

The Fiscal Year 2008 Department of Defense

Internet Protocol Version 6

Test and Evaluation Report



September 2008

**Assistant Secretary of Defense for Networks and Information Integration/
Department of Defense Chief Information Officer**

UNCLASSIFIED

THIS PAGE INTENTIONALLY LEFT BLANK.

The Fiscal Year 2008 Department of Defense
Internet Protocol Version 6
Test and Evaluation Report

This report is provided in response to Section 221 of the National Defense Authorization Act for Fiscal Year 2006 (Public Law 109-163). The report provides assessments of the test and evaluation results that the Department of Defense (DoD) Components have submitted to the DoD for the period July 2007 through June 2008, and integrates these assessments with the results previously reported by the DoD to Congress. The assessments follow the processes and methodologies of the test and evaluation strategy set forth in the DoD Internet Protocol Version 6 Master Test Plan Version 2.0.

Approved by: 

John G. Grimes
Assistant Secretary of Defense for
Networks and Information Integration/
DoD Chief Information Officer

Dated: 22 August 2008

Approved by: 

Dr. Charles E. McQueary
Director
Operational Test and Evaluation

Dated: 18 August 2008

THIS PAGE INTENTIONALLY LEFT BLANK.

Table of Contents

Executive Summary	1
1 Introduction	3
1.1 Purpose	3
1.2 Test and Evaluation Objectives	3
1.3 Scope	4
1.4 FY 2005 - FY 2007 Cumulative Results and Recommendations	4
2 FY 2008 IPv6 Test and Evaluation Results	5
2.1 Overview	5
2.2 Cumulative Analysis Methodology	5
2.3 Impact of FY 2008 Test and Evaluation Reports on Demonstration of Joint Staff IPv6 Operational Criteria	8
2.3.1 Criterion 1: Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IPsec, and integration with firewalls and intrusion detection systems	9
2.3.2 Criterion 2: Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment	11
2.3.3 Criterion 3: Demonstrate equivalent to, or better performance than, IPv4 based networks	13
2.3.4 Criterion 4: Demonstrate voice, data, and video integration	15
2.3.5 Criterion 5: Demonstrate effective operation in low-bandwidth environment	16
2.3.6 Criterion 6: Demonstrate scalability of IPv6 networks	17
2.3.7 Criterion 7: Demonstrate support for mobile terminals (voice, data, and video)	18
2.3.8 Criterion 8: Demonstrate transition techniques	19
2.3.9 Criterion 9: Demonstrate ability to provide network management of networks	21
2.3.10 Criterion 10: Demonstrate tactical deployability and ad hoc networking	22
3 FY 2008 Conclusions	23
4 Recommendations	29
5 Summary	33
Appendix A - References	35
Appendix B - Terms and Definitions	37
Appendix C - Acronym List	41
Appendix D - DoD IPv6 2008 Test Report Summaries	47

List of Tables

Table 2-1 Cumulative Test and Evaluation Matrix	7
Table 2-2 Joint Staff IPv6 Operational Criterion 1 Status.....	9
Table 2-3 Joint Staff IPv6 Operational Criterion 2 Status.....	11
Table 2-4 Joint Staff IPv6 Operational Criterion 3 Status.....	13
Table 2-5 Joint Staff IPv6 Operational Criterion 4 Status.....	15
Table 2-6 Joint Staff IPv6 Operational Criterion 5 Status.....	16
Table 2-7 Joint Staff IPv6 Operational Criterion 6 Status.....	17
Table 2-8 Joint Staff IPv6 Operational Criterion 7 Status.....	18
Table 2-9 Joint Staff IPv6 Operational Criterion 8 Status.....	19
Table 2-10 Joint Staff IPv6 Operational Criterion 9 Status.....	21
Table 2-11 Joint Staff IPv6 Operational Criterion 10 Status.....	22
Table D-1 2008 T&E Reports and Related Operational Criteria.....	48
Table D-2 JCS 4 Equipment Configuration.....	55
Table D-3 Cisco Test Equipment Configuration.....	58
Table D-4 Cisco Test Results.....	59
Table D-5 Equipment Configuration.....	63
Table D-6 Device Configuration.....	67
Table D-7 Cisco IPv4/IPv6 Combined Device Results.....	68
Table D-8 Performance Comparison Table.....	82
Table D-9 Test Equipment Configuration.....	84
Table D-10 Convergence Results.....	85
Table D-11 MP-BGP Results.....	85
Table D-12 Network Access Points Test Results.....	85
Table D-13 Network Equipment Configuration.....	87
Table D-14 Test Equipment Configuration.....	88
Table D-15 HTTP IP Network Ratio Comparison Results.....	89
Table D-16 Workstation and Server IPv4/IPv6 Comparison Results.....	90
Table D-17 Network Management Results.....	97
Table D-18 Ambriel Configuration.....	98
Table D-19 Ambriel Technologies Interoperability Status Summary.....	99

Table D-20	Certification of TechGuard PoliWall Test Configuration.....	101
Table D-21	Test Results for Functional Test Category.....	101
Table D-22	Quantum Configuration.....	103
Table D-23	Quantum Test Results for Functional Test Category.....	103
Table D-24	IBM Storage System Tape Library Configuration.....	105
Table D-25	IBM Storage System Tape Library Test Results for Functional Test Category.....	105
Table D-26	Cisco Catalyst 4510R Layer 3 Switch Configuration.....	107
Table D-27	Cisco Layer 3 Switch Test Results for Functional Test Category.....	107
Table D-28	Cisco Catalyst 6506-E Layer 3 Switch Configuration.....	108
Table D-29	Cisco Catalyst 6506-E Layer 3 Switch Results for Functional Test Category.....	109
Table D-30	Cisco 2800/7600 and 3800/7600 Integrated Services Router Configuration.....	111
Table D-31	Cisco 2800/7600 and 3800/7600 Test Results for Functional Test Category.....	112
Table D-32	Datatek Configuration.....	114
Table D-33	Datatek Technologies Test Results for Functional Test Category.....	114
Table D-34	Assessment Report for Evaluating Milestone Objective 2 IPv6 to IPv4 Architecture Enabled Applications and Services.....	119
Table D-35	Ethernet Switch Test Results.....	122
Table D-36	SuSE Test Results for Functional Test Category.....	137
Table D-37	Red Hat Test Results for Functional Test Category.....	139
Table D-38	Microprocessor Library Definition Required Tests.....	141
Table D-39	Test Configuration Hardware and Software.....	145
Table D-40	Sun Microsystems Test Results for Functional Test Category.....	145
Table D-41	IPv6 Transition Mechanism Test Report Equipment List.....	148
Table D-42	Hardware Software Configuration for Microsoft Windows ISATAP Test.....	160

List of Figures

Figure D-1	IPv6 Tunnel Broker Transition Test Diagram.....	155
------------	---	-----

THIS PAGE INTENTIONALLY LEFT BLANK.

Executive Summary

This report is provided in response to Section 221 of Public Law 109-163. It is based on field tests, exercises, demonstrations, experiments, simulations, and analyses conducted by Department of Defense (DoD) Components over the last five years, with emphasis on the most recent year test results (July 2007 through June 2008). This report is an update to last year's report submitted to Congress on September 14, 2007.

The DoD Internet Protocol Version 6 (IPv6) Transition Office (DITO) established a repository of IPv6 Test and Evaluation (T&E) reports provided by DoD Components in response to requests from the Assistant Secretary of Defense for Networks and Information Integration/DoD Chief Information Officer (ASD(NII)/DoD CIO). The data contained in these reports have been evaluated with respect to the principal T&E objectives of the DoD IPv6 Master Test Plan version 2.0 (MTP v2.0). The Army, Navy, Air Force, National Security Agency (NSA), and Defense Information Systems Agency (DISA), henceforth referred to as DoD Components, have provided 141 reports. For Fiscal Year (FY) 2008, 39 reports were received from the DoD Components.

Following the guidance set forth in the DoD IPv6 MTP v2.0, the DoD Components have developed, conducted, and reported on T&E for their specific Joint Staff IPv6 Operational Criteria. The DITO facilitates the sharing of IPv6 T&E results among DoD Components and other federal IPv6 working groups through the Defense Knowledge Online (DKO) web portal. Based on a cumulative analysis of all related reports, four of the 10 Joint Staff IPv6 Operational Criteria have been successfully demonstrated. The four completed criteria are Interoperability (Criterion 2), Performance (Criterion 3), Scalability (Criterion 6¹), and Transition Techniques (Criterion 8).

This year, Criterion 2, 3, and 8 testing was considered to be completed with the demonstration of all functional sub-elements other than security. With the concurrence of ASD(NII)/DoD CIO, Director Operational Test and Evaluation (DOT&E), and Joint Staff, the security sub-elements will be demonstrated under Criterion 1. NSA will ensure that the intent of the security related sub-elements will be incorporated into the reports prepared under Criterion 1. Application transition techniques (decomposition 8.2), though feasible, are currently prohibited under the existing DoD IPv6 Information Assurance (IA) Milestone Objective (MO) guidelines and were deleted from Criterion 8.

Significant progress was made in the demonstration of Network Management (Criterion 9). Testing of the available network management tools has been completed. No single tool has the necessary capabilities to monitor, configure, and account for IPv6 network resources. Multiple tools are required to meet all the threshold requirements and most tools do not provide the capability to use IPv6 communications paths to manage the devices. When new commercial tools become available, further testing will be necessary to ensure the DoD can manage network assets using both Internet Protocol Version 4 (IPv4) and IPv6 communication paths.

¹ Demonstrated in fiscal year 2007.

The development and availability of critical, fully functional IPv6 Capable Products lag in some areas that affect the DoD's schedule for IPv6 T&E and deployment. At present, commercial implementation of IA devices has not been certified for DoD use. NSA continues to assess requirements for IA devices such as Intrusion Detection Systems (IDS), Intrusion Prevention Systems (IPS), firewalls, and High Assurance IP Encryptors (HAIPE) in support of Joint Staff IPv6 Operational Criteria for Security (Criterion 1). The first certified HAIPE devices are anticipated to be available in FY 2010.

In FY 2008, the DoD successfully demonstrated IPv6 capability on its Unclassified but Sensitive IP Router Network (NIPRNet), which was configured with dual stack routers. This included the ability to pass and receive IPv6 packets on the core backbone network, satisfying the Office of Management and Budget (OMB) Memorandum M-05-22. While the successful demonstration of IPv6 on NIPRNet is an important milestone, further security implementation guidance and certified IA devices must be available before enabling the core network. The decision to enable the DoD core networks will be supported by the successful demonstration of the remaining Joint Staff IPv6 Operational Criteria. Additionally, Congress directed the Chairman of the Joint Chiefs of Staff to provide certification that conversion of the DoD networks to IPv6 would "provide equivalent or better performance and capabilities than that which would be provided by any other combination of available technologies and protocols." The successful demonstration or approved disposition of the Joint Staff IPv6 Operational Criteria will support this certification.

1 Introduction

1.1 Purpose

The publication of the Fiscal Year (FY) 2008 Department of Defense (DoD) Internet Protocol Version 6 (IPv6) Test and Evaluation (T&E) Report is in response to Section 221 of Public Law 109-163. This report provides an assessment of IPv6 T&E activities carried out by the DoD Components with respect to the T&E objectives of the DoD IPv6 Master Test Plan version 2.0 (MTP v2.0). This report is also an input to the congressionally directed IPv6 certification by the Chairman of the Joint Chiefs of Staff. Although this is the final report required under the public law, IPv6 T&E activities will continue.

1.2 Test and Evaluation Objectives

The DoD IPv6 T&E Report provides consolidated test results and assessments in support of the DoD transition to IPv6, and identifies what is completed and what T&E is still required. Assessment of the individual IPv6 T&E reports furnished by the DoD Components will address the progress in meeting the objective of demonstrating the functionality of IPv6 as delineated in the Joint Staff IPv6 Operational Criteria.

The Joint Staff enumerated 10 operational criteria to be demonstrated in support of the DoD's transition of its networks to IPv6. These criteria provide the top-level operational and technical capabilities necessary to verify that IPv6 fulfills the needs of the DoD. Each criterion was decomposed to provide two subordinate levels of measurable and verifiable functional elements that allow demonstration through T&E:

- Level 1 decomposition identifies capabilities required for each criterion.
- Level 2 decomposition identifies the specific technology, infrastructure, and/or functionality to demonstrate Level 1 decomposition.

Responsibility for Level 1 and Level 2 decomposition elements, as well as further decomposition levels associated with each Joint Staff IPv6 Operational Criteria, has been distributed among the DoD components, as outlined in the DoD IPv6 MTP v2.0.

Additionally, Congress directed the Chairman of the Joint Chiefs of Staff to provide certification that conversion of DoD networks to IPv6 would “provide equivalent or better performance and capabilities than that which would be provided by any other combination of available technologies and protocols.” The successful demonstration or approved disposition of the Joint Staff IPv6 Operational Criteria will support this certification.

1.3 Scope

The scope of analysis in this report is limited to T&E reports submitted by DoD Components in response to requests from the Assistant Secretary of Defense for Networks and Information Integration/DoD Chief Information Officer (ASD(NII)/DoD CIO) in a memorandum dated March 11, 2008. The DoD received 39 reports from the Components during FY2008. The evaluation team for this report consisted of the Defense Information Systems Agency (DISA) Joint Interoperability Test Command (JITC) and Director Operational Test and Evaluation (DOT&E) representatives. This report provides the results of analyses for the 39 reports and integrates the analyses with the 102 previously submitted reports to provide a cumulative status for IPv6 T&E.

1.4 FY 2005 - FY 2007 Cumulative Results and Recommendations

The FY 2007 report provided the cumulative results from all the previous reports. Based on the T&E results over the last three years, it was determined that Scalability (Criterion 6) had been fully demonstrated for transition to IPv6. Interoperability (Criterion 2) and Performance (Criterion 3) were expected to be completely demonstrated during FY 2008 as well as elements of Transition Techniques (Criterion 8). The 2007 T&E report recommended that more experience using mixed Internet Protocol Version 4 (IPv4)/IPv6 networks in an operationally realistic environment was needed. There was considerable T&E for Security (Criterion 1) during this reporting period; however, commercial development and implementation of security devices/applications are still needed to demonstrate this criterion. Voice, Data, and Video Integration (Criterion 4) and Operation in Low-bandwidth Environment (Criterion 5) need technical guidelines, defined standards, and products to further demonstrate these criteria. Lastly, the 2007 T&E Report stated Mobility (Criterion 7), Network Management (Criterion 9), and Tactical Deployability and Ad Hoc Networking (Criterion 10) lacked development and implementation, resulting in limited T&E.

2 FY 2008 IPv6 Test and Evaluation Results

2.1 Overview

This section provides the overall status of DoD IPv6 T&E in 2008, in support of the DoD's transition to IPv6 and summarizes IPv6 T&E results reported by DoD Components for the period July 2007 through June 2008. There were 39 T&E reports analyzed for the current reporting period. Appendix D contains a summary for each report. Reports submitted for the current year had a greater focus on the demonstration of the Joint Staff IPv6 Operational Criteria than in previous years. Results indicate the following criteria have been successfully demonstrated²: Interoperability (Criterion 2), Performance (Criterion 3), Scalability (Criterion 6), and Transition Techniques (Criterion 8). All reports used for this analysis can be found on the DoD IPv6 (Unrestricted) Knowledge Center on the Defense Knowledge Online (DKO): <https://www.us.army.mil/suite/folder/11731042>³.

2.2 Cumulative Analysis Methodology

The cumulative status of each Joint Staff IPv6 Operational Criterion is provided in the cumulative T&E matrix (Table 2-1). This matrix is based on analysis of all applicable tests conducted by DoD Components. The status of each Joint Staff Operational Criteria is represented by a pie chart with slices colored red, yellow, or green. Each slice of a criterion's pie represents one Level 2 decomposition element for that criterion. The status color for each Level 2 element is based on analysis and evaluation of the test results for the underlying decomposition elements. Underlying decomposition elements that need additional T&E are easily identified.

The color-coded rating scale for the Level 2 decomposition elements is as follows:

- Red - Limited progress has been made. A red slice indicates a Level 2 decomposition element that has had little or no T&E, or for which existing T&E results are inconclusive or unsatisfactory. Significant T&E and/or development is needed.

- Yellow - Significant progress has been made. A yellow slice indicates a Level 2 decomposition element that has had considerable T&E and for which multiple, independent T&E have provided substantially similar, positive results. Some combination of additional analysis, testing, or development is needed.

- Green - Successfully demonstrated. A green slice indicates a Level 2 decomposition element that has been successfully demonstrated or that the decomposition has an approved

² Test and evaluation confirmed equivalent performance and capability or an approved disposition is in place for decompositions not demonstrated.

³ Access to the DKO requires a DoD Common Access Card (CAC) and registration with the DoD IPv6 (Unrestricted) Knowledge Center.

disposition. The evaluation type, relevance, and scope (considered with the number of tests) provide enough data to yield a high confidence factor.

Table 2-1 presents the total number of T&E reports applicable to each criterion for the entire transition effort, categorized by the evaluation method (counts for this reporting period are in parentheses). A comparison of the cumulative pie chart for 2007 to the cumulative pie chart for 2008 provides an indication of the progress made in FY 2008 for each criterion. The cumulative pie charts provide the proportion of each criterion at each status level. A cumulative pie chart that includes red slices indicates that the demonstration of the underlying functional or technical elements is incomplete. A cumulative pie chart that includes yellow with no red slices indicates that the underlying elements have had considerable progress. A cumulative pie chart that is all green indicates that all underlying elements for that criterion were fully tested and the criterion has been demonstrated. The expected completion date to fully demonstrate each criterion is also provided.

Table 2-1 Cumulative Test and Evaluation Matrix

Joint Staff IPv6 Operational Criteria		Test Methods							Cumulative Status Thru		Expected Completion Date
		Engineering Analyses	Modeling & Simulation	Experiments	Demonstrations	Pilots	Exercises	Field Tests	2007	2008	
1	Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IP security (IPsec), and integration with firewalls and intrusion detection systems	22 (1)	1	19 (4)	12 (4)	2	12 (1)	2 (2)			4 th QTR FY 2010
2	Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment	11	2	17 (2)	11 (2)	1	19 (1)	8 (7)			4 th QTR FY 2008
3	Demonstrate equivalent to, or better performance than, IPv4 based networks	2	2	10 (3)	12 (4)		8				4 th QTR FY 2008
4	Demonstrate voice, data, and video integration	6		2	3 (1)		4	1			4 th QTR FY 2010
5	Demonstrate effective operation in low-bandwidth environment	2	2		3 (1)		5				4 th QTR FY 2010
6	Demonstrate scalability of IPv6 networks	1		1	1	1					1 st QTR FY 2008
7	Demonstrate support for mobile terminals (voice, data and video)	5	1	1	3 (1)		7	1			4 th QTR FY 2010
8	Demonstrate transition techniques	16	4	23	24 (8)	2	25 (4)	7 (7)			4 th QTR FY 2008
9	Demonstrate ability to provide network management of networks	3		6	6 (1)						4 th QTR FY 2010
10	Demonstrate tactical deployability and ad hoc networking	7	1	2	3 (1)			1			4 th QTR FY 2010

Key:
 Successfully demonstrated
 Significant progress has been made
 Limited progress has been made

Quarter (QTR), Fiscal Year (FY) Total Events (Current Fiscal Year Events)
 Note: The pie chart for Criterion 8 differs from 2007 due to the change of Level 2 decomposition elements. Refer to Section 2.3.8 for more detail

2.3 Impact of FY 2008 Test and Evaluation Reports on Demonstration of Joint Staff IPv6 Operational Criteria

This section provides the evaluation of each Joint Staff IPv6 operational criterion at Level 1 and Level 2 of the decomposed functional or technical elements. The DoD Components responsible for each criterion recommended status changes based on testing performed this year. The evaluation team used the recommendations and test reports to determine the decomposition status.

