# IPv6 TRANSITIONING MANAGEMENT – LAYING THE FOUNDATION FOR MANAGED IPv4/IPv6 INTEROPERATION

**Michael Mackay, Christopher Edwards**

Computing Dept, Lancaster University, Lancaster, LA1 4YW

Email: [m.mackay, ce]@comp.lancaster.ac.uk

## Abstract

*This paper highlights the need to supplement the work within the IPv6 community on IPv6 management with mechanisms specifically to support IPv6 transitioning. As a major feature of any IPv6 network for the considerable future, management support for the deployment and operation of a transitioning infrastructure is vital. We will propose a format for transitioning management information and examine how transitioning aspects of managed IPv6 networks can be handled within a transitioning management framework*

## 1 Introduction

The deployment of IPv6 will bring with it widespread changes to the Internet as it adjusts to cope with the introduction of a new Internet Protocol. During the deployment, there will be a lengthy period of IP co-existence that brings with it its own issues namely the area of IPv6 transitioning, this and its management is the focus of this paper. One aspect of most networks that will be affected by the introduction of IPv6 is network management, not only will there now be two protocols to manage within the same environment but also the transitioning process itself introduces significant disruption to the operation of the network and will place extra workload on the administrators. Of course, network management refers to a broad and complex field of computing that cannot be tackled in its entirety within this paper. Therefore, while we will provide a limited overview of the whole field from the perspective of IPv6 introduction, this paper focuses primarily on the issues relating to IPv6 transitioning management.

While IPv6 management has long been recognised as an important factor in the deployment of 'real world' IPv6 networks, certainly on a larger scale, viable solutions are now emerging as the development effort moves to mature IPv6 towards a commercial strength product. However, as an essential component of any IPv6 network for the foreseeable future, the issue of transitioning has generally been overlooked within the context of network management. Indeed, there is still a degree of uncertainty within the Internet standards bodies as to how to deploy real transitioning architectures in IPv6 networks and beyond the definition of some generic deployment scenarios, the finer points of transitioning, such as management has received little

attention. The aim of this paper is therefore to draw attention to the fact that work must be done on the definition of specific IPv6 transitioning deployment and management architectures as a precursor to large IPv6 deployments.

Section 2 of this paper will look at the IPv6 transitioning process, the mechanisms and deployment scenarios identified and then address how it might be managed. Thereafter, this paper will concentrate on three of the core management aspects for IPv6 and transitioning deployment. Fault and configuration management via the SNMP protocol is discussed in section 3 while policy management via SNMPConf (a functional extension) is discussed in section 4. Finally, section 5 introduces the STA, a management architecture for transitioning networks.

# 2 IPv6 Transitioning Management Architectures

This section discusses the issue of developing a transitioning management architecture and how it can integrate transitioning into management protocols. This primarily involves developing a prototype repository where transitioning management information can be represented and is suitably complete to contribute to the development effort. It will therefore be necessary to analyse how transitioning information is structured both individually and as a whole before putting forward a preliminary design. Initially however, we will present the transitioning mechanisms and generic scenarios developed within the IETF to establish both what constitutes the transitioning area and what work has been done towards defining transitioning deployment.

## 2.1 IPv6 Transitioning

Incorporating the range of transitioning mechanisms into a common framework presents a number of issues due to the disparate roles they perform. They can (and frequently are) however grouped according to 3 broad operational types, tunnelling, translators and dual stack that may also represent their common management characteristics. Each of the three types performs a specific task from a transitioning perspective and as such common information should be available for each mechanism according to the group to which they belong. While a detailed analysis of each transitioning tool is beyond the scope of this paper, this has been addressed in many other reports [1] and is briefly presented in table 1 below.

