Internet Engineering Task Force M. Morelli
Internet-Draft Telecom Italia Lab.
Expires: January 10, 2005 J. Palet
Consulintel
D. Fernandez
Technical Univ. of Madrid (UPM)
A. Gomez
University of Murcia (UMU)
July 12, 2004

Advanced IPv6 Internet Exchange model
draft-morelli-v6ops-ipv6-ix-00.txt

Status of this Memo

Copyright Notice

Abstract

Internet Exchanges (IX) have played a key role in the development of Internet, organizing and coordinating the traffic interchange among

Internet Service Providers (ISP). Traditionally, IXs have been based on layer 2 infrastructures, being the layer 3 services managed by the participant ISPs.

However, IPv6 hierarchical and aggregatable routing and addressing model comes to enhance the IX functionalities by proposing to directly assign addresses to IX customer's networks. The customers can connect with one or several upstream providers and have a separated addressing space, dependent on the IX instead of the providers, in order to facilitate multihoming or avoid renumbering procedures when changing the provider.

In addition, being an IX a central point where traffic is concentrated, several networks and application services benefit from the fact of being partial or totally offered from an IX, opening the IX to the world of new advanced services and functionalities like security, AAA, QoS, multicast, mobility, PKI, DNSSEC or policy provision, that could also facilitate the deployment of IPv6 and their required transition mechanisms.

This document describes an architectural model for an advanced IPv6 Internet Exchange that offers IPv6 address delegation services, as well as other advanced network and application services. The document discusses also about how these services can be offered from an IX and their associated requirements.

Table of Contents

1. Introduction

In recent years, Internet Exchanges (IX) and Network Access Points (NAP) have played a key role in the development of Internet.

An IX is commonly a neutral point where different Internet Service Providers (ISP) deploy their routers in order to exchange their traffic according to some kind of commercial and peering agreements (i.e. public peering agreement).

A NAP is basically an enhancement of the IX concept that, apart from a place to host bilateral peering arrangements between similar providers, it takes the role of a transit purchase venue, where regional ISPs can acquire transit services from long-haul or transit providers.

Although there are differences between an IX and a NAP, in the context of this document we will use the term IX to refer to both of them. As described later, the IX model presented here can be classified as an enhanced NAP, adding new services like IPv6 address delegation to a classical NAP.

An IX is usually based on a layer 2 infrastructure, fully redundant and made of high performance switches, where ISPs layer 3 equipment (routers) connect to.

Typically, IXs do not provide any layer 3 services to their customers, apart from route servers or other specific functions that help to organize and simplify routing in the IX. Other services, like AAA or multicast, are normally offered by each ISP.

However, being the IX a central point where the traffic is concentrated, several network and application services being offered nowadays would benefit from being partially or totally supported in the IX itself.

On the other hand, IPv6 proposes a strictly hierarchical routing and addressing model that essentially follows the principles stated in CIDR [1]: hierarchical assignment of addresses and routing based on aggregation. The addresses assigned to an organization depend on the point they connect to the Internet. As a consequence, if the site changes its provider, its global prefix must be changed according to its new location in the global topology.

In order to avoid renumbering when changing provider, IPv6 routing model [2] proposes a new way of assigning addresses based on IXs. It basically consists on delegating prefixes to IXs that are later sub-assigned to organizations connecting to the IX. In this way, the

address assignment is decoupled from the connectivity provision, taking the IX the role of address provider.

This new IPv6 IX based addressing model, as well as the advantages of locating network and application services inside the IXs, bring new possibilities for the design of new advanced IPv6 Internet Exchanges architectures, opening the providers market to new opportunities and actors.

Therefore, this document extensively describes an advanced IPv6 IX model and their requirements, providing details about the different ways customers can access IX services, the way routing is organized between customers and providers, as well as a list of services that benefit from being offered from the IX and the way they can be organized and operated.

2. Terminology

- Address Delegation (AD): The process by which an IX assigns address space to an IXC.

- Customer Exchange Routers (CER): The routers used to connect IX customers to the IX.

- Internet Exchange (IX): It refers to a generalized exchange point including basic peering functionalities, as well as transit purchase services found in NAPs

- IX Address Prefix: An IPv6 address prefix assigned by a Regional Registry to the IX with the objective to be later sub-delegated to IX customers.

- IX Administrator: The entity in charge of IX management.