The color-coded rating scale used in each criterion's decomposition status table is:

⊗ Red - Limited progress has been made. More T&E and/or development is needed to allow the decomposition item to be certified as having been demonstrated or T&E to date has not demonstrated satisfactory results.

⊕ Yellow - Significant progress has been made. Some portions of the decomposition element have not been successfully demonstrated or confidence in previous T&E results was low. Additional T&E and/or development is needed to allow the decomposition element to be certified as having been demonstrated.

✓ Green - The decomposition item has been successfully demonstrated or has an approved disposition. T&E has provided enough data to assure the decomposition element was demonstrated with a high confidence factor.

The rating symbol in the 2007 columns is the status reported in the 2007 T&E report for the Level 1 and Level 2 criterion decomposition. Rating symbols in the 2008 columns are the current status for each criterion. Specific T&E observations related to that criterion for 2008 follow each table.

The estimated completion date is the expected date that a Level 1 decomposition element will be satisfied. The responsible DoD Components provided the estimated completion dates.

2.3.1 Criterion 1: Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IPsec, and integration with firewalls and intrusion detection systems

Table 2-2 Joint Staff IPv6 Operational Criterion 1 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
1.1 Ensure that information is not disclosed to unauthorized persons, processes, or devices.	⊕	⊕	4 th Quarter FY 2010	1.1.1 Verify the implementation of IPsec with Encapsulating Security Protocol (ESP) in IPv6 hosts and routers. Verify integration with Public Key Infrastructure (PKI).	⊕	⊕
1.2 Ensure information received is the same as that which was sent (protect against unauthorized modification or destruction of information).	⊕	⊕	4 th Quarter FY 2010	1.2.1 Verify implementation of Authentication Header (AH) in IPv6 hosts and routers. Verify integration with PKI.	⊕	⊕
1.3 Ensure Authentication, Authorization, and Accounting (AAA) of persons and processes.	⊗	⊗	4 th Quarter FY 2010	1.3.1 Verify the implementation of an AAA server is able to ensure the Authentication, Authorization, and Accounting of persons, machines, and processes over an IPv6 network.	⊗	⊗
1.4 Ensure availability and mitigate denial of services (timely, reliable access to data, and information services for authorized users).	⊕	⊕	4 th Quarter FY 2010	1.4.1 Verify protection of the IPv6 stack of Hosts and Network Devices from intruders. (Note: Included in this are vulnerabilities that arise from errors in protocol specification or implementation or the associated device firmware).	⊕	⊕
				1.4.2 Demonstrate IPv6 traffic filtering capabilities of routers and firewalls according to security policies.	⊕	⊕
1.5 Ensure IPv6 traffic is interoperable with firewalls and Intrusion Detection Systems (IDS).	⊕	⊕	4 th Quarter FY 2010	1.5.1 Evaluate Firewalls and IDS functions that can be applied to IPv6 traffic. Evaluate Firewalls and IDS functions that can be applied to tunneled IPv6 traffic.	⊕	⊕
1.6 Ensure IPv6 traffic is interoperable with HAIPE devices.	⊗	⊗	4 th Quarter FY 2010	1.6.1 Evaluate HAIPE v3's ability to encrypt/decrypt IPv6 packets.	⊗	⊗

2008 T&E Observations Criterion 1

- Few products fully support IPv6 IPsec; however, vendors have implemented IPsec on some intermediate systems (i.e., routers).
(Test Reports D.8, D.9, D.16 through D.22; Decomposition 1.1.1)
- Client and server Certificate Authority (CA) certificates issued by two different Operating Systems (OS) proved the applicability of using PKI for both client and server based authentication over IPv6.
(Test Report D.8; Decomposition 1.1.1)
- The limited number of applications that underwent testing on the OS's (Red Hat Enterprise Linux 5.2 Server and Client, Novell SuSE Linux Enterprise Server 10, Microsoft Advanced Server 2008) met all required Request For Comments (RFCs) (4302, 4303, 4306, 4307) associated with ESP, AH and Internet Key Exchange version 2 (IKEv2) Protocol. Internet Key Exchange version 1 (IKEv1) is not interoperable with IKEv2 although some devices implement both standards for compatibility and interoperability.
(Test Reports D.16 through D.22; General Observations; Decompositions 1.1.1, 1.2.1)
- NSA testing of firewalls showed that the devices tested thus far do not provide IPv6 functionality when in transparent mode. Testing has been conducted on Juniper firewalls, and the results are favorable in supporting IPv6 functionality. However, results of that testing were not available for inclusion in this T&E report.
(Test Report D.5, General Observations; Decomposition 1.5.1)
- Although the High Assurance Internet Protocol Encryptor (HAIPE) version 3 specifications include IPv6 requirements, none were tested because DoD components are awaiting delivery of IPv6-capable HAIPE devices.
(Test Report D.5; Decomposition 1.6.1)
- Because of the lack of product availability, testing and certification of security products has been limited.
(General Observation)

2.3.2 Criterion 2: Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment

Table 2-3 Joint Staff IPv6 Operational Criterion 2 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
2.1 Demonstrate IPv4 application to IPv4 application over a mixed IPv4 and IPv6 network.	✓	✓	1 st Quarter FY 2008	2.1.1 Demonstrate core service interoperability: Domain Name System (DNS), directory services, File Transfer Protocol (FTP), email, web services, Network Time Protocol (NTP), and PKI.	✓	✓
				2.1.2 Demonstrate network core application interoperability: Voice over IP (VoIP) and video over IP.	✓	✓
				2.1.3 Demonstrate Commercial Off The Shelf (COTS) application interoperability (transaction, database access, and web services).	✓	✓
				2.1.4 Demonstrate Government Off The Shelf (GOTS) applications/systems interoperability.	✓	⊕ ⁴
2.2 Demonstrate IPv6 application to IPv4 application over a mixed IPv4 and IPv6 network.	⊕	✓	1 st Quarter FY 2008	2.2.1 Demonstrate core service interoperability: DNS, Directory, FTP, email, web services, NTP, and PKI.	⊕	✓
				2.2.2 Demonstrate network core application interoperability: VoIP and video over IP.	⊕	✓
				2.2.3 Demonstrate COTS application interoperability (transaction, database access, and web services).	⊕	✓
				2.2.4 Demonstrate GOTS application/system interoperability.	✓	⊕ ⁴
2.3 Demonstrate IPv6 application to IPv6 application over a mixed IPv4 and IPv6 network.	✓	✓	1 st Quarter FY 2008	2.3.1 Demonstrate core service interoperability: DNS, Directory, FTP, email, web services, NTP, and PKI.	✓	✓
				2.3.2 Demonstrate network core application interoperability: VoIP and video over IP.	✓	✓
				2.3.3 Demonstrate COTS application interoperability (transaction, database access, and web services).	✓	✓
				2.3.4 Demonstrate GOTS application/system interoperability.	⊕	⊕ ⁴

⁴ The change in status for Decomposition 2.1.4 and 2.2.4 is due to the change in the Department's approach to testing of GOTS applications and systems. Testing of GOTS applications and systems will be performed as IPv6 is implemented, vice in conjunction with COTS testing.

2008 T&E Observations Criterion 2

- Testing demonstrated the listed protocols as interoperable using Commercial Off-The-Shelf (COTS) equipment.
 - FTP (Get/Put)
 - Hypertext Transfer Protocol (HTTP)
 - Hypertext Transfer Protocol Secure (HTTPS)
 - Post Office Protocol version 3 (POP3)
 - Simple Mail Transfer Protocol (SMTP)
 - Simple Network Management Protocol (SNMP)
 - Lightweight Directory Access Protocol (LDAP)
 - Session Initiation Protocol (SIP)
 - DNS
 - G.711u VoIP
 - IP Television (IPTV) – Video, Audio(Test Reports D.11, D.13, D.15; Decompositions 2.1, 2.2, 2.3)

- Core services DNS, FTP, email, VoIP, and video over IP successfully interoperated in mixed IPv4 and IPv6 environments.
(Test Reports D.1, D.2, D.4, D.16, D.19, D.23; Decompositions 2.1, 2.2, 2.3)

- In a dual stack environment, a DNS server successfully responded to DNS queries from the host similar to the IPv4 cases. The server responded almost instantaneously to the DNS query (approximately 1ms).
(Test Report D.11; Decomposition 2.1)

- A router using Port Address Translation (PAT) translated incoming video packets' IPv6 source and destination addresses into IPv4 source and destination addresses. This resulted in reliable and high-quality video passing across the test network.
(Test Report D.4; Decomposition 2.2.)

- Responsibility for Information Assurance (IA) elements in each of the Level 2 decompositions is being transferred to the NSA.
(General Observation; Decomposition 2.1, 2.2, 2.3)

- No GOTS applications/systems were brought forward for interoperability testing during the nearly five years of testing.
(General Observation; Decompositions 2.1, 2.2, 2.3)

- All planned IPv6 interoperability T&E is considered complete.
(General Observation; Decompositions 2.1, 2.2, 2.3)

2.3.3 Criterion 3: Demonstrate equivalent to, or better performance than, IPv4 based networks

Table 2-4 Joint Staff IPv6 Operational Criterion 3 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
3.1 Demonstrate IPv6 throughput equivalent to or better than IPv4.			1 st Quarter FY 2008	3.1.1 Same as Level 1		
3.2 Demonstrate IPv6 latency equivalent to or better than IPv4.			1 st Quarter FY 2008	3.2.1 Same as Level 1		
3.3 Demonstrate IPv6 packet loss equivalent to or better than IPv4.			1 st Quarter FY 2008	3.3.1 Same as Level 1		
3.4 Demonstrate IPv6 service availability equivalent to or better than IPv4.			1 st Quarter FY 2008	3.4.1 Same as Level 1		

2008 T&E Observations Criterion 3

- Testing demonstrated equivalency between IPv4 and IPv6 combined throughput rates when the traffic was of a single protocol or when the traffic was split evenly between protocols. When traffic was split 90/10 or 10/90 IPv4/IPv6, inconsequential latency differences of an average of 1.43% were noted. (Test Report D.6; Decomposition 3.1.1)
- There was no appreciable difference between native IPv6 and dual stacked response times in networks with identical network configurations. (Test Report D.1; Decomposition 3.2.1)
- Throughput for the combined devices under test was identical for the two protocols with traffic levels evenly split between IPv4 and IPv6. Tests using identical frame sizes showed minor differences on specific devices that did not significantly affect intended operation of the device. (Test Report D.6; Decomposition 3.1.1)

- A single DNS server operated with performance degradation (~10% latency) in a dual stack network when responding to IPv6 DNS queries as compared to IPv4.
(Test Report D.11; Decomposition 3.1.)
- The results of the IPv6 over Multi-Protocol Label Switching (MPLS) test demonstrated that the usage of IPv6 in a dual stack environment does not affect performance when compared to the IPv4 baseline. When measuring maximum load and throughput, both protocols demonstrated nearly identical results. In each case, the throughput was close to line rate as expected.
(Test Report D.12; Decomposition 3.1.1)
- HTTP, SMTP, and Motion Picture Expert Group 2 (MPEG2) performance results demonstrated IPv4/IPv6 equivalency during end-to-end testing.
(Test Report D.13; Decompositions 3.1.1, 3.2.1, 3.3.1)
- In testing of three separate OS and hardware combinations, results indicate workstation and server performance parity between IPv4 and IPv6.
(Test Report D.13; Decomposition 3.3.1)
- Responsibility for IA elements in each of the Level 2 decompositions is being transferred to the NSA.
(General Observation; Decompositions 3.1, 3.2, 3.3, 3.4)
- All planned IPv6 performance T&E has been completed.
(General Observation; Decompositions 3.1, 3.2, 3.3, 3.4)

2.3.4 Criterion 4: Demonstrate voice, data, and video integration

Table 2-5 Joint Staff IPv6 Operational Criterion 4 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
4.1 Demonstrate simultaneous voice, data, and video (or any combination thereof) over shared IPv6 networks.	⊕	⊕	4 th Quarter FY 2010	4.1.1 Demonstrate Quality of Service (QoS) capabilities of IPv6 networks using Differentiated Services (DiffServ) and Resource Reservation Protocol (RSVP).	⊕	⊕
				4.1.2 Demonstrate transport control capabilities of IPv6 networks using Real Time Control Protocol (RTCP).	⊕	⊕
				4.1.3 Demonstrate session signaling capabilities of IPv6 networks using the Session Initiation Protocol (SIP).	⊕	⊕

2008 T&E Observations Criterion 4

- Testing included SIP, Reliable Transport Protocol (RTP) and Real Time Control Protocol (RTCP) protocols via an IPv6 network connection. It was found that these protocols are supported and they are effectively implemented for the support of real-time voice and video applications.
(Test Report D.2; Decompositions 4.1.2, 4.1.3)
- Testing demonstrated that voice transmission and video transmission using IPv6 was essentially equal to the same transmission via an IPv4 link in terms of quality and bandwidth consumption using RTCP and SIP.
(Test Report D.2; Decompositions 4.1.2, 4.1.3)
- Using RTCP and SIP, IPv6 achieved the same Mean Opinion Score (MOS)⁵ when compared to IPv4. The average voice and video scores were rated excellent.
(Test Report D.2; Decompositions 4.1.2, 4.1.3)

⁵ This system of testing calls for testers to watch or listen to and rate the transmission based on their opinion and was used as the scoring system for portions of this report.

2.3.5 Criterion 5: Demonstrate effective operation in low-bandwidth environment

Table 2-6 Joint Staff IPv6 Operational Criterion 5 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
5.1 Demonstrate ability to establish and maintain applications in low-bandwidth IPv6 environments.			4 th Quarter FY 2010	5.1.1 Demonstrate ability to establish and maintain applications (voice, data, video) in low-bandwidth IPv6 environments.		
				5.1.2 Demonstrate ability to maintain network operations (i.e., Network Management, DNS, Dynamic DNS, and Security) in low-bandwidth IPv6 environments.		

2008 T&E Observations Criterion 5

- Testing in specific low-bandwidth scenarios revealed an average increase of seven milliseconds (ms) packet latency between IPv4 only and dual stack enabled networks. Low-bandwidth data rates ranged from 64 to 1024 Kilobits per second (Kbps). (Test Report D.14; Decompositions 5.1, 8.1.1, 8.1.2, 8.1.3)
- Using automated test tools, testers completed 100% of the VoIP calls across the simulated Global Information Grid (GIG) network. The MOSs for this series of tests were determined to be identical. (Test Report D.13; Decomposition 5.1.1)
- Testing of transition techniques in a low bandwidth tactical environment included the evaluation of dual stack and various tunneling protocols. It was noted that as files sizes increased, throughput disparity between the protocols decreased. The difference in header size became less significant as packet sizes increased. (Test Report D.14; Decompositions 5.1.1, 3.1.1, 8.1.2)
- The bandwidth impact of the larger IPv6 header in low bandwidth environments has not been demonstrated. (General Observation)

2.3.6 Criterion 6: Demonstrate scalability of IPv6 networks

Table 2-7 Joint Staff IPv6 Operational Criterion 6 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
6.1 Demonstrate ability to add more network resources, services and users without negatively impacting existing users.	✓	✓	1 st Quarter FY 2008	6.1.1 Demonstrate the ability to build IPv6 networks comparable in size to existing IPv4 networks, with equal or better performance.	✓	✓
				6.1.2 Demonstrate the ability to populate IPv6 subnets with network elements of comparable numbers to existing IPv4 subnets, with equal or better performance.	✓	✓
				6.1.3 Demonstrate the ability to create IPv6 multicast sessions whose sizes are comparable to existing IPv4 multicast sessions, with equal or better performance.	✓	✓
				6.1.4 Demonstrate the ability to create IPv6 core services (DNS, Directory, FTP, email, Web, NTP, PKI) where the number of users are comparable to existing IPv4 core services, with equal or better performance.	✓	✓

2008 T&E Observations Criterion 6

- All planned IPv6 scalability T&E was completed and reported in the 2007 T&E Report.

2.3.7 Criterion 7: Demonstrate support for mobile terminals (voice, data, and video)

Table 2-8 Joint Staff IPv6 Operational Criterion 7 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
7.1 Demonstrate ability to establish and maintain IPv6 applications (voice, data, video) on the move.			4 th Quarter FY 2010	7.1.1 Demonstrate ability to initiate and maintain voice, data, or video applications using mobile terminals.		
				7.1.2 Demonstrate ability to maintain network operations of mobile terminals (i.e., Network Management, DNS, Dynamic DNS, and Security).		
				7.1.3 Demonstrate the ability to maintain connectivity of Mobile Nodes (MN) while On-The-Move (OTM) and network management of MN while OTM.		

2008 T&E Observations Criterion 7

- For a mobile airborne environment, testing showed minimal impact on throughput, latency, round trip time, and bit error rate shortly after the handover of the mobile host from one network to another. (Test Report D.14; Decomposition 7.1)
- Testing demonstrated increased capability in mobile node technology. Routers could incorporate Home Agent (HA) functionality and maintain connectivity during movement as well as while stationary. (Test Report D.14; Decomposition 7.1)
- Little operationally realistic testing has been attempted in tactical environments. (General Observation)

2.3.8 Criterion 8: Demonstrate transition techniques

Table 2-9 Joint Staff IPv6 Operational Criterion 8 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
8.1 Demonstrate DoD recommended network transition techniques.	⊕	✔	4 th Quarter FY 2010	8.1.1 Demonstrate the interoperability of IPv4 and IPv6 network transition techniques: <ul style="list-style-type: none"> Dual stack everywhere in an autonomous system Configured tunnels Tunnel Broker 	✔	✔
				8.1.2 Demonstrate the performance of IPv4 and IPv6 network transition techniques: <ul style="list-style-type: none"> Dual stack everywhere in an autonomous system Configured tunnels Tunnel Broker 	⊕	✔
				8.1.3 Demonstrate the security of IPv4 and IPv6 network transition techniques: <ul style="list-style-type: none"> Dual stack everywhere in an autonomous system Configured tunnels Tunnel Broker 	⊕	N/A ⁶

2008 T&E Observations Criterion 8

- Testing has demonstrated the interoperability and functionality of the dual stack, configured tunnels, and tunnel broker transition techniques. It has shown that these techniques are generally effective and secure. (Test Reports D30, D36, D39; Decompositions 8.1.1, 8.1.2, 8.1.3)
- Testing has shown that dual stacking creates the most flexible strategy. The coexistence of IPv6 with IPv4 is sufficiently stable to allow deployment of mixed networks. Performance degradation was minimal and should not affect the end user experience. (Test Reports D6, D30, D36; Decompositions 8.1.1, 8.1.2, 8.1.3)
- Not all transition techniques perform equally well in all circumstances. (Test Report D30; Decompositions 8.1.1, 8.1.2, 8.1.3)

⁶ Responsibility for Decomposition 8.1.3 is being transferred to the NSA.