| Device | Operational Summary |
|---|---|
| **Tunnels** | Automatic or Static, intra/inter site, manually configured and managed |
| **ISATAP** | Automatic intra-site tunnelling, uses IPv4 as a nonbroadcast multiple access (NBMA) link layer |
| **6TO4** | Automatic inter-site tunnelling, treats the IPv4 Internet as a unicast point-to-point link layer |
| **6OVER4** | *Historic- considered obselete* |
| **Terado** | Tunnelling between Terado-enabled devices via UDP for NAT traversal |
| **Tunnel Broker** | Inter-site tool manages IPv6 tunnel requests from isolated sites for dedicated Tunnel Servers |
| **DSTM** | Dual Stack mechanism, dynamically managed IPv4 address pool allocated to hosts |
| **NAT-PT** | Based on NAT, provides generic translator, incorporates ALGs including DNS |
| **BIS/BIA** | Host-based translator similar in operation to NAT-PT, operates at IP stack or API layer |
| **TRT / SOCKS** | Transport level IPv4/IPv6 TCP/UDP relay |
| **ALGs** | Dual stack gateway for protocols with embedded addresses (FTP, DNS, etc) |

Table 1 - Overview of Transitioning Tools

Within the IETF v6ops working group[1], 4 general transitioning scenarios have been identified to describe the issue of IPv6 deployment in these areas [2]. Again, a detailed analysis is outside the scope of this paper but they consist of unmanaged, enterprise, ISP and 3GPP network areas.

## 2.2 IPv6 Transitioning Management

Transitioning management describes the activity of incorporating transitioning functionality into network management tools, this is especially important for larger networks which are extensively managed. This deserves particular attention as the transitioning elements of IPv6 deployments are likely to not only be part of the network for some time, but will also be subject to quite radical change as the process continues, presenting a significant issue for administrators. Given the current work on transitioning scenarios within larger, managed environments, it is expected that a network will deploy a number of transitioning mechanisms over time as it migrates to IPv6 subject to its particular requirements. This may at consist of a variety of complex mechanisms such as 6TO4, NAT-PT, DSTM or Application Layer Gateways (ALGs), in addition to basic tunnelling and dual stack deployment (based on the existing IPv4 network). These mechanisms may be employed at various stages in order to provide specific functionality and removed when redundant, while others may be an integral part of the transitioning network for the duration of the deployment process. That such a group of disparate mechanisms may be expected to co-exist in a managed environment highlights the need for a method of managing this element of IPv6 migration in an organised way.

One of the main issues to be resolved in the use of IPv6 transitioning mechanisms is how they could be best deployed as part of the larger migration strategy, which necessitates the definition of a common transitioning architecture. In order to be managed effectively, it would be useful if they could be organised into a standard framework and managed under a common management protocol, SNMP being the *de facto* candidate. At present the development of a transitioning architecture is still at an early stage, certainly within the larger bodies such as the IETF, however at least one draft is currently in circulation [3]. In it, there is an attempt to establish some design goals for a common transitioning approach and present a guideline of how transitioning mechanisms fit into it given the current trends in IPv6 transitioning deployment. When considering a common transitioning architecture there are several key design principles that must be taken into consideration, it must as a basis be both robust and secure in order to cope with various network conditions and guard against misuse but also be simple to deploy and operate. Another issue that must be considered is that of service provision over the transitioning infrastructure and this is important as the type and capabilities of the traffic to be handled will affect quite radically the transitioning resources that must be deployed. That a transitioning architecture is expected to accommodate this variety of requirements implies a final, very important aspect, flexibility.

While IPv6-only deployment within managed environments itself still has many challenges, it is now becoming increasingly feasible. For example, before IPv6 deployment can begin, it is vital to first plan how it will be achieved. It will be necessary to introduce a suitable addressing architecture and IPv6 capable DNS in order to utilise IPv6 address space, also it will be

[1] http://www.6bone.net/v6ops/index.htm

necessary to plan a suitable site routing system and the issues of firewalls and NATs need to be addressed. Many of these issues are still the focus of discussion within the development community [4]. However, here we focus primarily on the issues of introducing complex transitioning architectures into a network managed environment. One of the most significant issues will be to model the transitioning mechanisms within a common format that will then be formalised in a structure such as the Management Information Base (MIB). In order to do this, it is important to first establish which elements should be modelled and how. To simplify this, the transitioning elements will be placed within three basic groups (as discussed above) for the purpose of modelling within the management base; tunnelling, translator and dual stack, this allows mechanisms with common functionality to be grouped together. Within each group, mechanism or implementation specific information is held as is common management information which can be collected and modelled for each group. There may also be a need for another 'general' group to hold non-specific transitioning information. Figure 1 shows how this group of elements might be represented.

