ATIS Internet Protocol version 6 (IPv6) Task Force Report on IPv6 Transition Challenges

July 2007

ATIS is a technical planning and standards development organization that is committed to rapidly developing and promoting technical and operations standards for the communications and related information technologies industry worldwide using a pragmatic, flexible and open approach. Over 1,100 participants from more than 350 communications companies are active in ATIS' 22 industry committees and its Incubator Solutions Program.

< http://www.atis.org/ >

ATIS Internet Protocol version 6 (IPv6) Report & Recommendation

This is an ***ATIS Report*** developed by the **IPv6 Task Force** for the **TOPS COUNCIL**.

This document is a *work in progress* and subject to change.

Published by
Alliance for Telecommunications Industry Solutions
1200 G Street, NW, Suite 500
Washington, DC 20005

TABLE OF CONTENTS

EXECUTIVE SUMMARY

In May 2006, the ATIS IPv6 Task Force (IPv6TF) issued a report and recommendation with respect to various aspects of IPv6 entitled "ATIS Internet Protocol version 6 (IPv6): Report & Recommendation, May 2006." Areas reviewed by the group included deployment, transition challenges and the market drivers behind deployment of IPv6. In response to that report, the ATIS Board of Directors requested that the ATIS Technical & Operations (TOPS) Council commission the extension of the IPv6TF to propose steps forward to the extent possible to address the numerous transition challenges identified. This second report of the IPv6TF attempts to address this objective.

Since the May 2006 IPv6TF report, work within the standards and policy arena (particularly the IETF and American Registry of Internet Numbers (ARIN)) has advanced to address several aspects of the challenges originally identified by the Task Force; however, additional work still remains. Of the items remaining to be resolved, most notable is the need for organizations to make key internal business decisions with respect to IPv6. More precisely, while IPv6 is unquestionably of importance to the industry, its wide-scale advancement has been preempted by more pressing technical and operational priorities demanding industry's focused attention. Given this reality in addition to the perception that market demands have yet to materialize to the point of driving IPv6 deployment beyond solutions presently available and supported today to enable IPv4-to-IPv6 interoperability (e.g., dual-stack with encapsulation), there is no sense of urgency to change this approach. However, as recently illustrated by ARIN's announcement with respect to the exhaustion of IPv4 address, certain market drivers for transitioning to IPv6 are starting to materialize.[1]

Notwithstanding the above, industry is keenly aware of the need to continue its efforts to address IPv6 to the extent possible. From the list of eighteen (18) transition challenges originally identified by the IPv6TF, Address Allocation Policies, Site Multi-homing, Quality of Service (QoS), Security, Interoperability between IPv4 and IPv6, Network Address Translators (NATs) and the impacts on existing network traffic and routing were quickly identified as high interest items when transitioning to IPv6.

Originally identified as an issue needing resolution shortly after the release of the IPv6TF report, ARIN approved a policy which granted service providers (SPs) the flexibility they requested in allocating addresses to their customers. This differs from the original policy whereby service providers were constrained by a fixed assignment standard and timeframe for allocating IPv6 addresses.

Conversely, site multi-homing continues to be a topic of high-interest and passionate debate. Every solution proposed to-date to address this gap has yet to advance to wide-scale deployment; with each option currently proposed having its advantages and disadvantages. Understanding the level of focus presently being placed on resolving

[1] http://www.arin.net/announcements/20070521.html

this gap and the resources allocated to this topic within the IETF, it should be expected that an agreement on an approach forward will eventually surface. In the interim, representation from the different interested organizations (IETF, Regional Internet Registers (RIRs), vendors and operators) should, to the extent possible, jointly compile the minimal and optimal technical and operational requirements to solve the problem and explore possible solutions forward.

A unified approach is also of value with respect to the interconnection and interoperability of IPv4 with IPv6 networks. The implementation of a dual-stack approach, which was recommended by the IPv6TF in its first report and reaffirmed within this report, effectively affords network providers the ability to minimize impacts on their existing core network by placing interoperability decisions on the edge devices (i.e., edge devices will decide which IP version is required to interoperate with the host). This approach however, as any approach to augment existing network-to-network interworkings, has known limitations, but steps towards mitigating any negative impacts can be properly dealt with through the development of acceptable operational best practices.

Many of the remaining items identified of high interest await internal business decisions before actionable steps forward can be proposed. For instance, the continued use of NATs in an IPv6 network is widely recognized as not the ideal approach. However, market realities mandate acceptance that NATs are deeply embedded within the communications networks today and will likely continue to be deployed. Consequently, NATs are here to stay for the foreseeable future and each organization must deal with these embedded devices in their own way; including their affects on network architectures when contemplating next generation topics such as wireline-wireless convergence, multicast and continuous media flows (i.e., VoIP).

Security issues --which seemingly change daily as IPv6 continues to be tested through laboratory evaluations and field trials -- as well as the impacts on existing traffic and routing tables as a result of transitioning to IPv6 will inherently call for industry caution and diligence for an extended period. Therefore, during this transition phase the support of standardization efforts and internal organizational decisions based on business needs and policy debates must continue.

In summary, findings contained in this report can be classified into three (3) categories:

- **Technical**: current standards activities, internal and external to ATIS affecting issues identified, are either available, under development or pending start-up (e.g., security, QoS, etc.)
- **Business Related**: internal company business decisions need to be made in order to build consensus for best practices, solutions or needs for solutions (e.g., site multi-homing, NATS, etc.)

- **Policy Related**: issues are outside the scope of the IPv6TF, but individual companies are actively addressing them in appropriate groups (e.g., privacy, numbering, etc.)

INTRODUCTION

ATIS, through its IPv6 Task Force (IPv6TF), has been actively reviewing and assessing various aspects of IPv6 for the past several years. Commissioned under the leadership of the *ATIS Technical & Operations (TOPS) Council* --a standing committee of the ATIS Board of Directors -- TOPS set forth the IPv6TF's objectives to include:

- review current standards activities internal and external to ATIS affecting those issues identified
- record consensus for best practices, solutions, or needs for solutions for those areas
- from the perspective of ATIS members, identify the areas in which ATIS committees or groups may contribute to, or anticipate activity in support of the deployment and ongoing management of IPv6 services
- to ascertain the readiness of IPv6 and identify technical and operational issues which must be brought to resolution; whereby "IPv6 readiness" is achieved when challenges identified in the ATIS IPv6 Report and Recommendation are addressed and IPv6 is minimally equal to current IPv4 offerings in the areas of security, Quality of Service (QoS) and operations management.

In May 2006, ATIS released its first report with respect to IPv6 entitled, "*ATIS Internet Protocol version 6 (IPv6) Report & Recommendation, May 2006*." In this report, the IPv6TF provided a survey of deployment drivers and challenges that service providers may face when considering their deployment of commercial IPv6 services. It also outlined transition strategies and challenges and recommended steps to mitigate deployment challenges.

On review of the challenges identified in the IPv6TF's report, the ATIS Board of Directors commissioned the extension of the IPv6TF to take a deeper dive into addressing the various challenges it identified in transitioning to IPv6. Of particular interest to the Board was attribution of the challenges to technical, operational or policy-related matters. The IPv6TF also was requested to ascertain the readiness of IPv6 and identify technical and operational issues which must be brought to resolution in order for the successful transition to and deployment of IPv6-based services. To the extent possible, this report attempts to addresses these objectives.

To initiate work, the IPv6TF prioritized the eighteen (18) identified challenges in its May 2006 report in order of importance to transitioning to IPv6.

Challenges identified as "HIGH" priority and selected to be addressed first included:

- Address Allocation Policies
- Site Multi-Homing
- Quality of Service
- Security
- Interoperability Between IPv4 & IPv6
- Network Address Translators (NATs)
- Impacts on Network Traffic & Routing

Challenges identified as "MEDIUM" priority and selected to be addressed secondly included:

- Impacts to Privacy/Legal Issues
- Management Tools (Dual-stack & IPv6 Networks)
- Impacts on Infrastructure Reliability
- Network Renumbering (Portability)
- Peering Evolution (Impacts to Settlements)
- Impacts to Access Networks

Challenges identified as "LOW" priority and selected to be addressed thirdly included:

- Separation of Locator & Identifier
- Vendor Availability
- Dual-Stack with Domain Name Server (DNS)
- Relationships with other Numbering Systems
- Cost

Other challenges were also identified during this exercise and are also included in this report.

1 ADDRESS ALLOCATION POLICIES

1.1 End Site Allocation Size

1.1.1 Issue

As outlined in the May 2006 IPv6TF report, the American Registry for Internet Numbers (ARIN) stated that within six (6) months of allocating IPv6 addresses, /32 must be announced to the global IPv6 networks and within five (5) years the Service Provider (SP) must have 200 /48 networks allocated to customers in order to retain its IPv6 allocation. As proposed by ARIN, service providers faced a major challenge in implementing this policy given that upon allocation of IPv6 address space, typically 6-12 months of internal testing is required prior to development of an implementation plan. Additionally, actual deployment of IPv6 services may take an additional two to three years.

The IPv6 allocation policy developed by ARIN stated the following for end-site allocations:

For fixed standard assignments:

a) /48 addresses should be assigned in the general case, except for very large subscribers
b) /64 addresses should be assigned when it is known that one and only one subnet is needed by design
c) /128 addresses should be assigned when it is absolutely known that one and only one device is connecting

A proposal to change the policy was submitted to ARIN in 2005 (2005-8). The proposal is as follows:

Flexible assignments between /64 and /48, with the following guidelines:

a) /64 addresses should be assigned when it is known that one and only one subnet is needed
b) /56 addresses should be assigned for small sites, those expected to need only a few subnets over the next 5 years
c) /48 addresses should be assigned for larger sites

1.1.2 Location of Work

ARIN is the body that sets IP addressing policies for the United States and Canada, and is where this policy and the proposal to change it reside.

1.1.3 Agreement

The main advantage of the proposal to allow for flexibility in assignments is that it would allow service providers flexibility in providing allocations to their customers,

instead of handing the same size allocation regardless of the size of the customer. Being able to go to smaller allocations may vastly increase the useable lifespan of the address space is another argument made.

There are two main disadvantages to the proposal. Flexibility in allocation size also introduces complexity. Having a fixed allocation size makes it easier to standardize filters and eliminates the need to take different allocation sizes into account when renumbering. Another problem with the proposal is that similar proposals have been not met with much support in other Regional Internet Registries (RIRs) in the world, so North America will have a different allocation policy than everywhere else.

1.1.4 Proposed Action

None. The proposal presented passed in ARIN on May 9, 2006. ATIS members are encouraged to abide by the current ARIN policies and participate in the ARIN Internet Resource Policy Evaluation Process (http://www.arin.net/policy/irpep.html). If ATIS members agree that the previous policy is 'correct,' members can use the ARIN process to attempt to have the existing policy changed.

1.2 Provider Independent IP Space

1.2.1 Issue

Provider Independent (PI) addressing was designed in the IETF to address the multi-homing problem by allowing for scalable Internet routing via the use of a prefix based on a geographic reference to the site. In one particular version (IETF internet-draft[2]) PI addresses are derived from a site's geographical location, which -- along with a site's longitude and latitude -- determines the global prefix portion of its address. With this prefix, the site can then connect to multiple service providers using the same address, as well as indirectly allowing sites to change providers without requiring networks to renumber.[3]

The ARIN IPv6 allocation policy however offered no possibility of PI IP space. Proposals have been made to allow for PI space for a number of reasons, including:

- Multi-homing (no viable solution yet)
- Elimination of renumbering
- Allowing for IPv4, but not IPv6
- Requiring organizations with PI IPv4 space to change way of business for IPv6

All of the above reasons cited by proposals for PI space may cause roadblock to IPv6 deployment.