- Dual IP stacks continue to exhibit stable coexistence and provide exceptional flexibility with acceptable impacts.
(General Observation; Decompositions 8.1.1, 8.1.2)
- As with IPv4 tunnels, IPv6 tunnel testing has shown the expected degradation in throughput, frame loss, and processor loading. The most noticeable difference was in the processor load on the routers.
(Test Report D.38; Decomposition 8.1.2)
- A report revealed that high volume traffic through tested routers could degrade performance, especially when using software-processing techniques. This performance degradation is nearly equal for IPv6 and IPv4 in dual stack networks. Processing in Application Specific Integrated Circuits performed better than software-based routing techniques.
(Test Report D.13; Decompositions 8.1.1, 8.1.2)
- For dual-stack traffic, IPv6 and IPv4 packets traversed the network at approximately the same rate, showing overall parity between the two protocols.
(Test Report D.13; Decomposition 8.1.2)
- Application transition techniques (Decomposition 8.2), though feasible, are currently prohibited under the existing DoD IPv6 IA MO guidelines. Consequently, the status of Decomposition 8.2., while reflected in the 2007 T&E report, has been deleted and is not reflected in the 2008 T&E report.
(Test Reports D.16, D.23; Decomposition 8.2)
- Responsibility for the IA element in the Level 2 decomposition is being transferred to the NSA, and therefore, 8.1.3 is no longer applicable to Criterion 8.
(General Observation; Decompositions 8.1.3)
- All planned IPv6 transition techniques T&E has been completed.
(General Observation; Decompositions 8.1.1, 8.1.2)

2.3.9 Criterion 9: Demonstrate ability to provide network management of networks

Table 2-10 Joint Staff IPv6 Operational Criterion 9 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
9.1 Demonstrate ability to monitor, configure, and account for IPv6 network resources.	✗	⊕	4 th Quarter FY 2010	9.1.1 Demonstrate that IPv6 devices can be monitored by Network Management Systems (NMS) commonly used by the DoD.	✗	✓
				9.1.2 Demonstrate that NMS commonly used by the DoD can configure IPv6 devices.	✗	✓
				9.1.3 Demonstrate that IPv6 devices can be accounted for by NMS commonly used by the DoD.	✗	⊕

2008 T&E Observations Criterion 9

- T&E reviewed a sampling of five network management tools and seven dual stack managed devices (network routers and server/client operating systems) commonly used in the DoD. Testing focused on the SNMP as the most widely used and accepted standard for network management. Testing revealed a heavy dependency on IPv4 network interfaces to communicate IPv4 and IPv6 related information. (Test Reports D.15, D.28; Decomposition 9.1)
- All of the tools tested supported the legacy SNMP protocols (SNMPv1 and SNMPv2), and most of the tools supported SNMPv3. (Test Report D.15; Decomposition 9.1)
- Only one of the tools tested could use IPv6 transport for SNMP communication. However, this tool could not send SNMP set requests. (Test Report D.15; Decomposition 9.1)
- Three of the five tools tested could use automatic discovery to identify clients. (Test Reports D.15, D.28; Decomposition 9.1.3)
- IPv6 polling consumed 30% more bandwidth than the comparable IPv4-only polling. This loss of efficiency is attributed to the greater IPv6 header size. (Test Reports D.15, D.28; Decomposition 9.1)

2.3.10 Criterion 10: Demonstrate tactical deployability and ad hoc networking

Table 2-11 Joint Staff IPv6 Operational Criterion 10 Status

Level 1 Decomposition (Capabilities to be demonstrated)	Cumulative Status Thru		Estimated Completion Date	Level 2 Decomposition (Specific technology/infrastructure/ functionality to be demonstrated)	Cumulative Status Thru	
	2007	2008			2007	2008
10.1 Demonstrate ability to move IPv6 networks as a whole, without reconfiguration.	✗	✗	4 th Quarter FY 2010	10.1.1 Demonstrate the ability to move networks to other locations while maintaining connectivity via the original IPv6 addresses, using Network Mobility (NEMO).	✗	✗
10.2 Demonstrate ability to support IPv6 networking without fixed router infrastructure.	✗	✗	4 th Quarter FY 2010	10.2.1 Demonstrate ability of IPv6 hosts to forward packets from peers, while on the move, using Mobile Ad hoc Networks (MANET) routing protocols.	✗	✗

2008 T&E Observations Criterion 10

- Testing successfully demonstrated the integration of Secure Neighbor Discovery with Network Mobility (NEMO). The autoconfiguration and neighbor discovery features in IPv6 enabled the warfighter to spread large numbers of sensors in the area of operations without manual configuration.
(Test Report D.10; Decomposition 10.1.1)
- Testing revealed difficulty assigning IPv6 addresses to nodes in a sensor network when using Mobile IPv6 (MIPv6). The IPv6 addresses were manually configured for nodes using the same HA server.
(Test Report D.10; Decomposition 10.1.1)
- The integration of NEMO with IP-enabled sensor networks provided a seamless integration in a Wide Area Network (WAN) infrastructure without a requirement for deploying proxies that convert between communication technologies.
(Test Report D.10; Decomposition 10.1.1)
- MANET as a technology was not included in any reported testing done this year.
(General Observation; Decomposition 10.2)

3 FY 2008 Conclusions

The following conclusions are based upon review and analysis of the 39 received reports for FY 2008. The DoD made significant progress in successfully demonstrating Joint Staff IPv6 Operational Criteria during this reporting period. Interoperability (Criterion 2), Performance (Criterion 3), and Transition Techniques (Criterion 8) are sufficiently mature and will support the Department's implementation of IPv6 and the Chairman's certification of equivalent performance and capability compared to other protocols. The lack of certified IPv6 Capable IA devices continues to hinder progress in Criterion 1 T&E, and affects the overall Department's planned transition and implementation of IPv6. The testing of IA requirements formerly in Criteria 2, 3, and 8 shall be transferred to NSA under Criterion 1.

Specific conclusions for the individual criterion are as follows:

Criterion 1: Demonstrate security of unclassified network operations, classified network operations, black backbone operations, integration of HAIPE, integration of IPsec, and integration with firewalls and intrusion detection systems.

- Testing of IPv6 IPsec attributes, AH, and ESP in network devices has demonstrated compliance with RFCs identified in Milestone Objective 2 version 2 (MO2v2), indicating that the technology is mature. Testing has demonstrated full IPsec capability in some routers, however IPsec attributes are not consistently applied across all products.
- PKI can be used for both client and server based authentication over IPv6.
- Testing successfully demonstrated IPv6 PKI implementation, however administrators must use IKEv2 in an IPv6 environment, due to the incompatibilities between IKE versions.
- IKEv1 and IKEv2 are not interoperable, but some devices can employ both versions. IKEv2 is the key exchange protocol of choice for any IPv6 enabled product requiring this attribute.
- Although more testing of IPv6 Capable firewalls is planned for this year, currently there are no firewalls that have passed NSA testing. Testing has been conducted on Juniper firewalls, and the results are favorable in supporting IPv6 functionality. However, results of that testing were not available for inclusion in this T&E report.
- Continued lack of IPv6 HAIPE devices is delaying demonstration of this criterion.
- The lack of IPsec implementation in vendor products indicates that IPsec has not been a high development priority, even though it is required by the applicable RFCs.

Criterion 2: Demonstrate end-to-end interoperability in a mixed IPv4 and IPv6 environment.

- Testing of COTS capabilities demonstrated that all the identified protocols are interoperable.
- Decompositions 2.1.4, 2.2.4, 2.3.4 call for GOTS testing. Further testing of GOTS applications and systems will be performed as IPv6 is implemented.
- IPv4 and IPv6 can coexist without adverse impact on network operations.
- T&E this reporting period demonstrated sufficient interoperability of network devices, services, and applications; hence, this criterion is considered satisfied.

Criterion 3: Demonstrate equivalent to, or better performance than, IPv4 based networks.

- Performance testing has indicated equivalence between IPv4 and IPv6.
 - Results demonstrate that throughput performance on native network configurations and combined protocol configurations are equivalent. Minor throughput differences were found on specific devices at specific frame sizes, but this did not significantly affect network operations. MPLS throughput measurements between IPv4 and IPv6 were identical.
 - Transmission Control Protocol (TCP) response times between the native IPv6 network and the dual stacked network demonstrated equivalency.
 - End-to-end testing showed that web page and email exchange, as well as video traffic, were equivalent.
- Latency testing revealed that differences in DNS response times between protocols were minimal and will be transparent to an end user. Using multiple DNS servers enhances performance.
- Testing has shown performance parity between IPv4 and IPv6; hence, this criterion is considered satisfied.

Criterion 4: Demonstrate voice, data, and video integration.

- Although some testing has been done, more testing is required of IPv6 applications and products using RTP, SIP, and specifically Assured Services SIP (AS-SIP) with the addition of RSVP.
- Interoperability testing of voice and video protocols has shown that IPv6 supports both technologies and that they can successfully transit dual stack networks.

- MOS testing on RTP and SIP voice and video transmissions resulted in “no” or “barely perceivable” differences.

Criterion 5: Demonstrate effective operation in low-bandwidth environment.

- Testing to date has demonstrated equivalent performance of IPv6 to IPv4 in low bandwidth environments. However, performance degradation was noted due to increased IPv6 header size. These effects are partially offset by an increase in processing efficiency due to the fixed-length IPv6 header (as opposed to the variable-length IPv4 header). Additionally, improved throughput performance was observed with the increased packet sizes available in IPv6.
- Multiple simulated VoIP calls were successfully completed, with network load traffic, on a limited bandwidth link.
- Voice, video, and data applications can successfully operate in low bandwidth IPv6 lab environments ranging from 64 bytes to 1500 Kbps.
- More technology development and testing are needed in low bandwidth environments.

Criterion 6: Demonstrate scalability of IPv6 networks.

- All planned T&E to support demonstration of this criterion was completed in FY 2007.

Criterion 7: Demonstrate support for mobile terminals (voice, data, and video).

- Testing revealed minimal performance impact during the handover of the mobile node from one network to another.
- Mobile node technologies are maturing, as demonstrated by a recent test that showed the improved capabilities of foreign agent and HA enabled routers.
- Tactical (battlefield) environments offer a number of challenges not commonly experienced by the standard industry application of this technology.
- More technology development and testing are needed with mobile terminals.

Criterion 8: Demonstrate transition techniques.

- Testing successfully demonstrated the interoperability and functionality of dual stack, configured tunnels, and tunnel broker transition techniques.
- Dual stacked transition technique: appears to create the most flexible strategy for the coexistence of IPv6 with IPv4; is sufficiently stable to allow deployment of mixed networks; and will enable legacy IPv4 dependent applications to continue operation.

- The network environment and mission requirements must be considered in selecting a transition mechanism. Not all mechanisms are expected to perform equally in all circumstances; regardless of performance, they may have certain advantages depending on the mission objectives.
- The application transition techniques (e.g., application translation, Bump in the Stack, Bump in the Application Programming Interface) outlined in the MO2v2 are not permitted on DoD networks by existing security guidance.
- Testing this reporting period demonstrated sufficient parity in transition techniques; hence, this criterion is considered satisfied.

Criterion 9: Demonstrate ability to provide network management of networks.

- Using available tools, it is possible to manage dual-stacked (mixed IPv4/IPv6) networks. However, a combination of two or more tools may be required and limitations may still exist.
- The most serious limitations of network management tools are:
 - Management of IPv6 devices must use IPv4 transport.
 - Available tools do not fully support IPv6 Management Information Bases (MIBs) as defined in RFC 2465 (in lieu of these MIBs, current tools use vendor- and device-specific MIBs to manage IPv6 functionality).
- Support for legacy SNMP protocols was shown in all tools tested.
- For those tools capable of using IPv6, reporting time for hosts was faster even though IPv6 used more bandwidth.
- In IPv6 native environments, successful performance of network management functions could not be consistently achieved.
- IPv4 is still required to provide full network management functionality.
- More technology development and testing are needed with network management tools and devices.

Criterion 10: Demonstrate tactical deployability and ad hoc networking.

- Improvements in mobile applications have been demonstrated (auto-configuration and multicasting protocols), but much work remains for development and T&E of the tactical deployability and ad hoc networking capabilities of IPv6.

- Mobility applications (NEMO and MANET) are in general an emerging technology. T&E for this criterion is dependent upon continued standards and mobile applications development.

THIS PAGE INTENTIONALLY LEFT BLANK.

4 Recommendations

The following recommendations are based on the T&E results, analyses, and DoD Components' input. These recommendations will support the Chairman, Joint Chiefs of Staff certification and assist in ensuring a smooth transition to IPv6 for the DoD.

Testing of IPv6 implementations in the areas of Interoperability (Criterion 2), Performance (Criterion 3), Scalability (Criterion 6), and Transition Techniques (Criterion 8) has shown that IPv6 protocol includes the required functionality and that some products are sufficiently mature to support limited operational use.

Recommendation 1: Sanction and resource operationally realistic use of IPv6 in large exercise environments. This will provide: visibility and experience with IPv6 for personnel outside the transition community; a venue for testing additional IPv6 functionality as it is developed; and a stable, long-term, easily accessible environment that can be used to test user-level applications.

Though the T&E of COTS applications does not indicate that there will be significant protocol issues, T&E of GOTS applications has not been accomplished. There have been no GOTS applications to date that require or use IPv6, nor have there been any stable, long-term, and easily accessible mixed network environments to use for testing such applications.⁷

Recommendation 2: Encourage deployment of IPv6 on operational networks in selected enclaves with operators who desire to experiment with IPv6 or who have a need that can be met by the base IPv6 protocol, such as a need for a larger address space or better aggregated hierarchical routing.

The DoD CIO has established a policy for requiring IPv6 Capable Products in acquisition programs. Adherence to this policy is evaluated in acquisition programs Information Support Plans (ISPs).

Recommendation 3: Enforce acquisition programs to include language in acquisition documentation and contracts for IPv6 capability.

Performance of IPv6 (with bandwidth of 1Mbs or higher) has been demonstrated to be equivalent to that of IPv4. Effective operation of IPv6 in low bandwidth environments (Criterion 5) has not yet been fully demonstrated below 1Mbs.

Recommendation 4: Concentrate future low-bandwidth performance testing on line-of-sight and satellite links. These links are an important part of the DoD's strategic and tactical networks, but remain largely untested.

⁷ This is excepting the Defense Research and Engineering Network (DREN) because it is not accessible within an MO2 enclave environment.

Recommendation 5: Continue testing in low-bandwidth environments representative of operational tactical networks.

To date vendor IPv6 implementations have focused on the basic functionality required to generate, accept, forward, and process IPv6 packets. Future development and T&E is required for network management tools, IA products, and devices. A full suite of IA products, tools, and policies is required before IPv6 can be implemented DoD-wide.

Recommendation 6: Require full IPsec functionality in all products procured by the DoD as appropriate to the individual product class.

Recommendation 7: Acquire pre-production HAIPEv3 devices, conduct beta T&E in mixed IPv4/IPv6 and native IPv6 environments, and provide performance and interoperability feedback to vendors.

Recommendation 8: Perform vulnerability analysis, and formulate mitigation and configuration guidance for IPv6 implementations (e.g., MO guidance).

Recommendation 9: Continue IPv6 IA, performance, and interoperability T&E efforts for routers, switches, and security products.

Recommendation 10: Develop and test IPv6-capable AAA and the PKI infrastructure within the DoD.

Recommendation 11: Encourage vendors to accelerate production of IPv6 Capable IA devices.

Recommendation 12: Transfer responsibility for assessing IA elements of criterion 2, 3, and 8 to NSA.

Recommendation 13: Transfer responsibility for interoperability and IA certification of IPv6 Capable security devices to DISA (JITC).

Network management functionality is gradually improving as vendors iterate through their products' lifecycles. However, current capabilities provide network management only through the dual stack phase of IPv6 transition, and IPv6-only management will eventually be necessary as IPv4 is eliminated from DoD networks.

Recommendation 14: Stress to vendors the need for greater IPv6 functionality in network management tools and in network devices, appliances, and software.

Recommendation 15: Network management testing should be a key objective during large exercises to demonstrate Network Management (Criterion 9). These exercises would allow testing in operational environments and expose the tools to systems that go beyond the challenges offered in a laboratory

setting. No further testing is directly required by the Air Force for Criterion 9.

The remaining criteria: Integration of Voice, Data, and Video (Criterion 4); Support for Mobile Terminals (Criterion 7); and Tactical Deployability and Ad-hoc Networking (Criterion 10) still require significant development and T&E.

Recommendation 16: Identify use cases and mission threads, and utilize large exercises to focus on these criteria as key testing objectives. These exercises would allow testing in operational environments with systems that go beyond the challenges offered in a laboratory setting.

Recommendation 17: Encourage vendors to develop and improve IPv6 functionality and performance for integrated voice, video, and data capabilities, and to support mobile terminals, tactical deployability, and ad-hoc networking.

The DoD continues to minimize duplicative testing. To conserve limited testing resources, DoD components should collaborate and identify objectives that can be satisfied in joint warfighter operational exercises to support the Chairman, Joint Chiefs of Staff certification.

Recommendation 18: Charter a tiger team led by the Joint Staff with support from ASD(NII)/DoD CIO, DOT&E, U.S. Joint Forces Command, and the DITO to identify and prioritize those criteria that require further testing. The tiger team should identify venues for operationally realistic testing to support the Chairman's certification of IPv6 performance and capability parity with IPv4.

THIS PAGE INTENTIONALLY LEFT BLANK.

5 Summary

The current state of the IPv6 products and services does not support full implementation DoD-wide at this time. T&E activities to date have demonstrated that vendor devices, operating systems, and network services do not fully support network requirements. Basic features required to enable information exchange using IPv6 are mature and suitable to enable basic connectivity, though many are not optimized. Advanced protocol features, where available, are inconsistently applied.

Important steps have been made in implementing IPv6 in the DoD. Four criteria (2, 3, 6, and 8) have been successfully demonstrated to date. The lack of IPv6 Capable IA products and HAIPE devices delays enterprise-wide implementation of IPv6. Although the IPv6 protocol is sufficiently mature, IPv6 implementations in software and hardware devices is lacking.

As new IPv6 Capable products and services are developed, further T&E will be required to assess interoperability, performance, and scalability. Successful implementation of IPv6 by DoD will require basic and advanced IPv6 protocol features and IA capabilities that do not currently exist. Further research, development, and testing are necessary to ensure that the DoD's networks can transition without affecting mission critical operations. Full implementation of IPv6 is dependent upon further development of standards, applications, services, and products by commercial industry.

