Internet Protocol Version 6 Migration Assessment

For OI&T

GENERAL DYNAMICS
Network Systems

Document Title: IP Version 6 Migration Assessment

Date	Filename/Version #	Author	Revision Description
6/20/2005	Ipv6-migration .doc	Tom Marsh/Andrew Delong	Initial Draft
8/25/2005	Ipv6-migration v1.1	Allen Heuvel	Added minor revision and document tracking page
1/9/2006	IPv6-migration v1.2	Allen Heuvel	Added input to testing section and minor updates

I. INTRODUCTION **4**
PURPOSE 4
PROBLEM 4
ADDRESS LIMITATIONS 4
TECHNOLOGY LIMITATIONS 4
II. PROPOSED SOLUTION **4**
IP VERSION 6 4
III. PROPOSED MIGRATION STRATEGY **5**
OVERVIEW 5
ARCHITECTURE 5
MIGRATION STAGES 6
Stage 1 - Core and Distribution *6*
Stage 2 - Regional *7*
Stage 3 - Campus and Data Center *8*
IV. LAB ENVIRONMENT **8**
MIGRATION PLAN 8
ADDRESSING SCHEME 8
V. TESTING PROCEDURE **11**
VI. CONCLUSION **11**
MIGRATION STRATEGY ADVANTAGES 11
MIGRATION STRATEGY DISADVANTAGES 11
APPENDIX A. **12**
TESTING AND CONFIGURATION STEPS 12
Enable IPv6 Routing and Configure Addresses *12*
Initiate BGP Process and Configure IPv6 Neighbors *12*
Initiate the OSPF version 3 Process and Configure IPv6 Neighbors *12*
Configure EIGRP (or OSPF) for IPv6 *13*

I. Introduction

Purpose

This document will touch on some of the inherent issues with the current Internet Protocol version 4 (IPv4) standards, examine the current state of the IPv4 VA WAN, and explore a potential successor and migration path. The document also details the IPv6 fielding in the TMP Lab on a replica of the One-VA Network backbone to move the IPv6 discussion beyond the theoretical and closer to implementation. This document will not address all possible solutions but rather will evaluate a single resolution and migration tactic.

Problem

As the growth of the Internet continues unabated, the addition of web integrated devices and internet hosts is increasing rapidly. This rapid expansion was not foreseen by the creators of the existing IPv4 standard, when it first appeared in documented form in 1981. They could not have predicted the massive integration of the Internet around the world and, as a result, their standard was not designed to accommodate this use.

Address Limitations

The existing IPv4 implementation allows for only 4,294,967,296 addresses, which may seem adequate to the casual observer however this has proven to not be the case. Mobile computing needs have blossomed in the past five years, with the surge in implementation of wireless computing, which has resulted in a mass of mobile users accessing the Internet from laptops, cell phones, and even cars. This impact is especially difficult within developing industrial areas, such as Asia, where a lack of IP addresses has the potential to cripple business operations and expansion. IP translation methods, such as NAT and PAT, have eased the demand substantially however they have only prolonged the inevitable address shortage.

Technology Limitations

Emerging technologies, such as voice over IP and secure computing, have further questioned the feasibility of sustaining IPv4. These systems often require the use of QoS and improved encryption functionality which is limited in the current IP implementation. Requirements such as these have been accounted for and are being integrated into the continued development of theIPv4 protocol standard.

II. Proposed Solution

IP Version 6

IP Version 6 (IPv6) was created to provide a solution to the many problems of IPv4 and provide for future expansion. First, the address size has been increased dramatically from 32-bit addresses to 128-bit. This provides enough addresses for each person on the planet to have one thousand thirty of their own. The restructured datagram (Fig 1) features no fragmentation, efficient auto configuration, improved header format, and improved support for mobile IP, integrated QoS, and IPSec support. The 128-bit addressing scheme will additionally provide for a hierarchical routing structure facilitating route aggregation.

