

Network Working Group  
Request for Comments: 4852  
Category: Informational

J. Bound  
Y. Pouffary  
Hewlett-Packard  
S. Klynsma  
MITRE  
T. Chown  
University of Southampton  
D. Green  
Command Information  
April 2007

## IPv6 Enterprise Network Analysis - IP Layer 3 Focus

### Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

### Copyright Notice

Copyright (C) The IETF Trust (2007).

### Abstract

This document analyzes the transition to IPv6 in enterprise networks focusing on IP Layer 3. These networks are characterized as having multiple internal links and one or more router connections to one or more Providers, and as being managed by a network operations entity. The analysis focuses on a base set of transition notational networks and requirements expanded from a previous document on enterprise scenarios. Discussion is provided on a focused set of transition analysis required for the enterprise to transition to IPv6, assuming a Dual-IP layer (IPv4 and IPv6) network and node environment within the enterprise. Then, a set of transition mechanisms are recommended for each notational network.

## Table of Contents

1. Introduction .....	3
2. Terminology .....	5
3. Enterprise Matrix Analysis for Transition .....	5
4. Wide-Scale Dual-Stack Deployment Analysis .....	10
4.1. Staged Dual-Stack Deployment .....	10
4.2. Routing Capability Analysis for Dual-IP Deployment .....	11
4.2.1. IPv6 Routing Capability .....	11
4.2.2. IPv6 Routing Non-Capability .....	11
4.2.2.1. Tunnel IPv6 over the IPv4 infrastructure ..	12
4.2.2.2. Deploy a Parallel IPv6 Infrastructure .....	12
4.3. Remote IPv6 Access to the Enterprise .....	12
4.4. Other Considerations .....	13
5. Sparse Dual-Stack Deployment Analysis .....	13
5.1. Internal versus External Tunnel Endpoint .....	13
5.2. Manual versus Autoconfigured .....	14
6. IPv6-Dominant Network Deployment Analysis .....	14
7. General Issues from Analysis .....	15
7.1. Staged Plan for IPv6 Deployment .....	15
7.2. Network Infrastructure Requirements .....	15
7.3. Stage 1: Initial Connectivity Steps .....	15
7.3.1. Obtaining External Connectivity .....	16
7.3.2. Obtaining Global IPv6 Address Space .....	16
7.4. Stage 2: Deploying Generic Basic Service Components .....	16
7.4.1. Developing an IPv6 Addressing Plan .....	16
7.4.2. IPv6 DNS .....	17
7.4.3. IPv6 Routing .....	17
7.4.4. Configuration of Hosts .....	18
7.4.5. Security .....	18
7.5. Stage 3: Widespread Dual-Stack Deployment On-Site .....	19
8. Applicable Transition Mechanisms .....	20
8.1. Recognizing Incompatible Network Touchpoints .....	20
8.2. Recognizing Application Incompatibilities .....	21
8.3. Using Multiple Mechanisms to Support IPv6 Transition .....	22
9. Security Considerations .....	22
10. References .....	22
10.1. Normative References .....	22
10.2. Informative References .....	24
11. Acknowledgments .....	25
Appendix A. Crisis Management Network Scenarios .....	26
A.1. Introduction .....	26
A.2. Scenarios for IPv6 Deployment in Crisis Management Networks .....	26
A.3. Description of a Generic Crisis Management Network .....	28
A.4. Stages of IPv6 Deployment .....	29

## 1. Introduction

This document analyzes the transition to IPv6 in enterprise networks focusing on IP Layer 3. These networks are characterized as having multiple internal links, and one or more router connections to one or more Providers, and as being managed by a network operations entity. The analysis focuses on a base set of transition notational networks and requirements expanded from a previous document on enterprise scenarios. Discussion is provided on a focused set of transition analysis required for the enterprise to transition to IPv6, assuming a Dual-IP layer (IPv4 and IPv6) network and node environment within the enterprise. Then, a set of transition mechanisms are recommended for each notational network.

The audience for this document is the enterprise network team considering deployment of IPv6. The document will be useful for enterprise teams that have to determine the IPv6 transition strategy for their enterprise. It is expected that those teams include members from management, network operations, and engineering. The analysis and notational networks presented provide an example set of cases the enterprise can use to build an IPv6 transition strategy.

The enterprise analysis begins by describing a matrix as a tool to be used to portray the different IPv4 and IPv6 possibilities for deployment. The document will then provide analysis to support enterprise-wide Dual-IP layer deployment strategy, to provide the reader with a view of how that can be planned and what options are available. The document then discusses the deployment of sparse IPv6 nodes within the enterprise and the requirements that need to be considered and implemented when the enterprise remains with an IPv4-only routing infrastructure for some time. The next discussion focuses on the use of IPv6 when it is determined to be dominant across or within parts of the enterprise network.

The document then discusses the general issues and applicability from the previous analysis. The document concludes by providing a set of current transition mechanism recommendations for the notational network scenarios to support an enterprise that is planning to deploy IPv6.

As stated, this document focuses only on the deployment cases where a Dual-IP Layer 3 is supported across the network and on the nodes in the enterprise. Additional deployment transition analysis will be required from the effects of an IPv6-only node or Provider deployments, and is beyond the scope of this document. In addition, this document does not attempt to define or discuss any use with network address translation [NATPT] or Provider Independent address space.

The following specific topics are currently out of scope for this document:

- Multihoming
- Application transition/porting (see [APPS]).
- IPv6 VPN, firewall, or intrusion detection deployment.
- IPv6 network management and QoS deployment.
- Detailed IT Department requirements.
- Deployment of novel IPv6 services, e.g., Mobile IPv6.
- Requirements or Transition at the Providers' network.
- Transport protocol selection for applications with IPv6.
- Application layer and configuration issues.
- IPv6 only future deployment scenarios.

This document focuses on IP Layer 3 deployment in the same way as the other IPv6 deployment analysis works have done [UMAN] [ISPA] [3GPA]. This document covers deployment of IPv6 "on the wire", including address management and DNS services.

We are also assuming that the enterprise deployment is being undertaken by the network administration team, i.e., this document does not discuss the case of an individual user gaining IPv6 connectivity (to some external IPv6 provider) from within an enterprise network. Much of the analysis is applicable to wireless networks, but there are additional considerations for wireless networks not contained within this document.

In Section 2, we introduce the terminology used in this document. In Section 3, we introduce and define a tools matrix and define the IP Layer 3 connectivity requirements. In Section 4, we discuss wide scale Dual-IP layer use within an enterprise. In Section 5, we discuss sparse Dual-IP layer deployment within an enterprise. In Section 6, we discuss IPv6-dominant network deployment within the enterprise. In Section 7, we discuss general issues and applicability. In Section 8, a set of transition mechanisms that can support the deployment of IPv6 with an enterprise are recommended.

