of assigning addresses centered on IXs. It basically consists on allocating prefixes to IXs that are later sub-assigned to organizations connecting to the IX (Figure 1-1).

This model decouples the address assignment from the connectivity provision. On the contrary to the classic provider assigned addressing model, where customers get addresses and connectivity from the same provider, in this model, the customer gets addresses from the IX and gets connectivity from one or more of the providers peering at the IX.

The advantages of this new IX based aggregation model are:

- Customers can change provider without having to renumber their network.
- Customers can multi-home with two or more providers on the IX without causing scalability problems, because providers do not need to propagate customers specific prefixes, only the whole prefix assigned to the IX.

However, this model is only roughly defined in IPv6 documents. There are a lot of details to be studied and experimented before being able to deploy it. For instance, the exact way in which customers connect to the IX, the way routing is organized between customers and providers, the business model behind, etc.

2. REQUIREMENTS

IP address delegation is the familiar process of assigning or allocating blocks of address space. In the IPv4 world this is a relatively simple process, handled by a hierarchy of address registries, starting with the IANA at the top, descending through the Regional Internet Registries (RIRs), in some cases to LIRs (Local Internet Registries) or ISPs.

IPv6 address delegation has a similar hierarchy, IANA, RIRs, local registries. However, address aggregation requirements mean that delegation is often tied to packet routing. Upstream network provision should come from the network provider that delegated the address space. However, for IXs, which can delegate IPv6 addresses, there may not be a backbone network to provide network connectivity with.

2.1 Requirements of IX Address Delegation Process

An IPv6 IX Address Delegation process should:

- Allow an IPv6 IX to delegate address blocks to other connecting networks.
- Allow the connecting networks to transit traffic via other connecting networks.
- Allow some means of control over routing policy.

3. IX BUSINESS MODELS

3.1 Business models

The new capabilities of IPv6 IXs could mean that they may have to operate using a new business model.

As the market is not yet mature these models are an open area of research, and it is not clear what will be commercially successful, or even commercially viable, in the future.

A future Euro6IX deliverable, will examine IPv6 IX business models.

4. A SOLUTION: LAYER 3 MEDIATION FUNCTION

4.1 Introduction to Layer 3 Mediation Function

One of the main goals to reach inside the Euro6IX project was to find a new model for the Internet Exchange (IX), starting from what is described in the RFC 2374 “An IPv6 Aggregatable Global Unicast Address Format”, where a new way to assign the IPv6 addresses was proposed, identifying the possibility for new entities, i.e. IXs, to assign IPv6 prefixes to its own customers.

This new way of thinking of the IX makes easier some operation like the renumbering of the providers’ network hosted by the Internet Exchange.

In fact, in IPv6 (as it happens also in IPv4), end users change their IPv6 prefix when they change connectivity provider. On the contrary, using this model, IX customers changing provider, do not need any renumbering procedure, because, as shown with more detail in the following, the prefix assigned to the customer belongs to that Internet Exchange and not to the Long Haul Provider. In particular, this is possible thanks to a mechanism decoupling the customer’s and the connectivity provider’s roles and called Layer Three Mediation Function (shortly, L3MF).

The work we are referring to is the result of the activity carried out inside the Work Package 2 during the first two years of the Euro6IX project and that can be summarized in this way:

- Functional definition of the architecture of the IX and of the L3MF functionality.
- L3MF Implementation.
- Basic L3MF tests.

The functionality we are referring to was called Layer 3 Mediation Function, because it plays an intermediation role between the ISPs and the customers getting access to the IX.

Figure 4-1 shows the model we are talking about, which is described in detail in Euro6IX Deliverable D2.2.

Briefly, the IX hosts the router belonging both to the “traditional” customers (Standard IX Customers, using the IX as a Neutral Access Point), and to the “new” customers (Next Generation IX Customers that uses the new functionality provided by the IX).

In other words, the IX becomes a meeting point between the customers and the ISPs (LH ISP1, LH ISP2 and LH ISP3). The customers can choose which ISP is more convenient for their aims (based on the fee they pay, time of the day, etc.) and changing it in a very easy way. In turn, the ISP finds, inside the IX, a point where the customers come to, taking advantage of this customer aggregation.