THIS PAGE INTENTIONALLY LEFT BLANK.

Appendix A - References

- Public Law 109-163 National Defense Authorization Act for Fiscal Year 2006, January 6, 2006.
<http://www.defenselink.mil/dodgc/olc/docs/PL109-163.pdf>
- Public Law 108-375 National Defense Authorization Act for Fiscal Year 2005, October 28, 2004.
<http://www.defenselink.mil/dodgc/olc/docs/PL108-375.pdf>
- Department of Defense (DoD) Internet Protocol Version 6 (IPv6) Master Test Plan version 2.0 (MTP v2.0), September 2006.
<https://www.us.army.mil/suite/doc/8958812>
- DoD IPv6 Generic Test Plan version 3 (GTPv3), August 2007.
<https://www.us.army.mil/suite/doc/9523305>
- DoD Deputy CIO Memorandum, DoD IPv6 Definitions, June 26, 2008.
<https://www.us.army.mil/suite/doc/11706660>
- DoD Information Technology Standards Registry.
<https://disonline.disa.mil/>
- DITO IA Guidebook Version 1-1.
<https://www.us.army.mil/suite/doc/7253350>

THIS PAGE INTENTIONALLY LEFT BLANK.

Appendix B - Terms and Definitions

Demonstration: Testing that is limited to a combination of related, perhaps interdependent, features or functions. It is usually an ordered sequence of tasks and is restricted from any operational network traffic.

Engineering Analysis: Category of testing based on engineers' previous experience with the technology, as well as use of equipment specifications to speculate about the performance or capability.

Exercise: Environment is a functional, operationally realistic network with controlled traffic and realistic loading. The test administrators and users are sympathetic to IPv6. Tests are focused on network and communications testing, perhaps with some training goals. This includes automated test generators running scripted test cases a large number of times. The test is well defined and of a limited duration.

Experiment: Testing that consists of a scope that is restricted to a single question or theory with a test network isolated from operational network traffic. Few repetitions of test cases and a limited number of participants are involved.

Field Test: Testing that uses an operationally-realistic network with common protocol traffic and assumed loading conditions. Focus is on the devices or systems operating within the environment in which it is deployed. A well-defined, limited duration is set for testing.

IPv6 Base Requirements: Requirements that are mandated for each specific device type in the IPv6 product profile in the DoD Information Technology Standards Registry (DISR).

IPv6 Capable Product: Products (whether developed by commercial vendor or the government) that can create or receive, process, and send or forward (as appropriate) IPv6 packets in mixed IPv4/v6 environments. IPv6 Capable Products shall be able to interoperate with other IPv6 Capable Products on networks supporting only IPv4, only IPv6, or both IPv4 and IPv6 and shall:

- Conform to the requirements for the DoD IPv6 Standards Profiles for IPv6 Capable Products document contained in the DISR.
- Possess a migration path and/or commitment to upgrade from the developer (company Vice President, or equivalent, letter) as the IPv6 standards evolve.
- Ensure product developer IPv6 technical support is available.
- Conform to National Security Agency (NSA) and/or Unified Cross Domain Management Office requirements for Information Assurance (IA) and products.

IPv6 Generic Test Plan Version 3 (GTPv3): A plan developed to specify conformance, interoperability, and performance procedures that IPv6 products must successfully complete

in order to be certified for interoperability by the Defense Information Systems Agency (DISA) Joint Interoperability Test Command (JITC).

<https://www.us.army.mil/suite/doc/5997794>

Joint Staff IPv6 Operational Criteria: Criteria that must be successfully demonstrated to support a decision to initiate DoD transition to IPv6 and identify key operational and technical capabilities at a high level.

Milestone Objective 1 (MO1): DoD Components are authorized to implement and operate IPv6 within an enclave. At MO1, the evaluation of the IPv6 protocol is sufficient, and the policy, procedures, and technical guidance have been developed to authorize DoD Components to operate in a single network domain or enclave environment within operational networks. The single domain or enclave requires strict access controls be maintained under a single administrative authority for IA and security policy. Information flow will be tightly controlled to prevent IPv6 packets from entering or leaving the domain. The border device shall not translate nor permit the transit of native or tunneled IPv6 packets. MO1 allows the use, familiarization, and testing of IPv6 protocol and applications to ascertain issues and derive migration strategies for this new protocol. MO1 was authorized as of October 1, 2005.

Milestone Objective 2 (MO2): DoD Components are authorized to implement and operate IPv6 across cooperative domain boundaries. At MO2, the policies, procedures, and technical guidance have been developed to expand the operation of IPv6 across cooperative domain boundaries, but limited to within DoD networks (no internet exchange of IPv6 packets, native or tunneled). MO2 will provide the ability to evaluate the scalability and further evaluate the IPv6 IA implications using tunneling and native IPv6 routing, as available. IPv6 traffic, which crosses cooperative domain boundaries, must be approved in accordance with the Defense Information Systems Network (DISN) connection-approval process to ensure compliance with IA policies. Multiple certification and accreditation authorities may be involved in MO2. MO2 permits applications to test IPv6-specific end-to-end capabilities and routing schema efficiencies. Limiting operation to within the DoD and only at approved locations reduces risk to IA and operational impacts on existing IPv4 networks. MO2 was authorized as of October 1, 2006.

Milestone Objective 3 (MO3): DoD Components are authorized to implement and operate IPv6 enterprise-wide. At MO3, policy, planning, and technical transition guidance will be provided to allow tunneled and native IPv6 traffic to exist on DoD operational networks. DISN and DoD Component core IP infrastructures are authorized to accept, route, and process IPv6 protocol traffic while maintaining interoperability with IPv4. Boundary protection and other security mechanisms to assure IA requirements shall be available and implemented to protect the DISN. MO3 permits applications and data owners to complete operational transition to IPv6 with at least the same functionality (parity) as currently found in IPv4.

Mixed IPv4 and IPv6 Environment: A mixed IPv4 and IPv6 environment includes the situations of tunneling IPv4 over IPv6 native network, tunneling IPv6 over an IPv4 native network, providing protocol translation at various points, and dual-stack operation.

Modeling and Simulation (M&S): Testing that uses a completely virtual environment to predict system or network performance. Software is used to simulate all involved devices and protocols.

Pilots (i.e., Pilot Testing): Testing that uses a functional, operational network with a limited number of administrators and users, but is realistic for the size of the network. There is no set time limit in conducting pilots, and all traffic is non-scripted (routine traffic).

THIS PAGE INTENTIONALLY LEFT BLANK.

Appendix C - Acronym List

A	DNS A record for an IPv4 Address
AAA	Authorization, Authentication, and Accounting
AAAA	DNS AAAA record for an IPv6 Address
ACL	Access Control List
AFATDS	Advanced Field Artillery Tactical Data System
AFB	Air Force Base
AFIOC	Air Force Information Operations Center
AFRL	Air Force Research Laboratory
AFSN	Air Force System Networking
AH	Authentication Header
AIPTL	Advanced IP Technology Laboratory
ARP	Address Resolution Protocol
ASA	Adaptive Security Appliance
ASD	Assistant Secretary of Defense
AS-SIP	Assured Services-SIP
AT&L	Acquisition Technology and Logistics
BER	Bit Error Rate
BGP	Border Gateway Protocol
BIND	Berkeley Internet Name Domain
CA	Certificate Authority
CAC	Common Access Card
CDS	Cross Domain Solutions
CIO	Chief Information Officer
CLI	Command Line Interface
CONUS	Continental United States
COTS	Commercial Off-The-Shelf
CPU	Computer Processor Unit
CJCS	Chairman of the Joint Chiefs of Staff
DAA	Data Acquisition Agent
DCP-ETSI	Distribution and Communication Protocol-European Telecommunications Standard Institute
DFS	Data Fusion Server
DHCP	Dynamic Host Configuration Protocol
DHCPv6	Dynamic Host Configuration Protocol version 6
DiffServ	Differentiated Services
DISA	Defense Information Systems Agency
DISN	Defense Information Systems Network
DISR	DoD IT Standards Registry
DITO	DoD IPv6 Transition Office
DKO	Defense Knowledge Online
DMZ	Demilitarized Zone

DNS	Domain Name System
DoD	Department of Defense
DoS	Denial of Service
DOT&E	Director, Operational Test and Evaluation
DREN	Defense Research and Engineering Network
DUT	Device Under Test
EIGRP	Enhanced Interior Gateway Routing Protocol
ERD	Electronic Report Distribution
ESP	Encapsulating Security Payload
FA	Foreign Agent
FTP	File Transfer Protocol
FW	Firewall
FY	Fiscal Year
GES	Ground Entry Sites
GIG	Global Information Grid
GN	Ground Node
GOTS	Government Off-The-Shelf
GRE	Generic Routing Encapsulation
GTP	Generic Test Plan
HA	Home Agent
HAIP	High Assurance Internet Protocol Encryptor
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
I3MP	Installation Information Infrastructure Modernization Program
IA	Information Assurance
ICMP	Internet Control Message Protocol
ICMPv6	Internet Control Message Protocol version 6
IDS	Intrusion Detection System
IE	Internet Explorer
IETF	Internet Engineering Task Force
IIAG	IPv6 Information Assurance Group
IIS	Internet Information Services
IKE	Internet Key Exchange
IOS	Internetwork Operating System
IP	Internet Protocol
IPS	Intrusion Prevention System
IPsec	IP security
IPTV	Internet Protocol Television
IPv4	Internet Protocol Version 4
IPv6	Internet Protocol Version 6
ISATAP	Intra-Site Automatic Tunneling Address Protocol

ISP	Information Support Plan
ISR	Integrated Services Router
IT	Information Technology
ITA	Information Technology Agency
JCAN	Joint Capability for Airborne Networking
JCS	Joint Chiefs of Staff
JIT	Joint Interoperability Tool
JITC	Joint Interoperability Test Command
JSTARS	Joint Surveillance Target Attack Radar Systems
JTEN	Joint Tactical Edge Networks
JUICE	Joint User Interoperability Communications Exercise
Kb	Kilobit
Kbps	Kilobits per second
L2	Layer 2
L3	Layer 3
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
M&S	Modeling and Simulation
MAC	Media Access Control
MANET	Mobile Ad hoc Networks
Mb	Megabit
Mbps	Megabits per second
μs	Microseconds
MIB	Management Information Base
MIP	Mobile IP
MN	Mobile Node
MO	Milestone Objective
MO1	Milestone Objective 1
MO2	Milestone Objective 2
MO2v2	Milestone Objective 2 version 2
MOB1	Main Operating Base 1
MOS	Mean Opinion Score
MP-BGP	Multiprotocol–Boarder Gateway Protocol
MPLS	Multi Protocol Label Switching
MPEG	Motion Picture Expert Group 2
MR	Mobile Router
MRD	Minimum Requirements Document
ms	milliseconds
MTP v2.0	Master Test Plan Version 2.0
MTU	Maximum Transmission Unit
NAP	Network Access Points

NAT-PT	Network Address Translation-Protocol Translation
NBMA	Non-Broadcast Multi-Access
NCOW	Net-Centric Operations Warfare
ND	Neighbor Discovery
NEMO	Network Mobility
NIDS	Network Intrusion Detection System
NII	Networks and Information Integration
NIPRNet	Unclassified but Sensitive Internet Protocol Router Network
NM	Network Management
NMI2	Network Management IPv6 Initiative
NM/OPS	NM Operations
NMS	Network Management Systems
NS	Name Server
NS	Neighbor Solicitation
NSA	National Security Agency
NTP	Network Time Protocol
OAM	Operation, Administration, and Maintenance
OC	Optical Carrier
OMB	Office of Management and Budget
OS	Operating System
OSPF	Open Shortest Path First
OSPFv3	Open Shortest Path First version 3
OTM	On The Move
PAT	Port Address Translation
PC	Personal Computer
PIC	Physical Interface Card
PKI	Public Key Infrastructure
PO	Participating Organization
POP3	Post Office Protocol version 3
PPP	Point-to-Point Protocol
PT	Port Translation
QFY	Quarter Fiscal Year
QoS	Quality of Service
RA	Router Advertisement
RF	Radio Frequency
RFC	Request for Comment
RHEL	Red Hat Enterprise Linux
RIM	Radio Interface Module
RIP	Routing Information Protocol
RO	Route Optimization
RSA	Rivest-Sharir-Adleman
RSVP	Resource Reservation Protocol

RTCP	Real Time Control Protocol
RTP	Reliable Transport Protocol
RTSP	Real Time Streaming Protocol
SATCOM	Satellite Communications
SDC	Standard Desktop Configuration
SDP	Service Delivery Points
SDP	Shelf Discovery Protocol
SEND	Secure Neighbor Discovery
SIIT	Stateless IP/Internet Control Message Protocol Translation
SIMR	Serial Interface to Military Radios
SIP	Session Initiation Protocol
SIPRNet	Secret Internet Protocol Router Network
SISTM	Simulator-Simulator
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SP	Service Pack
SPSS	Statistical Package for Social Sciences
STIG	Secure Technical Implementation Guide
STP	System Tracking Program
SUT	System Under Test
T&E	Test and Evaluation
TCP	Transmission Control Protocol
TDC	Theater Deployable Communications
TDM	Time Division Multiplexer
TEWG	Test and Evaluation Working Group
TIC	Technology Integration Center
TOC	Tactical Operation Center
TNT	Tactical Network Topology
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VLAN	Virtual Local Area Network
VoIP	Voice over IP
VPN	Virtual Private Network
WAN	Wide Area Network
WWW	World Wide Web

THIS PAGE INTENTIONALLY LEFT BLANK.

Appendix D - DoD IPv6 2008 Test Report Summaries

This appendix provides summaries for the 39 IPv6 Test and Evaluation (T&E) reports that DoD Components submitted for this reporting period (July 2007 through June 2008). The applicability of each report to the Joint Staff IPv6 operational criteria is summarized in Table D-1. The alphanumeric designator that precedes each report title in this table corresponds to the section number of the appendix that summarizes the report. Each report summary is comprised of the following eight elements: title, testing organization and publication date, summary, T&E method, relevant Joint Staff IPv6 operational criteria (including Level 1 and Level 2 decomposition relevancy), configuration, results, and conclusions/recommendations. Entries that summarize certifications contain a table that defines Requests For Comment (RFCs) found in the DoD Information Technology Standards Registry (DISR), which is available at <https://disronline.disa.mil>.

Table D-1 2008 T&E Reports and Related Operational Criteria

Section	Test Report Short Title	Joint Staff IPv6 Operational Criteria									
		1	2	3	4	5	6	7	8	9	10
D.1	IPv6 Dual Stack Transition Test Report		X	X					X		
D.2	JCS Criteria 4, Phase 2: 4.1.2.1, 4.1.2.2, 4.1.3.1, 4.1.3.3 Demonstration of the Real Time Protocol (RTP) and Session Initiation Protocol (SIP) Capabilities Over an IPv6 Network Test Report, v1.0			X	X						
D.3	Special Interoperability Test Certification of Cisco 1800, 2800, 3800, and 7200 Families of Routers	X	X						X		
D.4	TNT 07-4 AAR: IPv6 Testing with JITC		X								
D.5	Evaluation and Implementation of DISA IPv6 Information Assurance Guidance for Milestone Objective 2 version 2	X							X		
D.6	Cisco Networks Internet Protocol Version 6 Test Report			X					X		
D.7	Test of Internet Protocol Version 6 (IPv6) Configured Tunneling								X		
D.8	Net-Centric Operations Warfare (NCOW) IPv6 Demonstration: Security Features	X							X		
D.9	Net-Centric Operations Warfare (NCOW) IPv6 Demonstration: Security Features	X									
D.10	Net-Centric Operations Warfare (NCOW) IPv6 Demonstration	X							X		X
D.11	DNS IPv6 Test Plan and Report		X	X					X		
D.12	IPv6 Core Routing Test Plan and Report			X					X		
D.13	Joint Staff Internet Protocol Version 6 Operational Criterion 3 Test Report		X	X					X		
D.14	Demonstration of Operation of IPv6 in a Simulated Low Bandwidth Environment	X		X	X	X		X	X		
D.15	Technical Report For Network Management IPv6 Initiative (NMI2) (Tool Analysis)								X	X	
D.16	Special Interoperability Test Certification of Ambriel ATX-S Series IPv4/IPv6 Translator device		X						X		
D.17	Special Interoperability Test Certification of TechGuard PoliWall Version 1.21	X	X						X		
D.18	Special Interoperability Test Certification of Quantum Autoloader SuperLoader3 backup device		X						X		
D.19	Special Interoperability Test Certification of the IBM Storage System TS3100 Tape Library Express and IBM Storage System TS3200		X						X		
D.20	Special Interoperability Test Certification of Cisco Catalyst 4500 Family of Layer 3 Switches		X						X		

Table D-1 2008 T&E Reports and Related Operational Criteria (continued)

Section	Test Report Short Title	Joint Staff IPv6 Operational Criteria									
		1	2	3	4	5	6	7	8	9	10
D.21	Special Interoperability Test Certification of Cisco Catalyst 6500 Family of Layer 3		X						X		
D.22	Special Interoperability Test Certification of Cisco 2800 Integrated Services Router (ISR) Family of Routers	X	X						X		
D.23	Special Interoperability Test Certification of Datatek IPv4/IPv6 Translator		X						X		
D.24	Mobile IPv6 Implementation		X					X	X		
D.25	Assessment Report For Evaluating Milestone Objective 2 IPv6 To IPv4 Architecture		X						X		
D.26	2007 Ethernet Switch Comparison Report	X	X	X					X		
D.27	Transition Mechanisms Study AFATDS over IPv6								X		
D.28	Network Management IPv6 Initiative (NM12) (Client Analysis)		X						X	X	
D.29	Assessment Report For Evaluating Milestone Objective 2 Virtual Local Area Network Architecture	X	X						X		
D.30	Assessment Report for Evaluating MO2 ISATAP Architecture	X	X						X		
D.31	Assessment Report for Evaluating MO2 Microsoft Windows IPv6 to IPv4 Architecture	X	X						X		
D.32	Special Interoperability Test Certification of SuSE	X	X						X		
D.33	Special Interoperability Test Certification of Red Hat	X	X						X		
D.34	LOSSKNOT Section IV, Test Plan and Results	X	X						X		
D.35	Special Interoperability Test Certification of the Sun Microsystems SPARC T2000 and X86 V40z 32-bit and 64-bit Platforms Running Solaris 10	X	X						X		
D.36	IPv6 Transition Mechanism Test Report		X	X					X		
D.37	NIPRNet IPv6 Compliance Demonstration			X					X		
D.38	IPv6 Tunnel Broker Transition Test Report		X	X					X		
D.39	Assessment Report for Evaluating Milestone Objective 2 Microsoft Windows Intra-Site Automatic Tunnel Addressing Protocol		X						X		
Number of Test Reports Relevant to Each Joint Staff IPv6 Operational Criterion		16	27	11	2	1	0	2	36	2	1

D.1 IPv6 Dual Stack Transition Test Report

Testing Organization and Publication Date

Air Force Systems Networking (AFSN)
August 31, 2007

Summary

The AFSN conducted a study on the effects of using an IPv6 Dual Stack Transition mechanism on standard network equipment utilized to provide Wide Area Network (WAN) connectivity on the Unclassified Internet Protocol Router Network (NIPRNet) and Secret Internet Protocol Router Network (SIPRNet). The objective of this test is to evaluate the performance characteristics of a typical Air Force network architecture with dual stacked configurations.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.3, 2.3.1)

3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)

8 (8.1, 8.1.1, 8.1.2)

Configuration

The AFSN performed all tests in the Test and Integration Facility located in the same building. Testing included configurations for dual stack (IPv4 and IPv6) of both network equipment interfaces and routing protocols. Functionality and performance were evaluated by attempting to pass traffic over the test network set up in the Integration Facility. Traffic was generated using IPv6 and IPv4 addressing with the Spirent test device. Network equipment was evaluated for processor utilization, throughput, frame loss, latency, and average Transmission Control Protocol (TCP) Times, as well as functionality and other performance issues as pertinent to each respective type of equipment. After an initial baseline evaluation (with IPv4 traffic over an IPv4 network), traffic loads of 25% IPv6, 50% IPv6, 75% IPv6 and 100% IPv6 were tested. In addition, 100% IPv4 traffic over a stacked network and 100% IPv6 traffic over an all IPv6 network were evaluated.