This document then provides Appendix A for readers depicting a Crisis Management enterprise network to demonstrate an enterprise network example that requires all the properties as analyzed in Sections 3, 4, 5, 6, and 7. In addition, we recommend that readers of this document also read another use-case document to support an IPv6 Transition for a Campus Network [CAMP].

Readers should also be aware that a parallel effort for an enterprise to transition to IPv6 is training, but out of scope for this document.

## 2. Terminology

- Enterprise Network - A network that has multiple internal links, and one or more router connections to one or more Providers, and is actively managed by a network operations entity.
- Provider - An entity that provides services and connectivity to the Internet or other private external networks for the enterprise network.
- IPv6-capable - A node or network capable of supporting both IPv6 and IPv4.
- IPv4-only - A node or network capable of supporting only IPv4.
- IPv6-only - A node or network capable of supporting only IPv6. This does not imply an IPv6 only stack in this document.
- Dual-IP - A network or node that supports both IPv4 and IPv6.
- IP-capability - The ability to support IPv6 only, IPv4 only, or Dual-IP Layer
- IPv6-dominant - A network running IPv6 routing and control plane services that provides transport for both IPv4 and IPv6 protocol services
- Transition Implementation - The network strategy the enterprise uses to transition to IPv6.

## 3. Enterprise Matrix Analysis for Transition

In order to identify the best-suited transition mechanisms for an enterprise, it is recommended that the enterprise have an in-depth up-to-date understanding of its current IT environment. This understanding will help choose the best-suited transition mechanisms. It is important to note that one size does not fit all. Selection of mechanisms that reduce the impact on the existing environment is suggested. When selecting a transition mechanism, one must consider the functionality required, its scalability characteristic, and the security implications of each mechanism.

To provide context for an analysis of the transitioning enterprise at Layer 3, we have provided a matrix that describes various scenarios

which might be encountered during an IPv6 transition. The notional enterprise network is comprised of hosts attached to an enterprise-owned intranet(s) at two different global locations separated by the Internet. The enterprise owns, operates, and maintains its own intranetworks, but relies on an external provider organization that offers Internet Service. Both local and destination intranetworks are operated by different organizations within the same enterprise and consequently could have different IP-capability than other intranetworks at certain times in the transition period.

Addressing every possible combination of network IP-capability in this notional enterprise network is impractical; therefore, trivial notional networks (i.e., pure IPv4, pure IPv6, and ubiquitous Dual-IP) are not considered. In addition, the authors could not conceive of any scenarios involving IPv6-only ISPs or IPv6-only nodes in the near term and consequently have not addressed scenarios with IPv6-only ISPs or IPv6-only nodes. We assume all nodes that host IPv6 applications are Dual-IP. The matrix does not assume or suggest that network address translation is used. The authors recommend that network address translation not be used in these notional cases.

Future enterprise transitions that support IPv6-only nodes and IPv6-only ISPs will require separate analysis, which is beyond the scope of this document.

Table 1 below is a matrix of ten possible Transition Implementations that, being encountered in an enterprise, may require analysis and the selection of an IPv6 transition mechanism for that notional network. Each possible implementation is represented by the rows of the matrix. The matrix describes a set of notional networks as follows:

- The first column represents the protocol used by the application and, below, the IP-capability of the node originating the IP packets.  
(Application/Host 1 OS)
- The second column represents the IP-capability of the host network wherein the node originated the packet.  
(Host 1 Network)
- The third column represents the IP-capability of the service provider network.  
(Service Provider)

- The fourth column represents the IP-capability of the destination network wherein the originating IP packets are received.  
(Host 2 Network)
- The fifth column represents the protocol used by the application and, below, the IP-capability of the destination node receiving the originating IP packets.  
(Application/Host 2 OS)

As an example, notional network 1 is an IPv6 application residing on a Dual-IP layer host trying to establish a communications exchange with a destination IPv6 application. To complete the information exchange, the packets must first traverse the host's originating IPv4 network (intranet), then the service provider's and destination host's Dual-IP network.

Obviously, Table 1 does not describe every possible scenario. Trivial notional networks (such as pure IPv4, pure IPv6, and ubiquitous Dual-IP) are not addressed. However, the authors feel these ten scenarios represent the vast majority of transitional situations likely to be encountered in today's enterprise. Therefore, we will use these ten to address the analysis for enterprise deployment.

Table 1 - Enterprise Scenario Deployment Matrix

	Application ----- Host 1 OS	Host 1 Network	Service Provider	Host 2 Network	Application ----- Host 2 OS
A	IPv6 ---- Dual IP	IPv4	Dual IP or IPv4	Dual IP	IPv6 ---- Dual IP
B	IPv6 ---- Dual IP	IPv6	IPv4	IPv4	IPv6 ---- Dual IP
C	IPv4 ---- Dual IP	IPv4	Dual IP	IPv6	IPv4 ---- Dual IP
D	IPv4 ---- Dual IP	Dual IP or IPv6	IPv4	IPv6	IPv4 ---- Dual IP
E	IPv6 ---- Dual IP	Dual IP or IPv6	Dual IP	Dual IP or IPv6	IPv4 ---- Dual IP
F	IPv6 ---- Dual IP	IPv6	IPv4	IPv4	IPv4 ---- Dual IP
G	IPv4 ---- Dual IP	IPv6	Dual IP	IPv6	IPv6 ---- Dual IP
H	IPv4 ---- IPv4	IPv6	Dual IP	IPv4	IPv6 ---- Dual IP
I	IPv4 ---- IPv4	IPv6	IPv4	IPv6	IPv6 ---- Dual IP
J	IPv6 ---- Dual IP	IPv4	IPv4	IPv6	IPv4 ---- Dual IP



The reader should note that Scenarios A-C in Table 1 are variations of compatible hosts communicating across largely (but not entirely) homogenous networks. In each of the first three scenarios, the packet must traverse at least one incompatible network component. For example, Scenario B represents an enterprise that wishes to use IPv6 applications, but has yet to transition its internal networks; its Service Provider also lags, offering only a v4 IP-service. Conversely, Scenario C represents an enterprise that has completed transition to IPv6 in its core networks (as has its Service Provider), but continues to require a legacy IPv4-based application.

Scenario D represents the unusual situation where the enterprise has transitioned its core intranetworks to IPv6, but (like Scenario B) it's ISP provider has yet to transition. In addition, this enterprise continues to retain critical legacy IPv4-based applications that must communicate over this heterogeneous network environment.