- IX Customer (IXC): A customer network that uses an IPv6 address prefix sub-delegated from an IX Address Prefix.

- IX Local Customer (IXLC): An IXC which is directly connected to the IX through a router that can be either managed by the IX administrator or by the customer.

- IX Remote Customer (IXRC): An IXC which is not directly connected to the IX, but through a regional provider present on the IX. The regional provider manages the distribution of the customer IX sub-prefix through its network.

- IX Value Added Services: Network and application services offered by the IX to its customers and/or providers peering on the IX.

- Layer 3 Mediation Function (L3MF): An entity acting as the intermediary between IX customers and providers.

- Long Haul Provider (LHP): An Internet Service Provider (ISP) present on the IX that provides transit services to IX customers and regional providers.

- Regional Provider (RP): An Internet Service Provider that peers with other providers at the IX and optionally gets transit services from LHP.

3. Reference Scenario

The following figure describes the main reference scenario of the advanced IPv6 Internet Exchange model defined in this document.

```
                 Long Haul Providers
                 +------+   +------+
                 | LHP1 |...| LHPm |
                 +------+   +------+
                      |        |
   +------------------|--------|----------------+
   |               +----+   +----+              |
   |    IX         | R1 |...| Rm |              |
   |               +----+   +----+              |
   |                  |        |                | IX Customers
   |+----------+  +-----------------+  +------+ | +-------+
   ||          |  |                 |--| CER1 |-|-| IXLC1 |
   ||  Value   |  |   Layer Three   |  +------+ | +-------+
   ||  Added   |--|    Mediation    |      :    |      :
   || Services |  |    Function     |  +------+ | +-------+
   ||          |  |                 |--| CERn |-|-| IXLCn |
   |+----------+  +-----------------+  +------+ | +-------+
   |                  |       |                 |
   |               +----+   +----+              |
   |               | R1 |...| Rk |              |
   |               +----+   +----+              |
   +------------------|-------|-----------------+
                      |       |
                +--------+   +--------+
                | RP1    |   | RPk    |
                |        |   |        |
                | +-----+|...| +-----+|
                | |IXRC1||   | |IXRCk||
                | +-----+|   | +-----+|
                +--------+   +--------+
                  Regional Providers
```

The IX is made of:

- A classical L2 infrastructure (i.e, a high speed LAN), not represented on the figure.

- ISP routers (R), that connect ISP networks to the IX.

- Customer Exchange Routers (CER), that connect local customer networks to the IX.

- The L3MF that acts as an intermediary between IX customers and ISPs.

- Value Added Services, that refers to those additional network and application services that can be partially or totally offered from inside the IX.

Two types of IX customers are foreseen in the scenario:

1. IX Local Customers (IXLC), which are directly connected to the IX through a router that can be either managed by the IX administrator or the customer network administrator.

2. IX Remote Customers (IXRC), which are not directly connected to the IX but to a regional provider that is present on the IX.

Both type of customers use address prefixes sub-delegated from the IX addressing range.

The figure above is a functional scheme and it does not imply in which way the functionalities described here have to be implemented.

4. Relationship with the traditional IX model

From a theoretical point of view, there are no constraints in merging the traditional model with this new advanced model, that means that some customers may use the IX to exchange traffic among each others and simultaneously other customers (IXLC) may use the Internet Exchange according to this new functionality.

In the following, we refer only to the new advanced model so the term customer refers only to the IXLC.

5. Overview of the new advanced IPv6 IX model

In a classical model, an IX is not normally opened to the direct connection of customers (i.e. large corporate networks or small ISPs or whatever). Instead of that, customers are connected to ISP

networks, which are present on the IXs.

In those cases where customers are directly connected to an IX, they typically subscribe an agreement with one or several Long Haul Providers present on the IX to route their traffic and announce their prefix addresses. Customer addresses can be sub-assigned from the provider ranges, that is, customers can use Provider Aggregatable (PA) address space. Alternatively, customers can get Provider Independent (PI) addresses directly from the RIRs.

If the customer changes its Long Haul Provider and is using PA addresses, it will have to renumber its network. The only way to avoid renumbering is to use PI addresses. However, to preserve the scalability of the whole routing system, PI addresses do not longer exist in IPv6. Therefore, following a classical IX model, changing provider in IPv6 implies renumbering the customer network.