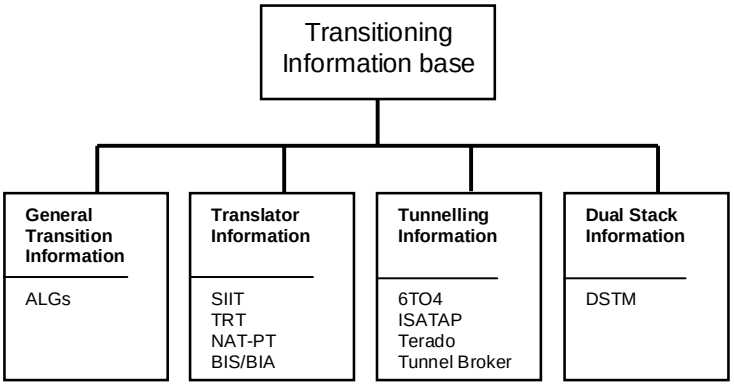


Figure 1 - Provisional grouping for a transitioning MIB

This layout presents a logical decomposition for transitioning information though it must be made clear that this is not the final format, rather it presents provisional design to motivate further development in this area. Given the above format, we must now discuss what data should be held under each category. The transitioning management format will not only store mechanism-specific data for the element but also any more general transitioning information that may be relevant. The information for each mechanism should be held inside the relevant table along with any more general transitioning information. Tables 2 and 3 presents an initial overview of what information could be held in each group and for each mechanism.

The tables show that some mechanisms contain similar management information which may simplify how they can be contained within a common format. By modelling common components together and only holding unique information separately, the framework can be greatly simplified. A common transitioning deployment will see devices dispersed throughout the network but in some cases components may be co-located, for example on gateway routers.

| Name | Group Summary |
|---|---|
| **Tunnelling mechanisms** | Traffic load, tunnel configuration, end-points, encapsulation statistics |
| **Translator devices** | Traffic load, packet translation statistics, session state information |
| **Dual-Stack mechanisms** | Traffic load, address binding information, packet transmission statistics |

Table 2 - Overview of transitioning groups

| Name | Modelling Summary |
|---|---|
| **Tunnels** | Tunnel configuration information, packet encapsulation statistics, end-points |
| **ISATAP** | Minimal - automatic |
| **6TO4** | mechanism configuration information, packet encapsulation statistics, |
| **6OVER4** | *Considered obsolete – not inckuded* |
| **Terado** | Server/relay confguration, packet encapsulation information |
| **Tunnel Broker** | Server states/loading, tunnel stats, throughput, end-point information |
| **DSTM** | Gateway / address pool statistics, packet stats, server information |
| **NAT-PT** | Translator information, address pool statistics, bindings information |
| **BIS/BIA** | Translator/ bindings information, internal address pool statistics |
| **TRT / SOCKS** | Translator information, session binding information, throughput |
| **ALGs** | Protocol specific session information |

Table 3 - Summary of specific tools

# 3 IPv6 developments for SNMP networks

One of the key areas for IPv6 and transitioning network management will be its incorporation into common management tools and protocols. As such, IPv6 development within the Simple Network Management Protocol (SNMP) [5] framework is important as this is regarded as the *de facto* standard for both fault and configuration management within IPv4 networks. IPv6 support in the Management Information Base (MIB) has been available since 1998 but has been under development since, changing once in 2000 and again more recently in order to better integrate its functionality with IPv4. SNMP IPv6 transport took longer to appear, possibly because managed networks were/are running dual stack and so are managed via SNMP over IPv4. Development implementations supporting IPv6 are now available, since March 2002 the net_snmp Open Source project[2] has supported IPv6 from version 5.0.3 with commercial implementations becoming available from multiple vendors including Cisco.