Scenarios E-J represent transitional situations wherein the enterprise has both IPv4 and IPv6 based instantiations of the same application that must continue to interoperate. In addition, these scenarios show that the enterprise has not completed transition to IPv6 in all its organic and/or Service Provider networks. Instead, it maintains a variety of heterogeneous network segments between the communicating applications. Scenarios E and J represent distinctly different extremes on either end of the spectrum. In Scenario E, the enterprise has largely transitioned to IPv6 in both its applications and networks. However, Scenario E shows that a few legacy IPv4-based applications may still be found in the enterprise. On the other hand, Scenario J shows an enterprise that has begun its transition in a very disjointed manner and, in which IPv6-based applications and network segments are relatively rare.

#### 4. Wide-Scale Dual-Stack Deployment Analysis

In this section, we address Scenario 1 as described in Section 3.1 of [BSCN]. The scenario, assumptions, and requirements are driven from the [BSCN] text. This analysis further corresponds to Scenario A in Section 3 above (although Scenario A shows a transitional situation wherein the enterprise has one network segment still lagging on transition to Dual-IP).

Within these IPv6 deployment scenarios the enterprise network administrator would introduce IPv6 by enabling IPv6 on the wire (i.e., within the network infrastructure) in a structured fashion with the existing IPv4 infrastructure. In such scenarios, a number of the existing IPv4 routers (and thus subnets) will be made Dual-IP, such that communications can run over either protocol.

Nodes on the Dual-IP links may themselves be IPv4-only or IPv6-capable. The driver for deploying IPv6 on the wire may not be for immediate wide-scale usage of IPv6, but rather to prepare an existing IPv4 infrastructure to support IPv6-capable nodes. Thus, while IPv6 is not used, Dual-IP nodes exist, and the enterprise can be transitioned to IPv6 on demand.

Analyzing this scenario against existing transition mechanisms for their applicability suggests a staged approach for IPv6 deployment in the enterprise.

##### 4.1. Staged Dual-Stack Deployment

Under these scenarios (as well as most others), the site administrator should formulate a staged plan for the introduction of a Dual-IP IPv6 network. We suggest that Section 7 of this document provides a good basis for such a plan.

In an enterprise network, the administrator will generally seek to deploy IPv6 in a structured, controlled manner, such that IPv6 can be enabled on specific links at various stages of deployment. There may be a requirement that some links remain IPv4 only, or some that specifically should not have IPv6 connectivity (e.g., Scenario A of Table 1). There may also be a requirement that aggregatable global IPv6 addresses, assigned by the enterprise's upstream provider from the address space allocated to them by the Regional Internet Registries (RIRs), be assigned.

In this document, we do not discuss the deployment of Unique Local IPv6 Unicast Addresses [ULA] because the address type and scope selected is orthogonal to the Layer 3 analysis of this document.

A typical deployment would initially involve the establishment of a single "testbed" Dual-IP subnet at the enterprise site prior to wider deployment. Such a testbed not only allows the IPv6 capability of specific platforms and applications to be evaluated and verified, but also permits the steps in Sections 7.3 and 7.4 of this document to be undertaken without (potential) adverse impact on the production elements of the enterprise.

Section 7.5 describes the stages for the widespread deployment in the enterprise, which could be undertaken after the basic building blocks for IPv6 deployment are in place.

#### 4.2. Routing Capability Analysis for Dual-IP Deployment

A critical part of Dual-IP deployment is the selection of the IPv6-capable routing infrastructure to be implemented. The path taken will depend on whether the enterprise has existing Layer 2/3 switch/router equipment that has an IPv6 (routing) capability, or that can be upgraded to have such capability.

In Section 4, we are not considering sparse IPv6 deployment; the goal of Dual-IP deployment is widespread use in the enterprise.

##### 4.2.1. IPv6 Routing Capability

Where IPv6 routing capability exists within the infrastructure, the network administrator can enable IPv6 on the same physical hardware as the existing IPv4 service. Enabling both is the end-goal of any enterprise to support Dual-IP deployment, when the capability, performance, and robustness of the Dual-IP operational deployment has been verified.

Ideally, the IPv6 capability will span the entire enterprise, allowing deployment on any link or subnet. If not, techniques from Section 4.4 may be required.

##### 4.2.2. IPv6 Routing Non-Capability

If the enterprise cannot provide IPv6 routing initially, there are alternative methods for transition. In this case, the enterprise administrator faces two basic choices, either to tunnel IPv6 over some or all of the existing IPv4 infrastructure, or to deploy a parallel IPv6 routing infrastructure providing IPv6 connectivity into existing IPv4 subnets.

It may thus be the case that a node's IPv4 and IPv6 default routes to reach other links (subnets) are through different routing platforms.

#### 4.2.2.1. Tunnel IPv6 over the IPv4 infrastructure

Consider the situation where there exists IPv6 edge routers that are IPv6-capable, while others, and perhaps the enterprise backbone itself, are not IPv6-capable (Scenario B of Table 1). Tunneling, as described in [BCNF], would be established between the Dual-IP capable routers on the enterprise, thus "bypassing" existing non IPv6-capable routers and platforms.

In the widespread Dual-IP scenario, a more structured, manageable method is required, where the administrator has control of the deployment per-link and (ideally) long-term, aggregatable global IPv6 addressing is obtained, planned, and used from the outset.

#### 4.2.2.2. Deploy a Parallel IPv6 Infrastructure

Alternatively, the administrator may deploy a new, separate IPv6-capable router (or set of routers). It is quite possible that such a parallel infrastructure would be IPv6-dominant.

Such an approach would likely require additional hardware, but it has the advantage that the existing IPv4 routing platforms are not disturbed by the introduction of IPv6.

To distribute IPv6 to existing IPv4 enterprise subnets, either dedicated physical infrastructure can be employed or, if available, IEEE 802.1q VLANs could be used, as described in [VLAN]. The latter has the significant advantage of not requiring any additional physical cabling/wiring and also offers all the advantages of VLANs for the new Dual-IP environment. Many router platforms can tag multiple VLAN IDs on a single physical interface based on the subnet/link the packet is destined for; thus, multiple IPv6 links can be collapsed for delivery on a single (or small number of) physical IPv6 router interface(s) in the early stages of deployment.

The parallel infrastructure should only be seen as an interim step towards full Dual-IP deployment on a unified infrastructure. The parallel infrastructure however allows all other aspects of the IPv6 enterprise services to be deployed, including IPv6 addressing, thus making the enterprise ready for that unifying step at a later date.

#### 4.3. Remote IPv6 Access to the Enterprise

When the enterprise's users are off-site, and using an ISP that does not support any native IPv6 service or IPv6 transition aids, the enterprise may consider deploying its own remote IPv6 access support. Such remote support might for example be offered by deployment of an IPv6 Tunnel Broker [TBRK].

#### 4.4. Other Considerations

There are some issues associated with turning IPv6 on by default, including application connection delays, poor connectivity, and network insecurity, as discussed in [V6DEF]. The issues can be worked around or mitigated by following the advice in [V6DEF].