To solve this problem, the advanced IX model presented in this document uses a different approach based on the possibility (new for an IX) to directly assign IPv6 addresses to the customers. Connectivity provision and IPv6 address assignment are now separated issues and they are no longer both linked to the LHP.

In this way, if an IX customer wants to change its service provider (e.g. because it gets a better service or price from another LHP), it does not need to change its own addresses, as they have been assigned by the IX and not by the service provider. It only will have to renumber if it changes the IX it is connected to (or from IX group, for instance, in case of distributed IX).

Consequently, in this new model the IX plays an important role in the Internet service provision: assigning addresses to customers and acting as a mediator between them and the LHPs. That is the reason why the main new IX functionality has been called Layer Three Mediation Function (L3MF), as it provides the decoupling among the customers and the LHPs.

Note that in the traditional IXs, agreements among the parties are usually taken among the ISPs connected to the IX itself. In fact, the IX is a neutral entity that does not have a particular role, since it basically provides the Layer 2 Connectivity infrastructure.

In the advanced model proposed here, three different entities are, in principle, involved: the IX itself, the IX customers and the LHPs. This means that different agreements could be taken, for example, according to the role of the IX.

Hence, it does not mean that IX customers will necessarily have to

negotiate with two or more different organizations (IX administrator and ISPs) in order to get Internet services. Instead, IXs could act as intermediaries between customers and ISPs, offering a one-stop-shop service. This will depend on business particularities/ models instead of technical ones and consequently are no further described in this document.

Another advantage is the possibility to use this model in order to provide value added services to the customers. The main concept is in fact that each IX can be considered as a place where services and ISP can be co-located and can be provided to a large amount of users taking advantage of the natural aggregation (in terms of Providers) that may happen inside an IX. These services will be discussed in the remaining sections.

The proposed model can be considered as the sum of different infrastructural layers. In fact, whereas the traditional IX model is mainly based on a layer 2 infrastructure used to exchange traffic in a quick, scalable and reliable way, on the other hand, the advanced IPv6 IX is to be considered like an extended one, where it is possible to find, together with the above mentioned layer 2 functionalities and the layer 3 infrastructure, other services and functionalities usually not present in the traditional model.

From a theoretical point of view, there are no constraints in merging the traditional IX model with this new advanced model. This means that some traditional customers may use the IX to exchange traffic among them while simultaneously other customers (so-called "next generation" customer) may use the Internet Exchange according to this new functionality. In the following, we will refer only to the new advanced model, so the term customer refers only to the "next generation" customers.

What is proposed here is essentially an architectural model, identifying the logical blocks to put inside the IX. This draft does not make reference to the real implementation model and its relationship with existing ISP architectural and commercial models.

5.1 Layer 2 infrastructure

It is basically the same as in the traditional IX. It consists on a switching platform (usually fully redundant and high performing) where the customers' routers are connected. In the new advanced model there are no particular changes related to this model.

5.2 Layer 3 infrastructure

This part of the structure depends on the implementation of the

so-called intermediation function between the connectivity provider and the customer. In a traditional IX, the layer 3 framework consists basically on the routers of the ISPs present in the IX.

Other L3 functionalities can be present, for example, a Route Server used to centralize and simplify the exchange of the routes between the peering partners. Additionally, other entities can be included to give support to services like multicast, mobility, etc.

5.3 Server Farm

The new model here proposed foresees that services are placed inside an IX. This is a revolutionary concept that permits the development of a different technical and business scenarios. Putting services inside an IX can take advantage because somehow it reduces the "distance" between who provides the service and who (the users) use that service. This may mean, for example, to reduce the response time of the service and to reduce the probability to have service unavailability. Obviously these considerations does not take into account the impact that this kind of model can have on the pre-existing ISP networks and a lot of attention should be paid on this aspect, not covered in this draft. Suitable IX models can be implemented in order to avoid conflicts with the pre-existing scenario and to take advantage of this solution.

6. Advanced IPv6 IX Services

This section describes and discusses about the services that can be offered from an IX based on the model introduced in this document according to the structure presented in the previous section. Services are grouped in network services, transition services, support and management services, security services and other services.

6.1 Network services

6.1.1 Address Delegation

Address delegation is one of the main functions to be offered by an advanced IPv6 IX. As mentioned, the new IX model decouples address assignment from connectivity provision, taking the IX the role of assigning addresses to its customers and the ISPs connected to the IX the provision of the connectivity to Internet.