Figure 4-1: D2.2 IX Model

Moreover, note that, in the traditional IXs, the relationships are normally between two entities (ISPs) that decide if making a peering or a transit agreement. On the contrary, the model here analyzed has got three or more entities involved: A customer, one (or more) ISP and the IX itself. These entities, from a commercial point of view, can interact in some different ways and this may open new business scenarios to be applied to real world.

From a theoretical point of view, this functionality hence has the main objective to solve the problem to assign the addresses provider independent to the ISPs accessing the IX.

In fact, usually, all IPv6 addresses are provider dependent, in the sense that, they are assigned by a provider and, as consequence, the structure of the IPv6 address reflects this hierarchical structure, having a section depending of the particular Provider they have subscribed to. So if a customer changes provider, the network of the customer will have to be renumbered, according the mechanism provided by the IPv6 protocol.

IX addressing based model makes easier this mechanism. In fact, in the best case, customers changing provider will not have to change the address they are using, because the addresses will be assigned by the Internet Exchange.

Additionally, another advantage related to this model is the possibility to manage the multi-homing more easily (each ISP accessing the IX is naturally multi-homed).

There are different ways to implement this mechanism. First one is to use a router connected to the Internet Exchange switch. This is the solution adopted in TOR6IX and is based on a BGP4+ approach.

Another possible solution could be to use a layer three switch, having also the L2 and L3 functionalities. This way, the switch could, in the same time, interconnect the different ISPs and to manage the routing between the different routers. This solution although seemed not viable in our particular case, above all from an economic point of view, because it will mean replacing all the switches in the Euro6IX network, but it could be useful in other situations.

Another solution could be that one of using virtual LANs to connect the customers to each LHP router, but it has not been adopted because it appeared not scalable. In fact, in the worst case, as many VLANs as the LHPs are needed, one for every customer connected to the switch.

4.2 IX Route Server based Architecture

The use of the L3MF, described above, to organize address delegation in an IPv6 IX can be combined with the use of a route server. Typically, route servers are used to centralize the peering between the different ISPs connected to the IX. In our case, apart from this important function, the route server will act as well as the mediator for IX clients. Figure 4-2 shows the IX model including the Enhanced Route Server.