#### 5. Sparse Dual-Stack Deployment Analysis

This section covers Scenario 2 as described in Section 3.1 of [BSCN]. This scenario assumes the requirements defined within the [BSCN] text.

IPv6 deployment within the enterprise network, with an existing IPv4 infrastructure, could be motivated by mission-critical or business applications or services that require IPv6. In this case, the prerequisite is that only the nodes using those IPv6 applications need to be upgraded to be IPv6-capable. The routing infrastructure will not be upgraded to support IPv6, nor does the enterprise wish to deploy a parallel IPv6 routing infrastructure at this point, since this is an option in Section 4.

There is a need for end-to-end communication with IPv6, but the infrastructure only supports IPv4 routing. Thus, the only viable method for end-to-end communication with IPv6 is to tunnel the traffic over the existing IPv4 infrastructure as defined in this analysis document.

The network team needs to decide which of the available transition tunneling mechanisms are the most efficient to deploy, so they can be used without disrupting the existing IPv4 infrastructure. Several conditions require analysis, as introduced in the following subsections.

##### 5.1. Internal versus External Tunnel Endpoint

Let's assume the upstream provider has deployed some IPv6 services, either native IPv6 in its backbone or in the access network, or some combination of both (Scenario B of Table 1). In this case, the provider will likely also deploy one or more transition mechanisms to support their IPv6 subscribers. Obviously, the enterprise could decide to take advantage of those transition services offered from the Provider. However, this will usually mean that individual nodes in the network require their own IPv6-in-IPv4 tunnel. The end result is somewhat inefficient IPv6 intranetworks communication, because all IPv6 traffic must be forwarded by the enterprise's IPv4 infrastructure to the Tunnel Endpoint offered by the Provider. Nevertheless, this may be acceptable, particularly if the IPv6

applications do not require intranetworks communication at all -- for example, when an application's server is located outside of the enterprise network, or on other intranetworks of the same enterprise.

Alternatively, the enterprise could decide to deploy its own transition mechanism node, possibly collocating it adjacent to the border router that connects to the upstream Provider. In this case, intranetworks communication using this tunnel endpoint is also possible.

## 5.2. Manual versus Autoconfigured

If the number of nodes to be using IPv6 is low, the first option is to use statically configured tunnels. However, automatically configured tunnels may be preferable, especially if the number is higher.

## 6. IPv6-Dominant Network Deployment Analysis

In this section we are covering Scenario 3 as described in Section 3.1 of [BSCN]. The scenario, assumptions, and requirements are driven from the [BSCN] text. Within this document, this situation is captured in Scenario C of Table 1.

Some enterprise networks may wish to employ an IPv6-dominant network deployment strategy. What this means essentially is that the network or specific sites within the enterprise network will transition to IPv6 using only IPv6 routing to transfer both IPv4 and IPv6 packets over the network, even though the network may be Dual-IP capable. IPv4 routing would not be turned on within an IPv6-dominant network, except if required to support edge IPv4 networks.

Under this scenario, communications between IPv6 nodes will use IPv6. When IPv6-capable nodes in the IPv6-dominant network need to communicate with IPv4 nodes, the IPv6 nodes will use their Dual-IP implementation to tunnel IPv4 packets in IPv6 [V6TUN]. An edge router within the IPv6-dominant network will decapsulate the IPv4 packet and route to the path of the IPv4 node on the network. This permits Dual-IP layer nodes to communicate with legacy IPv4 nodes within an IPv6-dominant network.

Scenarios E and F from Table 1 depict additional cases where an IPv6-dominant deployment strategy could be in place. In Scenario E, the entire network could be IPv6-dominant, but the Host OS 2 system is running an IPv4 application. In Scenario F, the Host OS 1 system network could be IPv6-dominant, but the rest of the networks are all IPv4.

In each case, communicating with an IPv4 end host or over an IPv4 network requires that a transition point exist within the network to support that operation. Furthermore, the node in the IPv6-dominant network must acquire an IPv4 address (to interoperate with the IPv4 end host), and locate a tunnel endpoint on their network which permits the IPv4 packet to be tunneled to the next-hop IPv6 router and eventually to a destination Dual-IP router.

While retaining interoperability with IPv4 is a noble goal for enterprise architects, it is an unfortunate fact that maintaining IPv4 services in an IPv6-dominant network slows and may even impede your ability to reap the maximum benefits of IPv6.

The decision whether or not to use an IPv6-dominant network deployment strategy is completely driven by the enterprise's business and operational objectives and guided by the enterprise's transition plan.

## 7. General Issues from Analysis

In this section, we describe generic enterprise IPv6 deployment issues, applicable to the analysis in Sections 4-6 of this document.

### 7.1. Staged Plan for IPv6 Deployment

The enterprise network administrator will need to follow a staged plan for IPv6 deployment. What this means is that a strategic identification of the enterprise network must be performed for all points and components of the transition.

### 7.2. Network Infrastructure Requirements

The considerations for the enterprise components are detailed in Section 3.2 of [BSCN]. We do not go into detail on all aspects of such components in this document. In this document, we focus on Layer 3 issues.

### 7.3. Stage 1: Initial Connectivity Steps

The first steps for IPv6 deployment do not involve technical aspects per se; the enterprise needs to select an external IPv6 provider and obtain globally routable IPv6 address space from that provider.

### 7.3.1. Obtaining External Connectivity

The enterprise service provider would typically be a topographically close IPv6 provider that is able to provide an IPv6 upstream link. It would be expected that the enterprise would use either native IPv6 upstream connectivity or, in its absence, a manually configured tunnel [BCNF] to the upstream provider.

### 7.3.2. Obtaining Global IPv6 Address Space

The enterprise will obtain global IPv6 address space from its selected upstream provider, as provider-assigned (PA) address space.

The enterprise should receive at least a /48 allocation from its provider, as described in [ALLOC].

Should an enterprise change their provider, a procedure for enterprise renumbering between providers is described in [RENUM].

## 7.4. Stage 2: Deploying Generic Basic Service Components

Most of these are discussed in Section 4 of [BSCN]. Here we comment on those aspects that we believe are in scope for this analysis document. Thus, we have not included network management, multihoming, multicast, or application transition analysis here; however, these aspects should be addressed in Stage 2.

### 7.4.1. Developing an IPv6 Addressing Plan

A site will need to formulate an IPv6 addressing plan, utilizing the globally aggregatable public IPv6 prefix allocated to it by its upstream connectivity provider.

In a Dual-IP deployment, the site will need to decide whether it wishes to deploy IPv6 links to be congruent with existing IPv4 subnets. In this case, nodes will fall into the same links or subnets for both protocols. Such a scheme could be followed, with IPv6 prefix allocations being made such that room for topological growth is provisioned (reducing the potential requirement for future renumbering due to restructuring).