As part of IPv6 transitioning management via SNMP, there are two aspects that must be addressed, transport and information storage. As an architecture that operates over transitioning areas of the network, the management protocol will require dual stack transport support and as IPv6 is starting to be supported, this is now available. The storage issue will, as is usually the case, be more problematic and in fact this aspect has yet to be addressed in any meaningful way. The logical method of storing transitioning information is within a MIB sub-tree specifically to hold transitioning information. With a wide variety of disparate mechanisms to model, this

[2] Net-SNMP home page, http://www.net-snmp.org

structure will require significant development within the standardisation bodies and as such, the aim here is to assess how this could be achieved and present some logical provisional results that can then motivate development in this area. As discussed in section 2.2, transitioning management information could be held according to the groups that roughly match their behaviour with general information held separately. This section builds upon the transitioning management information defined in the previous section to develop an initial MIB definition.

When considering a transitioning MIB, thought must be put into how the per mechanism information should be represented. As described in the previous section, there is some basic per mechanism information common to multiple mechanisms and to simplify the implementation these should be represented within a single table. For example, in the case of dual stack mechanisms, DSTM has elements of other classes and is insufficient to warrant a table of its own. Therefore, dual stack mechanisms will provisionally be omitted as basic deployments require no special transitioning management and DSTM will be included as a tunnelling mechanism. In addition to the basic information tables, there will be specific tables to handle management information for the specific tools. Information related to each aspect of the transitioning is held in the 'basic' tables while any mechanism-specific information is held within the mechanism tables. Clearly, no single node will be expected implement more than one or two transitioning mechanisms if only from a performance perspective and some work has been done to determine which mechanisms work together and which are incompatible [6]. As such, on any node only certain aspects of the MIB will be populated with information and certainly it will be difficult to imagine more than two mechanisms of the same class implemented on one node. The 'general transitioning' table has also been defined to hold any non-specific transitioning information that does not belong in any specific area of the MIB. In addition, the dual stack aspects of the table have not been included here on the basis that there isn't sufficient dual stack specific management information to justify a separate table. This is still subject to debate however and maybe included at a later date.

# 4 IPv6 Policy Management via SNMPConf

Policy based configuration management has been the focus of much recent work within the IETF [7] that has potentially great implications for configuration management and while it is aimed primarily at the service provisioning aspects it could be applied equally to other areas, such as IPv6 and transitioning. We will use this as a base to consider how transitioning management discussed in the previous section could benefit from this.

## 4.1 Policy based Configuration Management

One recent trend in network management is to provide configuration management via a policy based approach. To this end, an architecture was developed within the IETF to define a common approach. The system is composed of Policy Enforcement Points (PEP) managed by a Policy Decision Point (PDP) that uses a policy repository (PR) to store policies. The system outlines that the PEPs operate according to a set of polices, however if they encounter a situation it cannot handle with its current policy base it queries the PDP which then replies with an authoritative policy as seen in figure 2. Currently, two systems have been developed which implement this model; a new protocol, COPS [8] and an SNMP based system, SNMPConf.

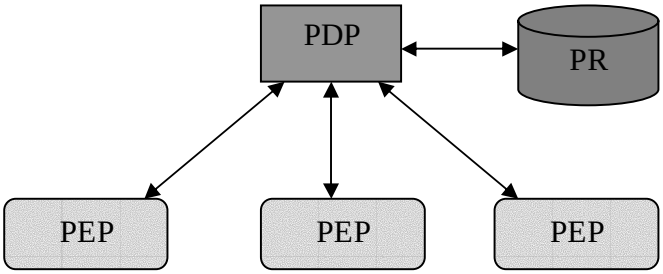

Figure 2 - IETF Policy-based Management Architecture

The SNMPConf working group[3] carried out work to apply SNMP to the IETF policy-based management architecture [9]. It outlines how a new notion called a template should be defined allowing configuration of multiple MIB instances and describes the new objects that should be defined for policy-based configuration. These objects describe the new notions that have not previously been incorporated into MIBs, these are roles, capabilities and time [10]. Using SNMPConf configuration management doesn't detract from using traditional SNMP allowing it to be used in combination with the more powerful policy based configuration. SNMPConf has already been applied to diffServ, but further uses include IPSec, mobileIP and transitioning.