IX administrators will request to their corresponding Regional Internet Registry (RIR) one or more address prefixes, basically following the same rules defined for ISPs or, as they are named, Local Internet Registries (LIRs). Technically, the IX will behave as

a LIR, being assigned by the registries the required prefixes.

Later, the prefixes assigned to the IX will be sub-delegated to IX customers, following the same rules a LIR uses to assign addresses to its customers [3].

The address delegation service is the basement of two basic services that can be offered to IX customers:

1. Provider Selection, that allows customers to choose the ISPs they want to get connectivity service from, as well as easily changing the selection without having to renumber their network.

2. Multihoming, which allows customers to simultaneously get connectivity services from two or more ISPs and define their routing policy. For example, customers can use one of the providers as the primary connection and using others as a fallback solution, or do load sharing among them.

L3MF will be the basic functional entity managing address delegation service in the IX. It will be in charge of:

1. Announcing the IX prefix/es to the ISPs peering on the IX.

2. Organizing the routing among IX customers and ISPs.

3. Implementing address delegation associated services like provider selection and multihoming.

In order to keep routing scalability, IX prefix/es must be announced aggregated to the IPv6 Internet through the ISPs peering on the IX. In principle, unaggregated prefixes assigned to IX customers must not be redistributed outside the IX (only to routers present on the IX). This fact imposes the limitation that all incoming traffic (that is, traffic destined to IX addresses) will follow the same path, independently of the IX customer it is directed to.

The only way to exert some control over the incoming path is to relax the rule mentioned above and allow the distribution of IX customer's unaggregated prefixes. However, that can only be done if they are distributed to a limited scoped, for example, though an ISP that has agreed with the IX to allow that or through a link that connects to another IX being part of an IX group or confederation. In any case, mechanisms must exist to guaranty that unaggregated prefixes are not freely redistributed (for example, filters based on community tags defined by the IX).

Typically, the L3MF will be implemented inside an IX managed router

using BGP protocol. L3MF will maintain external BGP peering with the ISPs and IX customer routers in order to direct each customer's traffic to its corresponding ISP. L3MF router will intelligently modify the NEXT_HOP BGP attribute to avoid that traffic passing through the L3MF router. However, there are other possible ways of implementing L3MF functionality, for example, using static routing or with the help of a route sever, as it is mentioned on the next section. The way the L3MF is implemented is out of the scope of this document.

As presented in section 3, two types of IX customers are defined:

1. IX Local Customers (IXLC), which have a direct layer 3 connection to a router in the IX (named CER), either managed by the customer or by the IX administrator. These customers can be connected to the IX using different technologies like point to point links, Frame relay, MPLS VPNs, etc. Customer routers (CER) can be dedicated to a customer or can be shared between several. However, in the shared case, all the preferences and routing policies will be common to all customers sharing the router.

2. IX Remote Customers (IXRC), which have not direct layer 3 connection with any router in the IX. They are customers of one of the providers present on the IX and so they are connected to the ISP network at some place different from the IX, but they use addresses form the IX prefix. In this case, the ISP has to cope with the distribution of the prefix assigned to the customer through its routing system, in order to have the customer accessible from the IX.

Both services defined (provider selection and multihoming) are available for IXLC. However, none of them are available for IXRC, as the customer is integrated in ISP network. Even that, the customer maintains the advantage that it is using IX assigned addresses, so in case he changes provider to another one present on the IX, it avoids having to renumber its network.

6.1.2 Route Server

Another service that it is commonly found on IXs is the route server. Roughly speaking, a Route Server is a modified BGP daemon designed to act as a central point for peering, changing the classical mesh topology of an IX into a star topology, where each ISP's border router maintains just one BGP peering with the route server.

By using a route server the scalability of the IX is greatly improved, because, being "n" the number of ISPs present in the IX, the number of BGP peering is O(n) and not O(n2) as it is with a mesh

topology. Moreover, the addition of new members to the IX is simplified, as only a new peering between new ISP router and the route server has to be configured.

The route server service can play an important role in the IPv6 IX model described in this document. Apart from the usual benefits that arise from the use of a route server mentioned above, it can help implementing the provider selection and multihoming services associated with IX based address delegation.