A beneficial property of IPv6 is that an administrator will not need to invest as much effort in address conservation. With IPv4, a site will likely allocate IPv4 subnets to be as small as possible for the number of hosts currently in the subnet (e.g., a /26 for 50 nodes) because IPv4 address conservation is required. This creates problems



when the number of nodes on a subnet grows, larger IPv4 prefixes are then required, and potentially time-consuming and disruptive renumbering events will follow.

With IPv6, a link can in effect have any number of nodes, allowing link growth without the need to adjust prefix allocations with the associated renumbering requirement. The size of the initial site allocation (currently recommended to be a /48) also is likely to allow room for site growth without a need to return to the connectivity provider to obtain more, potentially non-sequential, address space (as is the case for IPv4 today, with the associated paperwork and probable delays).

At the time of writing, best practice in IPv6 site address planning is restricted due to limited wide-scale deployments. Administrators should allocate /64 size prefixes for subnets, and do so in a way that has scope for growth within a site. The site should utilize a plan that reserves space for topological growth in the site, given that its initial IPv6 prefix allocation (currently recommended to be a /48) is likely to include such room for growth. Also see "IPv6 Unicast Address Assignment" [UNAD].

#### 7.4.2. IPv6 DNS

The enterprise site should deploy a DNS service that is capable of both serving IPv6 DNS records using the AAAA format [DNSV6R] and communicating over IPv6 transport.

Specific IPv6 DNS issues are reported in [DNSOP6].

#### 7.4.3. IPv6 Routing

The enterprise network will need to support methods for internal and external routing.

For a single-homed single-site network, a static route to a single upstream provider may be sufficient, although the site may choose to use an exterior routing protocol, especially where it has multiple upstream providers.

For internal routing, an appropriate interior routing protocol may be deployed. IPv6 routing protocols that can be used are as follows: BGP4+ [BGP4], IS-IS [ISIS], OSPFv3 [OSPF], and RIPng [RIPng].

#### 7.4.4. Configuration of Hosts

An enterprise network will have a number of tools available for the delegation and management of IPv4 addresses and other configuration information. These include manual configuration, NIS [NIS], and DHCP [DHCPv4].

In an IPv6 enterprise, Stateless Address Autoconfiguration [CONF] may be used to configure a host with a global IPv6 address, a default router, and on-link prefix information.

Where support for secure autoconfiguration is required, SEND [SEND] can be used. Readers should see the applicability statements to IPsec [IPSEC] within the SEND document.

A stateless configured node wishing to gain other configuration information (e.g., DNS, NTP servers) will likely need a Stateful DHCPv6 [DHCPv6] service available.

For nodes configuring using DHCPv6, where DHCPv6 servers are offlink, a DHCPv6 Relay Agent function will be required. Where DHCPv4 and DHCPv6 service are deployed together, dual-stack considerations need to be made, as discussed within current work on DHCP dual-stack issues [DHDS].

Hosts may also generate or request IPv6 Privacy Addresses [PRIVv6]; there is support for DHCPv6 to assign privacy addresses to nodes in managed environments.

#### 7.4.5. Security

When deploying IPv6 within a Dual-IP network, a site will need to implement its site security policy for IPv6-capable nodes as it does for IPv4-capable nodes. For example, a border firewall should be capable of filtering and controlling IPv6 traffic by enforcing the same policy as it already does for IPv4.

However, a site will also need to review its security policy in light of IPv6-specific functionality that will be deployed in the site, e.g., Mobile IPv6, stateless autoconfiguration (and SEND), IPv6 Privacy Extensions, and end-to-end IPsec. In addition, a site will need to review the use of globally aggregatable public address space where, for IPv4, private addressing and NAT may have been used.

An overview of how Network Architecture Protection (NAP) using IPv6 can provide the same or more benefits without the need for NAT can be found in [NAP]. This describes how the perceived security with IPv4 NAT can be achieved and surpassed with IPv6, i.e., how IPv6

technology can be used to provide the market-perceived benefits of IPv4 NAT.

Where deployed, intrusion detection systems will need to be enhanced to check IPv6 transport both for known application layer attack patterns and for new potential IPv6 threats, e.g., excessive hop-by-hop headers or errant IPv6 header options.

The deployment of specific transition mechanisms may also introduce threats, e.g., carrying IPv6 data tunneled in IPv4. The site security policy should embrace the transition mechanisms that are deployed.

An overview of IPv6 security issues can be found in [V6SEC]. This includes discussion of issues specific to the IPv6 protocol, to transition mechanisms, and to IPv6 deployment itself.

In addition, an enterprise should review all current host-based security requirements for their networks and verify support for IPv6.

#### 7.5. Stage 3: Widespread Dual-Stack Deployment On-Site

With the basic building blocks of external connectivity, interior IPv6 routing, an IPv6 DNS service, and address allocation management in place, the IPv6 capability can be rolled out to the wider enterprise. This involves putting IPv6 on the wire in the desired links, and enabling applications and other services to begin using an IPv6 transport.

In the Dual-IP deployment case, this means enabling IPv6 on existing IPv4 subnets. As described in Section 7.4.4, above, it is likely that IPv6 links will be congruent with IPv4 subnets because IPv4 subnets tend to be created for geographic, policy, or administrative reasons that would be IP version-independent.

While the use of IPv6 by some applications can be administratively controlled (e.g., in the case of open source software by compiling the application without IPv6 support enabled), the use of IPv6 transport, and preference over IPv4 transport, will vary per application based on the developer/author's implementation.

A Dual-IP deployment will often be made by sites wishing to support use of IPv6 within a site, even if IPv6 transport is not preferred by all applications. Putting support for IPv6 in all site infrastructure (DNS, email transport, etc.) allows IPv6 usage to be phased in over time. As nodes become IPv6 capable, and applications and services IPv6 enabled, the IPv6 capable infrastructure can be leveraged. For most networks, Dual-IP will be at the very least a

medium-term transition towards an IPv6-dominant future. However, the introduction of IPv6 support, with the potential benefits of globally aggregatable public address usage (with [NAP]) and other new IPv6 capabilities, can bring more immediate benefits for the site.

## 8. Applicable Transition Mechanisms

This section will provide general guidance for the use of specific transition mechanisms which in turn can be used by the enterprise to support the enterprise matrix notional networks (rows) in Section 3, and within the context of the analysis discussed in Sections 4, 5, and 6.

Table 1 provides a number of common scenarios that an enterprise architect might encounter as they consider how and where they should consider deploying transition mechanisms to support the network transition to IPv6. Selecting the most appropriate mechanism for each scenario is more of an art than a science and consequently making recommendations against each of the ten scenarios would be simply fodder for sharpshooters touting their favored product. However we can provide some high-level guidance that should benefit the architect's decision-making process.