[2] http://www.draft-hain-ipv6-pi-addr-10.txt
[3] ATIS Internet Protocol Version 6 (IPv6): Report & Recommendations, May 2006, Section 4.9.3

1.2.2 Location of Work

ARIN is the body that sets IP addressing policies for the United States and Canada, and is where this policy, and the proposal to change it, reside.

1.2.3 Agreement

The arguments for PI space are listed above. The main argument against PI space is that the very idea of PI space is contrary to the concept of strong aggregation and therefore PI space leads the IPv6 Internet towards some of the same routing issues that exist in the IPv4 network.

1.2.4 Proposed Action

None. The proposal passed in ARIN on May 9, 2006.

2 MULTI-HOMING

2.1 Issue

Site multi-homing is a scenario where a site chooses to connect to two or more SPs for the primary purpose of redundancy. If there is a failure on the link to a SP or with the SP provider edge router or if there is a total SP failure, the site can switch to use the alternate SP(s) to remain connected to the Internet. Other possible motivations for site multi-homing are load sharing and cost.

Currently there is no agreement on a solution for this very important issue facing the IPv6 community. There is a protocol based solution that is being developed within the IETF (Shim6); however, Shim6 is not widely accepted in the industry. In the absence of any accepted solution the RIRs have defined address allocation based policy solutions that allow large sites to be multi-homed but do nothing to solve the route explosion problem. Also, RIRs do not provide a multi-homing option for smaller sites. It is understood that these RIR policy based solutions will not scale as IPv6 becomes more widely deployed and that they are intended to be used only until a proper, scalable solution can be found.

Regardless of the approach eventually agreed upon for site multi-homing, it is generally agreed that a strong aggregation policy would be beneficial. An example for this approach is as follows:

- The default policy (set by ARIN and other RIRs) for IPv6 address allocation is for end users to be allocated a global prefix from their upstream service provider. This is known as *Provider Allocated (PA)* addressing.
- The SP then advertises an aggregate prefix to its peers which encompasses all of the end user prefixes which it allocated.
- An SP that has been allocated a /32 prefix can aggregate 64K /48 prefixes into a single advertised route.

- If access or availability of a site prefix changes, the change is not propagated beyond the upstream SP because it is part of the larger aggregate. This reduces the route processing in the DFZ and improves convergence times.
- The failure to enforce a relatively strong route aggregation mechanism would result in routing and forwarding table bloat and memory and processing power in core routers may not be able to keep up.

2.1.1 Site Multi-homing options

Over the last couple of years, a number of solutions for the multi-homing problem have been proposed. Provided in this section are a few options under consideration along with a list of their advantages (pros) and disadvantages (cons). The information provided is not intended to represent an exhaustive listing of the options or their pros and cons. Other options to include variants of those options listed may also exist.

2.1.1.1 Provider Allocated (PA) addresses re-advertised

End user PA prefix (/48) is advertised from the site into alternate (non-allocating) service providers.

- Pros:
  - This solution is simple and meets most of the requirements for the site and the immediate upstream provider.
- Cons:
  - The specific /48 address needs to be advertised throughout the Default-Free Zone (DFZ) defeating the goal of strong aggregation.
  - The primary SP also needs to advertise the specific /48 into the DFZ to avoid having all traffic flow through the alternate SP that advertised a more specific route.
  - Changes in site connectivity to either of the SPs need to be propagated and processed throughout the DFZ.
  - If the site chooses to not use the services of the primary SP, they have to renumber their network.

2.1.1.2 Shim6

Shim6 is a proposed solution from the IETF that is being actively developed. It is a host-to-host based solution where the site gets a PA from multiple SPs and all hosts form addresses based on the multiple prefixes. The hosts then negotiate with the far end of a connection to support the alternate address in the case of a failure. When there is a failure, existing connections will swap the alternate address for the original address for each packet. This swapping is handled by the Shim6 layer.

- Pros:
  - Maintains the goal of strong aggregation.
  - Changes to site connectivity are not propagated into the DFZ.

- Can maintain connections even after an SP failure (for long lived flows).

- Cons:
  - Multiple global addresses on a host would be difficult to manage and debug.
  - Both ends of a connection need to be Shim6 capable. A site has no control over far end of connection, therefore, no control over their network access resiliency.
  - Servers would have a huge amount of connection state to manage.
  - Traffic Engineering would be much more difficult to manage at the network or site level because hosts are making independent decisions on where to address and therefore route packets.
  - Establishing and maintaining filters and shapers is problematic when addresses change and filters include a combination of IP addresses and TCP/UDP ports.
  - Firewall and other security policies are very difficult to establish and maintain because addresses are changing on the fly for given hosts.
  - Shim6 adds additional complexity to the host. This can be a real support issue for low margin users that would have problems understanding the protocol (resulting in more support calls).
  - Shim6 only kicks in after a connection has been up for several seconds. It does not provide redundancy for short lived flows.

2.1.1.3 Provider Independent Addresses

Provider Independent (PI) Addresses are addresses allocated to a site directly from the RIR, not the SP. They are independent of any SP. As such, they can be advertised into any SP. The PI addresses will be from a distinctly identifiable prefix.

- Pros:
  - Meets all redundancy requirements of end sites.
  - Site not tied to any SP (never have to renumber).
- Cons:
  - Every PI address is advertised multiple times (once for each upstream SP) into the DFZ.
  - There is no way to aggregate these addresses in the DFZ.
  - There would be at least two specific (/48) routes for every site in the DFZ (same prefix, different destination/AS path).
  - Changes in connectivity from the site to any of the upstream SPs need to be advertised and processed throughout the DFZ.
  - Only large sites qualify for PI prefixes. This does not solve the site multi-homing issue for smaller sites.

2.1.1.3.1 Geographic-based Provider Independent Addressing

Geographic addressing is a form of the Provider Independent addressing solution proposed in a number of different IETF drafts, some of which are still current, while others are long expired.

In general, these drafts suggest allocated non-PA IP space in a manner that is tied to a geographic location. These addresses could then be aggregated in some area or zone, with additional aggregations as the geographical areas expanded. The idea is that only the routers (of multiple providers) within each of these zones need to know the specific IP allocations, and beyond these geographical zones, the aggregated addresses are sufficient for global reachability. The methods for determining the size of these zones, who allocates the addresses and how individual address assignments are derived, vary among the different drafts.

2.1.1.3.2 Location of Work

Currently this work and proposals exist as IETF drafts. However, actual policy is generally dictated in the individual RIRs, so ARIN would be the appropriate body for policy discussions relating to them.

2.1.1.3.3 Agreement

The pros and cons of PI space are already presented, so the items listed here are mainly with regards to PI space as it is being discussed today versus geographic PI allocations.

The biggest advantage of making geographic based PI allocations is that aggregation becomes a viable option. Current PI space proposals that are being adopted in RIRs set aside a block of IP space out of which these addresses are allocated, with no mechanism to provide for aggregation of these addresses, likened by some to the existing 'swamp' space of PI IPv4 addresses.

Complexity becomes the main disadvantage of this type of mechanism. There are many things that have to happen to make geographic PI addresses work as intended. The aggregation possibilities are diluted with each item that is not properly implemented.

In order to allow for providers to exchange specific routes inside each geographic zone, more exchanges may be required than currently exist. It also becomes much more important that providers exchange routes with peers at each area where they provide service to customers, increasing the financial burden on providers, even when capacity is not an issue. Multiple providers within each region also need to agree on the geographic boundaries where the aggregation would occur.

Authorities to assign addresses to the geographic regions need to be identified. For example, who assigns addresses in a given section of a given city? The RIR's may not be able to manage allocation at a minute geographic level.

Renumbering may also become an issue, because even though sites may no longer be tied to specific providers, they are now tied to geographic locations. This means the freedom from renumbering is not as much as an advantage as it is with non-geographic PI policies.

2.1.1.3.4 Proposed Action

Further discussion is required to determine the level of interest in and support of the idea of geographic based PI space in general, then more specifically which particular draft(s) are the best technical solutions for the problem.

Should the ATIS members believe this approach is a desirable solution to the multi-homing problem, then garnering support in the IETF Multi-homing Task Force is recommended; namely, aligned efforts will be needed to push one of the existing drafts in the IETF to a RFC status so there is a common methodology. Furthermore, due to the complexity of getting this type of solution to work on a global scale, additional focus, input (i.e., contributions), and time would be required to make this the multi-homing solution for IPv6; including:

- Gaining enough support amongst providers (not just tier-1) to push a policy proposal through ARIN and a similar, if not identical, proposal through the other RIRs and
- Working very closely with the rest of the industry to decide the specifics of running this type of architecture.

2.2 Location of Work

Standards/specifications work with respect to multi-homing is primarily done in IETF, while policy related issues are typically contributed to and resolved in ARIN and/or the multiple RIRs.

The reader may find documentation of additional "pros" and "cons" associated with the multiple approaches proposed for site multi-home from the Number Resource Organization (NRO)[4], an organization formed by the RIRs.

2.3 Agreement

All of the solutions described above for multi-homing are put forth by a single branch of the IPv6 community; although the problem has implications to the whole community (e.g., operators, protocol standards, vendors and RIRs). The solution, as such, should involve all of these groups (i.e., stakeholders).

2.4 Proposed Action

A cross organizational effort should be convened with representation from the different interested organizations (IETF, RIRs, vendors, operators) to jointly compile the minimal

[4] www.nro.net

and optimal technical and operational requirements to solve the problem and explore possible solutions.

3 NETWORK ADDRESS TRANSLATION (NAT)

3.1 Issue

This section describes the technical and business impacts to services and applications imposed by the current IPv4 Internet architecture, which is dominated by private addressing domains and the various middleware devices (NAT, ALG, Agents) required to maintain connectivity.

This section also describes the impact to services engineered and deployed into networks that predominantly use global addressing and as such do not require the deployment of NAT and other private addressing mitigation devices for IPv6 traffic.

3.1.1 Technical Issues

Services and applications that are engineered and deployed today need to be able to work in an environment where NATs, ALGs and other middleware devices are assumed to be part of the network infrastructure. The services need to be designed to allow their control and data flows to pass through these various devices with impunity. This often requires multiple service elements and multiple parties interacting to complete a service connection.

Application mediators have recently been developed to assist Peer-to-Peer services in establishing rendezvous points and to identify the specific middleware devices that are in the path of the connection.

In addition to this complexity, the nature of NAT devices imposes a constrained service interaction model that divides the space into clients and servers where the client is responsible for initiating transactions that can only be directed at servers.

Many services require the engineering and deployment of application specific middleware devices needed to traverse the different addressing domains and to coordinate multi-party rendezvous points specific to an application.

End-to-end security is difficult or impossible in today's IPv4 Internet. Encryption (IPsec ESP) does not work over NAT and Authentication (IPsec AH), which possible, is complicated and not widely deployed.

As services become more sophisticated, more complexity is being added to network and service infrastructures. The impact of this complexity increases the costs of development, deployment and operations of networks and services. The fragility and

costs of networks and services are increasing and this trend will continue as more peer-to-peer and collaborative services are deployed.