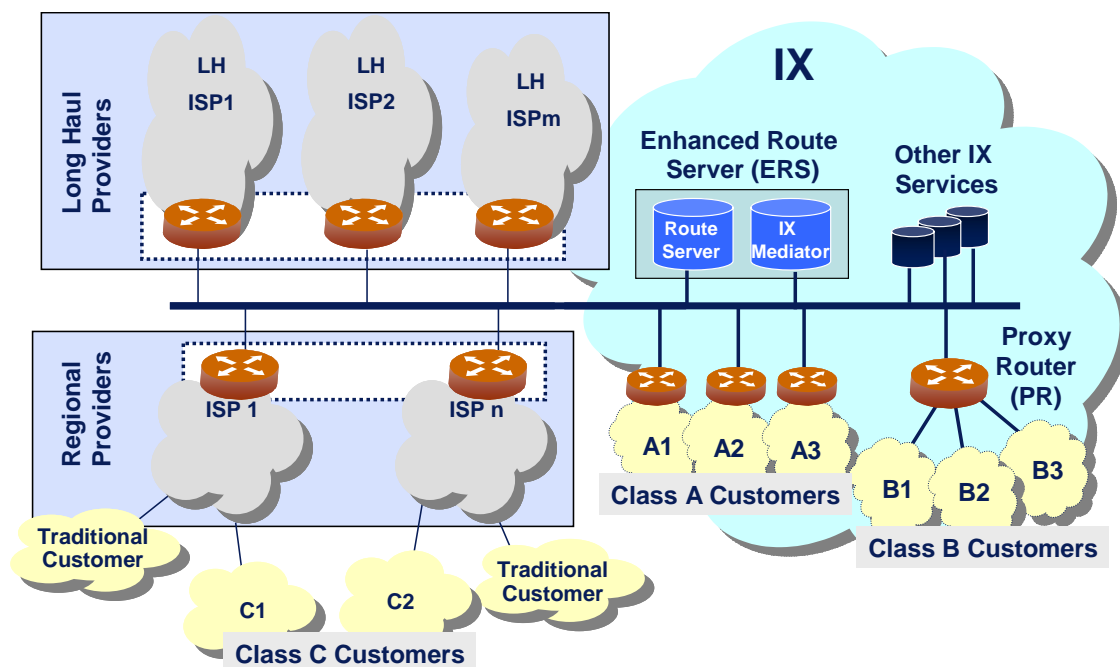


Figure 4-2: Route Server based IX Model

The core component of the IX is a layer-2 fast switching infrastructure that interconnects all layer-3 devices from:

- **Long Haul Providers**, which peer among them in the IX and offer transit services to Regional Providers.
- **Regional Providers**, which also peer among them and get transit services from Long Haul Providers.
- **IX equipment**, including an Enhanced Route Server (ERS) that, apart from the general route server functionality will include the “layer 3 mediation function” for the IX customers. IX could also include equipment to support additional services offered by the IX, not described in this document.

Traditional ISP customers connect directly to ISPs networks, as it is shown in the figure. They typically inherit their addresses (usually a /48 prefix¹) from the ranges assigned to their ISPs (one or more /32 prefixes). As the new IPv6 IX model provides the possibility to have customers with addresses inherited from the ranges assigned to the IX (one or more /32 prefixes as well), three types of IX customers have been identified:

- **Class A**, which are connected through their own router deployed in IX premises operated by them (A1, A2 and A3 in the figure).
- **Class B**, which are connected through an IX operated proxy router (PR), shared with other customers (B1, B2 and B3 in the figure).
- **Class C**, which are connected to an ISP, as the traditional ISP customers, but they use addresses from the IX range (C1 and C2 in the figure).

Note that the distinction between class A and class B customers is more administrative than technical, because a set of class B customers together with their proxy router can be considered as a class A customer.

The use of a proxy router for all or part of class B customers simplifies the IX management as it reduces the number of routers involved. However, class B customer's service choice is limited, due to the fact that all of them share the same routing policy in terms of ISPs selected or multi-home configurations. A possible solution to this problem could come from the use of virtual routers (that is, one physical router running several routing daemons in parallel, each one giving service to a different customer).

The objective of our IX model is to study and experiment how these different types of IX customers can benefit from the new services that the IX assigned addresses allow, that is:

- The choice of the provider they want to receive service from.
- The change of provider without renumbering.
- The possibility of multi-homing with two or more providers, in order to achieve, for example, fault tolerance, load sharing or more advanced services.

In our model, although Route Server and IX mediator are clearly two different functions, they will be implemented together using a modified route server implementation. Initially, we studied the possibility to separate them in two different devices, in order to keep the route server as it is nowadays and improve fault tolerance (if IX mediator fails, the route server could still be providing interconnection service for the ISPs connected to the IX). In that case, the two devices would interchange routing information through a BGP session.

However, this alternative shows a serious trouble, because the route server, as any conventional BGP entity, applies the BGP decision algorithm to routes received from the ISPs. For instance, if two or more ISPs advertise the same prefix route, the route server will choose one of them as the best route (by applying the standard BGP algorithm) and will propagate only this one. As a consequence, the mediator and the customers behind it will not receive the complete routing information they need.

Although the BGP algorithm could be modified in the route server in order to cope with this problem, to simplify the model and concentrate this study on how to offer the services demanded by IX customers, the decision was made to merge the mediation function and route server functionalities.

In the model proposed, all provider routers, class A customer routers and proxy routers servicing class B customers peer with the ERS using BGP. Therefore, the ERS will be responsible of implementing provider routing policies as well as customer choices.

Each ISP connected to the IX will have its own public Autonomous System Number (ASN), that will be included in the routes generated or propagated by it. In the case of the IX, private ASNs can be used, for example, by the route server when peering with ISPs. However, note that appropriate filters will have to be set up to remove these private ASN in the AS_PATH of the routes announced to others ISPs.

The use of an enhanced route server to implement the address delegation functionality allows a better integration of IX clients inside the IX architecture. As it will be show in the next section, IX client policies could be defined and controlled using the same tools used nowadays to manage route server based IXs.