### 8.1. Recognizing Incompatible Network Touchpoints

Mapping your specific situation into one of the ten scenarios of Table 1 is far less important than recognizing the critical touchpoints within the enterprise networks where incompatible networks interface. Unless a transition mechanism is being offered by the enterprise as a service, it is at these touchpoints that a mechanism must be considered.

A quick review of Table 1 reveals that the ten scenarios can be boiled down to variations of four major themes. The simplest, but also most favored (due to its flexibility), is widespread Dual-IP with compatible hosts at either end. This situation is illustrated in Scenario A, and transition mechanism considerations have already been described in some detail in Section 4.

In the second common theme (depicted in Scenarios B-D of Table 1), the enterprise is comprised of compatible hosts, with one or more incompatible network touchpoints in between. As described in Section 4.2.2.1, tunneling can be used to "bypass" the incompatible network segments. One tunneling option, manually configured tunnels [BCNF] could be used by the enterprise, but as the name implies, this mechanism provides no automated tunnel configuration.

"Connection of IPv6 Domains via IPv4 Clouds" [6TO4] can be used to support enterprises that do not have an assigned IPv6 prefix address.

Identifying the responsible device to perform the tunneling is driven by the position of the incompatible touchpoint. If a local network is incompatible, then host tunneling is appropriate. If the backbone (provider) network is incompatible, then gateway-to-gateway tunneling might be a better choice. By working to ensure tunnel endpoints are always configured at Dual-IP devices, end-to-end communication or services (IPv4 or IPv6) can be preserved.

Readers should review the current work regarding tunnels within the IETF Softwire working group and problem statement [SOFTW].

Having IPv6 applications on a Dual-IP host on a v4-only network requires some form of tunneling. Where configured tunnels are not sufficient, a more automatic solution may be appropriate. Available solutions include the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [ISTP] or Teredo [TRDO] to tunnel to a v6 end service. ISATAP [ISTP] can be used to provide end-node IPv6 connectivity from nodes on an isolated IPv4 network, through the use of automatic tunneling of IPv6 in IPv4. Teredo [TRDO] can be used when the enterprise network is behind a NAT.

Enterprise architects should consider providing a Tunnel Broker [TBRK] [TSPB] as a cost-effective service to local users or applications. Tunnel Brokers can be used to provide tunnel setup for an enterprise using manually configured tunnels and 6TO4 [6TO4]. Tunnel Brokers can automate the use of tunnels across an enterprise deploying IPv6.

Later in the transition process, after the enterprise has transitioned to a predominately IPv6 infrastructure, the architect will need to determine a network transition strategy to tunnel IPv4 within IPv6 [V6TUN] across IPv6-dominant links, or the enterprise Intranet. Or in the case of early deployment of IPv6-dominant networks, the architect will need to address this from the beginning of the required transition planning.

## 8.2. Recognizing Application Incompatibilities

Having recognized incompatible network touchpoints, it is also incumbent on the architect to identify application incompatibilities. During the transition period, particularly for large enterprises, it is to be expected that an application hosted at one location may lead (or lag) the IPv6-compatibility of its peer (or server) at some other location.

This leads us to the third theme (represented by Scenarios E and G): incompatible applications communicating across a homogenous network. Translation is an obvious solution, but not recommended except for legacy devices that are at the network edge and cannot or never will be upgraded to IPv6. A more scalable solution would be to use an Application Layer Gateway (ALG) between the incompatible hosts.

### 8.3. Using Multiple Mechanisms to Support IPv6 Transition

Inevitably, during the course of transitioning a large enterprise to IPv6, the architect will be faced with both incompatible hosts and simultaneously (at different parts of the enterprise) incompatible networks. These highly complex situations represent the fourth common theme in Table 1 (specifically depicted by Scenarios F, H, I, and J). Maintaining IP interoperability in these situations requires additional planning and may require multiple or even nested use of diverse transition mechanisms. For example, an ALG collocated with the application server may be required to service both IPv4 and IPv6 data streams that are simultaneously tunneled through incompatible network segment(s).

## 9. Security Considerations

Security considerations for IPv6 deployment in a Dual-IP environment are discussed above in Section 7.4.5, where external references to overview documents [V6SEC] [NAP] are also included.

## 10. References

### 10.1. Normative References

- [CONF] Thomson, S. and T. Narten, "IPv6 Stateless Address Autoconfiguration", RFC 2462, December 1998.
- [DHCPv6] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [6TO4] Carpenter, B. and K. Moore, "Connection of IPv6 Domains via IPv4 Clouds", RFC 3056, February 2001.
- [BSCN] Bound, J., Ed., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.
- [TRDO] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.

- [ISTP] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 4214, October 2005.
- [V6TUN] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, December 1998.
- [TBRK] Durand, A., Fasano, P., Guardini, I., and D. Lento, "IPv6 Tunnel Broker", RFC 3053, January 2001.
- [ALLOC] IAB and IESG, "IAB/IESG Recommendations on IPv6 Address Allocations to Sites", RFC 3177, September 2001.
- [NATPT] Tsirtsis, G. and P. Srisuresh, "Network Address Translation - Protocol Translation (NAT-PT)", RFC 2766, February 2000.
- [UMAN] Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks", RFC 3904, September 2004.
- [ISPA] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.
- [3GPA] Wiljakka, J., Ed., "Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks", RFC 4215, October 2005.
- [OSPF] Coltun, R., Ferguson, D., and J. Moy, "OSPF for IPv6", RFC 2740, December 1999.
- [BGP4] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, January 2007.
- [ISIS] Oran, D., Ed., "OSI IS-IS Intra-domain Routing Protocol", RFC 1142, February 1990.
- [RIPng] Malkin, G. and R. Minnear, "RIPng for IPv6", RFC 2080, January 1997.
- [APPS] Shin, M-K., Ed., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.
- [RENUM] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.

- [BCNF] Nordmark, E. and R. Gilligan, "Basic Transition Mechanisms for IPv6 Hosts and Routers", RFC 4213, October 2005.
- [ULA] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, October 2005.
- [DNSOP6] Durand, A., Ihren, J., and P. Savola, "Operational Considerations and Issues with IPv6 DNS", RFC 4472, April 2006.
- [DNSV6R] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", RFC 3596, October 2003.
- [NIS] Kalusivalingam, V., "Network Information Service (NIS) Configuration Options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3898, October 2004.
- [DHCPv4] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [IPSEC] Eastlake 3rd, D., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4305, December 2005.
- [SEND] Arkko, J., Ed., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", RFC 3971, March 2005.
- [PRIVv6] Narten, T. and R. Draves, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", RFC 3041, January 2001.