## 4.2 Policy-based Transitioning Management

Policy management shows potential if applied to the transitioning management, particularly for service provisioning as one of the key aspects of transitioning will be how to provide useable and manageable services over it. By deploying transitioning in a policy configured network, it will be much more feasible to provide manageable services especially when networks scale up to the size that would commonly be managed. This section outlines a policy managed transitioning architecture before looking in more detail at the types of transitioning that could be supporting according to a policy-driven approach. As described earlier a policy management system allows the administrator to define how the devices on the network behave in certain situations and if this were applied across the transitioning, greater control over the behaviour of the transitioning infrastructure could be achieved. The administrator can now define policies that specify how the transitioning architecture behaves under certain conditions, this is useful to handle situations of high resource contention such as that found in a transitioning network.

One example of how this might be applied to a transitioning context is in device/router based mechanisms such as 6TO4 or NAT-PT which would implement PEP functionality. More complex mechanisms such as DSTM or Terado are more difficult to define however and in these cases, both the server and relay mechanisms should implement PEPs. One outstanding issue is that of hosts incorporating transitioning complexity such as in DSTM or BIS/BIA. The overheads involved in host management raise significant scalability concerns but may be overcome if deployment is limited. Figure 3 shows how this system might be applied.

[3] http://ietf.org/html.charters/snmpconf-charter.html

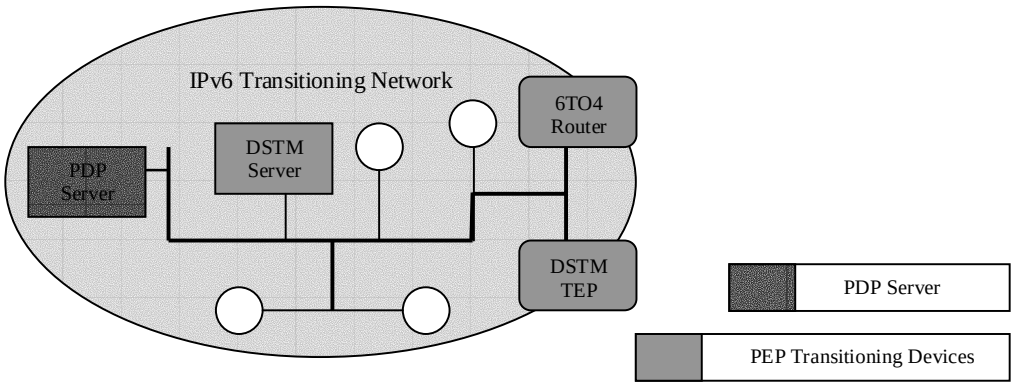


Figure 3 – Deployment of policy management over a transitioning network

This functionality allows the administrator to define configuration policies for all aspects of transitioning, controlling who may use the infrastructure and what performance they receive. In addition, the device can now be reconfigured to respond to network conditions (such as heavy load) or times of day. As with other services that have be adopted for the use of policy management such as DiffServ, the addition of transitioning support will involve the development and use of a specific transitioning policy MIB to define the relevant policy information in addition to the generic policy MIB. Service provisioning over the transitioning infrastructure could then be considerably simplified due to the increased control and flexibility it provides in terms of performance characteristics.

# 5 Site Transitioning Architecture (STA)

## 5.1 Overview

As we have seen, the need for IPv6 transitioning management will be of great importance in the near future. Moreover, a dedicated management system for IPv6 transitioning would be able to integrate the latest in network management to simplify the provision of services across the transitioning infrastructure in addition to providing standard fault management. This section will outline the possible uses for integrated transitioning management by introducing one such architecture, the STA and while anything beyond a simple outline is beyond the scope of this paper, further details can be found here [15]. The primary objective of the STA is to simplify the deployment and management of IPv6 transitioning during the migration process. By employing a flexible service-oriented architecture within which transitioning support can develop, the STA can simplify the process of transitioning for a large (Enterprise/ISP) site over the course of the IPv6 migration process from initial deployment to a prolonged co-existence and beyond. There follows a condensed list of the main objectives and benefits of the STA:

- Provide a unified framework suitable for the whole transitioning period
- Exhibit simplified management properties
- Introduce a service-based framework to the transitioning process

The STA is capable of deploying virtually any combination of transitioning mechanisms specified by the administrator and accommodate the changes that inevitably occur between sites and over time. For example, the functionality required in the early stages of IPv6 deployment will vary radically from that in the later stages of the process. Another key benefit of the STA is the management of the transitioning infrastructure as an entity as opposed to a set of individual components, allowing the administrator a wider degree of control over its operation. This makes it possible to define policies that govern the configuration of devices within it, ranging from fine-grained configuration of specific devices to more broad information about how the whole architecture will behave under certain conditions.

Another feature that will be possible with the STA is optimisation of the transitioning infrastructure itself, because each device is now part of a larger entity it is possible to optimise aspects of their operation to improve overall performance, one such example of this is address allocation. If a site deploys multiple interoperation devices, such as NAT-PT and DSTM for example, traditionally each device maintains its own IPv4 address pool. This is both wasteful and inefficient as each be individually configured and managed also, there is no guarantee that the address pool will be optimally utilised. If this could be unified within a single address allocation entity as we see with DCHPv6, it could be shared between all devices as necessary.

## 5.2 Design of the STA

The STA operates on two planes, operational and management. The operational plane composes the actual transitioning infrastructure deployed within the site while the management plane dictates how they are controlled and operated. The operational plane will consist of 3 abstract transitioning components; hosts, servers and border devices describing the layout of transitioning mechanisms used within the site, while the management plane is composed of a client/server model composed of agents and a management station that provides the administrator with remote access to the operational components deployed within the framework via a web-based GUI. While the components defined here will be sufficient to describe common transitioning devices, the STA must be capable of incorporating any other necessary network components such as switches or routers that deploy transitioning functionality. Only components that perform a transitioning role will be modelled within the STA, the layout of the STA is shown in figure 5.

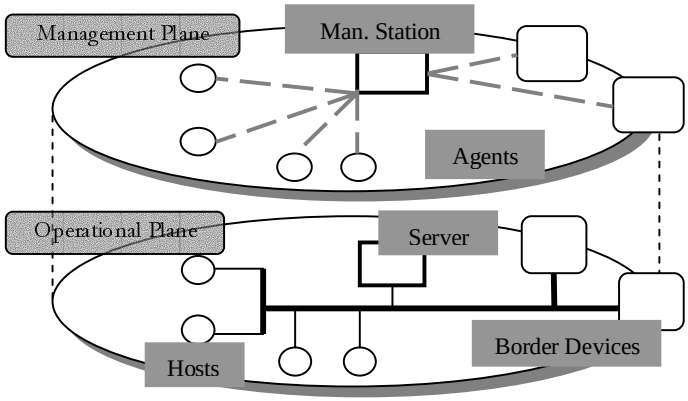


Figure 5 - Layout of the STA

**Border Devices -** Border devices will play a major role for of transitioning functionality in the network, both in terms of interoperation provisioning and inter site tunnelling. While the operation of many transitioning mechanisms revolves around a central border device, one of the effects of the STA will be the move away from this. In the STA, the transitioning complexity is still in the network, i.e. the border devices, but the administrative complexity is removed and unified inside the network, making the border devices 'dumb' nodes that perform transitioning functionality but are controlled remotely. This should have the effect of improving the scalability of the transitioning functionality making the deployment of more complex architectures possible.

**STA Server -** The STA Server is the only persistent component of the architecture, this will contain the management station element in addition to any necessary transitioning server functionality. Mechanisms that require a server (e.g. DSTM or tunnel brokers) when deployed within the STA will incorporate their servers within the STA server component. While performance is less critical in the server as it handles control not operational traffic, it may incorporate a number of critical services so it is still vital to be made scalable. Therefore it will be implemented as a 'cluster' allowing it to scale indefinitely (in the scope of handling control traffic) depending upon the complexity and size of the STA deployment.