In summary, the impact to enterprises for deploying and maintaining services in private addressing domains includes the following:

- Increased service and management costs
- Increased complexity
- Increased service fragility
- More network devices to manage
- Decreased end-to-end security
- Limited service model

In a network dominated by a single, transparent global addressing domain, services will be simpler, cheaper, more reliable, more diverse and more secure.

All of the negative impacts listed above go away when global addressing is dominant. The simplicity of deploying services in a global addressing domain plus the removal of the constrained service model will open the door to new service paradigms. Collaboration, sharing, distributed applications and processing will be the foundations of IPv6 enabled services.

3.1.2 Service Provider Business Issues

Internet Service Providers (ISPs) have the option of promoting private addressing domains or global addressing domains for their customer IPv6 deployments. This decision will be based on individual service provider business priorities and strategies. The consequences of one decision over the other should be examined.

The impact of promoting private IPv6 address domains to ISPs stems from the increased complexity of their maintenance and lack of new services. Customers will increasingly rely on the ISP to manage network elements and services, raising the amount of support rendered by a service provider in order to maintain its revenues. The range of offered/available services available from PI addressing will not rapidly change from current offered services which in the short term will allow an increased return on investment for current services but little growth of new services. Newer services that rely on global addressing cannot be offered to private addressing customers which may lead to customer dissatisfaction and the aforementioned loss of revenue. In addition, if a customer decides, after obtaining a PI address, that they want global addresses and their benefits, the customer will incur additional costs to transition, leading again to customer dissatisfaction and the potential loss of the customer.

The consequence of ISPs promoting global IPv6 address domains seems to offer more advantages, the first of which is simplicity. The use of global addresses may lead to customers relying less on SP for management services. ISPs may provide simplified service development and integrated services at reduced operational cost. Some new

services may replace existing SP offerings and services may be more distributed and not require central management. While new services that rely on global addresses can be offered (IPv6 mobility, network mobility, IMS services, Unencumbered P2P, etc) and increase revenues and future Return on Investment (ROI), those existing services which will be replaced may mean initial decreased revenue and decreased ROI on those existing services.

3.2 Location of Work

Not applicable.

3.3 Agreement

As the industry continues to implement a converged IP network and offer advanced services over IP, solutions become much more complicated with private addresses and/or a mix of private and public addresses. In addition to issues specific to service functionally with NATs, broader topics such as wireline-wireless convergence, multicast and continuous media flows (i.e.,VoIP) also need to be considered.

An approach should be taken that satisfies the business realities of SPs without hindering the ability of the SP or customer in moving towards a simplified network and service architecture. Individual ISPs also need to make business decisions based on their understanding and strategy.

3.4 Proposed Action

Individual ISPs need to make business decisions based on their understanding and strategy on the consequences of promoting public or private IPv6 addressing domains.

For SPs that feel the need to recommend private addressing, steps can be taken to transition to global addressing. Provided are two examples:

- Example 1: An enterprise is allocated a global prefix which is used for numbering their IPv6 subnets. A global prefix/address is then translated into an alternate global prefix/address within the NAT device. The enterprise can then transition to global addressing by turning off NAT or selectively electing to not NAT certain subnets or addresses.
- Exmaple 2: An enterprise uses Unique Local addresses which are NAT'ed to the public Internet. If the enterprise wants global addressing, the SP allocates a global prefix in which the enterprise adds global addresses along side Unique Locals. They can then turn off NAT or selectively allow certain addresses to pass through.

4 QUALITY OF SERVICE (QoS)

This section gives a brief summary of IPv6 QoS functionality, including DiffServ, flow label, Multi-protocol Label Switching (MPLS) support of DiffServ, DiffServ-aware MPLS

traffic engineering (DSTE), aggregation of Resource Reservation Protocol (RSVP) reservations over DSTE tunnels and vendor support of IPv6 features. The IPv6 specification developed by the IETF is given in RFC 2460.[5]

4.1 DiffServ

4.1.1 Issue

The meaning, purpose and use of the IPv6 traffic class bits for DiffServ is consistent with the analogous type of service (TOS) bits in IPv4. The 8-bit "Traffic Class" field in the IPv6 header, as specified in RFC 2460 "IPv6 Specification,"[6] is used to carry the DiffServ code point (DSCP), as specified in RFC 2474 "Definition of the DiffServ Field in the IPv4 and IPv6 Headers"[7]:

The DS field structure for the IPv6 Traffic Class is as follows:

```
 0  1  2  3  4  5  6  7
+---+---+---+---+---+---+---+---+
|       DSCP        | ECN |
+---+---+---+---+---+---+---+---+
```

DSCP: Differentiated Services Codepoint
ECN: Explicit Congestion Notification (under discussion in the Transport Area Working Group)

Use of the ECN field is specified in RFC 3168 "The Addition of ECN to IP."[8]

4.1.2 Location of work

IETF, specifically the Transport Area Working Group (TSWG), is addressing the remaining Diffserv-related work items, since the Diffserv WG itself is long concluded.

4.1.3 Agreement

Not applicable.

4.1.4 Proposed Action

Continue to support and utilize this work.

[5] http://www.ietf.org/rfc/rfc2460.txt?number=2460
[6] http://www.ietf.org/rfc/rfc2460.txt?number=2460
[7] http://www.ietf.org/rfc/rfc2474.txt?number=2474
[8] http://www.ietf.org/rfc/rfc3168.txt?number=3168.

4.2 Use of Flow Label

4.2.1 Issue

Various efforts have been made to define the flow label field but so far no agreement has been reached in the IETF. Various proposals include:

1. "A proposal for the IPv6 Flow Label Specification"[9]
2. "A Model for DiffServ use of the IPv6 Flow Label Specification"[10]
3. "A Modified Specification for use of the IPv6 Flow Label for providing An efficient Quality of Service using a hybrid approach"[11]
4. "A Radical Approach in providing Quality-of-Service over the Internet using the 20-bit IPv6 Flow Label field"[12]

4.2.2 Location of Work

Current work is being done in the IETF on RFC 3697.[13] The Security Section of this RFC indicates that there is significant potential for abuse of the flow label feature. Therefore there is some doubt that this feature will be pursued further.

4.2.3 Agreement

RFC 2460 "IPv6 Specification" specifies a 20-bit flow label in the IPv6 header.[14] The flow label may be used by a source to label sequences of packets ("flows") for which it requests special handling by the IPv6 routers, such as non-default QoS or real-time service. The nature of that special handling might be conveyed to the routers by a control protocol, such as RSVP or NSIS, or by information within the flow's packets themselves, e.g., in a hop-by-hop option. There may be multiple active flows from a source to a destination, as well as traffic that is not associated with any flow. A flow is uniquely identified by the combination of a source address and a non-zero flow label. Packets that do not belong to a flow carry a flow label of zero. A flow label is assigned to a flow by the flow's source node. New flow labels must be chosen pseudo-randomly and uniformly from the range 1 to FFFFF hex. All packets belonging to the same flow must be sent with the same source address, destination address and flow label.

RFC 3697 "IPv6 Flow Label Specification" specifies the following:[15]

a) The 3-tuple source, destination, and flow label is used to identify the flow.
b) Packets arriving with the same flow label value after 120 seconds should not be treated as belonging to the same old flow.

[9] http://www.watersprings.org/pub/id/draft-conta-ipv6-flow-label-02.txt
[10] http://www.watersprings.org/pub/id/draft-conta-diffserv-ipv6-fl-classifier-01.txt
[11] http://www.watersprings.org/pub/id/draft-banerjee-flowlabel-ipv6-qos-03.txt
[12] http://www.watersprings.org/pub/id/draft-jagadeesan-rad-approach-service-01.txt
[13] http://www.ietf.org/rfc/rfc3697.txt?number=3697
[14] http://www.ietf.org/rfc/rfc2460.txt?number=2460
[15] http://www.ietf.org/rfc/rfc3697.txt?number=3697

4.2.4 Proposed Action

Industry must formulate an agreement on the assignment and use of the flow label field in IPv6.

4.3 MPLS Support of DiffServ, DiffServ-aware MPLS Traffic Engineering (DSTE) and Aggregation of RSVP Reservations over DSTE Tunnels

4.3.1 Issue

QOS must be maintained across an MPLS network, even over traffic engineering tunnels.

4.3.2 Location of Work

Current work is being done in the IETF on RFC 3697.

4.3.3 Agreement

The following IETF RFCs and Internet drafts should be used for their various areas of implementation:

- RFC 3270 "MPLS Support of DiffServ"[16] allows support of DiffServ for both IPv4 and IPv6 traffic transported over an MPLS network.

- MPLS signaling & MPLS fast reroute (FRR): RFC 3036 "LDP Specification" specifies support for IPv6.[17] RFC 3209 "RSVP-TE: Extensions to RSVP for LSP Tunnels "specifies support for IPv6.[18] RFC 4090 "Fast Reroute Extensions to RSVP-TE for LSP Tunnels" also specifies support for IPv6.[19]

- RFC 4124 "Protocol Extensions for Support of DiffServ-aware MPLS Traffic Engineering".[20]

- IETF draft "Aggregation of RSVP Reservations over MPLS TE/DSTE Tunnels".[21] The procedures defined in this draft are applicable to these cases:

  - Aggregation of E2E IPv4 RSVP reservations over IPv4 TE Tunnels.
  - Aggregation of E2E IPv6 RSVP reservations over IPv6 TE Tunnels.
  - Aggregation of E2E IPv6 RSVP reservations over IPv4 TE tunnels, provided a mechanism such as [6PE] is used by the Aggregator and Deaggregator for routing of IPv6 traffic over an IPv4 MPLS core.

[16] http://www.ietf.org/rfc/rfc3270.txt
[17] http://www.ietf.org/rfc/rfc3036.txt
[18] http://www.ietf.org/rfc/rfc3209.txt
[19] http://www.ietf.org/rfc/rfc4090.txt
[20] http://www.ietf.org/rfc/rfc4124.txt?number=4124
[21] http://www.ietf.org/internet-drafts/draft-ietf-tsvwg-rsvp-dste-03.txt

- Aggregation of E2E IPv4 RSVP reservations over IPv6 TE tunnels, provided a mechanism is used by the Aggregator and Deaggregator for routing IPv4 traffic over IPv6 MPLS.

4.3.4 Proposed Action

RFC 4124 makes no mention of support for IPv6. This should be addressed with a minor RFC "Patch."

RSVP-TE provides more features including fast-reroute, traffic engineering and DiffServ-TE, which a native IPv6 only core does not support. The implication is that vendors do not support the IPv6 features since the standards appear to specify IPv6 support for FRR and TE, as discussed above.

4.4 IPv6 End-to-End QoS Signaling

4.4.1 Issue

For signaling of individual flows, the IETF Next Steps in Signaling (NSIS) working group is progressing QoS signaling.[22] NSIS QoS signaling addresses the shortcomings of the RSVP protocol and defines a two-layer QoS signaling model:

- NSIS QoS signaling layer protocol (QoS-NSLP)[23]
- NSIS transport layer protocol (NTLP) [24]

4.4.2 Location of Work

IETF Next Steps in Signaling (NSIS) working group.

4.4.3 Agreement

Support IPv6 and allows for the population of the flow label in identifying an IPv6 flow.

A QoS specification (QSPEC) object which contains the QoS parameters required by a QoS model (QOSM) should be defined.[25] A QOSM specifies the QoS parameters and procedures that govern the resource management functions in a QoS-aware router. Multiple QOSMs can be supported by the QoS-NSLP the QoS-NSLP allows stacking of QSPEC parameters to accommodate different QOSMs being used in different domains. As such, NSIS provides a mechanism for interdomain QoS signaling and interworking.