The test network or System Under Test (SUT) consisted of three 7206 VXR Core routers representing a simulated DISA WAN and three base networks (Eglin, Tyndall and MacDill Air Force Base). Eglin had a setup resembling a future Block 30 or dual diversity/path configuration (with two Service Delivery Point (SDP) routers). The other two bases had architectures more closely resembling today's NIPRNet architecture (one SDP router and an External router).

Results

IPv4 Baseline Test

For the baseline test, the network was configured for IPv4 traffic/configuration only. Multiple IPv4 traffic flows were used across the network. No IPv6 configuration was used. Test results indicated negligible losses. Specifically, there was almost no loss during the throughput test for 128 byte frame sizes up to about 76% load. For loads of 76%-81%, there was some loss, but still less than 10%. For loads above 86% (still, for the 128 byte frame size), losses increased, but never went much above 30%. When the frame size increased to 256 bytes, there was virtually no loss with any load. This pattern of no loss continued with increasing frame sizes.

Dual Stack Baseline Test

For the baseline over a dual stacked network test, only IPv4 traffic over a network consisting of equipment running dual stacks of IPv4 and IPv6 (addresses on interfaces, routing protocol stacks, etc.) was used. Multiple IPv4 traffic flows were used across the network. The results were similar to the baseline testing on the all IPv4 network tested above during the initial baseline. While no or minimal losses were seen up to about 81% load, at this point there were some higher losses noted than in the initial baseline. For instance, in the original baseline test, no losses were seen over approximately 30%. However, in this stacked configuration test, some paths experienced losses much higher than 30%, even some as high as 80-85% for the higher loads on smaller frame sizes. However, as was the case above, for frame sizes over 256 bytes, there was virtually no loss with any load. Running 100% IPv4 over the stacked network configuration (with no IPv6 traffic present) caused little or no additional loss compared to the original baseline test.

In the first test over the dual stacked network, using 100% IPv6 traffic resulted in decreased throughput and frame loss numbers increased slightly more than the results that were seen when 75% IPv6 traffic was tested during other testing. At 128 byte frames with 91-96% loading, roughly one-fourth of the traffic paths experienced 100% loss. Again, as frame sizes increased, loss decreased. As was the case in all the other tests, once frame sizes of 512 bytes were reached, no more significant loss was experienced. Even at smaller frame sizes, losses were only significant with larger loads (for example, above 81% loading, or 81 Mbps on a 100 Mbps interface).

25% IPv6 Test

During this test, increased loss was noted in the smaller frame sizes. In particular, in the 128 byte frame size, some loss started occurring at about 50% loading. The loss continued to grow through 96% loading, with slightly increased loss compared to previous tests with no IPv6 traffic. Likewise, some additional loss was seen at 256 and 384 byte frame sizes that were not seen in previous tests. However, it can be summarized that these losses were negligible when compared to the baseline testing. Again, no significant loss occurred with larger frame sizes. The Avalanche testing showed no noticeable difference in the response times from the IPv4 baseline test. The 2000 and 2500 simulated users test was similar in response times. The first

Hypertext Transfer Protocol (HTTP) get request arrived around 220 ms; the first Acknowledgement (ACK) of response data arrived around 400 ms; and the connection closed around 6 secs. It was noticed that the test device could not handle anything over 2400 simulated users. At approximately 2400 simulated users, the device stopped running the test and the errors increased exponentially. This caused the average response times to be higher for the 5000 simulated user test.

50% IPv6 Test

Increasing traffic to 50% IPv6 produced minimal increases in loss. For instance, some loss was noticed at 46% loading for the 128 byte frame size. The increase in loss was minimal, and no increase in loss was seen in larger frame sizes. The Avalanche testing showed a slight difference in the response times (about 50 ms) from the IPv4 baseline test. The 2000 and 2500 simulated users test was similar in response times. The first HTTP get request arrived around 275 ms; the first ACK of response data arrived around 460 ms; and the connection closed around 8 secs. It was noticed that the test device could not handle anything over 2400 simulated users. At around 2400 simulated users, the device stopped running the test and the errors increased exponentially. This caused the average response times to be higher for the 5000 simulated user test.

75% IPv6 Test

Slightly higher throughput and frame loss were noticed for this test. With higher loads (91-96%) on the small 128 byte frame size, loss on some paths reached 100%. Also, on the next pass of the test, with only 1% loading on a 256 byte frame size, some loss continued to be observed. This seemed to indicate that processor utilization on some routers in the network had not recovered from the 100% loss (and near 100% utilization) seen on the higher loaded smaller frame size immediately before this pass of the test. As was consistently the case in all tests, no loss was seen in larger frame sizes. The Avalanche testing provided a slight difference in the response times of about 50 ms from the IPv4 baseline test. The 2000 and 2500 simulated users test was similar in response times. The first HTTP get arrived around 275 ms; the first ACK of response data arrived around 460 ms; and the connection close around 8 secs. Again, it was noticed that the test device could not handle anything over 2400 simulated users. At around 2400 simulated users, the device stopped running the test and the errors increased exponentially. This caused the average response times to be higher for the 5000 simulated user test.

100% IPv6 Test

When the network architecture was configured to native IPv6 (no dual stacks), there was less frame loss and throughput loss than in the previous testing over a dual stacked network. For instance, no losses were seen until loading for 128 byte frame sizes reached greater than 51% (compared to losses beginning at 41-46% loading for the stacked configuration). At no time did any paths experience 100% loss, no matter how high the loading was (compared to numerous traffic paths experiencing 100% loss for 128 byte frame sizes in the dual stacked configuration). Again, with larger frame sizes, no loss was seen. These results are consistent with the knowledge that running an IPv6 only configuration should have lower processor utilization than running a dual stacked IPv4 and IPv6 configuration.

During the Avalanche testing, there was no appreciable difference in the response times from the all IPv6 network and the dual stacked network. The 2000 and 2500 simulated users test was similar in response times. The first HTTP get arrived around 300 ms; the first ACK of response data arrived around 500 ms; and the connection close around 16 secs. It was noticed that the test device could not handle anything over 2400 simulated users. At around 2400 simulated users, the device stopped running the test and the errors increased exponentially.

Conclusions/Recommendations

Selection of an appropriate Internetwork Operating System (IOS) for operation of routers will be critical to successful implementation of IPv6. Many current IOS versions are IPv6 Capable, to some extent. However, to do some functions like running Open Shortest Path First (OSPF) or Enhanced Interior Gateway Protocol (EIGRP) instead of Routing Information Protocol (RIP), running IPv6 tunnels (IPv4 traffic passing over IPv6 networks) require some of the larger, more recent versions of IOS. For Cisco routers, a good starting point as of the time of this study would be to use IOS version 12.4 or higher. It is also recommended not only to look at routers and switches, but also at the servers that run on the network. All the servers will have to start implementing a dual stack architecture, which will require more processing time.

The dual stack configuration, probably the most likely implementation for the Air Force, is more demanding on the network and the network's equipment. Comparing preliminary baseline results with IPv6 configuration test results reveals that there is additional throughput loss and frame loss (and increased latency) on systems under test processing IPv6 packets.

Recommend using larger frame sizes whenever possible. Larger frame sizes are more efficient (with less overhead) and therefore produce better performance and higher throughput. On dual stack transition testing, almost no significant losses were ever experienced on frame sizes larger than 512 bytes. Recommend using routers with higher processing capabilities, where possible, to eliminate throughput and frame loss.

D.2 JCS Criteria 4, Phase 2: 4.1.2.1, 4.1.2.2, 4.1.3.1, 4.1.3.3 Demonstration of the Real Time Protocol (RTP) and Session Initiation Protocol (SIP) Capabilities Over an IPv6 Network Test Report, v1.0

Testing Organization and Publication Date

Space and Naval Warfare Systems Center
September 27, 2007

Summary

This laboratory testing was designed to successfully demonstrate the following segments of Joint Chiefs of Staff (JCS) Criteria 4: Level 3 decomposition items 4.1.2.1, 4.1.2.2, 4.1.3.1, and 4.1.3.3. Specifically, the objectives of this demonstration were to demonstrate that IPv6 supports SIP, RTP, and Real Time Control Protocol (RTCP) to transport voice, video, and data traffic over independent and shared IPv6 environments, and to compare the performance of SIP, RTP, and RTCP over IPv4 and IPv6 based environments.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

3 (3.1, 3.1.1, 3.2, 3.2.1)

4 (4.1, 4.1.2, 4.1.3)

Configuration

This testing used the Counterpoint Eyebeam Softphone application. The Eyebeam Version 1.5 Beta is a telephony client that runs under Microsoft Windows and MAC operating systems. The performance was compared to that of the same data via an IPv4 network connection. The tests were performed in IPv6 mode or IPv4 mode only, but not in dual stack mode. Three situations were tested:

- Voice Transfers (one-to-one) using the Voice Speedex Wideband FEC (64K) codec
- Video Transfers (one-to-one) using the High Quality H.326 codec
- Data transfers (Instant Messaging using the SIMPLE protocol).

The Softphone application was installed on three laptops. The three nodes were connected in an IPv4 and IPv6 (not dual stack) network by two routers via two satellite simulators. The equipment used in testing is listed in Table D-2.

Table D-2 JCS 4 Equipment Configuration

Vendor (number of devices)	Equipment	Model	Operating System
Sony (2)	PC	PCG-Z1RA	XP-Pro
Dell (1)	PC	Optiplex-GX620	XP-Pro
Cisco (3)	Router	3825	12.4(4)T1
AdTech (2)	Link Simulators	SX/12	N/A
Ixia (1)	Automated Test Device	IxChariot	N/A
Logitech (1)	Camera	3000	N/A
Logitech (1)	Headset	350	N/A

Results

Voice Testing Comparison Summary

The quality of the voice transmission for IPv6 compared to IPv4 showed little, if any, perceptible difference. The consumed bandwidth for these test conditions was only 11% more than that consumed by the same voice transmission via an IPv6 network connection, which is not significant. The jitter and latency were well within acceptable limits set for the pass/fail criteria. The voice transmission for this test was essentially equal to the same transmission via an IPv4 link in terms of quality and bandwidth consumption.

Video Testing Comparison Summary

The quality of the video transmission for IPv6 compared to IPv4 showed “no perceptible difference”. The consumed bandwidth for these test conditions were 12% less for IPv6 than that consumed by the same voice transmission via an IPv4 network connection. It is unclear why the transmission consumed less bandwidth; it may be attributed to application differences in compression algorithms. In any case, the difference is not significant. The jitter and latency were well within acceptable limits set for the pass/fail criteria. The video transmission for this test was essentially equal to the same transmission via an IPv4 link in terms of quality and bandwidth consumption.

Data Testing Comparison Summary

The network efficiency was somewhat lower for IPv4, due to the 20 byte longer length of IPv6 headers; but this is not significant for a chat-like application. All performance measures were well within the pass/fail criteria. Short message transfers via a non-congested link were equivalent to IPv4.

Conclusions/Recommendations

IPv6 appears to provide adequate support for the SIP, RTP, and RTCP protocols. On a small scale and in a controlled lab environment, the end users perception of the application’s quality over an IPv6 network was approximately equivalent to that of the same application via an IPv4 network. It was noted that while comparing the protocol’s head-to-head performance measurements (i.e., jitter and latency) between IPv4 and IPv6, there were significant differences; however, the performance variations were imperceptible to the end user. While the basic

operational functionality of SIP, RTP, and RTCP were successfully demonstrated, and their comparable performance and usability confirmed, scalability (Decomposition 4.1.3.4) and end-to-end security (Decomposition 4.1.3.2) tests are still required before issuing a complete endorsement.

D.3 Special Interoperability Test Certification of Cisco 1800, 2800, 3800, and 7200 Families of Routers Running Internetwork Operating System Version 12.4(11)T For Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
July 24, 2007

Summary

This report presents the results of the Special Interoperability Test Certification of the Cisco 1800, 2800, 3800, and 7200 Families of Routers running IOS Version 12.4(11)T as IPv6 Capable routers. This special certification is based on IPv6 Capable testing conducted at the Joint Interoperability Test Command's (JITC's) Advanced IP Technology Laboratory from April 2, 2007 through June 8, 2007.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.4, 1.4.1)
2 (2.1, 2.1.1, 2.3, 2.3.1)
8 (8.1, 8.1.1)

Configuration

The routers were tested as part of a simulated (DISN) IP Core Node test architecture managed by the Advanced IP Technology Laboratory (AIPTL) at JITC. The Devices Under Test (DUTs) and equipment used during testing is listed in Table D-3.

Table D-3 Cisco Test Equipment Configuration

Equipment Name	Model Number	IOS/OS/Version(s)
Hardware		
Cisco Router – Device Under Test (DUT)	Cisco 1841	12.4(11)T
Cisco Router – DUT	Cisco 2811	12.4(11)T
2 Cisco Routers – 1 DUT	Cisco 3845	12.4(11)T
Cisco Router – DUT	Cisco 7200	12.4(11)T
2 Juniper Routers	Juniper M40e	V 7.4R2.6/V 7.6R3.6
2 Juniper Routers	Juniper T320	V 7.4R2.6
Juniper Router	Juniper T640	V 7.1R3.3/V 7.4R2.6
5 Dell Power Edge Servers	2850	MS 2003 Server
2 Gateway Notebooks	450ROG	Windows XP Professional
Gateway Workstation	E Series	Windows XP Professional
Software		
Windows XP Professional	Not Applicable (N/A)	Build 5.1.2600 SP2
Windows Server 2003	N/A	Build 5.2.3790 SP1
SimpleTesterPro	N/A	V11.0.1
VLC Media Player	N/A	V0.8.6b
Wireshark	N/A	V.0.99.2

Results

JITC distributes interoperability information via the JITC Electronic Report Distribution (ERD) system, which uses Unclassified-But-Sensitive Internet Protocol Router Network (NIPRNet) e-mail. More comprehensive interoperability status information is available via the JITC System Tracking Program (STP). The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC Joint Interoperability Tool (JIT) at <http://jit.fhu.disa.mil> (NIPRNet) or <http://199.208.204.125> (Secret Internet Protocol Router Network[SIPRNet]). Information related to IPv6 Capable testing is at <http://jitc.fhu.disa.mil/apl/>.

Table D-4 presents a condensed test results table. More tests were conducted than reported in this appendix. This table provides the RFC, RFC title, testing that was completed (conformance, performance, and interoperability), and whether the router met the requirements.

Table D-4 Cisco Test Results

Cisco 1800, 2800, 3800, and 7200 Family of Routers						
RFC	RFC Title	Testing Completed			Router	
		Conformance	Performance	Interoperability	Requirement	Met/Not Met
2408	Internet Security Association and Key Management Protocol	Stated in Letter of Conformance (LoC)	No Performance Test Required	Yes	Required (R)	Met
2409	Internet Key Exchange (IKE)	Stated in LoC	No Performance Test Required	Yes	R	Met
4301	Security Architecture for Internet Protocol	Stated in LoC	No Performance Test Required	Yes	R	Met
4302	IP Authentication Header	Stated in LoC	No Performance Test Required	Yes	R	Met
4303	IP Encapsulating Security Payload (ESP)	Stated in LoC	No Performance Test Required	Yes	R	Met
4305	Cryptographic Algorithm Implementation Requirements for ESP and Authentication Header (AH)	Stated in LoC	No Performance Test Required	Yes	R	Met
4306	Internet Key Exchange version 2 (IKEv2) Protocol	Not Listed	Not Tested	Not Tested	Optional (O)	Not Tested
4307	Cryptographic Algorithms for Use in the IKEv2	Not Listed	Not Tested	Not Tested	O	Not Tested
4308	Cryptographic Suites for IPsec	Stated in LoC	No Performance Test Required	Yes	R	Met
4213	Transition Mechanisms for IPv6 Host and Routers	Stated in LoC	No Performance Test Required	Yes	R	Met
2474	Definition of the DiffServ Field in the IPv4 and IPv6 Headers	Stated in LoC	No Performance Test Required	Yes	R	Met
3413	Simple Network Management Protocol (SNMP) Applications	Stated in LoC	No Performance Test Required	Yes	R	Met
2460	Internet Protocol version 6 (IPv6) Specification	Stated in LoC	No Performance Test Required	Yes	R	Met
2461	Neighbor Discovery for IP version 6 (IPv6)	Stated in LoC	No Performance Test Required	Yes	R	Met
2462	IPv6 Stateless Address Auto configuration	Stated in LoC	No Performance Test Required	Yes	R	Met
2464	Transmission of IPv6 Packets over Ethernet Networks	Stated in LoC	No Performance Test Required	Yes	R	Met
2710	Multicast Listener Discovery (MLD)	Stated in LoC	No Performance Test Required	Yes	R	Met

Conclusions/Recommendations

The Cisco 1800, 2800, 3800, and 7200 Families of Routers running IOS Version 12.4(11)T are certified for listing as IPv6 Capable routers.

D.4 TNT 07-4 AAR: IPv6 Testing with JITC

Testing Organization and Publication Date

Naval Postgraduate School (NPS) Center for Network Innovation and Experimentation
September 2007

Summary

The Tactical Network Topology (TNT) 07-4 exercise conducted an initial “connectivity” test that ascertained the feasibility of connecting an IPv6 internetwork (e.g., the Defense Research and Engineering Network (DREN)) to an IPv4 tactical edge network (e.g., TNT) and successfully passed data one-way between the two.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

2 (2.2, 2.2.2)

Configuration

JITC operates several servers connected to the DREN through a dual-stack IPv4/IPv6 connection. The TNT network is an IPv4-only, isolated testbed network, connected to outside sites only via Virtual Private Network (VPN) connections. For this test, a direct connection was established between the TNT network and JITC via the NPS DREN connection. A Cisco 2821 router running IOS version 12.4(3f) on the NPS side of the DREN link was utilized to provide IPv6 to IPv4 address translation services.