5. IMPLEMENTATIONS

This section describes the various possible implementations that have taken place within the project. The various IX partners within the Euro6IX consortium all have slightly differing requirements, both technical and business, and therefore all the IXs are different. This diversity allows a useful variety of implementations to potentially exist and interconnect.

5.1 TILAB TOR6IX Implementation

As already said, the solution adopted in TOR6IX is based on the introduction of an “entity” (i.e. a router) inside the IX managing the routing and having the following features:

- It is a neutral entity, not belonging to any of the ISPs. The Layer 3 Mediation Function has to be hence placed inside the IX building and has to be managed by the IX technical staff.
- It has, at least, to manage the routing between the different entities coming to the IX.
- It does not have to be involved in forwarding mechanism, because the traffic has to flow directly from the customer to the ISP and it has not to transit on any other external equipment (because in a Layer 2 IX, all the traffic flows through the switch).
- It has to be possible to modify its configuration in a dynamic way.
- It could be involved in the IPv6 address delegation.

This mechanism was implemented through a router that is located in the same building as the Internet Exchange switch and it will have the same AS number as that one owned by the Telco owning the IX.

Following scheme reports the test bed developed in TOR6IX.

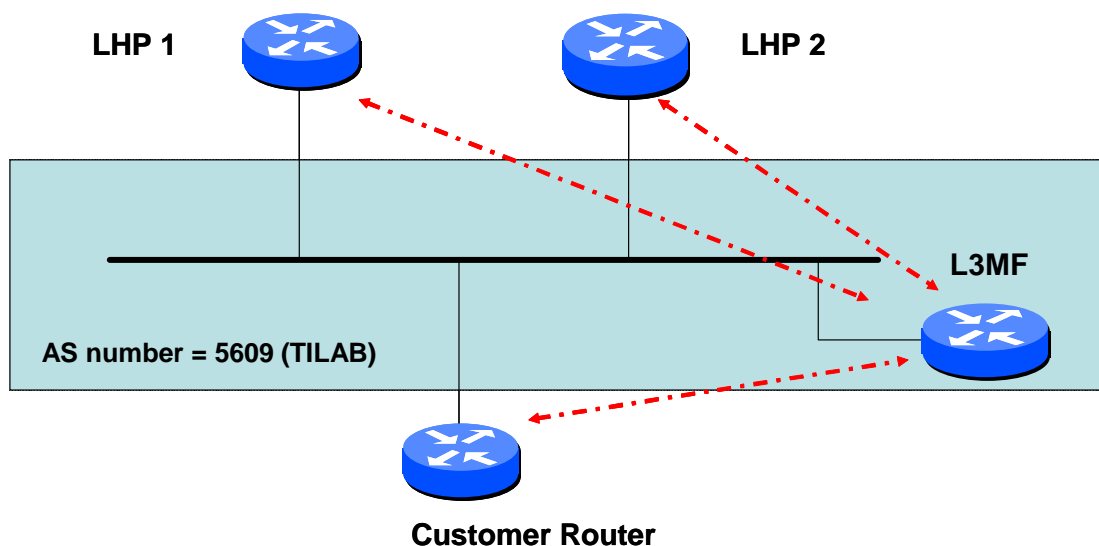


Figure 5-1: TOR6IX Test-bed Network

The main aim is to show that, working on the routing inside the IX, is possible that a customer connected to an IX can change the provider chosen for the long haul connectivity. Consequently it was needed to emulate at least two Long Haul Providers and one customer accessing the IX.

In the particular case, a Juniper router (emulating Euro6IX Long Haul Provider) and a Cisco Router (emulating the 6Bone Long Haul Network Provider) were used. These routers are connected to the layer 2 switch that represents the core of the IX network (represented in the picture as a simple LAN segment).

The router that provides the Layer 3 Mediation function is a Cisco router that has been connected to the IX switch. The customer router is always a Cisco router connected to the switch.

According the TOR6IX model, L3MF router does not have to forward the traffic but it has only to exchange the routes between the parties involved in the process.