In particular, L3MF can be integrated with the route server functionality, giving rise to an Enhanced Route Server (ERS) that centralizes the peering among ISPs and IX customers routing related functions.

The use of a route server on an IX must not change the way routing is organized among ISPs. In other words, the route server must be transparent, meaning that the routes learned and installed by each border router must be the same when the route server is present than when it is not (mesh topology).

With this premise in mind, two types of route servers arise, depending on how much "intelligence" we relay on them:

1. Transparent Route Server, which propagates all the routes received from any router to the rest of routers, without modifying route attributes. In this way, routers acquire the same routing information as they would do via direct peering ("full mesh" topology), allowing them to apply their selection criteria according to their local policies.

2. Smart Route Server, which is an "intelligent" BGP daemon which performs best path selection on behalf of client routers. For this purpose, it must maintain a separate RIB for each of the clients, and it must also know the routing policies of all the clients. When the Route Server receives some announce it applies the appropriate policies before inserting them into the Loc-RIB corresponding to each client.

The main drawback of the first type of Route Server is that it needs to introduce some modifications to BGP protocol, because in BGP-4 it is not possible to send more than one announcement corresponding to the same prefix through a single peering. There are some proposals for modifying the protocol in order to make that possible, for example, the mechanism proposed in [4] or the new attribute proposed in [5].

However, the second type seems the most appropriate for the advanced

IX model purposes, due to the following reasons:

- o It does not need any modification to route server clients BGP implementations (any currently available BGP implementation will suffice for ISP and customer routers).

- o The provision of the new services associated to IX based address delegation suggests keeping the client-side BGP configuration as simple as possible (at least, for IX customer routers), and this means moving all the policies and best path selection process into the Route Server. This is precisely the idea of the smart route server.

In summary, the use of an ERS based on a smart route server to implement L3MF greatly simplifies the provision of complex services like load sharing multihoming. All the complex functions (filtering, policies, etc) that in other case should be implemented on IX customer routers are moved into the route server, who performs them on behalf of the customers and under the management of the IX administrator. In this way, requirements for IX customer routers are reduced and their management greatly simplified.

It is important to note that the use of an ERS to give service to IX customers does not necessarily imply that ISP peering sessions must be integrated on that. ISPs can peer directly among them over the layer 2 infrastructure or even the ERS can coexist with another route server that centralizes peering among ISPs.

6.1.3 QoS

The management of QoS can be efficiently done by means of policy provision architectures. With this approach, the functionality and flexibility increases notably. The utility of this kind of solution can be improved by advanced IX, as they can be a key element of such architectures.

The QoS provision usually is made by means of Policy Based Networks (PBNs) architectures based on the one proposed at [6], although it is strongly oriented to intra-domains enviroments.

The IX can provide solutions more scalable and more powerful in order to achieve end-to-end QoS allocation among different domains. Being the IX the common point where different domains are connected to, they might store information about the network status, user rights and so on, in order to decide if a given QoS request can be successfully granted.

Even if domains physically attached to other IX have to participate

to provide QoS resources, communication between the IX involved could be done to share the required information in order to decide whether or not the QoS resources allocation is possible.

Consequently, the main functionalities that the IXs could provide regarding QoS provision are:

- Provision of policies among QoS edge routers belonging to different domains.

- Define QoS policies for classifying data streams traversing the network.

- Store the policies edited by network administrators.

One of the main keys to succeed with this solution is the implementation of the policies. In general, PBN enforces the usage of network resources based on policies derived from criteria defined by network administrators, particularly, QoS PBNs allocate QoS reservations based on such policies.

However, policy is a general and abstract concept, which needs to be specified in particular actions. On the other hand, such policies should be defined by using high-level languages to let administrators easily define conditions that influence on the resource reservations and at the same time, easily understand the policies defined by other administrators. Thus, the more flexible are the policies, the more powerful is the PBN.

Given the fact that an IX is a network point where a number of different technologies can be present, the policies used on this type of environments should be as more general as possible in order to do not exclude any type of network.

6.1.4 AAA Services

AAA infrastructure can be deployed in an IX service network to offer authentication, authorization and accounting services in a wide variety of ways and to different kind of users and purposes. The main advantage to use AAA services is that they can provide support to mobile and roaming end users that roam between different ISPs or IXs. The recommended protocol defined to transport AAA information is DIAMETER [7].