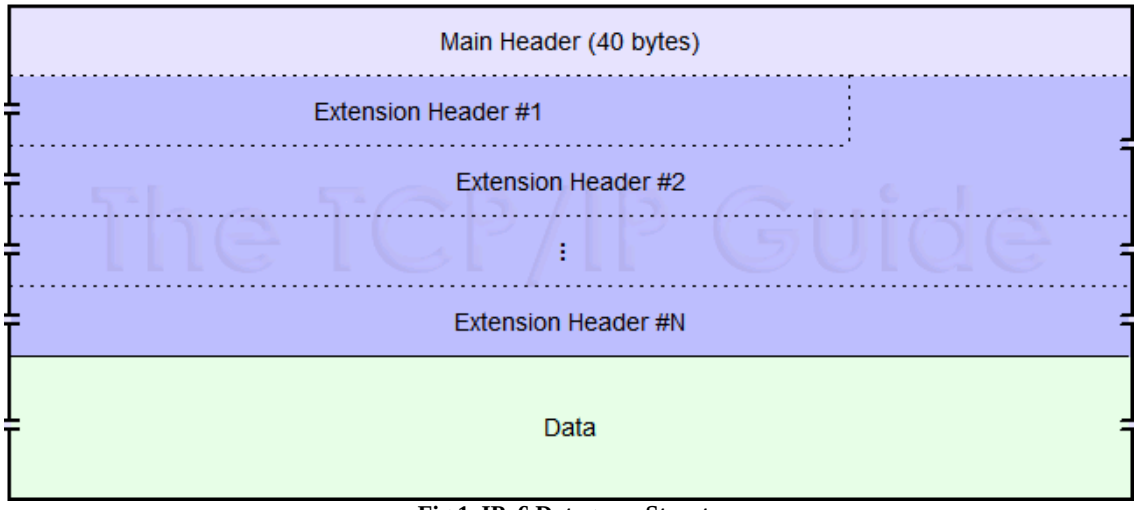

Fig 1. IPv6 Datagram Structure

III. Proposed Migration Strategy

The purpose of this section outlines a proposed IPv6 methodology for the migration of the VA WAN. Further testing and research will be necessary to ensure full interoperability with productions systems before performing actual migration steps. This document covers the basics of the logical process steps and device input for VA Backbone migration.

Overview

The proposed solution to the limitation of IPv4 is to perform a gradual migration to a pure IPv6 environment. Cisco routers, operating with the correct IOS version, will support IPv4 and IPv6 protocol stacks concurrently over the same interfaces. Due to the homogeneous nature of the VA WAN, an excellent opportunity exists for a simple migration strategy having minimal impact on the user community and requiring little administrative overhead.

Architecture

The existing VA network's logical and physical architecture will remain intact. IOS upgrades may be required for a number of devices however these can be scheduled off peak hours to minimize any potential impact to the end user.

Applying the IPv6 process was found to slightly increase the CPU utilization of a device and should be taken into consideration before attempting deployment in the VA WAN. Initial results during the testing phase of this assessment, CPU utilization raised less than 1% after process activation. The TMP lab hardware is not identical to VA production nodes however the CPU utilization needs to be examined on an actual production devices in future testing and after activation. The process increase is not significant and should not be considered as a potential setback unless a node is operating with a high CPU level prior to activation. Minimal network hardware would be required for the first stages of

this migration. Upgrades to memory and flash would be the more common requirement to handle the higher CPU levels and larger code revisions to run the dual protocols concurrently.

Migration Stages

The proposed migration method is a multi-staged conversion which will permit a gradual integration of new hardware, and IOS software if necessary. This will enable the VA to perform extensive testing and troubleshooting gradually at every stage during the integration of this relatively new and still evolving technology.