Between the routers, the routing protocol is:

- eBGP between the customer and the Layer 3 Mediation Function router.
- eBGP between the Layer 3 Mediation Function router and the router emulating the Long Haul Provider 6Bone (indicated as LHP1 in the previous picture) and between the L3MF router and the Euro6IX router (indicated as LHP2).
- No routing protocol running between the customer router and the LHPs routers.

The prefix associated to the IX is that one owned by TILAB, i.e. 2001:6b8::/32. From this prefix, the customer prefix is picked up. The prefix used for Euro6IX network is the 3ffe:4011::/32.

The administrator of the L3MF (i.e. the IX owner that provides the service) can modify, upon request of the customer, the BGP rules to make possible that the customer itself uses a different provider and this can be made using the BGP protocol.

Basically, this solution is based on the possibility to influence the direction of the incoming and outgoing traffic using the well-know Next-Hop BGP attribute (for the upstream direction) and the No Export community (for the downstream direction).

In particular:

Upstream direction

The traffic in the upstream direction comes from the customer to the long haul provider. In order to do that, L3MF router receives by eBGP routes received from the routers LHP1 and LHP2, re-advertising to the customer only those routes coming from the LHP selected for the long haul connectivity (based on the next-hop attribute).

Downstream direction

Traffic in the reverse direction can follow every time different paths depending on the routing outside the LHP. It means that, even if the direct path is always fixed, this is not true for the reverse path that can change, causing asymmetrical routing (direct and reverse path are different).

The reverse path is not under control because it depends on the routing outside the LHPs network. This problem could be avoided if it was possible to advertise the more specific IX routes (i.e. /48 prefix associated to the IX customer) outside the Euro6IX network, but this is not possible because of the aggregation policies (to avoid the increase of the routing table of the Internet routers).

A possible solution, currently used in TOR6IX L3MF implementation and anyway not fully solving the problem, is to let the customer prefix circulate inside the LHP network (but not outside) using eBGP no-export community to tag the routes from L3MF router to the LHP routers. The more specific routes (/48 prefix) would be visible inside the chosen LHP network (towards the internal BGP peers) while, outside this network, this unaggregated prefix would not be visible.

This is the problem commonly faced in all multi-homing scenarios, where a trade-off between the unaggregation of the routes and the asymmetric routing has to be reached. Consequently, the upstream path (from the customer to the LHP) is fixed, but not the downstream path. This could not be a problem (most of the routing in Internet is not symmetric) and in any case this is not the problem we are going to solve.

Basic tests, carried-out up to now, cover the basic functionality in order to verify the possibility to change the provider, when requested. These first tests showed that L3MF functionality implemented using eBGP is able to do this change.

The solution practically adopted by TOR6IX uses a route-map that filters the routes directed to the customer router based on the next-hop attribute.

In the tests carried out up to now, long Haul Provider are 6Bone and Euro6IX. When changing next-hop based filter, test packets (ICMPv6 traceroute packets) flow directly through the selected LHP routers. As said, only the forward path is selected and the reverse path is not dependent by LHP or L3MF routing choices.

L3MF router makes a job similar to that one made by a Route Reflector router. In fact, it is not involved in the forwarding process and it is used only for exchanging the routes between the routers (customers and providers).

For this reason, it could be interesting to merge together two functionalities, as already indicated in some papers produced by UPMⁱⁱ. This could be done considering a single box inside the IX that works as Route Reflector when two or more ISPs interconnect between each other and as L3MF router when a customer wants to connect to an ISP.

5.2 UPM's Route Server based IX Address Delegation Architecture

UPM's effort has been focused on the implementation of the tools needed to emulate a complete IX scenario. Basically, a general purpose emulation tool has been created for the emulation of real IX scenarios. This tool, named Virtual Network User Mode Linux (VNUML)ⁱⁱⁱ is able to emulate complex scenarios made of tenths of peering routers interchanging an important amount of prefixes, with or without a route server, using few physical devices (ideally a single PC based box).

Besides, as the emulated scenarios are based on linux operating system, an IPv6 route server implementation was developed for the open source "quagga"^{iv} routing daemon (derived from the well-known zebra project). Apart from the basic BGP route server functionalities, this implementation will include some specific modifications for the management of IX clients.