The objective was to share video between the JITC and TNT, providing each site with video feeds from the other. To do this, two separate sets of translation rules were established: one to expose a TNT IPv4 video source to the IPv6 network, and one to expose a JITC IPv6 video source to the IPv4 network. The former was done via a static Network Address Translation (NAT) mapping (i.e., translating between a specific IPv4 address and a corresponding IPv6 proxy address). The latter used Port Address Translation (PAT) to transpose any incoming IPv6 address onto a single IPv4 address that acted as its proxy within the IPv4 network.

Results

Video was successfully passed from the JITC to the TNT network using a standard User Datagram Protocol (UDP) unicast connection. The Video LAN Client (VLC) media server and client software were used to serve and view the video. In this mode of operation, the video viewing computer had a native IPv4 address, and the Cisco router provided a static NAT mapping onto a reserved IPv6 address, to which the native-IPv6 video server sent the video. The

Cisco PAT translated the incoming video packets from IPv6 source and destination addresses into IPv4 source and destination addresses by using the router's IPv4 address as a proxy source address. This configuration resulted in reliable and high-quality video being passing across the DREN and through the TNT network to the Tactical Operation Center (TOC) in Camp Roberts, California.

Establishing video in the opposite direction was a significant challenge, which was never accomplished. This was possibly due to limitations in Cisco's implementation of IPv4 to IPv6 translation; although passing application traffic from IPv6 to IPv4 worked, traffic from IPv4 to IPv6 was unreliable in the best case.

Conclusions/Recommendations

Initial testing demonstrated that both vendor protocol support and known best practices are still maturing, thus requiring the DoD to establish standard operating procedures for these scenarios before they become commonplace in operational settings.

Extending IPv6 to edge devices will provide valuable insight into interoperability issues, especially as IPv6 is carried over existing switches, wireless data links, and other underlying (Layer 2) devices. It may result that certain equipment is not compatible with IPv6 traffic, even in cases where compatibility is claimed or anticipated. Testing IPv6 over existing equipment and upgrading to include newer, IPv6-compliant equipment may result in better knowledge and understanding for future DoD network applications.

D.5 Evaluation and Implementation of DISA IPv6 Information Assurance Guidance for Milestone Objective 2 version 2 (MO2v2)

Testing Organization and Publication Date

National Security Agency (NSA), Network Infrastructure Division Systems, and Network Analysis Center
September 30, 2007

Summary

This document contains an analysis of each MO2v2 architecture, the functional and security requirements, recommendations, and configuration guidance to implement those requirements. MO2v2 describes network architectures that allow IPv4 and IPv6 traffic to pass between participating enclaves and the network core. NSA evaluators built and evaluated these architectures to determine if the MO2v2 architectures were functional and secure.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.4, 1.4.1, 1.4.2, 1.5, 1.5.1)
8 (8.1, 8.1.3)

Configuration

The evaluators built four enclaves corresponding to the four MO2v2 architectures: split domain, dual-stack, Intra-Site Automatic Tunneling Address Protocol (ISATAP), and NAT-Port Translation (PT). A network core connects the four enclaves. The test bed allows communications between the enclaves subject to the filtering at the enclave boundaries.

The network core consisted of four routers, each of which connects to an enclave representing one of the MO2v2 architectures. These routers were dual-stacked, so they could not route native IPv4 and native IPv6 traffic. Each enclave has an IPv6 connection to the core, so that each enclave can communicate with any other enclave using native IPv6. The split domain, dual-stack, and ISATAP enclaves also support a native IPv4 connection to the core, allowing these three enclaves to communicate using IPv4. The network core does not enforce any of the security requirements of the enclaves.

Table D-5 lists the equipment used during testing.

Table D-5 Equipment Configuration

Enclave	Vendor	Model	Type	Software Version
Network Core	Cisco	3825	Router	C3825-ADVENTERPRISEK9-M, v 12.3(14)T1
	Cisco	3825	Router	C3825-ADVENTERPRISEK9-M, v 12.4(16)
	Cisco	3845	Router	C3845-ADVENTERPRISEK9-M, v 12.4(16)T
	Cisco	3845	Router	C3845-ADVENTERPRISEK9-M, v 12.4(16)XT
Split Domain Architecture	Cisco	3660	Router	C3660-JK9S-M, v 12.4(6)XT
	Cisco	3725	Router	C3725-ADVENTERPRISEK9-M, v 12.4(6)XT
	Cisco	3660	Router	C3660-JK9S-M, v 12.4(6)XT
	Cisco	ASA 5510/ AIP-SSM-10	Firewall/IPS	7.2(2)/6.0(2)
	Cisco	ASA 5520/ AIP-SSM-10	Firewall/IPS	7.2(2)/5.0(2)
	Juniper	Netscreen 204	Firewall/IPS	5.40R5.0
	Cisco	3725	Router	C3725-ADVENTERPRISEK9-M, v 12.4(6)T
Dual Stack Architecture	Cisco	3725	Router	C3725-ADVENTERPRISEK9-M, v 12.4(16)
	Cisco	PIX 515	Firewall	7.2(1)
	Cisco	4215	IDS	6.0(3)E1
	Cisco	PIX 515E	Firewall	7.2(2)
	Cisco	3725	Router	C3725-ADVENTERPRISEK9-M, v 12.4(6)T
	Cisco	4215	IDS	6.0(3)E1
ISATAP Architecture	Cisco	3725	Router	C3725-ADVIPSERVICESK9-m, v 12.3(23)
	Cisco	3725	Router	C3725-ADVIPSERVICESK9-m, v 12.3(23)
	Juniper	ISG 1000	Firewall/IPS	6.0.0r1.0
	Cisco	PIX 515E	Firewall	7.2(2)
	Cisco	4215	IDS	6.0(3)E1
	Cisco	2851	Router	C2800NM-ADVANTERPRISEK9-M, v12.4(6)XT
NAT-PT Architecture	Cisco	3725	Router	C3725NM-ADVANTERPRISEK9-M, v12.4(6)XT
	Cisco	2851	Router	C2800NM-ADVANTERPRISEK9-M, v12.4(4)T
	Cisco	3725	Router	C3725NM-ADVANTERPRISEK9-M, v12.4(6)T
	Cisco	ASA5520/ AIP-SSM-10	Firewall/IPS	7.2(2)/5.0(2)
	Cisco	ASA 5510/ AIP-SSM-10	Firewall/IPS	7.2(2)/6.0(2)

Results

No specific results were provided in this report. The report describes the configuration and provides detailed recommendations derived from testing.

Conclusions/Recommendations

During the evaluation, it became apparent that current perimeter security devices (i.e., firewalls and Intrusion Detection System [IDS]/Intrusion Prevention System [IPS]) lack the necessary security features proposed in the architectures. Evaluators proposed techniques to compensate for some of the devices' shortcomings. Recommendations from this evaluation include secure configuration examples for enclave routing and perimeter security devices. During the evaluation, multiple functionality deficiencies were found while configuring the firewall and IDS/IPS devices. System administrators should use the DoD IPv6 Transition Office (DITO) Information Assurance (IA) guidance in conjunction with this document to provide a secure transition to MO2v2.

The following headings summarize the test recommendations.

Split Domain Architecture

Requirements

- The IPv4 firewall and IPv4 Network Intrusion Detection Systems (NIDS) must support the current IPv4 deployment.
- The IPv6 firewall and IPv6 NIDS must provide equivalent or better support than the current IPv4 deployment for IPv6 traffic.

Recommendations

- The firewall and NIDS on the IPv4 path through the intra-enclave security zone support IPv4 while the firewall and NIDS on the IPv6 path support IPv6.

Dual-Stack Architecture

Requirements

- No other IPv6 transition mechanisms may cross the enclave boundary.
- The enclave firewalls and NIDS must support current IPv4 deployment for IPv4 and IPv6 traffic.

Recommendations

- Dual-stack is the only transition mechanism implemented as enforced by the router interface rules.
- The firewalls and NIDS support IPv4 and IPv6.

ISATAP Architecture

Requirement

- The enclave firewall and NIDS must support current IPv4 deployment for IPv4 and IPv6 in IPv4 tunneled traffic.

Recommendations

- In the evaluators' implementation, the ISATAP tunnel is terminated at the ISATAP router, so tunnels are not allowed in the intra-enclave security zone as described above. The enclave firewall and NIDS support both native IPv4 and native IPv6.
- All IPv6 traffic unencapsulated from the ISATAP tunnel traffic that traversed outside of the enclave contained the IPv6 ISATAP address assigned to the Windows host. Since these addresses are formatted specifically for ISATAP, an external entity will know that the enclave host is using ISATAP when any IPv6 communication occurs. Additionally, an external entity will be able to extract the internal host's IPv4 address embedded as the last 32

bits of the ISATAP address. It may be undesirable for this information to be available to external entities.

NAT-PT Architecture

Requirements

- The enclave firewall and NIDS must support current IPv4 deployment for IPv4 and IPv6 traffic.
- IPv4 to IPv6 address mappings must be one-to-one. Specifically, only one native IPv4 address and translated IPv6 address may be associated with each layer-two address supported on the IPv4 device.

Recommendations

- The firewall and NIDS support IPv6. IPv4 is not supported in the intra-enclave security zone according to the router interface rules below.
- The evaluators configured one-to-one IPv4-IPv6 mappings via the commands shown in the implementation section. However, the NAT-PT router cannot associate an IP address with a layer-two address. MAC address filtering cannot be performed on a layer-three interface; filtering on a layer-three router interface is done using IP addresses. Therefore, although one native IPv4 address and translated IPv6 address are logically associated with a layer-two address on the IPv4 device, this layer-two binding is not enforced by the NAT-PT router.

Firewall Issues

The Cisco firewalls do not provide IPv6 functionality when in transparent mode. In addition, the evaluators determined that the Cisco Adaptive Security Appliance (ASA) firewall could not block IPv6 traffic encapsulated in IPv6. IPv6-in-IPv6 traffic is not permitted by any of the architectures, but writing a rule to block this type of traffic resulted in the firewall blocking all IPv6 traffic.

IDS/IPS Issues

All traffic entering its interfaces is inspected by the IPS blade on a back channel interface before it can pass through the ASA. The IPS blade can deny packets based on vendor-supplied signatures or custom-defined signatures. Unfortunately, the evaluators were unable to utilize the blade for native IPv6 or any traffic tunneled in IPv6, due to a software issue with the ASA. The ASA could not group IPv6 traffic, which is a requirement that must be met for the ASA to send traffic to the IPS module.

D.6 Cisco Networks Internet Protocol Version 6 Test Report

Testing Organization and Publication Date

Joint Interoperability Test Command
September 2007

Summary

This test compared IPv6 performance to IPv4 performance using six types of Internet routers manufactured by Cisco Networks. Those routers were 1841, 2811, 3825, 7200, 7301, and 7600 series routers. These represent the Cisco routers typically found in the DISN. They were tested using the Cisco IOS Versions 12.4(11)T and 12.2(33)SRB, which are expected to be used in the DISN during the IPv6 transition. The Advanced IP Technology Laboratory personnel witnessed testing of the Cisco routers for IPv6 performance from April 8-12, 2007, at Cisco Networks Laboratory in Research Triangle Park, North Carolina.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

3 (3.1, 3.1.1, 3.2, 3.2.2)

8 (8.1, 8.1.2)

Configuration

Tests were conducted on individual devices separate from any network architecture and therefore are not representative of that device's performance on a network. The impact of the tested devices on a DISN replica network will be determined in later testing. Two connections were made from the automated test chassis to the device. The device interfaces were chosen to ensure that the offered traffic would constitute 100% of a single port's capacity. The Spirent Test Center network/device performance tester was the automated test chassis used in this test.

When evaluating throughput and latency of the device, several IPv4/IPv6 ratios were used. These ratios were 100% IPv4, 100% IPv6, and the following IPv6/IPv4 percent ratios: 10/90, 50/50, and 90/10.

Table D-6 lists the devices and the IOS versions that were tested.

Table D-6 Device Configuration

Device Platform	Interface Model Number	Module Location #/Firmware Version	IOS	Processing Engine
Cisco 1841	100 Mb Ethernet	fastethernet0/0 fastethernet0/1	12.4(11)T	N/A
Cisco 2811	100 Mb Ethernet	fastethernet0/0 fastethernet0/1	12.4(11)T	N/A
Cisco 3825	1 GE	gigabitethernet0/0 gigabitethernet0/1	12.4(11)T	N/A
Cisco 7200	1 GE	gigabitethernet0/1 gigabitethernet0/2	12.4(11)T	NPE-G2
Cisco 7600	10 GE	Tengigabitethernet2/0/0 Tengigabitethernet3/0/0	12.2(33)SRB	SUP720-3BXL
Cisco 7301	1 GE	gigabitethernet0/0 gigabitethernet0/1	12.4(11)T	N/A

Results

Slight differences were found between IPv6 and IPv4 combined throughput rates when the devices were running one protocol exclusively or when the traffic was split evenly between protocols. Differences noted when traffic was 90% one protocol and 10% of the other were small enough to be within measurement and rounding error. It was also noted that IPv6 introduced additional latency to the devices, which is unlikely to have significant impact on network operations. The routers with these discrepancies were always the lower capacity customer edge routers.

Throughput for the combined devices was identical for the two protocols with evenly split traffic levels. Minor differences were found on specific devices at same frame sizes, but not enough to significantly impact operations.

Table D-7 presents the IPv4/IPv6 combined device results.

Table D-7 Cisco IPv4/IPv6 Combined Device Results

Frame Size (Bytes)	IPv4/IPv6 Frame Latency (µs)	IPv4/IPv6 Throughput (Mbps)
100% IPv4/IPv6		
86	27/35	35/35
128	32/40	38/38
256	37/59	67/67
512	53/66	70/70
768	69/75	90/90
1024	86/92	90/90
1280	104/107	90/90
1518	121/126	90/90
10% IPv4/IPv6		
86	29/40	N/A
128	34/48	N/A
256	53/52	N/A
512	59/63	N/A
768	72/81	N/A
1024	89/98	N/A
1280	107/116	N/A
1518	122/111	N/A
50% IPv4/IPv6		
86	36/38	35/35
128	36/38	45/45
256	54/57	68/68
512	59/61	87/87
768	73/76	90/90
1024	90/94	90/90
1280	106/111	90/90
1518	121/126	90/90
90% IPv4/IPv6		
86	37/34	N/A
128	47/38	N/A
256	51/55	N/A
512	57/62	N/A
768	73/75	N/A
1024	90/93	N/A
1280	106/111	N/A
1518	103/126	N/A

Conclusions/Recommendations

Test results indicated parallels in the frame throughput. While minor differences were found in frame latency, these differences will have no operational impact. Therefore, the performance of IPv4 and IPv6 in the tested Cisco routers is considered equivalent.

D.7 Test of Internet Protocol Version 6 (IPv6) Configured Tunneling

Testing Organization and Publication Date

Air Force Communications Agency
September 18, 2007

Summary

This report presents the results of testing configured tunneling of IPv6 packets through an IPv4 network. The testing demonstrated the general feasibility and performance of tunneling scenarios within a laboratory environment and is indicative of what can be expected in an operational environment. The test objective was to demonstrate the establishment and performance of configured IPv6 tunnels. Tunnel types included protocol 41 (IPv6 in IPv4) and protocol 47 Generic Routing Encapsulation (GRE). Tunnel performance was compared against IPv4-only performance. Tunnel scenarios included router-to-router, host-to-host, and router-to-host.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

8 (8.1, 8.1.1, 8.1.2)

Configuration

An emulated network enterprise was configured with two Air Force bases interconnected via emulated DISN connectivity. All connections between devices are 100 Mbps Fast Ethernet. Using the Smartbits 600 performance analyzer with Smartflow software, test frames were generated between Main Operating Base 1 (MOB1) and MOB2. Network throughput, frame loss, latency, and processor utilization measurements were recorded for various frame sizes and connection loading percentages. The two gateway routers used were Cisco 7206 VXR's with IOS 12.4(11)T1 on a Network Processor Engine 400. The two host computers were running Microsoft Windows XP with Service Pack 2 (SP2).

Results

Router-to-Router

Protocol 41 Encapsulation

Interoperability:

Testing demonstrated the ability to establish configured protocol 41 IPv6 tunnels through an IPv4-only network between two Cisco 7206 VXR routers running IOS 12.4(11)T1. Configuration of the routers was easy, with no issues. The only impact on the IPv4-only network was the need to configure the bases' firewalls to permit protocol 41 UDP port 9000 packets.

Performance:

When comparing the performance of protocol 41 encapsulated IPv6 traffic against native IPv4 traffic, there was significant degradation. The biggest difference was in the processor load on the routers. For the evaluation, throughput was compared at a frame size of 1280 bytes; frame loss was compared at a frame size of 1280 bytes and a load of 51 Mbps; latency was compared at a frame size of 1280 bytes and a load of 51 Mbps; and router processor utilization was compared at a frame size of 512 bytes and a load of 50 Mbps. These frame sizes and loading were used based on analysis of the data and determining the range of frame sizes and loading that provide the most consistent range of results. Smaller frame sizes and lower loading showed greater fluctuation in data that could be attributed to spurious bit errors.

Generic Routing Encapsulation

Interoperability:

The testing demonstrated the ability to establish configured GRE IPv6 tunnels through an IPv4-only network between two Cisco 7206 VXR routers running IOS 12.4(11)T1. Configuration of the routers was easy, with no issues found. The only impact on the IPv4-only network was the need to configure the bases' firewalls to permit GRE UDP port 9000 packets

Performance:

When comparing the performance of GRE IPv6 traffic against native IPv4 traffic, there was significant degradation. The biggest difference was the processor load on the routers. For the evaluation, throughput was compared at a frame size of 1280 bytes; frame loss was compared at a frame size of 1280 bytes and a load of 51 Mbps; latency was compared at a frame size of 1280 bytes and a load of 51 Mbps; and router processor utilization was compared at a frame size of 512 bytes and a load of 50 Mbps. These frame sizes and loading were used based on analysis of the data and determining the range of frame sizes and loading that provide the most consistent range of results. Smaller frame sizes and lower loading showed greater fluctuation in data that could be attributed to spurious bit errors.

In contrast to native IPv4 and protocol 41 performance measurement results, the router processor showed severe degradation when using GRE to tunnel IPv6 packets through an IPv4 network. This overloading of the processor resulted in significant frame loss and high latency. The router processor utilization exceeded 98% between 20 and 30 Mbps for GRE while the protocol 41 measurements did not exceed 98% until a load of 90 Mbps was applied.

Host-to-Router

Protocol 41 Encapsulation

Interoperability:

The testing demonstrated the ability to establish configured protocol 41 IPv6 tunnels through an IPv4-only network between a Cisco 7206 VXR router running IOS 12.4(11)T1 and a Windows XP SP2 host. Configuration of the router and host was easy, with no issues found. The only impact on the IPv4-only network was the need to configure the bases' firewalls to permit protocol 41 packets.