**Hosts -** In order to reduce overheads, hosts will not be specifically modelled by the STA unless their modification is specified as part of the mechanism being used, as in the case of DSTM for example. Also, even if host modification is specified as part of the tool, every effort will be made to minimise this effect even to the extent of modifying the operation of the tool to shift complexity into the network although at this stage this must be firmly classified as further work, beyond the scope of this report. It is in the interest of the STA to minimise the functionality required in the host as this will be replicated across the network, increasing the management load. Hosts which do implement transitioning mechanisms, BIS/BIA for example are obviously therefore undesirable in the STA. However, this may not apply to all hosts on the network and those that are may be managed according to a 'group policy' which applies to all managed devices of that type, thereby relieving the issue significantly.

Management of the STA will implement an SNMP structure with agents and MIBs located on transitioning devices with a management station that handles access and configuration of the devices through their agents as shown in figure 6.

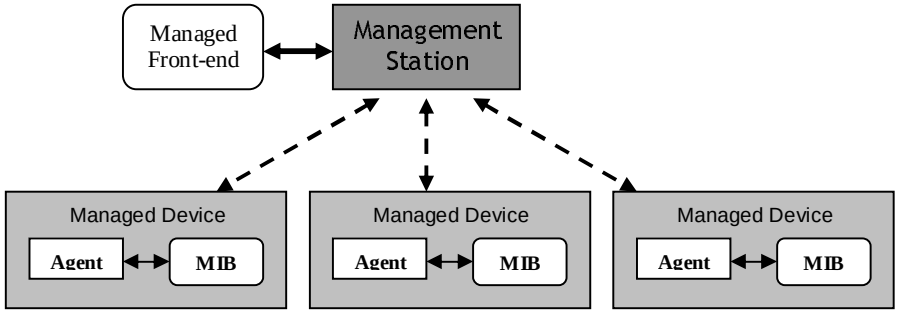


Figure 6 – Design of the management framework

In addition to basic SNMP management information for transitioning devices, in the STA the management plane will include policy information based on SNMPConf, these policies may be defined as site-level or specific to the type of mechanism they represent as per policy management. Operational policies are the key to tailoring the transitioning architecture to offer certain service characteristics and are the building blocks to offering a service-based infrastructure supporting QoS, security and others.

## 5.3 STA Deployment Case Study

While the STA design describes it's the core components, it is also essential to show how it can be applied to various transitioning scenarios. It is important to stress that in each case the basic format of the STA does not alter, rather it is applied to meet certain transitioning requirements.

**Case Study 1 – A Large Enterprise Deployment**

Consider an Enterprise deployment such as a University campus. Here, IPv6 deployment may require a mix of interoperation and tunnelling functionality to be supported. Initially, existing IPv6 deployments are likely to be relatively limited and so the level of support required will be extensive, for example there may be a need for both internal and external tunnelling. Internal tunnelling might be supported by ISATAP, but in the case of lab space or residence networks where connectivity is provided via a NAT box, Terado is a stronger candidate. Initial IPv6 deployment may also be based extensively on dual stack. Whilst any dual stack deployment inherently impacts on all of the devices in the 'dual stacked' path, it is certainly realistic to think that this might take place within the main campus backbone (where the number of devices that require both IPv4 and IPv6 addresses will at least be limited). A dual stacked backbone will provide support for IPv6-only networks that might begin to appear within the campus (perhaps being limited to specific subnets initially). Over time, the tunnelling mechanisms will be removed as native IPv6 becomes available, translator type mechanisms will also be required in scenarios where IPv4-only nodes talk to IPv6-only nodes. As IPv6 deployment becomes more widespread, these interoperation mechanisms will of course be required to scale accordingly.

**Case Study 2 – An IPv6-only Deployment**

This case considers a sizeable IPv6-only deployment within a large corporate network as part of a departmental upgrade. The system administrators for the department have decided that the upgraded infrastructure is entirely IPv6-only with external connectivity to the wider corporate network via dual stack. In the absence of dual stack internally to the department (either through the lack of desire to deploy it or the plain lack of resources), a significant transitioning deployment will be necessary to support extensive IPv6-only operation. In this case, a variety of mechanisms may be employed including protocol relays (TRT, FTP/DNS-ALG), IP translators (NAT-PT) or mechanisms such as DSTM that provide limited dual stack by implementing some functionality within the host. The decision as to which tools to use will be made based on the size of the deployment and the functionality to be provided. As such, the majority of the interoperation complexity will be deployed initially at the subnet boundary and over time will be moved outwards and extended as the rest of the corporate network upgrades to IPv6.