The QSPEC parameters include, among others:

- TRAFFIC DESCRIPTION Parameters:
  - <Token Bucket> Parameters

[22] http://www.ietf.org/html.charters/nsis-charter.html
[23] http://www.ietf.org/internet-drafts/draft-ietf-nsis-qos-nslp-11.txt
[24] http://www.ietf.org/internet-drafts/draft-ietf-nsis-ntlp-09.txt
[25] http://www.ietf.org/internet-drafts/draft-ietf-nsis-qspec-10.txt

- <Bandwidth> Parameter

- CONSTRAINTS Parameters:
  - <PHB Class> Parameter
  - <Y.1541 QoS Class> Parameter
  - <DSTE Class Type> Parameter
  - <Reservation Priority> Parameter
  - <Preemption Priority> & <Defending Priority> Parameters
  - <Path Latency> Parameter

Definitions of various QOSMs in progress, are as follows:

- 'NSIS QOSM for Networks Using Y.1541 QoS Classes' based on ITU-T Recommendations [Y.1541] 'Network Performance Objectives for IP-Based Services' and 'Signaling Requirements for IP-QoS',[26]
- 'NSIS QOSM for Networks Using Resource Management in Diffserv (RMD)''[27]

The ability to achieve end-to-end QoS through multiple Internet domains is also an important requirement. End-to-end QoS signaling is used to insure end-to-end QoS, and an example of using NSIS QoS signaling and the Y.1541-QOSM to achieve this is provided.

The QoS originating node (ON) initiates an end-to-end, inter-domain QoS RESERVE message containing the QoS parameters, including for example, <Y.1541 QoS Class>, <Token Bucket>, <Reservation Priority>, <Path Latency>, and perhaps other parameters for the flow. The RESERVE message may cross multiple domains perhaps supporting different QOSMs. At the ingress edge node of the local-QOSM domain, the end-to-end, inter-domain RESERVE message triggers the generation of local QoS parameters derived from the RESERVE message and consistent with the local QOSM. The local parameters and QOSM are used for QoS processing in the local-QOSM domain, and then the original RESERVE message sent by the ON is used for QoS processing at the QoS destination node (DN).

Each node on the data path checks the availability of resources and accumulating the delay, delay variation and loss ratio parameters as described below. If an intermediate node cannot accommodate the new request, it indicates this by marking a single bit, and the reservation is denied. If no intermediate node has denied the reservation, the ON RESERVE message is forwarded to the next domain. If any node cannot meet the requirements designated by the ON RESERVE message to support a QoS parameter, for example, it cannot support the accumulation of end-to-end delay with the <Path Latency> parameter, the node sets a flag that will deny the reservation. Also, parameter negotiation can be done, for example, by setting the <Y.1541 QoS Class> to a lower class than specified in the ON RESERVE message. When the available <Y.1541 QoS Class> must be reduced from the desired <Y.1541 QoS Class>, say because the delay objective

[26] http://www.ietf.org/internet-drafts/draft-ietf-nsis-y1541-qosm-02.txt
[27] http://www.ietf.org/internet-drafts/draft-ietf-nsis-rmd-07.txt

has been exceeded, then there is an incentive to respond to the ON with an available value for delay in the <Path Latency> parameter. For example, if the available <Path Latency> is 150 ms (still useful for many applications) and the desired QoS is 100 ms (according to the desired <Y.1541 QoS Class> Class 0), then the response would be that Class 0 cannot be achieved and Class 1 is available (with its 400 ms objective). In addition, the response includes an available <Path Latency> = 150 ms, making acceptance of the available <Y.1541 QoS Class> more likely.

4.4.4 Proposed Action

Continue to support the standards work as identified.

4.5 IPv6 QoS with Flow State Networking

4.5.1 Issue

QOS Issue 4: Flow State Networking (FSN) Issues:

a) FSN approach inconsistent with end-to-end QoS approach to signal Y.1541 IP QoS classes end-to-end; the FSN approach does not signal Y.1541 classes.
b) Service provider network evolution may not be toward implementing FSN routers.

4.5.2 Location of work

ITU-T Study Group 12 and Study Group 13 are currently working on a new Recommendation, Y.flowreq, which is relevant in scope to flow state networking.

4.5.3 Agreement

Flow state networking is an approach to IPv6 QoS that embeds the QoS signaling in the IPv6 header itself. An IPv6 flow state router uses the 3-tuple (source address, destination address and flow label) to identify the flow (an IPv4 flow state router uses the 5-tuple to classify the flow). Each flow is assigned a flow state that identifies the QoS treatment while the flow is in progress. QoS is applied at the flow level with a per-flow scheduler to allow scheduling of flows independent of each other. The advantage of flow state networking is that per-flow reservations are not required, and the solution may be more scalable in that respect. However, identifying and classifying flows, and retaining flow state are large computational tasks, and are done in hardware/firmware in current implementations.

It has been proposed that an IPv6 QoS signaling standard to permit the specification of the QoS of a flow (or group of flows) in a QoS Extension Header be created. This permits the QoS to be setup in real time without a separate signaling message structure such as RSVP. The QoS request and response are incorporated into the data flow packet headers so that the QoS can be setup during the first round trip. The QoS extension header can be used to specify rate, burst tolerance, precedence and delay priority for a

flow and to determine the available rate the network path can support. Each router in the path examines the QoS Header and agrees or adjusts the rates requested to the rates it can support.

4.5.4 Proposed Action

None. The path forward proposed by SGs 12 and 13 does present challenges. Routers and hosts in the Internet will need to be modified for this scheme to work. Based on these challenges, the creation of necessary standards will likely be delayed or rejected outright.

5 INTEROPERABILTIY BETWEEN IPv4 AND IPv6

5.1 Translation of IPv6-IPv4

5.1.1 Issue

As industry transitions from IPv4 to IPv6, there is a need to assure service continuity and hence network interoperability. To this end, it is anticipated that industry's general transition strategy is to achieve network (IP) protocol independence (application to application interoperability throughout) in the network by means of evolving, to the extent possible, the network edges into dual-stack capabilities. The aim is to not to deal with network interoperability but rather move the selection of which version of IP to use to the (dual-stack) edge devices.

Known challenges with this approach need to be considered and appropriately addressed. Issues needing consideration include (1) the lack of wide-scale dual-stack deployment whereby one edge device is IPv6 only (e.g., mobile handsets) and the other remains IPv4 only, and (2) utilization of gateway devices is not ideal and could potentially create problems such as single point failure, congestion, added deployment expense and redundancy; all while also potentially limiting services and service features (via use of NAT devices).

An example of an 'end-tailing-off' interoperability scenario, as described above, would be the United States Department of Defense (DOD) eventual transition to native IPv6. Once this happens, problems will inevitably arise for their partners who remain solely IPv4.

5.1.2 Location of Work

NAT Protocol Translator (PT) specifications (which translate IP protocol from one to another) have been developed by the IETF IPv6 Operations Working Group (IETF RFC (RFC 2766). These specifications, however, have been demoted as a standard.

5.1.3 Agreement

The general agreement amongst the ATIS IPv6 Task Force members is to achieve an IP version independent network by means of evolving the edge devices to dual-stack. It has also been agreed that industry guidelines (e.g., best practices) should be developed to ensure continued interoperability between and amongst carriers during industry's transition to dual-stack, while also minimizing the potential impacts of incompatibility between carriers that elect to deploy only IPv6 or remain IPv4 only.

5.1.4 Proposed Action

Industry should undertake steps to: (1) have carriers transitioning to IPv6 collect valuable lessons learned; (2) conduct more research to determine the standard operational practices of IPv6; (3) utilize the data-points collected to describe some of the best practices and their consequences.

5.2 Partitioned Internet

5.2.1 Issue

As IPv4 addresses become more limited and IPv6 services become more common, IPv6 only host networks and services will become viable and start to emerge. During this stage of industry's migration to IPv6, many IPv4 only host networks and services will remain. This will create window of time where a partitioned Internet exists; for example, an IPv4 only Internet and an IPv6 only Internet.

During this window, it will not be impossible for emerging IPv6 host networks and services to access the existing IPv4 Internet; however, it will be increasingly difficult and costly to do so. Also during this period, it will not be impossible for IPv4 only hosts and services to reach the IPv6 network; but again, there will be a certain cost and complexity involved.

5.2.2 Location of Work

Individual organizations will need to assess these conditions and ascertain their plans forward.

5.2.3 Agreement

Avoiding a partitioned Internet will be based on an individual organization's requirements. If the value of solving the problem for a particular organization is high enough, the investment will be made in various proxy, ALG and translation devices.

5.2.4 Proposed Action

None. The pace in which industry transitions to IPv6 will be decided by the marketplace. Accordingly, industry should ascertain at what point they believe critical mass is achieved (e.g., a particular service offering or feature, the deployment of IPv6 by a collection of large enterprises, and/or overwhelming user/subscriber demands) and continually assess these market forces and conditions.

6 IMPACT ON TRAFFIC AND ROUTING

6.1 Issue

Transition strategies inherently call for an extended period of operation where both IPv4 and IPv6 protocols are either explicitly supported or indirectly supported by various tunneling mechanisms. While these tunnels are in effect, the network operator now has to manage essentially two different networks with different topologies. Any traffic engineering activities (e.g., manipulation of routing administrative weights, use of MPLS traffic engineering (TE) tunnels), also need to be duplicated to include changes in traffic patterns; since not all peering relationships will transition to supporting IPv6 at the same time.

Routing policies at peering points will also need to be developed for IPv6, as well as the existing IPv4. It may be expected that these peering arrangements will have to also act as tunneling/conversion points in the case where operators have chosen anything other than native IPv6 support.

One of the values of IPv6 is advanced multicast support. Efforts need to be made to ensure there are standards/best practices in place on how to support multicast on a MPLS network. The use of the flow-label field in IPv6 also presents yet another challenge in offering a feature-rich end state.

6.2 Location of Work

Multicast support is being developed for MPLS in the IETF MPLS Working Group. [28]

Several related specifications include:

- RFC 4461 "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)" [29]
- RFC 4687 "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks". [30]
- RFC 3353" Framework for IP Multicast in MPLS [31]

6.3 Agreement

The above standards should be assessed as possible when IPv6 is deployed, even if in a trial state.

6.4 Proposed Action

On completion of the specifications, efforts need to be undertaken to determine best practices to minimize impact on traffic and routing. Tunneling issues needing

[28] See the MPLS WG charter at http://www.ietf.org/html.charters/mpls-charter.html
[29] http://www.ietf.org/rfc/rfc4461.txt
[30] http://www.ietf.org/rfc/rfc4687.txt
[31] http://www.ietf.org/rfc/rfc3353.txt

consideration include (1) traffic management, (2) IPv6 over IPv4, (3) encapsulation of IPv6 in IPv4 and (4) NAT transversal.

7 SECURITY

7.1 IPv6 Overall Security Issues

7.1.1 Issue

The argument can be made that end-to-end IPv6 will provide better security through its elimination of NATs and implicit use of IPsec, a generic security architecture in the IP level, although it should be noted in practice not all implementations of IPv6 include IPsec. The use of basic security mechanisms at the network level will alleviate some of the security issues previously found in the IPv4 environment, because network level security mechanisms are available to all layered applications. Due to the simplicity of their implementation, a large percentage of network attacks are low level, so their prevention in the IP level seems an appropriate function. However, IPsec does not provide security to user-oriented applications, a reminder that the security of a system is based on the sum of all its elements.