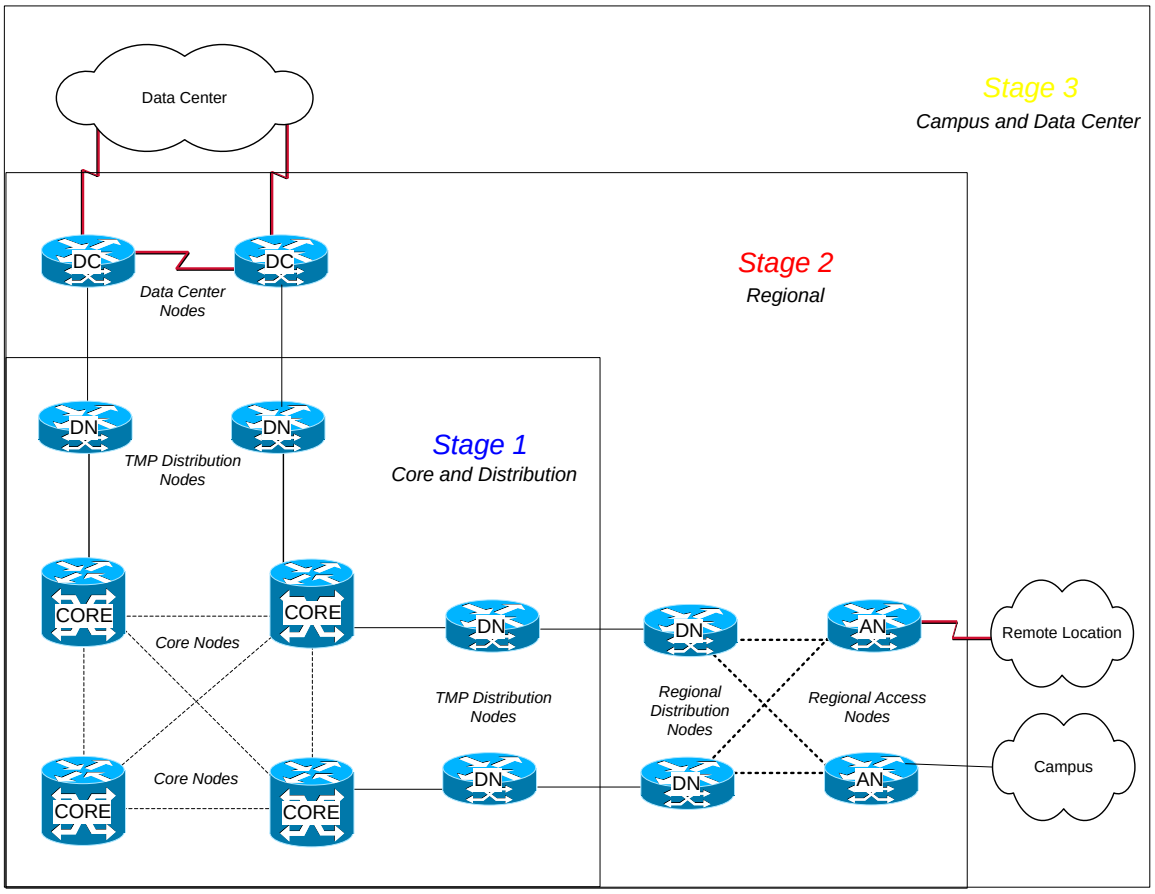

Fig 2. IPv6 Migration Stages

Stage 1 - Core and Distribution

The initial stage consists of an IPv6 dual stack implementation at the core and TMP distribution nodes and is roughly equivalent in scope to the TMP Phase III project. This will enable network administrators to thoroughly test connectivity and BGP/EIGRP communications before proceeding and provide the initial opportunity to examine multiple EIGRP processes redistributing between differing BGP autonomous systems. Existing VA Core and Distribution devices at this stage are currently 100% capable of operating in a dual stack environment and could be migrated immediately. This stage meets the requirements outlined in the OMB memo M-05-22 to be accomplished by June

2008. Details of this mandate will be further clarified with a NIST Federal standards compliance document which is scheduled to be published in the near future.

Stage 2 - Regional

The second phase will consist of spreading the IPv6 conversion to the regional and data center nodes. This configuration will be similar to Stage 1, however only approximately 80% of the Access Node (AN) devices are currently capable of running in the dual stack environment. Legacy devices must be upgraded and IOS versions must be examined to confirm IPv6 compatibility. It is anticipated that routine technical refresh following the OMB memo will increase the percentage of devices that will be capable of IPv6 over the next 18 months.

Stage 3 - Campus and Data Center
The third stage will complete the migration and integrate the individual remote locations, campuses, and data centers into the IPv6 backbone. This will be the most lengthy and labor intensive of all stages as it will require a thorough examination, design, and upgrade, of the existing operating systems, applications, and IPv4 hardware devices. This stage will require the most extensive training for local support personnel and should only be conducted after a detailed migration strategy has been developed in order to prevent communications disruptions and ensure business continuity.