The detailed study of the different scenarios that arise from the IX model presented in the previous section is now underway. The precise scenarios for each type of customer are being defined, as well as their emulation using the VNUML tool.

The Figure 5-2 presents the general peering schema used for the scenarios. As mentioned before, all routers, either from providers or IX customers, peer using external BGP sessions with the ERS.

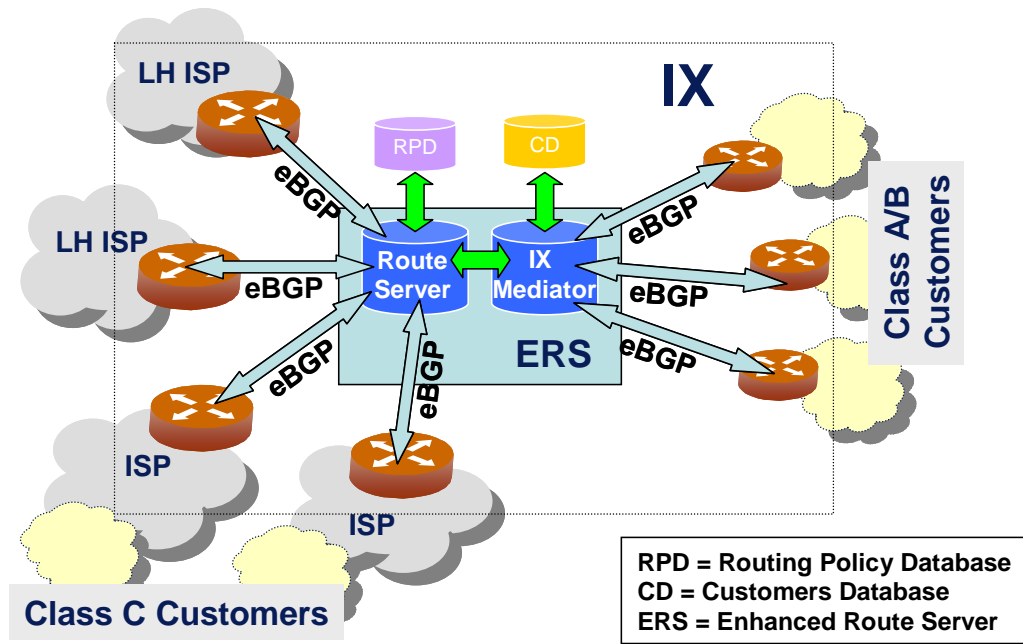


Figure 5-2: BGP Peering Scheme

Depending on the service demanded by each customer (stored in the Customers Database, CD), different information is sent by the IX Mediator through the customer's router BGP session. For instance, in the simpler case where the customer is single-homed, it will only announce a default route with the next-hop field pointing to the selected ISP router. In case of customers that demand to be multi-homed with two ISPs for fault tolerance reasons, the IX Mediator could announce a default route to one of the providers, changing to the other in case of failure. Finally, if the customer demands a strict control over its routes, it could receive the full BGP tables from the mediator.

Although route control is centralized by the ERS, data traffic goes directly between IX routers using the layer-2 switching infrastructure. ERS uses BGP's next-hop attribute to direct customer's traffic to its provider router. In case the provider changes, the CD has to be modified and the ERS must adjust the next-hop attribute according to the change.

When the traffic sent by a customer has its destination on a different provider from the one it gets service from, the traffic will traverse the layer-2 infrastructure twice: One to reach customer's provider router and another one to reach the router of the next provider in the way to the destination. This is clearly inefficient. It could be solved using redirecting techniques; however, it is not a major problem, as nowadays layer-2 switching devices are powerful enough to cope with that traffic. Besides, security reasons could preclude direct connections between customer and provider's routers (for example, a provider could include layer-2 filters in its router to only allow traffic from routers it is peering with).

The main limitation of the scheme proposed is that it only controls the choice of ISPs for customer's outgoing traffic. This is an intrinsic limitation, because the path followed by incoming traffic must be the same for all customers, as the whole IX prefix (typically a /32) is announced to ISPs. In order to have different paths for each customer, the IX prefix should be announced in a disaggregated way, breaking the hierarchical routing principle.