Many security issues that will occur in the IPv6 space are not unlike those that appear in the IPv4 environment. Some are unique to or could be more likely in the IPv6 environment and in the transitory phases of IPv6 deployment and IPv4/v6 coexistence.

Those areas of security risk affecting IPv6 overall include:

- Attack
- Transparency/Peering
- Addressing
- Routing Headers

7.1.1.1 Similarities to IPv4

Similar to the IPv4 space, sniffing, application layer attacks, rogue devices, man in the middle (MITM) attacks and flooding will occur in the IPv6 space. Without IPsec, IPv6 as a protocol is no less likely to experience a sniffing attack than IPv4 and any attacks using MITM will also have similar likelihood. Regardless of IPsec, flooding will occur in IPv6 as in IPv4. Rouge devices will be as easy to insert into an IPv6 network as an IPv4 network. Most importantly it should be noted that the use of IPv6, even with IPsec, will not prevent or alleviate the vulnerabilities for attack found in the application layer.

7.1.1.2 Reconnaissance

While reconnaissance in IPv4 is relatively easy, in the IPv6 space the difference in subnet size makes the process all but impossible. The default subnets in IPv6 have 2(64) addresses, making the time required to scan each address extraordinary and unlikely to

take place. Additionally New Mobility Anchor Point (NMAP) does not support ping sweeps on IPv6 networks. In IPv6, public servers will require Domain Name Server (DNS) reachability, creating a point at which attack may occur. A decrease in NAT usage will increase dynamic DNS usage, making DNS a likely target for attack.

The most likely solution to reconnaissance in IPv6 is to change scanning methods. The use of addresses from Organizationally Unique Identifier (OUI) designations allows scanning focus on popular Network Information Centre (NIC) vendor's ranges. Administrators can adopt easy to remember addresses and scanning should occur at compromised routers in key transit points.

Other practices to lessen the impact of reconnaissance attack are privacy extensions, filtering, standard non-obvious addressing, and host and application security. While using privacy extensions may lessen the opportunity for outside attack, they should be implemented carefully as their use can also complicate trackback and troubleshooting within a network. Internal IPv6 addresses should be filtered at the organizational border routers. Unneeded services should be filtered at the firewall and Internet Control Message Protocol (ICMP) should be selectively filtered. Standard non-obvious static addresses should be used for critical systems.

7.1.1.3 Unauthorized Access

Implementation of authorized access control is a policy decision that occurs most often in layers 3 and 4. In IPv6 the use of privacy extensions offer some limit to the exposure of a host address, but as noted above, also makes it harder for a network administrator to identify an end host. The use of local unicast addressing in an enterprise network can be used to automatically deny inbound and outbound access for enterprise only services. Bogon filtering is another option since only the three Top-Level Aggregation identifiers (TLAs) have been allocated; Access Control Lists (ACLs) can permit those ranges exclusively and block other IPv6 traffic.

7.1.1.4 Maladies

While the scanning qualities of most hybrids and worms will prevent their effectiveness in IPv6, viruses and some adapted worm development will still require the detection and mitigation practices developed in IPv4 to maintain network security and throughput.

7.1.1.5 Transparency/Peering

IPsec with encryption makes current network based firewalls blind to the upper layer information. Distributed firewalls (personal) can see the packet after decryption. IPsec when used with authentication headers does not provide encryption but does provide integrity.

ICMP is more heavily relied upon in IPv6 than in IPv4. ICMP policy on firewalls needs to be changed to accommodate this reliance. Firewalls at a minimum should allow link-

local multicast traffic to a device on the network. In transparent mode a firewall must understand the new uses of ICMPv6 and allow filters to be defined for each case. Network administrators should determine which ICMPv6 messages are required though the access control device and apply filters appropriately. In addition ICMPv6 policies should be mapped as closely to the equivalent ICMPv4 policies as possible.

Neighbor discovery (ND) protocol is used in the IPv6 space to replicate many of the functions performed by address resolution protocol, router discovery and IMCPv4 in the v4 space. ND allows for different nodes on the same link to advertise their existence to their neighbors, and to learn about the existence of their neighbors for the purposes of forwarding and redirecting packets. Unsecured ND can make networks and node vulnerable to DoS (i.e. a malicious node prevents communication between nodes) or redirect attacks (i.e. an attacker redirects packets to another node instead of the intended node). During a redirect attack, information can be snooped and intercepted, or in the case where the redirected information is large, flood the network interface and log files of the victim host, rendering it inoperable.

In the IPv6 space, efforts to secure networks against neighbor discovery attacks (SEND) initially espoused the use of IPsec authentication headers and automatic keying. However, automatic key management does not provide adequate provision for mobility and roaming devices. Manual keying is also not a scalable solution and may cause administration hardships if enacted on a large scale. More recently the use of cryptographic generated addressing has been investigated and may provide more mitigation of ND attacks without the effort needed to manually key authentication.

7.1.1.6 Addressing

The use of multicast addresses can enable an adversary to identify key resources on a network and attack them. These addresses must be filtered at the edge in order to make them unreachable from the outside.

The use of Cryptographically Generated Addresses (CGA) does provide some relief from Neighbor Discovery. The use of the CGA specific id (hash of the sender public key) allows certificate functionality without requiring a key managed infrastructure.

7.1.1.7 Routing Headers, Manipulation and Fragmentation

IPv6 endpoints may accept IPv6 packets with a routing header. As such, it is possible that an end-host will process a routing header for redirection or forwarding of a packet.

The unlimited size of the header chain can make filtering difficult due to the larger number of boundary conditions to exploit, and denial of service attacks may result with poor IPv6 stack implementations. Header manipulation may also occur if IPv6 header fragments are not denied to inter-network devices.

The network designer should validate that the operating systems within their network do not forward packets that include a routing header. A specific set of nodes should be

designated as Mobile IPv6 (MIPv6) home agents. If MIPv6 is not needed then the network manager can filter IPv6 packets that contain routing headers at an access control device. IPv6 extension header policies should also be designed to closely follow the IPv4 header policy.

7.1.1.8 Mobility

The growth of wireless and mobile services is expected to dynamically increase during the next decade, and these services and devices are expected to use IPv6 to network. The IETF has begun work on a Mobile IP solution in which mobile devices have a permanent IP address which is broadcast to anyone and a temporary address in connection with its placement on a network (home or roaming). Information is directed to the device via the devices permanent address which then forwards to the temporary address. This system creates a duplicate number of communications, each of which increases a node's vulnerability to attack. The aforementioned use of IPsec and CGA are likely solutions to vulnerabilities faced by mobile IPv6, however the broadcast of a mobile device's permanent IP address has required further examination. ARMOR and HMIPv6 (Hierarchical Mobile IPv6) were developed for the ability to move without revealing a device's permanent address (ARMOR), and for local movements without requiring a new temporary address to be established thus hiding the device from other users (HMIPv6).

7.1.2 Location of Work

Some work has begun in IETF IPv6 Operations (v6ops) and IP Next Generation (IPNG) working groups.

7.1.3 Agreement

A "best practices" document for all communications service providers and users of services to combine the IETF RFCs addressing security against attack should be developed.

7.1.4 Proposed Action

- Addressing solutions (e.g., use of unique local addresses whenever applicable) and filtering schemes should be implemented.
- RIRs should implement policies for address allocation which allows for the selective filtering of IPv6 traffic on a regional basis.
- Remain vigilant in the development and implementation of application security.
- Encourage inclusion of IPsec in all IPv6 implementations.
- Review the RFC list of IPv6 Security issues and threats on an ongoing basis:
  - IPv6 Transition and Coexistence Security Considerations
  - Recommendations for Filtering ICMPv6 Messages and Firewalls
  - Network Architecture Protection

7.2 Security during Enterprise Testing/Service Development

7.2.1 Issue

During the evaluation phrase of IPv6, there exists the possibility of unsecured service piloting. Accordingly, network architects must be cognizant of the possible immaturity of software applications for implementing IPv6 security when in the testing phase.

7.2.2 Location of Work

Some work begun in the IETF v6ops working group.

7.2.3 Agreement

A "best practices" document for service providers should be developed.

7.2.4 Proposed Action

During the testing phase of IPv6, it is recommended that service providers do not connect to the IPv6 Internet until it is absolutely critical for the organization to do so; since there are many threats. Even then, connections should be limited to those areas that are necessary.

7.3 Security During v4/v6 Co-Existence

The main techniques and mechanisms for the transition to IPv6 are tunneling, translation and dual stack operation. Tunneling and translation are more often associated with security impact; however dual stack scenarios also have some security considerations. Tunneling (IMCP, 6to4, 4to6) and translation (NAT) are security issues regardless of the protocol in which they are working and have long since been the avenue for covert network attacks.

7.3.1 6to4 Tunneling

7.3.1.1 Requirement

The architecture for automatic 6to4 tunneling includes 6to4 routers and 6to4 relay routers, which accept and de-capsulate IPv4 protocol-41 ("IPv6-in-IPv4") traffic from any node in the IPv4 Internet. This characteristic enables a number of security threats including DoS and spoofing. The means for a "meta-threat", traffic laundering, is also possible and can be used in conjunction with other threats, whether specific to 6to4 or not.

The main vulnerabilities of the 6to4 tunneling scenarios as outlined in RFC 3964 are:

- 6to4 routers have to consider all 6to4 relays, and other 6to4 routers, as "on-link"
- 6to4 relays have to consider all 6to4 routers as "on-link", and
- the discovery that at least a couple of major 6to4 implementations do not implement all the security checks.

Generally, Layer 4 (L4) spoofing remains the same in IPv6 as the global aggregation of addresses makes spoof mitigation at aggregation point easy to deploy. However, L3-L4 spoofing may occur during 6to4 tunneling, the tunneling mechanisms acting as a conduit because there is no way for any 6to4 router to confirm the identity of the IPv4 node from which it receives traffic. In the case where an IPv4 address is spoofed, 6to4 ACLs are rendered ineffective. The IPv4 header is stripped by the 6to4 relay, making traceback difficult, and the inner packet is forwarded on. An attacker may even choose the relay as the spoofed source address. There is hope for mitigation of harm however, due to the 1:1 ratio of packets which prohibits amplification of an attack and the possibility of checks within the 6to4 relay not to accept 6to4 addresses.

Solutions for 6to4 security concerns are to only allow authorized endpoints to establish tunnels. Static tunnels may be a more secure way to tunnel, but since they are not scalable, may not work in every scenario.

7.3.1.2 Location of Work

IETF v6ops working group.

- RFC 1933/2893 - Transition Mechanisms for IPv6 Hosts and Routers
- RFC 2401 - Security Architecture for IP November
- RFC 2473 - Generic Packet Tunneling in IPv6 Specification
- RFC 2529 - Transmission of IPv6 over IPv4 Domains without Explicit Tunnels
- RFC 3964 - Security Considerations for 6to4
- RFC 3056 - Connection of IPv6 Domains via IPv4 Clouds

7.3.1.3 Agreement

A "best practices" document for all communications service providers and users of services to combine the IETF RFCs addressing security against attack should be developed.