IV. Lab Environment

The TMP lab network consists of multiple Cisco 2600 series routers simulating a single AS location, distribution node, and data center configuration. All devices are configured as replicas of the actual nodes on the production VA network Backbone.

Migration Plan
All Cisco routers running IOS version 12.3x and higher will support a dual-stack configuration running both IPv4 and IPv6 simultaneously. The migration will consist of addressing the required interfaces with IPv6 addresses, configuring BGP, OSPF, and EIGRP to redistribute the new protocol, and finally disabling IPv4 and operating within a pure IPv6 environment at the AS level. The staged migration, described above, will be followed in its entirety.

Addressing Scheme
All existing IPv4 addresses were converted into an IPv6 equivalent (Figs 2-5) for the purposes of this experiment. This addressing structure has been designed for the purposes of this exercise only and is not intended for actual use upon the VA WAN. If, and when, the decision is made to proceed with an IPv6 migration, the addressing structure allotted and eventually obtained by the VA must be examined and designed in much greater depth accordingly to accommodate the existing WAN architecture and to be able to provide for future computing needs. For the purposes of this experiment we have followed the Cisco recommended IPv6 addressing scheme, a summary of which can be retrieved at:
http://www.cisco.com/application/pdf/en/us/guest/tech/tk872/c1482/cdccont_0900aecd8026003d.pdf

Cores	IPv4 Address	Mask	IPv6 Address	Mask
Reston				
ATM1/0.5	20.252.5.1	255.255.255.252	2005::14FC:0501	/126
ATM1/0.23	20.252.23.133	255.255.255.252	2005::14FC:1785	/126
ATM1/0.101	20.252.255.18	255.255.255.252	2005::14FC:FF12	/126
ATM1/0.102	20.252.255.1	255.255.255.252	2005::14FC:FF01	/126
ATM1/0.103	20.252.255.5	255.255.255.252	2005::14FC:FF05	/126
Loopback0	199.12.113.128	255.255.255.255	2005::C70C:7180	/128
Lees Summit				
ATM1/0.5	20.252.5.129	255.255.255.252	2005::14FC:0581	/126
ATM1/0.23	20.252.23.1	255.255.255.252	2005::14FC:1701	/126
ATM1/0.101	20.252.255.13	255.255.255.252	2005::14FC:FF0D	/126
ATM1/0.102	20.252.255.2	255.255.255.252	2005::14FC:FF02	/126
ATM1/0.103	20.252.255.9	255.255.255.252	2005::14FC:FF09	/126
Loopback0	199.12.113.160	255.255.255.255	2005::C70C:71A0	/128
Richardson				
ATM1/0.23	20.252.23.129	255.255.255.252	2005::14FC:1781	/126
ATM1/0.101	20.252.255.17	255.255.255.252	2005::14FC:FF11	/126
ATM1/0.102	20.252.255.21	255.255.255.252	2005::14FC:FF15	/126
ATM1/0.103	20.252.255.14	255.255.255.252	2005::14FC:FF0E	/126
Loopback0	199.12.113.192	255.255.255.255	2005::C70C:71C0	/128
Rancho Cordova				
ATM1/0.26	20.252.23.13	255.255.255.252	2005::14FC:170D	/126
ATM1/0.101	20.252.255.6	255.255.255.252	2005::14FC:FF06	/126
ATM1/0.102	20.252.255.22	255.255.255.252	2005::14FC:FF16	/126
ATM1/0.103	20.252.255.10	255.255.255.252	2005::14FC:FF0A	/126
Loopback0	199.12.113.224	255.255.255.255	2005::C70C:71E0	/128

Fig 3. Core IPv4 to IPv6 Address Conversion

DC	IPv4 Address	Mask	IPv6 Address	Mask
Austin-dc1				
FastEthernet0/0	20.254.128.1	255.255.255.0	2005::14FE:8001	/120
ATM1/0.23	20.252.23.2	255.255.255.252	2005::14FC:1702	/126
Loopback0	199.12.113.237	255.255.255.255	2005::C70C:71ED	/128
Austin-dc2				
FastEthernet0/0	20.254.128.2	255.255.255.0	2005::14FE:8002	/120
ATM1/0.23	20.252.23.130	255.255.255.252	2005::14FC:1782	/126
Loopback0	199.12.113.160	255.255.255.255	2005::C70C:71A0	/128