Performance:

In comparing the performance of tunneled protocol 41 host-generated traffic against native IPv4 traffic, there was a significant reduction of throughput. Native IPv4 throughput was measured at 93.542 Mbps while tunneled protocol 41 throughput was measured at 30.064 Mbps. This is a 67.9% reduction in performance.

GRE

Configured GRE tunneling between Host 1 and the MOB1 gateway router was not available, due to lack of the feature on Host 1.

Host-to-Host Tunneling

Protocol 41 Encapsulation

Interoperability:

The testing demonstrated the ability to establish configured protocol 41 IPv6 tunnels through an IPv4-only network between two Windows XP SP2 hosts. Configuration of the hosts was easy, with no issues found. The only impact on the IPv4-only network was the need to configure the bases firewalls to permit protocol 41 packets.

Performance:

When comparing tunneled protocol 41 host-generated traffic against native IPv4 traffic, there was a significant reduction in throughput. Native IPv4 throughput was measured at 93.542 Mbps while tunneled protocol 41 throughput was measured at 31.073 Mbps. This is a 62.5% reduction in performance.

GRE

Configured GRE tunneling between the two hosts was not available due to lack of the feature on the hosts.

Conclusions/Recommendations

Testing demonstrated the technical feasibility of all three tunneling scenarios. All were easy to configure. The only configuration change to the two emulated bases was to allow tunneling through the firewalls. When compared to IPv4-only performance, tunnel performance was degraded. Data collected during the router-to-router scenario testing indicated the use of GRE to encapsulate the IPv6 packets for transport through an IPv4 network resulted in significant performance degradation. The load on the router processor during GRE tunneling quickly surpassed the capability of the processor as the test traffic rate was increased.

The testing proved tunneling could be applied in the router-to-router, host-to-router, and host-to-host scenarios using Cisco 7206 VXR routers and Windows XP SP2 hosts. The performance results were only applicable to the specific scenarios, hardware, and software used during this testing. A significant finding was severe degradation of performance when using GRE to encapsulate IPv6 packets for transport through an IPv4 network. Other hardware may support GRE tunneling with less degradation than the Cisco 7206 VXR; however, there is no known advantage to the use of GRE versus protocol 41.

D.8 Net-Centric Operations Warfare (NCOW) IPv6 Demonstration: Security Features

Testing Organization and Publication Date

Defense Information Systems Agency
September 30, 2007

Summary

This paper describes an IPv6 security demonstration in a NCOW setting. The goal of the demonstration is to highlight IPv6 security capabilities for NCOW that cannot be easily or cost effectively realized using other technologies. The exercise demonstrates the net-centric IA goals of edge-to-edge non-repudiation, authentication, and encryption support for data transport and Operation, Administration and Maintenance (OAM) for authorized network administrators using IPv6 and IPv6 end-to-end security.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.3, 1.3.1)

8 (8.1, 8.1.1, 8.1.3)

Configuration

Most test scenarios were implemented with two Microsoft Windows Server 2008 machines. Although both machines have Microsoft Windows Server 2008, only one machine had server roles. The server was configured to be dual-stack; the client was IPv6-only.

For Linux testing, Gentoo Linux was used as the base distribution for the majority of testing with Gentoo base version 2007.0 and the current Linux kernel 2.6.20.

Results

Microsoft Windows Server 2008 Scenarios

Scenario 1: The client does not have the Certificate Authority's (CA's) root certificate or a client authentication certificate from the CA.

Result: When the client attempted to access <https://demo.ncowpki.com/>, the web browser reported that the site was not trusted because it had a server certificate that could not be verified. In addition, the client was denied access to the site because it could not supply

the proper credentials. Internet Information Services (IIS) was configured to trace failed requests. A request trace was logged when an error status code was generated.

Scenario 2: The client has the CA's root certificate in its Trusted Root Certificate Authorities store, but does not have a client authentication certificate issued by the CA.

Result: When the client attempted to access <https://demo.ncowpki.com/>, the server certificate was considered valid and trusted. The client was denied access to the site because it could not supply the proper credentials.

Scenario 3: The client has the CA's root certificate in its Trusted Root Certificate Authorities store and holds a client authentication certificate issued by the CA.

Result: When the client attempted to access <https://demo.ncowpki.com/>, the server certificate was considered valid and trusted, and the client was allowed access to the content displayed by the web server.

Scenario 4: The client has the CA's root certificate in its Trusted Root Certificate Authorities store, but its client authentication certificate has been revoked by the CA and the CA has published the Certificate Revocation List.

Result: When the client attempted to access <https://demo.ncowpki.com/>, the server certificate was considered valid and trusted. The client was denied access to the site because it did not hold a valid client authentication certificate.

Public Key Infrastructure (PKI)

PKI (client and server authentication) over IPv6 was demonstrated using the Apache (version 2.0.58) based Hypertext Transfer Protocol Secure (HTTPS) Data Fusion Server (DFS). Client and server CA certificates issued by Windows and Linux were used to prove the applicability of using PKI for both client and server based authentication over IPv6.

Secure Sockets Layer and Transport Layer Security over IPv6

HTTPS over IPv6 testing demonstrated full interoperability with Windows and Linux client systems showing no negative test results.

IP Security (IPsec)

IPsec over IPv6 was demonstrated between multiple Linux client systems. The IPsec endpoints were configured using manual text files to instruct the IPsec enabled endpoints and the appropriate keys that the endpoints should offer during Internet Key Exchange Version 2 (IKEv2) key exchange. IKEv2 and Internet Security Association and Key Management Protocol were used to dynamically instantiate the IPsec transport sessions. The IPsec over IPv6 testing used Encapsulating Security Payload (ESP) (IP/50) transport mode tunnels with Authentication Header (AH) authentication to demonstrate end-to-end authentication and encryption. IPsec over

IPv6 capabilities used the Openswan library (version 2.4.9) using Rivest-Shamir-Adleman (RSA) cryptographic keys and Microsoft CA generated certificates

Conclusions/Recommendations

The exercise successfully demonstrated the net-centric IA goals of edge-to-edge non-repudiation, authentication, and encryption support for data transport. Using IPv6 during this demonstration provided:

- Authentication, Authorization, and Accounting (AAA) services for user and administrator privilege groups
- Secure configuration of the standalone sensor network
- Authentication between the sensor network elements
- Privacy of the standalone sensor network
- Filtering between sensor network elements and IPv6 network elements
- Authentication and encryption for IPv6 network elements.

D.9 Net-Centric Operations Warfare (NCOW) IPv6 Demonstration: Security Features

Testing Organization and Publication Date

Defense Information Systems Agency
October 1, 2007

Summary

This report analyzes the requirements that are needed to attain the overarching goal of the Global Information Grid (GIG)-IA. These requirements include:

- Network core encryption
- Edge-to-edge non-repudiation, authentication and encryption
- Support of Cross Domain Solutions (CDS)
- Network stability
- Hardened against Denial of Service (DoS)
- Ability for authorized users to manage and operate the network.

Test and Evaluation Method

Engineering Analysis

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.4, 1.4.1, 1.6, 1.6.1)

Configuration

The Unclassified and Secret Network configuration of the GIG is used to hypothesize the impacts of implementing IPv6 solutions and installing the goals mentioned above in the summary section.

Results

Encryption and High Assurance Internet Protocol Encryptor (HAIPE)

The GIG-IA Increment 1 environment requires “system high” security domains. All enclaves will connect through an IP packet edge-edge encryption (E3) device such as a HAIPE or commercial IPsec device. These devices enables the creation of an encrypted IP core which supports the secure, shared transport of all classification levels of data, ranging from Unclassified through Top Secret. There should be no difference between IPv4 and IPv6 network core encryption since HAIPE V3 specification supports both IPv4 and IPv6 packet encryption.

Identity Management, Authentication, and Privileges

The GIG-IA Increment 1 environment defines IPsec as the primary solution for data non-repudiation, authentication, and encryption. The security architecture for IP defines IPsec, which provides security services at the IPv4 or IPv6 layers. It comprises the use of AH, ESP, and IKEv2. The IPv6 base protocol specification requires that all implementations of IPv6 must support ESP and may include AH extension headers. This requirement, along with secure transmission of keys using IKEv2, provides an end-to-end secure channel for communication. It has been determined that IPv6 provides better integration of IPsec through the use of the modular AH and ESP header extensions, where the nested header approach may enable better router and firewall (FW) processing of IPv6 header extensions based on the NSA analysis.

Mediate Security Assertions and Cross Security Domains Exchange

This design tenet involves the development and deployment of a combination of technical solutions including:

- Firewalls (FWs)
- Access Control Lists (ACLs)
- IPsec
- Secure VPNs
- PKI
- Demilitarized Zones (DMZ) Enterprise Architecture

Current CDS mechanisms include FWs, ACLs, and DMZ enterprise architecture based on the DoD Enclave and Access Control Secure Technical Implementation Guides (STIGs). The FWs are topological defense mechanisms that rely on a well-defined boundary between the good “inside” and the bad “outside” of the enclave, with the FW mediating the passage of information between them. ACLs, generally implemented in core and edge systems such as screening routers, operate on a security policy to accept or deny packets based on protocol address, IP protocol type, and/or port.

Network Stability

The technical assessment of network stability can be summarized in the following manner:

- Network stability and availability will be increased by implementing native IPv6.
 - Vendors and network operators are unwilling to implement new security features (e.g., Domain Name System Security) in large deployments of IPv4 products and networks.
 - Network Management/Operations (NM/OPS) of a transitioning IPv4 network to IPv6 will be most cost effective if the NM/OPS infrastructure is running as a single stack network.
- New authenticated NM/OPS paradigm
- Secure control plane
- Automatic rerouting and reconfiguration
- Dynamic addressing structure.

Failure to implement IPv6 will result in continued manual intervention to restore routing links, create security associations, and manage NM/OPS and IA devices in the tactical environment.

The DoD IPv6 requirement to implement IPsec, AH, and ESP will assist in the development of a secure routing and secure NM/OPS.

Hardened against DoS

The technical assessment of DoS can be summarized in the following manner.

- DoS is the most difficult security issue to mitigate
- DoS mitigation requires multiple IP-based and non IP-based security solutions.

The migration of IA network services to enclaves, DMZs and host-systems will aid in the protection against DoS by implementing IPv6 IPsec. The integration of AH and ESP header extensions provides better routing and switch processing characteristics than IPv4 IPsec.

Ability for Authorized Users to Manage and Operate the Network

The technical assessment of NM/OPS can be summarized in the following manner.

- Secure NM/OPS may be increased implementing native NM/OPS IPv6 infrastructure
- New authenticated NM/OPS paradigm
- Secure, authenticated access to NM/OPS Systems.

Conclusions/Recommendations

The GIG-IA Increment 1 environment only should be implemented using IPv6, since it would be difficult to implement the GIG-IA using the current IPv4 architecture. The technical assessment of implementing the GIG-IA Increment 1 environment with IPv4 or IPv6 can be summarized in the following manner:

- IPv6 IPsec, through better integration of AH and ESP header extensions, provides better routing and switch processing characteristics than IPv4 IPsec.
- IP layer does not have an effect on CDS certification or improvement. CDS processes operate above layer three at the data level; thereby they are not affected by the transport to and from the CDS solution.
- Network stability and availability will be increased by implementing native IPv6.
- Failure to implement IPv6 will result in continued manual intervention to restore routing links, create security associations, and manage NM/OPS and IA devices in the tactical environment.
- The migration of IA network services to enclaves, DMZs and host-systems will aid in the protection against DoS by implementing IPv6 IPsec.

The GIG-IA Increment 1 environment solutions may require additional analysis and development based on vendor implementation and operational experience. This additional development may include IETF IPv6 protocol standards redevelopment. This is the principle strength in implementing IPv6 for the GIG-IA, since the IETF IPv4 protocol standards cannot be redeveloped.

D.10 Net-Centric Operations Warfare (NCOW) IPv6 Demonstration

Testing Organization and Publication Date

Defense Information Systems Agency
September 30, 2007

Summary

This document describes an IPv6 demonstration in a NCOW setting. The goal of the demonstration was to highlight IPv6 capabilities for NCOW that cannot be easily or cost effectively realized using other technologies. This demonstration created a typical NCOW setting by integrating IPv6 enabled sensor networks at various locations in the country. The sensors are used to remotely monitor DoD assets. Based on the real-time information gathered in this process, sensors were remotely tasked to fulfill the requirements of the mission.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1, 1.1.1, 1.2, 1.2.1, 1.3, 1.3.1)

8 (8.1, 8.1.1)

10 (10.1, 10.1.1)

Configuration

This demonstration monitored an area using a set of networked sensors that are deployed in a remote area with no direct connectivity to the outside world. A Data Acquisition Agent (DAA) relayed sensor data to a DFS via satellite. The DFS gathered and processed the data from all the sources, and took actions (triggers the video camera). An end user monitored the situation.

The IPv6 Camera (Panasonic KX-HCM110) was connected to a dual stack subnet where it had IPv4 and IPv6 connectivity. A motion detector was connected to the IPv6 camera so that a window with the camera's view of the IPv6 lab room would pop up once motion was detected in the room. When deployed, the sensor nodes automatically formed a secured area using the IPv6 attributes of neighbor discovery, link-local addressing, stateless auto-configuration, and the bridge node bridges between the sensor network and the local area network. In this demonstration, a web-enabled wireless sensor network application was used. Arch Rock manufactured the deployment platform.

Results

This demonstration attained the overarching net-centric IA goals of edge-to-edge non-repudiation through authentication and encryption, as well as the ability for authorized users to manage and operate the network by providing:

- Secure configuration of stand-alone sensor networks
- Authentication between the sensor network elements
- Privacy of stand-alone sensor networks
- Authentication and encryption for IPv6 network elements
- Filtering between sensor network elements and IPv6 network elements.

This demonstration also incorporated two major IA technical efforts, the implementation of PKI over native IPv6 systems, and the implementation of IPsec.

Conclusions/Recommendations

There were issues in how the sensor networks get their IPv6 addresses. For example, are they manually assigned or do they come from the NEMO Client router or bridge? Also, how are the external interfaces allocated to the NEMO client for its home address?

Manually configuring IPv6 addresses in the Home Agent to match the IPv6 addresses that were provisioned for Mobile IPv6 nodes and NEMO clients, as well as, the individual sensors was difficult. These issues were related to spending too much time manually configuring and numbering these moving components. Today, no automation tool exists within the industry to do this for NEMO or the NEMO network. This will pose a major problem when the network is deployed.

It was noted that the demand for route optimization between the sensor decision-making entities will require route optimized paths rather than always going through the NEMO home agent. Currently, the base NEMO specification does not support route optimization, but the IETF NEMO group is working on such a standard. The implementation and pitfalls of various route optimization approaches for NEMO must be examined before DoD deploys these methods.

This demonstration project showed that there is a requirement from the DoD that simultaneous access from different access networks to the sensor network is desirable. This would mean that a NEMO client that servers as a bridge to the sensor network would be extended to allow the connection from different DoD networks to dynamically associate with the NEMO client.

D.11 Domain Name System (DNS) IPv6 Test Plan and Report

Testing Organization and Publication Date

Army, Information Technology Agency (ITA)
March 16, 2006

Summary

The overall objective of the ITA IPv6 test effort is to have the Pentagon ready to support IPv6 communication in the network. These tests were conducted to verify proper DNS operation within the ITA network infrastructure.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.3, 2.3.1)

3 (3.2, 3.2.1)

8 (8.1, 8.1.1)

Configuration

The test environment utilized one master server and two public caching servers for the trusted (internal) and untrusted (external) zones. This was sufficient to test the functionality, availability, and performance of the IPv6 DNS architecture. The DNS master server operated on a SunFire v440 server, and the public caching servers ran on SunFire v240 servers.

A baseline test captured statistics for IPv4 DNS functionality and performance. IPv6 DNS testing was then conducted and compared to results from the baseline test. These tests included various dual-environment scenarios, such as a dual stack (IPv4/IPv6) client requesting a DNS record that contains both IPv4 and IPv6 addresses, to identify potential issues during the IPv6 transitioning period. The scope of functional testing was to verify basic forward and reverse name resolution for DNS clients in IPv4 and dual-stacked environments (querying for 'A' and 'AAAA' records). Performance testing was conducted by measuring server response times to client DNS queries, and by load testing the DNS servers in IPv4-only and dual-stack scenarios.

Results

IPv4 Baseline Test

This test provided a baseline of the network and DNS services to ensure proper functionality of the DNS servers on the test network. Once the master server was loaded with all the zone files, the caching servers obtained all the zone files through zone transfers.

Dual-stack DNS Test

The DNS server successfully responded to DNS queries from the host in the Wedge. Similar to the IPv4 case, the server responded almost instantaneously to the DNS query (approximately 1ms).

When testing dynamic DNS updates using IPv6, the DNS server could be reached with pings, but the nsupdate utility could not communicate with the DNS server to update the record in the test zone. Researching this error confirmed that this is a problem with the Berkeley Internet Name Domain (BIND) version. Beginning with 9.3.5, this error was corrected.

Performance

Table D-8 presents the comparison table between IPv4 and IPv6 querying A and AAAA records.

Table D-8 Performance Comparison Table

DNS Transport	Record Type	Performance (queries/sec)
IPv4	A	10,300
IPv4	AAAA	9,300
IPv6	AAAA	9,000
IPv6	A	9,500

Conclusions/Recommendations

The functional tests proved that the DNS server could respond to any DNS queries from host machines located in Wedges, regardless of protocol. Response times from the server were the same whether using IPv4 or IPv6 to query the DNS server. Performance tests showed that a single Sunfire V240 server running BIND 9.3.4 took a small performance hit (~10%) when responding to IPv6 DNS queries. The ITA DNS Anycast configuration is such that the end users will not notice the decreased performance since multiple DNS servers will be available to respond to DNS queries.

D.12 IPv6 Core Routing Test Plan and Report

Testing Organization and Publication Date

Army, Information Technology Agency
July 16, 2007

Summary

As part of this test effort, IPv6 routing must be successfully tested to guarantee that the ITA network infrastructure can efficiently read, process, and forward IPv6 packets reliably. The IPv6 Core Routing test report describes the results of the test by using the test procedures outlined in the IPv6 Core Routing test plan. These tests were conducted to verify proper IPv6 routing operation within the ITA network infrastructure.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

3 (3.1, 3.1.1, 3.3, 3.3.1)

8 (8.1, 8.1.2)

Configuration

Performance testing involved testing reliability of the network under varied loads. Performance was evaluated by comparing the throughput and frame loss of a network utilizing only IPv4 routing protocols versus both protocols. It should be noted that performance testing was conducted on lab equipment similar to, but not identical to the operational environment. Therefore, performance data such as maximum routes, convergence, and failover times will likely be improved for both protocols. The devices used within the test network are listed in Table D-9.