7.3.1.4 Proposed Action

Participate in and contribute to the IETF work. Review the RFC list of IPv6 Security issues and threats on an ongoing basis:

- IPv6 Transition and Coexistence Security Considerations
- Network Architecture Protection

7.3.2 Dual Stack

7.3.2.1 Issue

Host security in the dual stack scenario is of the utmost importance. Applications can be subject to attack from both IPv4 and IPv6. It is recommended that host security controls block and inspect traffic from both IPv4 and IPv6 through the use of host intrusion prevention, personal firewalls and Virtual Private Network (VPN) clients. Also for

consideration in the dual stack environment is the timeliness of the development of an IPv6 backbone, and the native support of IPv6 security technologies.

7.3.2.2 Location of Work

The IETF v6ops working group.

7.3.2.3 Agreement

A "best practices" document for all communications service providers and users of services to combine the IETF RFCs addressing security against attack should be developed.

7.3.2.4 Proposed Action

Participation in and contribution to the IETF work is recommended. The RFC list of IPv6 Security issues and threats should be reviewed on an ongoing basis, including those for:

- IPv6 Transition and Coexistence Security Considerations
- Network Architecture Protection

7.3.3 NAT Gateways

7.3.3.1 Issue

The perceived use of NATs in the IPv4 space as a security feature is one of the most perplexing concerns of IPv6 security. Even in the IPv4 space, NATs have been known to break applications and security by modifying IP packets. (This is essentially why IPsec and IPv4 don't play nice.) They prevent the trace-back of an attack to a source machine.

The continued use of NATs in the IPv6 space may cause security issues because they prevent end-to-end security in the network layer. Workarounds are complex, costly and sometimes not entirely possible.

NATs in an IPv6 network would also limit many of the main perceived values of IPv6. Ultimately, the deployment of NAT in IPv6 networks however will depend on the following:

a) Maturity and field hardening of the IPv6 security mechanisms that are meant to replace the security provided by NAT. If these mechanisms are deployed and proven to be effective, the value and arguments for NAT deployment will be diminished
b) Development of services, applications and security mechanisms that provide a clear value add when deployed without NAT. The promise of IPv6 is end-to-end type services and security. The availability of such applications that are easier, cheaper or provide superior value when deployed without NAT will cause organizations to explore security options other than NAT.

c) Development of an IPv6 multi-homing capability. NATs allow a multi-homed network to change service providers more easily. For some, this is a strong argument for NAT deployment.

7.3.3.2 Location of Work

IPv6 advocacy organizations such as IPv6 Forum.

7.3.3.3 Agreement

- The potential value of IPv6 is greater when deployed without NAT.
- If organizations are to move away from NAT for IPv6, the risks of such deployments must be diminished and demonstrated, and the values of networks without NATs must be clear and demonstrated.
- "Proof" that the security of networks using IPv6 without NAT is equivalent or greater than IPv4 with NAT is paramount. This may be achieved in a number of ways, one of which may be a demonstration of an IPv6 network's abilities to deal with attack.

7.3.3.4 Proposed Action

Solicit IPv6 advocacy organizations to provide proof of the NAT replacement security mechanisms and values of IPv6 without NAT.

8 PRIVACY IMPLICATIONS OF IPv6 ADDRESSES

8.1 Issues

Privacy concerns for IPv6 and the ability to track the identity of a user across the network based on an IP address are basically similar to that of IPv4. Once an IP address is known to be associated with a particular user, that address can be tracked and a history can be recorded to gather information about the individual's net behavior. While IPv6 does not introduce any new concerns to privacy over IPv4 once the association between IP address and individual is known, IPv6 addresses generated by stateless address autoconfiguration have the potential to facilitate providing information about that association. The technical detail that explains how this addressing works is described in RFC 2373 relating to the IPv6 addressing architecture, RFC 2642 entitled'IPv6 Stateless Address Autoconfiguration' and RFC 2374 titled 'IPv6 Aggregatable Global Unicast Address Format.'

The rationale behind stateless address autoconfiguration is to generate a global unique address without the need for a Dynamic Host Configuration Protocol (DHCP) server. The IPv6 address is generated from a combination of information about the public topology and site topology. The site topology portion of the address is made up of 16 bits of a Site-Level Aggregation Identifier (SLA-ID) which is used by an individual organization to create its own local addressing hierarchy and to identify subnets and a 64-bit interface ID. Interface identifiers are unique serial numbers or addresses that are

link dependent, used to identify interfaces on a link, and are required to be unique on that link. Some IPv6 systems use the right 64 bits of the address to store an IEEE defined global identifier (EUI64). This identifier is composed of a company id value assigned to a manufacturer by the IEEE Registration Authority. The 64 bit identifier is a concatenation of the 24-bit company identification value and a 40-bit extension identifier assigned by the organization with that company identification assignment. The 48-bit MAC address of a network interface card may also be used to make up the EUI64. It is this portion of the address that is the basis for the privacy concerns in that it is generated based on the ID of the hardware interface of the device and could be used to identify each machine individually.

RFC 3041 titled 'Privacy Extensions for Stateless Address Autoconfiguration in IPv6' describes the potential problem with this type of IPv6 address and highlights the fact that any communication system which has a constant address or identifier for incoming and outgoing communication has potential privacy concerns. As mentioned above, this concern would be the same for IPv4 or IPv6. Not all nodes and interfaces contain IEEE identifiers, and in many cases an interface identifier is generated through some other means (e.g., at random), and the resultant interface identifier is not globally unique and may also change over time. The focus of RFC 3041 is on addresses derived from IEEE identifiers, as the privacy concerns are only in those cases where the interface identifier is globally unique and non-changing.

8.2 Location of Work

While RFC 3041 was originally drafted under the Network Working Group of the IETF, there is a current internet draft under the IPv6 Working Group titled 'Privacy Extensions for Stateless Address Autoconfiguration in IPv6'[32] that, if approved, will obsolete RFC 3041. This draft describes an extension to IPv6 stateless address autoconfiguration for interfaces whose interface identifier is derived from an IEEE identifier. Use of the extension causes nodes to generate global scope addresses from interface identifiers that change over time, even in cases where the interface contains an embedded IEEE identifier.

8.3 Agreement

In order to alleviate privacy concern associated with the machine generated portion of the IPv6 address derived from EUI64 identifier, IPv6 privacy extension methods should be used in conjunction with stateless address autoconfiguration.

8.4 Proposed Action

None.

[32] http://www.ietf.org/internet-drafts/draft-ietf-ipv6-privacy-addrs-v2-05.txt

9 MANAGEMENT TOOLS (DEVELOPMENT) – MANAGING DUAL STACK AND IPv6 NETWORKS

9.1 Issue

With transitioning to IPv6 (either via deployment of dual-stack or native IPv6), the tools developed and utilized today for management of IPv4 networks will need to be reevaluated and likely revised. For instance, network management models and the databases used to manage the devices in a communications network will need to be adapted to IPv6 (e.g., the Management Information Base (MIB)) which includes objects such as an output queue length and Address Translation table (like ARP tables). Tools for extracting addresses that are auto-derived from IPv6 will also need to be developed and implemented for IP address management.

9.2 Location of Work

IETF is trying to keep the MIB's up to date, while vendors are trying to follow requirements.

IETF RFCs of interest are:

- RFC 1155, "Structure and Identification of Management Information for TCP/IP based Internets",
- RFC 1213, "Management Information Base for Network Management of TCP/IP-based Internets", and
- RFC 1157, "A Simple Network Management Protocol".

9.3 Agreement

Not applicable.

9.4 Proposed Action

Service providers should communicate internally with their Operations department to gauge the level and types of tools used that are company proprietary and consider if those tools should continue to be proprietary.

Efforts should also be given to ensuring MIB objects are appropriately adapted and/or adopted for dual-stack and native IPv6.

10 IMPACT ON INFRASTRUCTURE RELIABILITY

10.1 Issue

Maintaining a reliable and secure network is of paramount importance to service providers. Any movement away from existing configurations and additions or changes

to operationally tested and field-proven embedded infrastructure devices must be thoroughly and comprehensively evaluated for impacts.
Examples of issues needing consideration include:

- The maturity of IPv6, as it has yet to be field-proven or field-hardened.
- The deployment of new software developed and hardware deployed for IPv6; as potential instabilities with IPv6 increases the risks to existing (IPv4) networks.
- The lack of sufficient memory capabilities embedded in infrastructure devices to support IPv6.
- Increased vulnerability to DOS attacks in and with IPv6, which can also impact existing infrastructure.
- Potential for network disruptions, delay and overall quality due to overtaxed routing tables and overflowed memory buffers.

10.2 Location of Work

Not applicable.

10.3 Agreement

Service providers should continue to look at the individual vendors' infrastructure devices to better understand current capabilities versus future needs (e.g., memory capacity) to meet the additional memory requirements of IPv6. It is also important to be mindful that simple software upgrades to existing devices might not be sufficient. It is important therefore that organizations transitioning to IPv6 evaluate their immediate and long-term needs and communicate these requirements with their vendors.

10.4 Proposed Action

Continue to evaluate IPv6 and the impacts of IPv6 on existing infrastructures through laboratory evaluations, deployment of a closed IPv6 network and field trials. Share experiences of deploying IPv6 in live and trial networks.

11 IMPACT TO ACCESS NETWORKS (OPERATING AT LAYER 3)

11.1 Issue

With service providers actively engaged in capital intensive, multiyear investments to upgrade facilities for consumer Internet access from "dial-up" to broadband access, the premature introduction of IPv6 could introduce additional project risks for current deployment programs and potentially premature obsolescence of newly deployed equipment.

Although standardization efforts and activities anticipate no adverse affects to the access networks resulting from IPv6, operational aspects need to be given due consideration and factored into a carrier's transition strategy.

11.2 Location of Work

DSL Forum; most notably WT-144 which is a revision document TR58 "Multi-Service Architecture and Framework Requirements".

11.3 Agreement

The DSL Forum should continue to evolve work affecting this area.

11.4 Proposed Action

Continue to monitor/interact with the DSL Forum to promote continued/completed work.

12 PEERING EVOLUTION (IPv6 OVER IPv4)

12.1 Issue

IPv6 peering works by exchanging IPv6 routes at the peering session and allowing IPv6 traffic to transit between two Autonomous Systems. In the United States, the Moon6 backbone currently provides a native IPv6 backbone to support peering. The development of a native IPv6 backbone and environment will allow the development of peering based services and applications.

From early IPv6 peering investigations, the preferred method is through an exchange point, however direct peering to the backbone router is possible in some cases as is peering through GigaPoPs, through tunnels or BGP multihop.

For the purposes of BGP routing design for peering sessions, IPv6 peering and IPv4 peering can be separated so that the impacts of IPv6 deployment are lessened on the existing IPv4 network. In this situation the same protocol should be used for both IPv4 and IPv6.

For the purposes of Open Shortest Path First (OSPF) routing (using version 2 for IPv4 and version 3 for IPv6), the versions used for each are essentially different protocols, enabling the deployment of IPv6 peering without effect on the IPv4 network.

For IS-IS routing, there is a single topology for IPv4 and IPv6 though there is a multi-topology extension. Without the use of the multi-topology extension there is need for an x-day. The single topology also makes it more difficult to enable IPv6 peering to gradually co-exist with IPv4 peering as all IS-IS node must have IPv6 enabled. Once IPv6 is enabled, it is suggested that the same peering protocol be used as in IPv4.

12.2 Location of Work

To be determined.

12.3 Proposed Action

None.