In the case of Class C customers, each ISP will have to announce through their peering session with the ERS the sub-prefixes assigned to its class C customers (typically a /48), as well as to distribute that prefix inside its intra-domain routing. In this case, the IX prefix is being distributed in a disaggregated way. However, that disaggregation is limited in scope, as each ISP only has to carry its class C customer prefixes.

As mention before, the detailed study of the different scenarios presented is already underway and will be finished during the third year of the project. Apart from that, the use of the RPSLng language to define the IX client policies will be evaluated. This will allow the inclusion of IX client's routing policies in a routing policy database like the ones being used nowadays for IPv4, and its automatic processing using new management tools.

5.3 BT LON6IX Implementation

5.3.1 LON6IX summary

The LON6IX is a layer 2 and layer 3 IPv6 IX. It provides layer 2 traditional IX functions to connecting ISPs and companies (using ATM and Ethernet switches), but it also provides the new layer 3 addressing services. Connecting ISPs and companies can be allocated address blocks and can obtain automatic multi-homed transit connectivity.

The traditional layer 2 IX service provided by LON6IX will not be described in detail, as it is fairly simple and similar to IPv4 IX services. Customers interconnect to each other via ATM or Fast Ethernet. Their routers are allocated an IPv6 address from the IX infrastructure prefix.

Most LON6IX customers make use of more advanced layer 3 services. The LON6IX solution uses a layer 3 router and BGP4+, and provides connecting entities with a choice of methods for obtaining their connectivity.

For customers who already have a RIR-allocated address prefix, transit services can be provided by BGP peering with the LON6IX router. The customer advertises, as a minimum, their own prefix, and if they are a transit-capable network they may advertise other prefixes they have reachability for too. Thus the LON6IX router receives connectivity to many other networks.

The LON6IX router advertises either the LON6IX prefixes or some transit routes to the customer, depending on the customer requirements. Thus the customer gains connectivity to many other networks.

For customers who do not have a RIR-allocated address prefix, the LON6IX can allocate them an address block. This address block is part of the LON6IX prefix, and can be further sub-allocated by the customer. Using a BGP peering, the customer advertises this address block to the LON6IX router. The LON6IX router advertises the full transit table to the customer router. Thus the customer gains connectivity to all networks that the LON6IX can reach. The customer does not need to worry about obtaining transit from backbone networks, or negotiating connectivity.