Fig 4. DC IPv4 to IPv6 Address Conversion

AS	IPv4 Address	Mask	IPv6 Address	Mask
Martinsburg-AS5				
FastEthernet0/0	20.40.189.13	255.255.255.0	2005::1428:BD0D	/120
ATM1/0.5	20.252.5.2	255.255.255.252	2005::14FC:0502	/126
Loopback0	199.12.113.133	255.255.255.255	2005::C70C:7185	/128
silver-spring-as5				
FastEthernet0/0	20.40.189.33	255.255.255.0	2005::1428:BD21	/120
ATM1/0.5	20.252.5.130	255.255.255.252	2005::14FC:0582	/126
Loopback0	199.12.113.165	255.255.255.255	2005::C70C:71A5	/128

Fig 5. AS IPv4 to IPv6 Address Conversion

DW	IPv4 Address	Mask	IPv6 Address	Mask
Martinsburg-AS5-DW				
FastEthernet0/0	20.40.189.14	255.255.255.252	2005::1428:BD0E	/126
Serial1/0	20.40.189.49	255.255.255.252	2005::1428:BD31	/126
Serial1/1	20.40.189.65	255.255.255.252	2005::1428:BD41	/126
Serial1/2	20.40.189.54	255.255.255.252	2005::1428:BD36	/126
Loopback0	199.12.113.133	255.255.255.255	2005::C712:7185	/128
Loopback3	199.12.112.49	255.255.255.255	2005::C712:7131	/128
Silver_Spring-AS5-DW				
FastEthernet0/0	20.40.189.34	255.255.255.252	2005::1428:BD22	/126
Serial1/0	20.40.189.41	255.255.255.252	2005::1428:BD29	/126
Serial1/1	20.40.189.45	255.255.255.252	2005::1428:BD2D	/126
Serial1/2	20.40.189.53	255.255.255.252	2005::1428:BD35	/126
Loopback0	20.40.189.200	255.255.255.255	2005::1428:BDC8	/128
Loopback3	199.12.112.50	255.255.255.255	2005::C70C:7032	/128

Fig 6. DW IPv4 to IPv6 Address Conversion

AR	IPv4 Address	Mask	IPv6 Address	Mask
BAL-AS5-AR				
FastEthernet0/0	20.42.254.28	255.255.255.0	2005::142A:FE1C	/112
Serial1/0	20.40.189.50	255.255.255.252	2005::1428:BD32	/126
Serial1/1	20.40.189.46	255.255.255.252	2005::1428:BD2E	/126
Loopback0	20.40.189.203	255.255.255.255	2005::1428:BDCB	/128
Loopback3	199.12.112.52	255.255.255.255	2005::C70C:7034	/128
WAS-AS5-AR				
FastEthernet0/0	20.42.1.1	255.255.255.0	2005::142A:0101	/126
Serial1/0	20.40.189.42	255.255.255.252	2005::1428:BD2A	/126
Serial1/1	20.40.189.66	255.255.255.252	2005::1428:BD42	/126
Loopback0	20.40.189.202	255.255.255.255	2005::1428:BDCA	/128
Loopback3	199.12.112.53	255.255.255.255	2005::C70C:7035	/128

Fig 7. AR IPv4 to IPv6 Address Conversion

V. Testing Procedure

See Appendix A regarding details of the TMP lab testing

VI. Conclusion

Based on this assessment, the option of running a dual stack IPv4 and IPv6 network (also recommended by Cisco) will provide the smoothest, least hazardous, and most efficient method of transitioning the VA WAN in its entirety to comply with the initial OMB deadline. This plan puts emphasis on a seamless transparent path to the end-users during the course of the initial migration providing time to properly design, layout and provide training for the next generation architecture of the VA Enterprise. This will fully depend upon the required timetable for an IPv6 conversion, set by the VA. The network migration in Stages 1 and 2 could potentially be completed with minimal additional cost and effort however the majority of this project will take place in Stage 3, as many of the existing systems and applications must be modified or replaced at each campus, remote location, and data center.