13 IPv6 IMPACTS TO BILLING/ACCOUNTING

13.1 Issue

In the IPv6 space, billing and accounting will continue to become integrated with features such as authentication and authorization. The billing models used for IPv4 may not encompass the needs for billing for IPv6 access and services. In IPv4, billing models have evolved from a strict volume based charge, to a fixed tariff, to a combination of volume based differential charges for local value additions and content based services in addition to a fixed tariff for general services access.

In IPv6, billing systems will have the ability to do flow-based accounting through addressing schemes and interface to flow policing devices. Systems will also include interfaces to quality monitoring and content monitoring with the ability for session change on demand. These abilities are based on a real-time integrated system with monitoring and authentication mechanisms, online service provisioning, and support for the implementation of provider agreements and roaming policies.

Many IPv6 billing impacts will be associated with the need and demand for next generation services. These services and accounting mechanisms will be dependent upon the use of user profiles and payment/credit information. The billing function thus becomes much more complex than in an IPv4 environment, especially with the addition of converged and/or bundled services, multi-provider agreements and revenue sharing.

Mobile Internet, applications and content delivery as well as the use of peer-to-peer applications will be a driving factor behind the evolution from the current billing models used by ISPs to the complex IPv6 and next generation service ready models.

13.2 Location of Work

Some standardization work has been undertaken in the IETF on authentication, authorization and accounting (AAA). However, much of the work affecting IPv6 billing is greatly affected by the development of NGN and converged services which will be offered by both traditional ISPs, telcos and mobile service providers.

13.3 Agreement

Not applicable.

13.4 Proposed Action

It is generally recommended that this topic should be monitored and development of IPv6 ready billing and accounting models started in those forums where billing and accounting standardization work already takes place.

14 NETWORK RENUMBERING

14.1 Issue

Network renumbering may be required for a variety of reasons. Although network management and address allocation practices are being developed to deal with many scenarios, network renumbering introduces challenges, including those associated with the transition from IPv4 to IPv6.

14.1.1 Change of uplink prefix

One cause for renumbering is a change in the site's upstream provider. With IPv6, sites are highly unlikely to be able to obtain provider independent (PI) address space, as have in some cases been obtained in the past with IPv4. Rather, sites use provider assigned (PA) addressing. As a result, if a site changes provider, it must also change its IPv6 PA prefix. This change can be triggered by customer migration to new provider, intermittent connectivity to an upstream provider, provider migration and upstream renumbering or IPv6 transition.

14.1.2 Change of internal topology

A site may need to renumber all or part of its internal network due to a change of topology, such as creating more or less specific subnets, or acquiring a larger IPv6 address allocation. Motivations for splitting a link into separate subnets may be to meet security demands on a particular link (policy for link-based access control rules), or for link load management by shuffling popular services to more appropriate locations in the local topology. Link-merging may be due to restructuring within the hosting organization, for example, acquisition or merger, network growth or network mobility.

14.2 Location of Work

IETF:

- "Procedures for Renumbering an IPv6 Network without a Flag Day", IETF RFC 4192.[33]

- "Things to think about when Renumbering an IPv6 network", IETF draft-chown-v6ops-renumber-thinkabout-05, September 2006.[34]

[33] http://www.ietf.org/rfc/rfc4192.txt
[34] http://www.ietf.org/internet-drafts/draft-chown-v6ops-renumber-thinkabout-05.txt

14.3 Agreement

Not applicable.

14.4 Proposed Action

Additional work within the IETF is needed in the area of address architecture. Formal hierarchy was defined within the IPv6 addressing architecture, providing a separation between the global prefix and the site prefix and helping to facilitate the automated renumbering of an end site.

15 SEPARATION OF LOCATOR AND IDENTIFIER

15.1 Separating routing from addressing

15.1.1 Issue

The IETF has identified Internet routing and addressing as a serious issue that needs to be resolved. The current IPv6 addressing and routing architectures do not scale over the long term. Therefore, it needs to be urgently investigated.

15.1.2 Location of Work

The IETF. Current discussions on Locator/ID seperation is occuring on the Routing and Addressing Mailing list (RAM list) which is an out growth from the Routing and Addressing WorkShop (RAWS).[35]

15.1.3 Agreement

The solution will involve Service Providers and standards bodies in particular the IETF. ATIS should be involved.

15.1.4 Proposed Action

ATIS members should participate in and contribute to the IETF on this topic.

16 VENDOR AVAILABILITY

16.1 Issue

Many of the network routers, switches, servers and embedded system devices deployed today support dual-stack capabilities in order to "future proof" product lines. However, not all vendors support IPv6 in their full range of products and IPv6 forwarding is often times done using software manipulations versus hardware/software configurations. This method lends itself to potential performance problems and service degradation as

[35] For more infor on RAWS:
http://www.iab.org/about/workshops/routingandaddressing/index.html

this method does not scale. In addition, from a client/customer perspective, their existing operating systems may not support DHCPv6 which is required for IPv6.

The availability of devices/services as noted below also need to be considered:

- Software applications and services have not been ported on a large scale to IPv6
- Media entertainment applications (e.g. gaming)
- Peer-to-Peer file sharing
- Home network cable/xDSL modems and routers need to support the encapsulation of IPv6 within IPv4.

While concerns remain over product availability, progress is being made in identifying IPv6 capable products. The IPv6 Forum has implemented an IPv6 ready logo program to assist others in identifying IPv6 products and services. A list of IPv6-ready devices/services is available at http://www.ipv6ready.org/frames.html.

16.2 Location of Work

IETF is addressing some of these technical issues related to equipment availability. Additional work is required.

The IPv6 Forum IPv6 Ready Logo program is identifying IPv6 devices/services.

16.3 Agreement

Broad IPv6 deployment will not occur until vendors are able to resolve the needs for graceful restart, non-stop forwarding and service software upgrades to implement IPv6. These issues are the key requirements that are prerequisites to service providers. Most organizations would be willing to take a hit to obtain these features, but would not do the same thing to implement IPv6.

16.4 Proposed Action

Service providers and their suppliers (i.e., equipment, software and systems vendors) should jointly discuss IPv6 capabilities. The continued or increased delivery of IPv6 capable devices, services and/or applications will be based on market forces.

17 RELATIONSHIP TO OTHER NUMBERING SYSTEMS

The relationship of IPv6 to other numbering systems is important to assess in the context of converged next generation services. Unlike other existing numbering systems, IP addresses are not by and large given attributes via geographical borders, nor instrumented through national authorities. In the United States, traditional phone numbering administration takes place through the North American Numbering Council (NANC), a federal advisory committee to the FCC. Other issues with numbering in the North American Numbering Plan (NANP) area are worked to resolution at the ATIS

Industry Numbering Committee (INC). Also inherently related to IPv6 numbering is the use of ENUM.

At this time it does not appear that the NANC or ATIS INC are undertaking any consideration or work on the relationship of the E.164 numbering system to IPv6. The IETF, ENUM Forum, and ENUM LLC have recognized IPv6 as a related system and have included its consideration in their work. The point of intersection with ENUM is of primary importance as the state of telephone numbers seems to be more akin to acting as a domain name, rather than a location and equipment identifier. The portability of numbers and permutation of VoIP alternatives to POTS will drive the need for ENUM/IPv6 interoperation.

17.1 ENUM

ENUM facilitates the integration of telephone numbers (15 digits) to IP addresses and has largely been deployed as a VoIP enabler via RFC2916 and ITU E.164. The usefulness of the relationship between ENUM and IPv6 will depend upon the continued development and availability of IPv6 and ENUM capable equipment and ENUM aware applications, the coordination of national ENUM authorities via the ITU-T (such as the case for SIP-ENUM) and coordination of protocols for DNS (e.g. DNSSEC.)
Work on the adjustments to address the requirements of IPv6 addressing for ENUM is done through the IETF DNS Extensions (DNSEXT) WG.

17.1.1 Location of Work

IETF DNS Extensions (DNSEXT) WG.

17.1.2 Agreement

Not applicable.

17.1.3 Proposed Action

It is generally recommended that this topic should be monitored and the awareness of IPv6 to affected numbering systems be raised as necessary in those forums where numbering related work takes place.

18 COST

18.1 Issue

As presented in ATIS' first report and recommendation with respect to IPv6, "ATIS Internet Protocol version 6 (IPv6) Report & Recommendation, May 2006," transition costs have many variables and each operator will have different ratios of costs. Each service provider must assess transitioning to IPv6 based on their own internal economic and operational benefits -- e.g., their return on investment (ROI).

Notwithstanding the above however, it is generally agreed that the cost of maintaining an existing IPv4 infrastructure, NAT devices, ALGs, and Proxies over time will increasingly outweigh existing benefits.

18.2 Location of Work

Several reports have been issued that assess anticipated costs associated with IPv6.

U.S. Department of Commerce (NIST/NTIA):

- "IPv6 Economic Impact Assessment: Final Report", Planning Report 05-2, National Institute of Standards and Technology, U.S. Department of Commerce, October 2005. (http://www.nist.gov/director/prog-ofc/report05-2.pdf)

- "Technical and Economic Assessment of Internet Protocol Version 6 (IPv6)", National Institute of Standards and Technology, National Telecommunications and Information Administration, U.S. Department of Commerce, Chapter 2.2 – Stakeholder Costs of Adopting IPv6, pp. 23-31, January 2006. (http://www.nist.gov/director/prog-ofc/IPv6-final.pdf)

18.3 Agreement

It is generally agreed that with the exception of achieving economy of scale by ensuring interoperability through standardization activities, costs associated with the transitioning to IPv6 is beyond the scope of this exercise.

18.4 Proposed Action

None.

APPENDIX A: REFERENCES

The following references were identified during the course of this work effort. As it is known that specifications, standards and documents are revised, updated and/or deleted over time, readers are cautioned that the following specific references may have changed. Accordingly, readers are encouraged to refer to the specific groups and their website for the most recent documents.