Depending of the functionality provided by the IX to the clients (ISPs or end users), the AAA service has different requirements. For example, the IX could need to offer secure access control to end users connected directly to the network IX, also, the AAA services

can be used by the security services defined in the clients ISP service network. Some of the possible alternatives are listed below.

6.1.4.1 IX provides only internal AAA service

An IPv6 IX can provide AAA services to directly connected AAA end users, that is, users that do not arrive from a local ISP.

When the IX network receives connections from end users through network access services (NAS), users can be authenticated and authorized to obtain the network connection using the local AAA server sited in the IX service network.

If the number of NAS in the IX network becomes unmanageable, intermediate elements can be used, like Relay or Proxy AAA agents.

If the user's authentication or authorization process requires the exchange of information between external AAA servers, then a Redirect AAA service (described below) can be deployed locally.

6.1.4.2 IX provides accounting service to ISPs

The AAA service can be used only to get accounting information about the connected ISPs. For example, an IX charges to ISP by the bandwidth consumed in the link between them, the IX can deploy an AAA service used to obtain accounting information from the link to the ISP. In this case could not have authentication or authorization process.

6.1.4.3 IX provides AAA Routing service

Assuming an IX with several local ISPs, when an AAA server (AAA-A) sited on a local ISP network (ISP-A) needs to exchange information to another AAA server (AAA-B) sited on another local ISP network (ISP-B), the IX can provide Relay and Proxy AAA services to allow AAA-A to reach easily AAA-B and vice-versa.

6.1.4.4 IX provides AAA Redirect service

If AAA servers sited in ISPs local to an IX need to interact with external to the IX AAA servers. The own IX can provide a common AAA Redirect service to allow locate and provide information about these remote servers.

6.1.4.5 IX provides AAA Translation Service

Supposing the common situation, in which the ISPs have a RADIUS [8] or TACACS+ [9] system used to authenticate users, the IX can offer a

common point to translate the RADIUS domains to DIAMETER domains, adding additional functionalities to the RADIUS technology. This translation can be done placing an AAA translation service in the IX network.

To locate an AAA infrastructure inside an IX network implies to provide authentication and authorization services. In base of the AAA specification, the AAA server could use ASM modules to interact with external entities, which offer advanced authentication and authorization services to help the AAA server. That is, advanced authentication and authorization entities could be deployed in the IX. An authentication system can be deployed using a Public Key Infrastructure (PKI) and an authorization system can be deployed using an Authorization Authority, based, for example, on SAML or PMI technologies.

Moreover, to locate an AAA infrastructure inside an IX may imply that the AAA services need to establish authorization and authentication data exchanges between external AAA servers. These external servers sometimes have a trusted relationship with the original IX, but in other occasions, the AAA servers belong to not previously known organizations. In any case the exchange between servers must be protected establishing secure channels, commonly using IPsec or TLS protocols. The secure channels should be established using public key cryptography to avoid the problem of distribute secret keys between organizations.

If the AAA servers belong to a well-known organization, a cross-certification relationship can be established between the servers to protect the AAA exchanges. If the AAA servers belong to a not known organization the exchange only could be possible if a certification path can be discovered between the peers CAs.

Additionally in large networks, which may include millions of users, the management of mobility services and as a consequence their Home Agents, become operationally and administratively a complex scenario. Then a solution where the AAA infrastructure can help to solve this situation and then the integration of a AAA services like this within the IX is needed.

6.1.5 Multicast

Multicast Transmission Services are one of the services that can be provided in this advanced IX scenario. With reference to the ASM approach (PIM-SM particularly), the IX could be in fact the right location where to place the RendezVous Point, since it is the focal point where the IX customers and the ISP could meet. The introduction of RP inside the IX could take some advantage as for

example the optimisation of latency in transmission with a better-perceived quality by the users.

6.1.6 Mobility

TBD.

6.1.7 Multihoming

Moreover, this model could facilitate a multihoming solution since a customer is naturally multihomed (connected to more than one LHP/ISP at the same time). ??? TBD.

6.1.8 DNS

TBD.

6.2 Transition services

Most of transition mechanisms are tunnel-based and/or require a relay, server, tunnel-end-point (TEP) or other similar functionalities.