Migration Strategy Advantages

By implementing a gradual transition from IPv4 to IPv6, the VA will mitigate numerous issues related to rapid migration (including excessive administrative demands, hardware cost, and data disruption potential). A blossoming transition, beginning at the core and spreading to the remote locations, will enable the VA to gradually replace outdated hardware and troubleshoot a live IPv6 environment without disrupting the existing IPv4 data flow. Once the IPv6 stack has been initiated at the Core, Distribution, and Access layers, VA WAN engineers can gradually “grandfather in” the IPv6 technology at each “tech refresh” interval for the campus, remote, and data center locations. This will minimize the immediate administrative burden and permit the VA to spread the hardware cost over numerous fiscal years.

Migration Strategy Disadvantages

The major disadvantage to a dual stack migration strategy is the lack of a driving force behind this implementation. Currently, no specific business or technology requirement is pressing the IPv6 migration and the potential exists for a lengthy or even stalled migration. Specific guidelines and timetables must be implemented in a future detailed design plan to force network modification goals and require new hardware to operate with a dual stack upon installation. If this policy is in place and adequately enforced, the potential exists for a complete migration of the VA Enterprise within the next three to five years.

Appendix A.

The first step in converting to IPv6 is to enable the IPv6 routing process on the devices which will carry IPv6 traffic. Since the stacks will be running concurrently, and consequently IPv4 functionality will remain unaffected, IPv6 configurations may be applied to devices in any predetermined order.

Testing and Configuration Steps

All lab configurations are deigned for simplicity and testing purposes only. Specific security and management options, such as route maps and prefix lists, are omitted throughout this test scenario. These options must be reexamined prior to an actual migration on the VA WAN.

Enable IPv6 Routing and Configure Addresses

To enable the IPv6 routing and address each interface, the following actions were performed on all routers, in each phase:

1. Enable IPv6 routing globally
    a. (config)# **ipv6 unicast-routing**
2. Address each interface (from table in Fig 2)
    a. (config-if)# **ipv6 address** *ipv6-address/mask*

Initiate BGP Process and Configure IPv6 Neighbors

The second step is to activate the BGP process for IPv6 and configure neighbors. This is performed in Stage 1 (on each core and distribution node). Route summarization has been performed where possible to facilitate connectivity and ease configuration complexity within the lab environment. Specific aggregate routes must be determined prior to an actual deployment.

1. Disable the IPv4 unicast address family default for the BGP routing process
    a. (config-router)# **no bgp default ipv4-unicast**
2. Configure IPv6 neighbors
    a. (config-router)# **neighbor** *ipv6-address remote-as autonomous-system-number*
3. Define neighbor source interface
    a. (config-router)# **neighbor** *ipv6-address* **update-source** *interface*
4. Specify the IPv6 address family and enter address family configuration mode.
    a. (config-router)# **address-family ipv6 unicast**
5. Activate each neighbor to exchange prefixes for the IPv6 address family with the local router.
    a. (config-router-af)# **neighbor** *ipv6-address* **activate**
    b. (config-router-af)# **neighbor** *ipv6-address* **next-hop-self**
    c. (config-router-af)# **neighbor** *ipv6-address* **send-community**
    d. (config-router-af)# **neighbor** *ipv6-address* **soft-reconfiguration inbound**
6. Define networks to be injected into the IPv6 BGP database (repeat this step for each network to be advertised).
    a. (config-router-af)# **network** *ipv6-address/prefix-length*

Initiate the OSPF version 3 Process and Configure IPv6 Neighbors

The third step is to activate the OSPF v3 process for IPv6 and configure networks to advertise. OSPF version 3 functions slightly different then the previous version in that the OSPF process must be configured on each interface to be advertised from. This is configured in Stage 1 (on each core and distribution node).