A1 QoS

References are available at
http://njc240srvr02.ugd.att.com/home1/gash/public/ipv6.qos/

A2 Security - RFC list of IPv6 Security issues and threats:

Some security drafts that should be reviewed by an organization deploying IPv6:

- http://www.ietf.org/internet-drafts/draft-ietf-v6ops-security-overview-04.txt
- http://www.ietf.org/internet-drafts/draft-ietf-v6ops-icmpv6-filtering-recs-02.txt
- http://www.ietf.org/internet-drafts/draft-ietf-v6ops-natpt-to-exprmntl-03.txt

IETF Reference Documents:

Recommendations for IPv6 in 3GPP Standards (RFC 3314) (48168 bytes)
Default Address Selection for Internet Protocol version 6 (IPv6) (RFC 3484) (55076 bytes)
Basic Socket Interface Extensions for IPv6 (RFC 3493) (82570 bytes) obsoletes RFC 2553
A Flexible Method for Managing the Assignment of Bits of an IPv6 Address Block (RFC 3531) (11314 bytes)
IPv6 for Some Second and Third Generation Cellular Hosts (RFC 3316) (48741 bytes)
Advanced Sockets Application Program Interface (API) for IPv6 (RFC 3542) (173028 bytes) obsoletes RFC 2292
IPv6 Global Unicast Address Format (RFC 3587) (8783 bytes) obsoletes RFC 2374
IPv6 Flow Label Specification (RFC 3697) (21296 bytes)
Optimistic Duplicate Address Detection (DAD) for IPv6 (RFC 4429) (33123 bytes)
Requirements for IPv6 prefix delegation (RFC 3769) (10287 bytes)
Deprecating Site Local Addresses (RFC 3879) (24142 bytes)
Management Information Base for the Transmission Control Protocol (TCP) (RFC 4022) (47360 bytes) obsoletes RFC 2012,RFC 2452
IPv6 Scoped Address Architecture (RFC 4007) (53555 bytes)
IP Tunnel MIB (RFC 4087) (53206 bytes) obsoletes RFC 2667
Management Information Base for the User Datagram Protocol (UDP) (RFC 4113) (40323 bytes) obsoletes RFC 2013,RFC 2454
Unique Local IPv6 Unicast Addresses (RFC 4193) (35908 bytes)
Default Router Preferences and More-Specific Routes (RFC 4191) (34672 bytes)
IPv6 Host-to-Router Load Sharing (RFC 4311) (10156 bytes) updates RFC 2461
IP Version 6 Addressing Architecture (RFC 4291) (52897 bytes) obsoletes RFC 3513
Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification (RFC 4443) (48969 bytes)

IPv6 Node Requirements (RFC 4294) (39125 bytes)
Management Information Base for the Internet Protocol (IP) (RFC 4293) (242243 bytes) obsoletes RFC 2011,RFC 2465,RFC 2466
IP Forwarding Table MIB (RFC 4292) (69321 bytes) obsoletes RFC 2096
Neighbor Discovery Proxies (ND Proxy) (RFC 4389) (38124 bytes)
A Method for Generating Link Scoped IPv6 Multicast Addresses (RFC 4489) (12224 bytes) updates RFC 3306
IPv6 Node Information Queries (RFC 4620) (31134 bytes)

A5 Mobility for IPv6 (mip6)

Mobile IPv6 (MIPv6) specifies routing support to permit an IPv6 host to continue using its "permanent" home address as it moves around the Internet. Mobile IPv6 supports transparency above the IP layer, including maintenance of active TCP connections and UDP port bindings. The primary goal of the MIP6 working group will be to enhance base IPv6 mobility by continuing work on developments that are required for wide-scale deployments. Additionally the working group will ensure that any issues identified by the interop testing of the MIPv6 specifications are addressed quickly.

Internet-Drafts:

Mobility management for Dual stack mobile nodes A Problem Statement (16271 bytes)
Why Authentication Data suboption is needed for MIP6 (37182 bytes)
IP Address Location Privacy and Mobile IPv6: Problem Statement (19611 bytes)
Mobile IPv6 support for dual stack Hosts and Routers (DSMIPv6) (60457 bytes)
MIP6-bootstrapping via DHCPv6 for the Integrated Scenario (40833 bytes)
Mobility Header Home Agent Switch Message (21629 bytes)
Home Agent Reliability Protocol (93477 bytes)
DHCP Option for Home Information Discovery in MIPv6 (24194 bytes)

Request For Comments:

Using IPsec to Protect Mobile IPv6 Signaling between Mobile Nodes and Home Agents (RFC 3776) (87076 bytes)
Mobility Support in IPv6 (RFC 3775) (393514 bytes)
Mobile Node Identifier Option for Mobile IPv6 (MIPv6) (RFC 4283) (14653 bytes)
Mobile IP version 6 Route Optimization Security Design Background (RFC 4225) (98584 bytes)
Authentication Protocol for Mobile IPv6 (RFC 4285) (40874 bytes)
Mobile IPv6 Management Information Base (RFC 4295) (209038 bytes)
Mobile IPv6 and Firewalls: Problem Statement (RFC 4487) (32022 bytes)
Securing Mobile IPv6 Route Optimization Using a Static Shared Key (RFC 4449) (15080 bytes)
Extension to Sockets API for Mobile IPv6 (RFC 4584) (53995 bytes)

A6 Mobility for IP: Performance, Signaling and Handoff Optimization (mipshop)

Mobile IPv6 enables IPv6 mobile nodes to continue using a given "home address" in spite of changes in its point of attachment to the network. These changes may cause

packet loss, and also represent overhead traffic on the network. Hierarchical Mobile IPv6 (HMIPv6, RFC 4140) reduces the amount and latency of signaling between a MN, its Home Agent and one or more correspondent nodes. Fast Handovers for Mobile IPv6 (FMIPv6, RFC 4068) reduces packet loss by providing fast IP connectivity as soon as the mobile node establishes a new point of attachment at a new link.

Internet-Drafts:
Mobile IPv6 Fast Handovers over IEEE 802.16e Networks (35857 bytes)
Fast Handovers for Mobile IPv6 (94466 bytes)
Mobile IPv6 Fast Handovers for 3G CDMA Networks (72329 bytes)
Fast Handovers for Mobile IPv6 (94430 bytes)
Applying Cryptographically Generated Addresses and Credit-Based Authorization to Mobile IPv6 (75926 bytes)

Request For Comments:
Fast Handovers for Mobile IPv6 (RFC 4068) (93591 bytes)
Hierarchical Mobile IPv6 mobility management (HMIPv6) (RFC 4140) (71503 bytes)
Mobile IPv6 Fast Handovers for 802.11 Networks (RFC 4260) (35277 bytes)

A7 Mobile Nodes and Multiple Interfaces in IPv6 (monami6)

There is currently rapid development in the area of new wireless standards (802.11*, 802.16, 802.20, UMTS, Bluetooth and others). At the same time, terminals which have radio and protocol support for two, three or even more standards are appearing. This opens the possibility of using multiple access types simultaneously, with each access used to transport the traffic for which it is most appropriate. The objective of the Monami6 WG is to produce a clear problem statement and to produce standard track specifications to the straightforward problems associated with the simultaneous use of multiple addresses for either mobile hosts using Mobile IPv6 or mobile routers using NEMO Basic Support and their variants (FMIPv6, HMIPv6, etc). Where the effects of having multiple prefixes on a single interface is identical to the effects of having multiple interfaces each with a single prefix, Monami6 will consider a generalized approach to cater for multiple prefixes available to a mobile host/router.

Internet-Drafts:
Analysis of Multihoming in Mobile IPv6 (62282 bytes)
Multiple Care-of Addresses Registration (66005 bytes)

A8 Site Multihoming by IPv6 Intermediation (shim6)

The background documents to be considered by the WG include:
RFC 3582
draft-ietf-multi6-architecture-04.txt
draft-ietf-multi6-things-to-think-about-01.txt
draft-ietf-multi6-multihoming-threats-03.txt

The input documents that the WG will use as the basis for its design are:

draft-huston-l3shim-arch-00.txt
draft-ietf-multi6-functional-dec-00.txt
draft-ietf-multi6-l3shim-00.txt
draft-ietf-multi6-failure-detection-00.txt
draft-ietf-multi6-hba-00.txt
draft-ietf-multi6-app-refer-00.txt

Internet-Drafts:
Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming (80672 bytes)
Applicability Statement for the Level 3 Multihoming Shim Protocol (Shim6) (38079 bytes)
Hash Based Addresses (HBA) (52444 bytes)
Level 3 multihoming shim protocol (282751 bytes)
Default Locator-pair selection algorithm for the SHIM6 protocol (21214 bytes)

A9 Site Multihoming in IPv6 (multi6)

For the purposes of redundancy, load sharing, operational policy or cost, a site may be multi-homed, with the site's network having connections to multiple IP service providers. The current Internet routing infrastructure permits multi-homing using provider independent addressing, and adapts to changes in the availability of these connections. However if the site uses multiple provider-assigned address prefixes for every host within the site, host application associations cannot use alternate paths, such as for surviving the changes or for creating new associations, when one or more of the site's address prefixes becomes unreachable. This working group will produce specifications for an IPv6-based site multi-homing solution that inserts a new sub-layer (shim) into the IP stack of end-system hosts. It will enable hosts on multi-homed sites to use a set of provider-assigned IP address prefixes and switch between them without upsetting transport protocols or applications.

Request For Comments:
Goals for IPv6 Site-Multihoming Architectures (RFC 3582) (17045 bytes)
IPv4 Multihoming Practices and Limitations (RFC 4116) (26598 bytes)
Architectural Approaches to Multi-Homing for IPv6 (RFC 4177) (95374 bytes)
Threats relating to IPv6 Multihoming Solutions (RFC 4218) (75969 bytes)
Things Multihoming in IPv6 (MULTI6) Developers Should Think About (RFC 4219) (25141 bytes)

A10 IPv6 Operations (v6ops)

The global deployment of IPv6 is underway, creating an IPv4/IPv6 Internet consisting of IPv4-only, IPv6-only and IPv4/IPv6 networks and nodes. This deployment must be properly handled to avoid the division of the Internet into separate IPv4 and IPv6 networks while ensuring addressing and connectivity for all IPv4 and IPv6 nodes.

The IPv6 Operations Working Group (v6ops) develops guidelines for the operation of a

shared IPv4/IPv6 Internet and provides operational guidance on how to deploy IPv6 into existing IPv4-only networks, as well as into new network installations.

The main focus of the v6ops WG is to look at the immediate deployment issues; more advanced stages of deployment and transition are a lower priority.

Request For Comments:
Transition Scenarios for 3GPP Networks (RFC 3574) (23359 bytes)
Unmanaged Networks IPv6 Transition Scenarios (RFC 3750) (48153 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Sub-IP Area Standards (RFC 3793) (11624 bytes)
Introduction to the Survey of IPv4 Addresses in Currently Deployed IETF Standards (RFC 3789) (22842 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Internet Area Standards (RFC 3790) (102694 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Routing Area Standards (RFC 3791) (27567 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Security Area Standards (RFC 3792) (46398 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Transport Area Standards (RFC 3794) (60001 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Application Area Standards (RFC 3795) (92584 bytes)
Survey of IPv4 Addresses in Currently Deployed IETF Operations & Management Area Standards (RFC 3796) (78400 bytes)
Evaluation of Transition Mechanisms for Unmanaged Networks (RFC 3904) (46844 bytes)
Security Considerations for 6to4 (RFC 3964) (83360 bytes)
Application Aspects of IPv6 Transition (RFC 4038) (69727 bytes)
Scenarios and Analysis for Introducing IPv6 into ISP Networks (RFC 4029) (64388 bytes)
IPv6 Enterprise Network Scenarios (RFC 4057) (33454 bytes)
Procedures for Renumbering an IPv6 Network without a Flag Day (RFC 4192) (52110 bytes) updates RFC 2072
Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks (RFC 4215) (52903 bytes)
Basic Transition Mechanisms for IPv6 Hosts and Routers (RFC 4213) (58575 bytes) obsoletes RFC 2893
Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks (RFC 4554) (23355 bytes)

A11 Location of Work re v6–v4 translation

NAT-PT work is being done in the IETF v6ops working group (http://www.ietf.org/html.charters/v6ops-charter.html); the v6ops working group draft that proposes to deprecate NAT-PT to experimental status is http://www.ietf.org/internet-drafts/draft-ietf-v6ops-natpt-to-exprmntl-03.txt; a slide

summary of this work is at http://www3.ietf.org/proceedings/04nov/slides/v6ops-6.pdf

APPENDIX B: TASK FORCE MEMBERS

Name	Company
Convener	
Tim Jeffries	ATIS
Members	
Joe Houle	AT&T
Steven Wright	AT&T (BellSouth, Inc.)
Wayne Zeuch	Alcatel-Lucent
Anik Sane	Bell Canada
Asok Chatterjee	Ericsson, Inc.
Dwight Jamieson	Nortel, Inc.
Mike Fargano	Qwest, Inc.
Mark Desterdick	Verizon, Inc.
ATIS Support	
Martha Ciske	ATIS