Some of them, such as 6to4 [10][11], use mechanism such as anycast, to discover the TEP. But others require a manual configuration (users have to know where the TEP is located or a method to find it), while already automatic discovery procedures had been proposed [12].

Furthermore users without technical knowledge don't know in general, how to choose the best transition mechanism (the one that provides de "best performance").

The IXs can play an important role to help to overcome such barriers, so users do not need to know anything about which are the best mechanisms and/or TEPs or other configuration parameters (i.e. tunnel brokers/servers). Therefore, the following advantage can be obtained by using IX with auto-discovery capabilities:

- Simplicity to client's configuration because they only would need to know one destination to get the best IPv6 connectivity, and this can be automatically discovered.

- Transparency in the communication in case that a server gets down because datagrams would be automatically redirected to the nearest alternative TEP.

- Load balancing in order to have a uniform resource share.

- Facility for the scalability of new TEPs.

In general, the IX seems to be a good place to provide transition mechanisms, and even not just one, but a set of them which can be used by different hosts, in different scenarios connected to ISPs collocated in the IX.

In addition to that, the IX can be used as a policy distribution service for the managed auto-transition [13], as a kind of helper to increase the performance of the transition at a large. For instance given the fact that a broker located into the IX has real-time information about the associated TEPs implemented on different ISP, this information should be utilized by the auto-transition mechanisms in order to select the best one.

The combination of services like DNS, AAA, anycast, within the IX, together with transition, provide a perfect framework and facilitates:

- The scalability of the transition mechanisms.

- The automation of the discovery and selection of the "best performing" transition mechanism and its parameters.

- The management of the transition service.

- Avoid deploying transition mechanism in every ISP, but providing the service to all the users (depending on the defined ISP policy).

6.3 Support and policy-based management services

The goal of policy-based management is to enable network, service and application control and management on a high abstraction layer. The administrator specifies rules that describe domain-wide policies independent of the implementation of the particular network node, service or application. It is then the policy management architecture that provides support to transform and distribute the policies to each node and thus to enforce a consistent configuration in all involved elements, which is a prerequisite for achieving end to end security services, for example.

Use of policies is an intrinsically layered approach allowing several levels of abstraction. There can be, for example, general policies expressing an abstract business goal, and on the other end there can be policies that express a rather specific, device or service dependent rule. In fact, one of the current research issues in

policy management is the refinement of high-level goals into implementable policies.

Policy rules are independent of a specific device and implementation, but define in abstract terms a desired behaviour. They are stored and interpreted by the policy framework, which intent to provide a consistent behaviour in all affected policy enforcement points.

As IX models become more complex with increased functionalities, will required the more management and then the support of Policy Based Management Network will certainly be fundamental.

6.4 Security services

6.4.1 PKI

One central component in the Security Architecture is the provision/distribution of keys. In order to do that one of main component for the support of the security services is the PKIs. They are needed in the IPv6 world to offer cryptographic features to security services like HTTPS, IPsec, etc. and also can be used for the securization of the different elements appearing in the infrastructure.

6.4.2 DNSSEC

Services like DNSSEC are being used to secure DNS transactions between clients and servers and among servers, as well as to store cryptographic information about the entities stored in it. Therefore, DNSSEC can be naturally used by a PKI to store the security information of the PKI clients, offering a worldwide distributed cryptographic storage mechanism. Using DNSSEC to store Public Key Certificates (PKC) and Certificate Revocation Lists (CRL), an IPv6 user can easily find the public certificate associated to other person/entity by means of a simple DNS request, in order to exchange information in a secure way.

6.4.3 Distributed security

With the deployment of IPv6 networks a revision of existent security models and architectures is a must. As new paradigms and applications enabled by IPv6 end-to-end appears the security infrastructure must be ready.

The IPv6 Distributed Security model [14][15] goal is to move the Security Policy Enforcement Point (PEP) to the end device to be protected. This is accomplished by the distribution of the security policy to the PEP from a central Policy Decision Point (PDP). Three main elements are needed:

1. Security Policy definition language.

2. Security Policy distribution protocol.

3. Unique and secure identification of entities.

As the Distributed Security involves policy distribution the IX could be used as a repository for Security Policies. Even different repositories could be located in an IX, acting as the PDP, and share some security information and alerts.

The fact that an IX will also have IPv4 connectivity could enhance the collaboration between IPv6 and IPv4 Security Policies repositories.