1. Enable OSPF for IPv6 on the interface (repeat this step on each interface to be advertised from).
    a. (config-if)# **ipv6 ospf** *process-id area area-id*
2. Enable OSPF for IPv6 routing.
    a. (config)# **ipv6 router ospf** *process-id*
3. Consolidate and summarize routes at an area boundary.
    a. (config-rtr)# **area** *area-id* **range** *ipv6-address/prefix-length*
4. Redistribute BGP through OSPF.
    a. (config-rtr)# **redistribute bgp** *process-id*

Configure EIGRP (or OSPF) for IPv6

The final step is to configure EIGRP, or OSPF, for IPv6 on the remaining AS, DW, and AR nodes. Because EIGRP was not yet available for IPv6 at the creation of this document, IPv6 for OSPF (OSPFv3) will be addressed only. The configuration steps taken will be similar to EIGRP; however not exact, and must therefore be examined prior to any actual migration plan. This is configured in Stage 1 and 2 (on the distribution and regional nodes).

This implementation could provide an appropriate window for eliminating EIGRP on the network entirely; greatly reducing the potential for a future WAN reconfiguration. This would enable to VA to host a fully open standard routing protocol backbone.

The following configuration is similar to the one applied above for the core, DC, and AS nodes however it will utilize a different OSPF area.

1. Enable OSPF for IPv6 routing.
    a. (config)# **ipv6 router ospf** *process-id*
2. Consolidate and summarize routes at an area boundary.
    a. (config-rtr)# **area** *area-id* **range** *ipv6-address/prefix-length*
3. Enable OSPF for IPv6 on the interface (repeat this step on each interface to be advertised from).
    a. (config-if)# **ipv6 ospf** *process-id area area-id*

For the purposes of connectivity testing, a Windows XP laptop (already IPv6 compatible) was integrated on the lab network. This device was installed behind the Access Layer to simulate a client machine, inside a medical center, traversing the VA WAN. End-to-end IPv6 ICMP echo requests were sent from the end host to IPv6 router interfaces. Sample results are displayed below:

```
C:\WINDOWS\System32\cmd.exe
Microsoft Windows XP [Version 5.1.2600]
(C) Copyright 1985-2001 Microsoft Corp.

C:\Documents and Settings\Lab02>ipconfig

Windows IP Configuration


Ethernet adapter Local Area Connection:

        Connection-specific DNS Suffix  . :
        Autoconfiguration IP Address. . . : 169.254.109.146
        Subnet Mask . . . . . . . . . . . : 255.255.0.0
        IP Address. . . . . . . . . . . . : fe80::20d:60ff:fed0:609c%4
        Default Gateway . . . . . . . . . : fe80::213:80ff:fe50:8370%4
                                            fe80::20f:24ff:fe27:4e00%4

Tunnel adapter Automatic Tunneling Pseudo-Interface:

        Connection-specific DNS Suffix  . :
        IP Address. . . . . . . . . . . . : fe80::5efe:169.254.109.146%2
        Default Gateway . . . . . . . . . :

C:\Documents and Settings\Lab02>
```

Fig 8. Client IPv6 Configuration

```
C:\WINDOWS\System32\cmd.exe

C:\Documents and Settings\Lab02>ping 2005::142a:fe1c

Pinging 2005::142a:fe1c with 32 bytes of data:

Reply from 2005::142a:fe1c: time<1ms
Reply from 2005::142a:fe1c: time<1ms
Reply from 2005::142a:fe1c: time<1ms
Reply from 2005::142a:fe1c: time<1ms

Ping statistics for 2005::142a:fe1c:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 0ms, Average = 0ms

C:\Documents and Settings\Lab02>
```

Fig 9. ICMPv6 Echo Request Sample Result

Testing conducted in this lab environment is not suitable to final implementation in a production environment. These tests were intended as a baseline tests for IPv6 IOS versions on the tested platforms and to perform limited connectivity and routing tests from a single client workstation. Further interoperability testing will be required covering far greater depth in associated technologies, protocols, hardware platforms and applications. Specific testing tools that test against the existing and developing IPv6 related RFC standards have been widely used in early IPv6 research network testing (i.e. DREN, Abilene, and MoonV6). These tests have been supported and made possible from vendors such as Agilent (www.agilent.com) and Ixia (www.ixiacom.com) whose testing products cater to this evolving technology. Further required testing initiatives will be covered in detail in the One-VA IPv6 Test Plan.