## 10.2. Informative References

- [TSPB] Blanchet, M., and F. Parent, "IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)", Work in Progress, August 2005.
- [V6SEC] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/Co-existence Security Considerations", Work in Progress, October 2006.
- [NAP] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", Work in Progress, January 2007.
- [CAMP] Chown, T., "IPv6 Campus Transition Scenario Description and Analysis", Work in Progress, March 2007.



- [DHDS] Chown, T., Venaas, S., and C. Strauf, "Dynamic Host Configuration Protocol (DHCP): IPv4 and IPv6 Dual-Stack Issues", RFC 4477, May 2006.
- [UNAD] Van de Velde, G., Popoviciu, C., and T. Chown, "IPv6 Unicast Address Assignment", Work in Progress, March 2007.
- [VLAN] Chown, T., "Use of VLANs for IPv4-IPv6 Coexistence in Enterprise Networks", RFC 4554, June 2006.
- [V6DEF] Roy, S., Durand, A., and J. Paugh, "IPv6 Neighbor Discovery On-Link Assumption Considered Harmful", Work in Progress, January 2006.
- [SOFTW] Dawkins, S., Ed., "Softwire Problem Statement", Work in Progress, March 2007.

## 11. Acknowledgments

The authors would like to acknowledge contributions from the following: IETF v6ops Working Group members, Fred Baker, Pekka Savola, and Jordi Palet

## Appendix A. Crisis Management Network Scenarios

### A.1. Introduction

This appendix first describes different scenarios for the introduction of IPv6 into a crisis management network for emergency services, defense, or security forces that are currently running IPv4 service. Then, the scenarios for introducing IPv6 are analyzed, and the relevance of already defined transition mechanisms are evaluated. Known challenges are also identified.

When a crisis management enterprise deploys IPv6, its goal is to provide IPv6 connectivity on its institutional fixed networks and on the mobile wireless services that are deployed to a crisis area. The new IPv6 service must be added to an already existing IPv4 service, the introduction of IPv6 must not interrupt this IPv4 service, and the IPv6 services must be interoperable with existing IPv4 services.

Crisis management enterprises accessing IPv4 service across mobile ground networks, airborne networks, and satellites will find different ways to add IPv6 to this service based on their network architecture, funding, and institutional goals. This document discusses a small set of scenarios representing the architectures for IPv6 expected to be dominant in crisis management networks during the next decade. This document evaluates the relevance of the existing transition mechanisms in the context of these deployment scenarios, and points out the lack of essential functionality within these methods for a provider to support IPv6 services for these scenarios.

The document focuses on services that include both IPv6 and IPv4 and does cover issues surrounding accessing IPv4 services across IPv6-only networks. It is outside the scope of this document to describe detailed implementation plans for IPv6 in defense networks.

### A.2. Scenarios for IPv6 Deployment in Crisis Management Networks

#### Scenario 1: Limited IPv6 Deployment Network

Sparse IPv6 dual-stack deployment in an existing IPv4 network infrastructure. Enterprise with an existing IPv4 network wants to deploy a set of particular IPv6 "applications" and have some ability to interoperate with other institutions that are using IPv6 services. The IPv6 deployment is limited to the minimum required to operate this set of applications.

Assumptions: IPv6 software/hardware components for the application are available, and platforms for the application are IPv6 capable.

Requirements: Do not disrupt IPv4 infrastructure.

#### Scenario 2: Dual-Stack Network

Wide-scale/total dual-stack deployment of IPv4 and IPv6 capable hosts and network infrastructure. Enterprise with an existing IPv4 network wants to deploy IPv6 in conjunction with their IPv4 network in order to take advantage of emerging IPv6 network-centric capabilities and to be interoperable with other agencies, international partners, and commercial enterprises that are deploying an IPv6 architecture.

Assumptions: The IPv4 network infrastructure used has an equivalent capability in IPv6.

Requirements: Do not disrupt existing IPv4 network infrastructure with IPv6. IPv6 should be equivalent or "better" than the network infrastructure in IPv4. It may not be feasible to deploy IPv6 on all parts of the network immediately. Dual-stacked defense enterprise network must be interoperable with both IPv4 and IPv6 networks and applications.

#### Scenario 3: IPv6-Dominant Network

Enterprise has some limited IPv4-capable/only nodes/applications needing to communicate over the IPv6 infrastructure. Crisis management enterprise re-structuring an existing network, decides to pursue aggressive IPv6 transition as an enabler for network-centric services and wants to run some native IPv6-only networks to eliminate cost/complexity of supporting a dual stack. Some legacy IPv4 capable nodes/applications within the enterprise will have slow technical refresh/replacement paths and will need to communicate over the IPv6 dominant infrastructure for years until they are replaced. The IPv6-dominant enterprise network will need to be interoperable with its own legacy networks, commercial networks, and the legacy networks of similar organizations that will remain IPv4-dominant during a long transition period. Reserve units, contractors, other agencies, and international partners may need IPv4 service across this enterprise's IPv6-dominant backbone.

Assumptions: Required IPv6 network infrastructure is available, or available over some defined timeline, supporting the aggressive transition plan.

Requirements: Reduce operation and maintenance requirements and increase net-centricity through aggressive IPv6 transition. Interoperation and coexistence with legacy IPv4 networks and applications is required. Legacy IPv4 nodes/applications/networks will need to be able to operate across the IPv6 backbone and need to

be able to interoperate with the IPv6-dominant network's nodes/applications.

A.3. Description of a Generic Crisis Management Network

A generic network topology for crisis management reflects the various ways a crisis management network can connect customers through their network infrastructure. Because the institution's existing wired and fixed-site wireless infrastructure can be destroyed or unavailable in a crisis, the crisis management network must be able to deploy its own mobile wireless network or connect through external wired and wireless networks provided by ISPs or partner organizations. This infrastructure lets us divide the basic areas for IPv4/IPv6 interoperability into three main areas: the customer applications, the local network, and the network backbone.

The basic components in a crisis management network are depicted in Figure 1.