Even some business case could be foreseen if the up-dated last-minute security policy distribution service is charged to some users.

6.5 Other services

TBD.

7. Conclusions

This advanced IPv6 IX model could be described closer to an advanced Point of Presence (PoP) instead of just a traditional IX, when considering that it provides services and addresses to the customers.

A traditional PoP can also provide addresses, being the difference that in the advanced IPv6 IX model, the addresses are assigned by the IX itself and consequently do not change even if the customer of the IX changes its provider.

The model could be further extended to groups of advanced IX and/or distributed IX.

8. Security Considerations

This memo does not add any new security implication.

TBD.

9. IANA Considerations

This document requests no action for IANA.

[[note to RFC-editor: this section can be removed upon publication.]]

10. Acknowledgements

This memo is result of the work done in the Euro6IX project. The authors would also like to acknowledge the inputs from Ivano Guardini, Miguel A. Diaz, Alvaro Vives, Gabriel Lopez, Jose A. Garcia, Jose L. Rubio and the European Commission support in the co-funding of the Euro6IX project, where this work is being developed.

11 Informative References

[1] Fuller, V., Li, T., Yu, J. and K. Varadhan, "Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy", RFC 1519, September 1993.

[2] Hinden, R. and S. Deering, "An IPv6 Aggregatable Global Unicast Address Format", RFC 2374, July 1998.

[3] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.

[4] Haskin, D., "A BGP/IDRP Route Server alternative to a full mesh routing", RFC 1863, October 1995.

[5] Bhatia, M., "Advertising Equal Cost Multi-Path (ECMP) routes in BGP", draft-bhatia-ecmp-routes-in-bgp-00 (work in progress), May 2003.

[6] Yavatkar, R., Pendarakis, D. and R. Guerin, "A Framework for Policy-based Admission Control", RFC 2753, January 2000.

[7] Calhoun, P., Loughney, J., Guttman, E., Zorn, G. and J. Arkko, "Diameter Base Protocol", RFC 3588, September 2003.

[8] Rigney, C., Willens, S., Rubens, A. and W. Simpson, "Remote Authentication Dial In User Service (RADIUS)", RFC 2865, June 2000.

[9] Finseth, C., "An Access Control Protocol, Sometimes Called TACACS", RFC 1492, July 1993.

[10] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.

[11] Huitema, C., "An Anycast Prefix for 6to4 Relay Routers", RFC 3068, June 2001.

[12] Palet, J. and M. Diaz, "Evaluation of v6ops Auto-discovery for

Tunneling Mechanisms", draft-palet-v6ops-tun-auto-disc-01 (work in progress), June 2004.

[13] Palet, J. and M. Diaz, "Evaluation of IPv6 Auto-Transition Mechanism", draft-palet-v6ops-auto-trans-00 (work in progress), April 2004.

[14] Vives, A. and J. Palet, "IPv6 Security Problem Statement", draft-vives-v6ops-ipv6-security-ps-00 (work in progress), April 2004.

[15] Palet, J., Vives, A., Martinez, G. and A. Gomez, "IPv6 distributed security requirements", draft-palet-v6ops-ipv6security-00 (work in progress), March 2004.

[16] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.

[17] Case, J., Fedor, M., Schoffstall, M. and J. Davin, "Simple Network Management Protocol (SNMP)", STD 15, RFC 1157, May 1990.

[18] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M. and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

Authors' Addresses

Mario Morelli
Telecom Italia Lab.
???
Torino
IT-????? - Italy

Phone: +????
Fax: +???
EMail: mario.morelli@tilab.com

Jordi Palet Martinez
Consulintel
San Jose Artesano, 1
Alcobendas - Madrid
E-28108 - Spain

Phone: +34 91 151 81 99
Fax: +34 91 151 81 98
EMail: jordi.palet@consulintel.es

David Fernandez Cambronero
Technical Univ. of Madrid (UPM)
Ciudad Universitaria s/n
Madrid
E-28040 - Spain

Phone: +34 91 549 57 00
Fax: +34 91 336 73 33
EMail: david@dit.upm.es

Antonio F. Gomez Skarmeta
University of Murcia (UMU)
Campus de Espinardo s/n
Murcia
E-30071 - Spain

Phone: +34 968 364 607
Fax: +34 968 364 151
EMail: skarmeta@um.es

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.