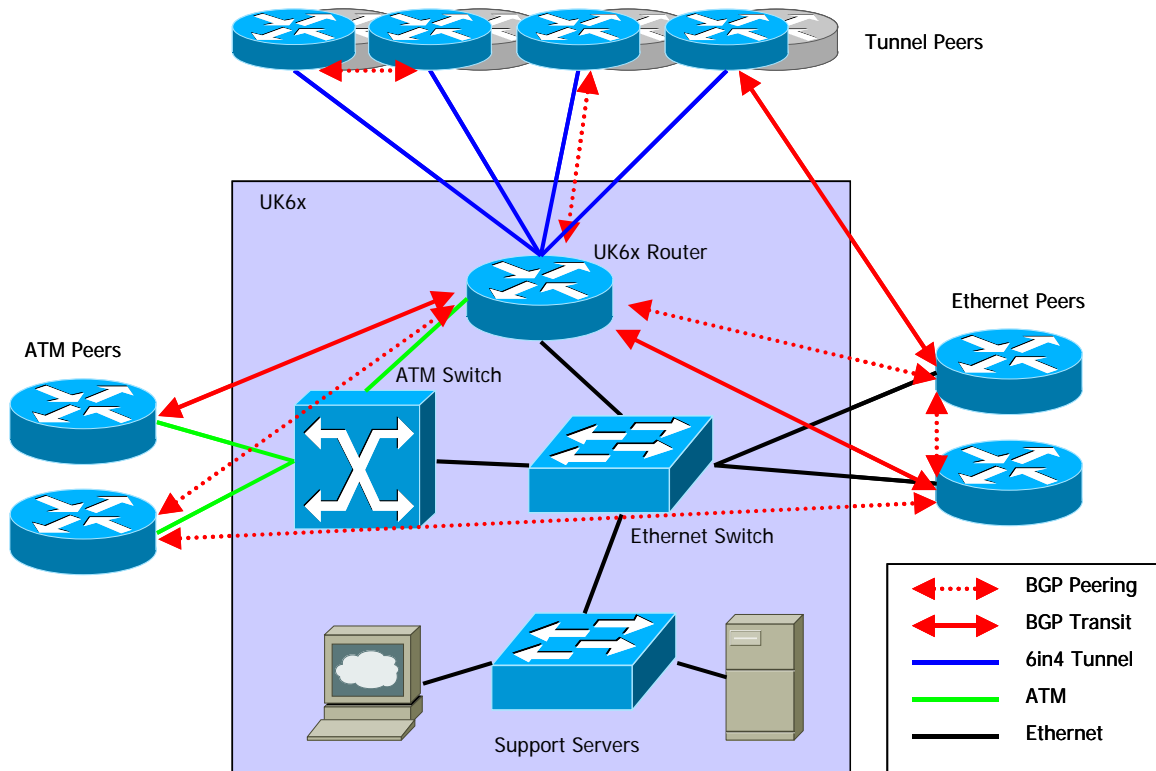


Figure 5-3: LON6IX Architecture

5.3.2 Summary of variations to L3MF

The LON6IX solution is simpler than the L3MF solution. The customers have no choice about the long-haul providers that their traffic uses. These transit decisions are made by the router and are based purely on the BGP policies of the router.

This means that the LON6IX router does forward packets between some customer routers. This happens when both customers do not peer with each other, but do peer with the LON6IX. The customer is relying on the LON6IX to have optimal routing.

5.3.3 Summary of issues with variations to L3MF

This solution is not capable of scaling beyond the level of traffic that can be forwarded by a single router. The actual forwarding of traffic between customers is more scalable when the customer routers send traffic directly over layer 2. It is expected that the LON6IX will implement a solution similar to L3MF for those customer wishing to make choices about the provision of their transit.

6. ADDRESS DELEGATION TOOLS

A simple tool to help manage IPv6 address delegation is a spreadsheet. The LON6IX addressing spreadsheet tool uses macros to automatically convert sequences of binary digits into hexadecimal representations of addresses. This allows sequences of address blocks to be created easily.

Each line represents an address block, and contains details about the allocation of that address block to a customer.

It is perhaps more appropriate to perform this task using a database application, to prevent the accidental duplication of records.

7. SUMMARY AND CONCLUSIONS

The main goal of the design of the Euro6IX project was to study a model of IPv6 Internet where the Internet Exchanges can take a role in the deployment since the early beginning, instead of appearing later when Internet was already wide spread as it was the case for IPv4.

This concept is based on the RFC 2374 “An IPv6 Aggregatable Global Unicast Address Format”, where a new way to assign the IPv6 addresses was proposed, identifying the possibility for new entities, i.e. IXs, to assign IPv6 prefixes to its own customers.

This model decouples the address assignment from the connectivity provision, so the customer can obtain addresses directly from the IX and the connectivity from one or more of the providers present in that IX.

This new IX based aggregation model has to main advantages:

- Customers can change the provider without having to renumber their network.
- Customers can multi-home with two or more providers on the IX without causing scalability problems.

The document describes the requirements for the address delegation process of the IPv6 IX and also a possible solution for providing this function, named Layer 3 Mediation Function (L3MF), together with tools that could facilitate the implementation, evaluation and operation.

The new capabilities of IPv6 IXs could provide a new business model, which might include additional services.

8. REFERENCES

-
- ⁱ APNIC, ARIN, RIPE NCC, “IPv6 Address Allocation and Assignment Policy”, ripe-267, January 2003.
- ⁱⁱ Fernandez, D., Galán F. and de Miguel, T., “Study and Emulation of IPv6 Internet Exchange based Addressing Models”. IEEE Communications Magazine, January 2004.
- ⁱⁱⁱ Virtual Network User Mode Linux Project. Available at: <http://www.dit.upm.es/vnuml>.
- ^{iv} Quagga Routing Suite. Available at: <http://www.quagga.net>.