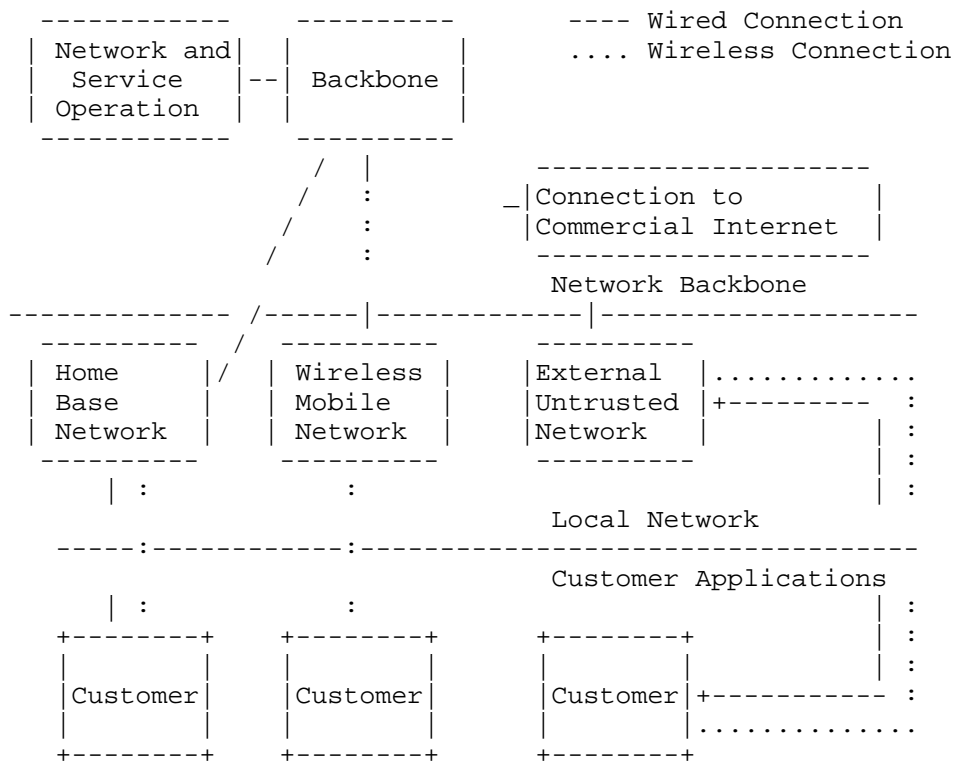


Figure 1: Crisis Management Network Topology.

#### A.4. Stages of IPv6 Deployment

The stages are derived from the generic description of scenarios for crisis management networks in Section 2. Combinations of different building blocks that constitute a crisis network environment lead to a number of scenarios from which the network engineers can choose. The scenarios most relevant to this document are those that maximize the network's ability to offer IPv6 to its customers in the most efficient and feasible way. In the first three stages, the goal is to offer both IPv4 and IPv6 to the customer, and it is assumed that in the distant future, all IPv4 services will be eventually switched to IPv6. This document will cover engineering the first four stages.

The four most probable stages are:

- o Stage 1            Limited Launch
- o Stage 2            Dual-Stack Dominance
- o Stage 3            IPv6 Dominance
- o Stage 4            IPv6 Transition Complete

Generally, a crisis management network is able to entirely upgrade a current IPv4 network to provide IPv6 services via a dual-stack network in Stage 2 and then slowly progress to Stages 3 and 4 as indicated in Figure 2. During Stage 2, when most applications are IPv6 dominant, operational and maintenance costs can be reduced on some networks by moving to Stage 3 and running backbone networks entirely on IPv6, while adding IPv4 backwards compatibility via v4 in v6 tunneling or translation mechanisms to the existing configuration from Stage 2. When designing a new network, if a new IPv6-only service is required, it can be implemented at a lower cost by jumping directly to Stage 3/4 if there are only limited or no legacy concerns.

##### Stage 1: Limited Launch

The first stage begins with an IPv4-only network and IPv4 customers. This is the most common case today and the natural starting point for the introduction of IPv6. During this stage, the enterprise begins to connect individual IPv6 applications run on dual-stacked hosts through host-based tunneling using Tunnel Broker, ISATAP, or Teredo. Some early adopter networks are created for pilot studies and networked together through configured tunnels and 6to4.

The immediate first step consists of obtaining a prefix allocation (typically a /32) from the appropriate RIR (e.g., AfriNIC, APNIC, ARIN, LACNIC, RIPE) according to allocation procedures.

The crisis management enterprise will also need to establish IPv6 connectivity between its home base networks and mobile wireless networks over its backbone. It will need to negotiate IPv6 service with its service providers and with peer organizations; it is of utmost importance to require IPv6 capability or an upgrade plan when negotiating purchases of network applications and infrastructure. In the short term, network connections, especially legacy wireless networks that cannot provide IPv6 services, can provide IPv6 services through the use of tunnels. However, the longer-term goal must be requiring and obtaining IPv6 native connectivity from the transit networks. Otherwise, the quality of IPv6 connectivity will likely be poor and the transition to Stage 2 will be delayed.

#### Stage 2: Dual-Stack Dominance

Stage 2 occurs when most applications, local networks, and network backbones become dual-stacked so that native IPv6 connections are enabled. At this point there is a mix of IPv4 and IPv6 applications and services in use across the enterprise. The enterprise may be made IPv6-capable through either software upgrades, hardware upgrades, or a combination of both. Generally IPv6 is added during normal technical refresh as the enterprise buys new equipment that is IPv6 ready.

Specialty legacy applications and wireless/satellite networks may be especially slow to transition to IPv6 capability due to upgrade costs, so plans must be made for backwards compatibility for these systems. Since some new IPv6 services cannot be provided through IPv4, and some legacy network connections may not yet be upgraded, tunneling mechanisms have to be provided on the backbone to provide IPv6 connectivity through to customer IPv6 applications still relying on legacy IPv4-only networks. The tunnels may provide host-based tunneling for individual customers or site-to-site tunnels to connect small IPv6 domains through IPv4-only networks. If any new applications are IPv6-only rather than dual-stacked, and need to interact with IPv4-only legacy applications, translators will be used as a transition mechanism of last resort during this stage.

#### Stage 3: IPv6 Dominance

Applications are deployed specifically to use IPv6 as benefit; thus, network backbone and nodes use IPv6 and not IPv4, except where IPv4 is legacy.

## Authors' Addresses

Jim Bound  
HP  
110 Spitbrook Road  
Nashua, NH 03062  
USA  
Phone: 603.465.3130  
EMail: jim.bound@hp.com

Yanick Pouffary  
HP Competency Center  
950, Route des Colles, BP027,  
06901 Sophia Antipolis CEDEX  
FRANCE  
Phone: + 33492956285  
EMail: Yanick.pouffary@hp.com

Tim Chown  
School of Electronics and Computer Science  
University of Southampton  
Southampton SO17 1BJ  
United Kingdom  
EMail: tjc@ecs.soton.ac.uk

David Green  
Command Information  
13655 Dulles Technology Drive  
Suite 500  
Herndon, VA 20171  
USA  
Phone: 703.561.5937  
EMail: green@commandinformation.com

Steve Klynsma  
The MITRE Corporation  
7515 Colshire Drive  
McLean, VA 22102-5708  
USA  
Phone: 703-883-6469  
EMail: sklynsma@mitre.org

## Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

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

## Intellectual Property

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

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

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

## Acknowledgement

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