**Case Study 3 – A Provider/ ISP Deployment**

Consider an ISP that wishes to transition to IPv6, with the aim to provide IPv6 connectivity, eventually natively, to its customers. Such a deployment might involve upgrading both the core and access network devices, and has the potential to be both costly and time consuming. Therefore, initial IPv6 connectivity may be provided by configured tunnels or via a tunnel broker. Native IPv6 connectivity will eventually be provided either by upgrading the existing

components or deploying a new IPv6-only infrastructure alongside the existing network. In addition to providing IPv6 connectivity, an ISP may choose to provide extra transitioning service to its customers in the form of interoperation for IPv6-only customers. This would require a considerable interoperation deployment based on translators and protocol relays but as existing IPv4 resources are gradually exhausted, it is a service that may be increasingly utilised.

# 6 Conclusions

It is positive to note from the work presented here that the management infrastructure necessary for IPv6 deployment is for the most part either in place or under development. Somewhat less impressively however, the transitioning management aspect is still in need of significant work. This paper explores the issues of transitioning management architectures and puts forward an initial design for a transitioning management structure that shows how this can be deployed in future managed networks. Also, in the STA, we put forward the design of an advanced transitioning management architecture that is not only capable of managing IPv6 transitioning deployments but also of handling the entire requirements of a site throughout the migration process from initial deployment to long-term interoperation to gradual phase-out.

# Acknowledgments

This paper was written with support from the 6Net project, a European IPv6 research project.

# References

[1] Christian Schild, Tina Strauf (ed.), et al, “D2.3.2: Initial IPv4 to IPv6 transition cookbook for end site networks/universities”, 6NET project deliverable, February 2003

[2] http://ietf.org/html.charters/v6ops-charter.html

[3] P. Savola, “A View on IPv6 Transition Architecture”, draft-savola-v6ops-transarch-02.txt, October 2003, work in progress

[4] Tim Chown (ed.) et al, “D2.5.2: Updated IPv6 Deployment Issues”, 6NET project deliverable, September 2003

[5] J.D. Case, M. Fedor, M.L. Schoffstall, and C. Davin, “Simple Network Management Protocol (snmp)”, May 1990, RFC 1157

[6] A. Baudot, G. Egeland, C. Hahn, P. Kyheroinen, A. Zehl, “Interaction of transition mechanisms”, draft-ietf-ngtrans-interaction-01.txt, June 2002

[7] R. Rajan et al., “A Policy Framework for Integrated and Differentiated Services in the Internet”, IEEE Network, Vol. 13 No. 5, October 1999

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

[9] M. MacFaden, D. Partain, J. Saperia, W. Tackabury, “Configuring Networks and Devices with Simple Network Management Protocol (SNMP)”, April 2003, RFC 3512

[10] Steve Waldbusser, Jon Saperia, Thippanna Hongal, “Policy Based Management MIB”, draft-ietf-snmpconf-pm-14.txt, September 2003, work in progress

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

[12] P.R. Calhoum, J. Loughney, E. Guttman, G. Zorn, and J. Arkko, “Diameter Base Protocol”, draft-ietf-aaa-diameter-15.txt, October 2002, work in progress

[13] Isabelle Astic, Olivier Festor, “Current Status of IPv6 Management”, http://www.inria.fr/rrrt/rt-0274.html, December 2002

[14] R. Droms, (ed.), J. Bound, B. Volz, T. Lemon, C. Perkins, M. Carney, “Dynamic Host Configuration Protocol for IPv6 (DHCPv6)”, July 2003, RFC 3315

[15] M. Mackay, C Edwards, “IPv6 migration implications for Network Management – Introducing the Site Transitioning Framework (STF)”, October 2003, work in progress