Table D-9 Test Equipment Configuration

Device	Software Version
Extreme 6804	Extremeware 7.6.3.3
Extreme 5i	Extremeware 7.6.3.3
Cisco 2691	IOS 12.3(21)
Cisco 3550	IOS 12.1(22) EA8a
Cisco 6503	IOS 12.2(18)SXF7
Cisco 6506	IOS 12.2(18)SXD7b
Juniper M10i	JunOS 8.2R2.4
Juniper M20 (W1U/L)	JunOS 7.5 / JunOS 8.2 R2.4
Juniper M20 (W4U/L)	JunOS 8.0 R2.8 / JunOS 8.3 R1.5
Juniper M20 (Vcomp)	JunOS 7.5 R4.4
Netscreen 5200	ScreenOS 6.0.0b3.0
HP Laptop	Windows XP Service Pack 2
Dell Optiplex GX270	Windows XP Service Pack 2
Spirent Smartbits	SmartFlow 4.70.022.1
Spirent Smartbits	TeraRouter Tester 5.00.150

Results

IPv6 over MPLS Test

The results of the IPv6 over Multi Protocol Label Switching (MPLS) test demonstrated that the usage of IPv6 in a dual stack environment does not affect performance when compared to the IPv4 baseline. When measuring maximum load and throughput, both protocols had almost identical results. In each case, the performance was close to the theoretical maximum load of 100%.

Open Shortest Path First (OSPF) Version 3 (OSPFv3) Routing Test

The OSPF scalability test was performed by determining the maximum number of IPv4 and IPv6 OSPF routes a Wedge Router could support. To determine the OSPF route maximums, routes were injected until the memory utilization reached approximately 100%. As the number of routes increased, traffic was sent to the advertised routes verifying the router could continue to forward traffic.

Overall, OSPFv3 (IPv6) performed slightly better than OSPF Version 2 (OSPFv2) (IPv4). At 100% memory utilization, the Wedge router handled 1 million IPv6 routes using OSPFv3, compared to 960k IPv4 routes using OSPFv2. The traffic rate of the streams remained consistent for all of the routes tested.

Route Flapping

The convergence times of OSPFv2 and v3 were tested using three route-flapping cases: the stopping of Hello messages, withdrawing routes, and breaking the physical link. The convergence time was then measured from the point when the routes stopped flapping, to the point when traffic flow resumed to previous levels.

For each of the three route flapping cases tested, IPv6 performed similar to the IPv4 baseline. These results indicate that IPv6 has no performance impact on the OSPF convergence times. Table D-10 compares the convergence results collected for IPv4 and IPv6.

Table D-10 Convergence Results

Route Flapping (OSPF)	Convergence Time (Seconds)	
	IPv4 Baseline	IPv6 Dual Stack
Stop Sending Hellos	15	14
Withdrawing Routes	18	22
Break Physical Link	47	46

Multiprotocol-Border Gateway Protocol (MP-BGP) Routing Test

The results collected show that advertising IPv4 routes using an IPv4 BGP session allows for the highest number of routes to be loaded into the router’s routing table. The two cases that used IPv6 routes performed similar to the IPv4 baseline until it reached about 500,000 routes. At this point, the memory utilization and number of routes held began to deviate. Table D-11 lists the route flapping results.

Table D-11 MP-BGP Results

Route Flapping (MP-BGP)	Convergence Time (Seconds)	
	IPv4 Baseline	IPv6 Dual Stack
Withdraw Routes	11	10
Break TCP Session	10	6
Break Physical Link	89	77

Network Access Point (NAP) Failover Tests

The ITA network currently has three Network Access Points (NAP) referred to as NAP A, NAP B, and NAP C. This test ensured that IPv6 traffic would properly flow through the NAPs and test failover of the NAPs. This was accomplished by enabling IPv6 functionality on the NAP firewalls and routers. The primary NAP that a customer may use converges to a secondary and tertiary path upon failure in the primary NAP. Customers can fall into three NAP routing categories: ABC, BAC, or CAB.

The results of the failover testing showed that IPv6 failover performance was similar to the IPv4 baseline for customers with a NAP preference of ABC and BAC. The failover results are shown in Table D-12.

Table D-12 Network Access Points Test Results

NAP Preference Failover	Convergence Time (Seconds)	
	IPv4 Baseline	IPv6 Dual Stack
Preference ABC A->B	36	36
Preference BAC->A	38	35

Conclusions/Recommendations

The Core Routing tests verified that the ITA network infrastructure could reliably handle an IPv4-IPv6 dual stacked environment. For the test cases evaluated, the routing performance and functionality remained similar for IPv4 and IPv6. Therefore, a dual-stack routing infrastructure will allow ITA to transition its network and provide IPv6 services to its customers.

D.13 Joint Staff Internet Protocol Version 6 Operational Criterion 3 Test Report

Testing Organization and Publication Date

Joint Interoperability Test Command
November 2007

Summary

This test was designed to compare the end-to-end network performance characteristics of IPv6 in relation to IPv4. The Joint Interoperability Test Command's Advanced IP Technology Facility personnel at Fort Huachuca, Arizona, conducted testing of IPv6 in end-to-end networks from July 3 through September 14, 2007. Tests characterized response time across an IP-based network.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.1.1, 2.1.2, 2.3, 2.3.1, 2.3.2)

3 (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)

8 (8.1, 8.1.1, 8.1.2)

Configuration

Testing was done on an end-to-end network in a dual stack environment. End-to-end testing included the protocols that made up the most commonly used applications on the NIPRNet. These protocols were HTTP, Simple Mail Transfer Protocol (SMTP), Motion Picture Expert Group 2 (MPEG2), and SIP/Voice over Internet Protocol (VoIP). The VoIP was tested; while not commonly used on the NIPRNet today, it offered an effective user experience evaluation environment, due to the delay-sensitive nature of VoIP packets.

The network equipment used in testing is presented in Table D-13.

Table D-13 Network Equipment Configuration

Device	Software Version
Cisco 3745 Router	12.4(11)T
Cisco 3845 Router	12.4(11)T
Juniper T640 Router	JUNOS 7.6R3.6
Juniper T320 Router	JUNOS 7.6R3.6
Juniper M40e	JUNOS 7.6R3.6

Test equipment configuration during testing is listed in Table D-14.

Table D-14 Test Equipment Configuration

Component	Operating System	Application Software	Functionality
Dell PowerEdge 2950 Server	Windows 2003 Server R2	Agilent N4190B/ NetPressure 3.7.73	Protocol Traffic Loading
Gateway 4100E Desktop	Windows Vista Enterprise 6.0	IBM NetBIOS 3.0	Client Traffic
Dell PowerEdge 2950 Server	Windows 2003 Server Standard Edition	Network General InfiniStream 4.0.237	Protocol Traffic Analysis
Gateway 4100E Desktop	Red Hat Enterprise Linux ES release 4 (Nahant Update 5) Linux Version 2.6.9-55.0.2.EL NMON utility	N/A	Client Traffic
Gateway 4100E Desktop	Windows XP Professional SP2	Spirent Smartbits 600B 2.80 LAN-3325A Tera-Metrics 6.00 Avalanche 7.56	Bit-level Loading
Gateway 4100E Desktop	Windows XP Professional SP2	Spirent Smartbits 600B 2.80 LAN-3325A Tera-Metrics 6.00 Smartbits 5.50	Bit-level Loading
Gateway 4100E Desktop	Windows XP Professional SP2	Spirent ClearSight 3.2.0.25	Protocol Traffic Analysis
Gateway 4100E Desktop	Windows XP Professional SP2	WireShark 0.99.5 Microsoft Office Excel SPSS V. 15.0	Protocol Traffic Analysis
Gateway 450ROG Laptop	Windows XP Professional V.2002 SP1	Spirent Smartbits 6000C 2.80 POS-3519A Tera-Metrics Smartbits 5.50	Bit-level Loading

Results

HTTP

The average response times for the IPv4/IPv6 ratios shown in Table D-15 were compared for equality in Statistical Package for the Social Sciences (SPSS) using a *t*-test. These results showed no significant difference in response time averages at the 99% confidence interval. The average response times for the 90% ratio were compared using the same *t*-test. These results showed a statistically significant difference of one millisecond (ms) at the 99% confidence interval. While this statistical difference was found with the 90% IPv4/IPv6 ratio, the absolute value of the difference is so small that no operational impact is expected.

Table D-15 HTTP IP Network Ratio Comparison Results

HTTP				
	100% IPv4/IPv6 Ratio	90% IPv4/IPv6 Ratio	50% IPv4/IPv6 Ratio	10% IPv4/IPv6 Ratio
Metrics				
Average Response Time (ms)	68 / 68	106 / 105	106 / 106	105 / 105
Average Response Time Standard Deviation (ms)	5.40 / 2.10	4.70 / 4.62	5.22 / 4.21	5.44 / 9.41
Error +/- (ms) for Average Response Time	.889	1.40	1.16	1.73
Sample Size (sessions)	1106 / 1214	870 / 1123	893 / 896	1124 / 1035
Throughput (Mbps)	35 / 35	29 / 31	31 / 32	28 / 33
Packet Loss	0 / 0	0 / 0	0 / 0	0 / 0
Packet Re-Ordering	0 / 0	0 / 0	0 / 0	0 / 0
Packet Size Distribution				
<65	0 / 0	0 / 0	0 / 0	0 / 0
65-127	1238327 / 673054	1526727 / 1918452	928554 / 1213612	213756 / 217374
128-255	91818 / 18078	117230 / 51712	71746 / 36549	14902 / 7696
256-511	27300 / 6394	34860 / 18348	21000 / 12927	4200 / 2780
512-1518	33900 / 0	43159 / 0	26000 / 0	5200 / 0
>1518	186060 / 21022	232974 / 59867	141186 / 42501	30898 / 9140

SMTP

The average response times for the 100, 90, 50, and 10% IPv4/IPv6 ratios shown in Table D-15 were compared for equality in SPSS using a *t*-test. These results showed no significant difference in response time averages at the 99% confidence interval.

MPEG 2

Based on this limited data set of five MPEG2 sessions, no differences were observed in the Mean Opinion Score (MOS). Additional T&E is necessary to effectively streaming video performance over IPv4/IPv6 effectively.

SIP/VoIP

With no suitable IPv6 enabled VoIP products available from vendors, automated test equipment was substituted as the only available means to conduct testing. These results should be seen as representing an immature IPv6 SIP/VoIP environment. When vendor implementations become available, additional T&E will be necessary to characterize SIP/VoIP performance properly. While no firm conclusions regarding the relative performance of IPv4 and IPv6 should be drawn from this limited data set, call completion rates and MOS are the same for IPv4 and IPv6.

Workstations and Server

The results in Table D-16 show IPv4 and IPv6 performance comparison results for each of the three separate Operating Systems (OS) and hardware combinations. These results indicate workstation and server performance parity between IPv4 and IPv6.

Table D-16 Workstation and Server IPv4/IPv6 Comparison Results

Protocol: HTTP	
Windows Vista Workstation	
Metrics	100% IPv4/IPv6 Ratio
CPU utilization %	5 / 5
Memory utilization %	52 / 52
Network utilization %	1 / 1
Packets per second	113 / 113
Windows 2003 Server	
CPU utilization %	5.9 / 4.4
Packets per second	102 / 107
Red Hat Server	
CPU utilization %	2.6 / 2.4
Memory utilization %	39 / 39

Conclusions/Recommendations

The most critical measure for this test, HTTP response time performance, was operationally equivalent for all traffic ratios. The SMTP performance was equivalent for all traffic ratios. The MPEG2 performance results showed IPv4/IPv6 equivalency but additional T&E is necessary to develop a statistically reliable sample. While using automated test equipment to simulate a SIP/VoIP system, IPv4/IPv6 equivalency was shown. To characterize SIP/VoIP properly, additional T&E is necessary once vendor implementations become available. The combination of testing on Windows Vista, Windows 2003, and Red Hat server showed performance parity.

D.14 Demonstration of Operation of IPv6 in a Simulated Low Bandwidth Environment

Testing Organization and Publication Date

Air Force, MITRE
September 28, 2007

Summary

The primary objective of this test was to quantify the operational impact of IPv6 traffic versus equivalent IPv4 traffic. The secondary objective was to quantify critical metrics for new features of the IPv6 protocol where no IPv4 equivalent exists in the low-bandwidth environment. The final objective was to characterize the operational impact of transition mechanisms specific to the low-bandwidth tactical environment. This work considers the operational parameters of the Theater Deployable Communications (TDC) and is restricted to testing the operational effect of IPv6 in below T1 Satellite Communication (SATCOM) links. This work is also bounded by consideration of NIPRNet topology.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

- 1** (1.1, 1.1.1, 1.2, 1.2.1)
- 3** (3.1, 3.1.1, 3.2, 3.2.1, 3.3, 3.3.1)
- 4** (4.1, 4.1.1)
- 5** (5.1, 5.1.1)
- 7** (7.1, 7.1.3)
- 8** (8.1, 8.1.1, 8.1.2)

Configuration

A baseline of IPv4 traffic was quantified and compared against the IPv6 equivalent, where possible. For new IPv6 features without equivalent in IPv4, the same metrics were taken and analyzed. TDC staff was interviewed and the topology scenarios were detailed and built. To imitate a SATCOM link, an SX-14 simulator was used with the standard latency of 250 ms per uplink, and a bit error rate (BER) of varying 10^{-7} to 10^{-9} as implemented for each scenario. From the Agilent N2X traffic generator, packet sizes were varied from 64 bytes to 1500 bytes for each bandwidth tested (e.g., 64, 128, 256, 512, 768, and 1500 Kbps). Bandwidth for the SATCOM simulator was varied from 128 Kbps to 1024 Kbps. The maximum percentage of traffic stream was pushed from the Agilent without packet loss so that Mbps rates could be accurately measured for the return traffic without subtraction for frame drop.

Results

Testing encompassed numerous scenarios. An IPv4 baseline was established first for Quality of Service (QoS), mobility, Anycast Multicast, and IPsec. The test network was then configured for the use of transition mechanisms; e.g., dual stack, Network Address Translation-Protocol Translation (NAT-PT), Protocol 41 (IPv6inIP), and GRE were tested as transition mechanisms in the low bandwidth tactical environment. The pattern seen consistently when comparing the IPv6 throughput per packet size to the IPv4 was that increasing the packet size closed the gap between the IPv4 baseline throughput and the IPv6 total throughput. This was expected because the larger the packet size, the less the impact of the additional header on a per packet basis. This finding was consistent for all transition mechanisms.

Dual Stack

Throughput - Averaged 87.22% total throughput of the IPv4 baseline for each SATCOM bandwidth and packet size.

Latency - Averaged 7 ms higher than the IPv4 baseline for any given packet size and bandwidth for dual stack.

Round Trip Time (RTT) - Average RTT for dual stack for all bandwidths and packet sizes was 549.9 ms.

Bit Error Rate (BER) - This group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

NAT-PT

Throughput - NAT-PT during this testing was the second best performer of the four mechanisms for both throughput and latency. Throughput was on average 83.9% of the dual stack throughput results for each SATCOM bandwidth in each packet size.

Latency - Average latency for all bandwidths and packet sizes was 5 ms higher than that of dual stack.

RTT - Average RTT was 573.9 ms for all bandwidths and packet sizes.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

IPv6 GRE

Throughput - Average throughput with GRE protocol 47 statically configured tunnels had 26.7% less throughput than the base dual measurement stack for the bandwidths measured.

Latency - Averaged 21.2 ms average latency and 81.9 ms maximum average latency relative to dual stack.

RTT - Average RTT for this group for all bandwidths and packet sizes is 592.3 ms.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

IPv6inIPv4 (Protocol 41)

Throughput - Average throughput with IPv6inIPv4 was 17.5% lower than that of dual stack for all measured SATCOM bandwidths.

Latency - For all bandwidths and packet lengths, latency for this transition mechanism was 15 ms greater than the average latency for dual stack.

RTT - The RTT for this group was 579 ms average for all bandwidths and packet sizes.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

Multicast

Throughput - Average throughput for the group under test was 5% higher on average than that found for the IPv4 Anycast baseline for all packet sizes and bandwidths measured.

Latency - Average latency was 281.95 ms for all packet sizes and bandwidths measured.

RTT - Average RTT for this group was 563.9 ms for all packet sizes and bandwidths measured.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER for all packet sizes and bandwidths measured.

IPsec

Throughput - For the use of the ESP transform set alone, the throughput was decreased across all bandwidths (for respective packet sizes) an average of 6%. For the use of AH as the transform set alone, the throughput was decreased across all bandwidths (for respective packet sizes) an average of 4%.

Latency - Comparisons revealed no more than 1% difference in IPv4 IPsec average latency for AH and ESP. This 1% difference was averaged across all packet sizes and bandwidths measured in this study. This will add a penalty of the entire IPsec header to every fragmented packet in IPv6, as opposed to the IPv4 IPsec overhead of only the initial fragment.

RTT - Averaged no more than a 1% difference in the RTT between the IPv4 baseline and each of the respective IPv6 averages taken for the ESP and AH transform sets for all packet sizes and bandwidths measured.

BER - The BER for this group showed a 2% degradation of throughput and a 1% effect on latency for a BER of 10^{-7} over the IPv4 baseline. This finding was the same for ESP and AH transform sets for all packet sizes and bandwidths measured.

QoS

Throughput - The throughput for this group varied less than 1% on average from that of the dual stack group. The same methodology was used for the marked traffic with the non-guaranteed DSCP 0, best effort, and traffic. The ultimate throughput of the marked and unmarked traffic was the same as dual stack.

Latency - The latency for this group was within 2% on average for all packet sizes and bandwidths of the latency for the dual stack group.

RTT - The RTT for this group was within 2% on average for all packet sizes and bandwidths of the latency for the dual stack group.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

Mobility

Throughput - The throughput for this group was the same as the dual stack group.

Latency - Average latency was 124 ms for the handover of the home node and identical to that of the dual stack environment for all other cases (once the handover is complete).

RTT - This group was identical to that of the dual stack group for all packet sizes and bandwidths measured.

BER - The BER for this group showed no statistical difference in terms of latency or throughput for the packets transmitted from the same traffic streams without BER.

Conclusions/Recommendations

This demonstration served to quantify some critical methodologies as put forth in the Joint Staff Operational Criterion 5 Level 4 decomposition. It was successful in quantifying for the NIPRNET-simulated architecture in lab. The effects of varying SATCOM bandwidths and differing packet sizes streaming through this architecture are likely to have significant operational impact on the tactical environment. These findings can be used to aid not only architectural planning in a low bandwidth environment but also to serve as a catalyst and initial recommendation for which transition methodologies most appropriately need to be implemented for lessening the overhead penalty in a dual stack environment.

There is a recognized need for IPv6 testing in a low bandwidth environment. Overall, the findings in this report provide an initial, preliminary level of information that will assist the DoD in its transition to IPv6. Additional study is required in this area and more testing in a larger scale, operational network environment should occur.

D.15 Technical Report For Network Management IPv6 Initiative (NMI2) (Tool Analysis)

Testing Organization and Publication Date

Air Force Research Laboratory (AFRL)
October 2007

Summary

This document explores the initial results and conclusions of the effort called NMI2. NMI2 is follow on work by AFRL to obtain a current network management “snapshot” and further investigate the effects of network management within a dual stack (IPv4/IPv6) environment. Application functionality and transition capabilities were explored primarily from the network management server-side of the equation with a follow on document to further detail the client interactions with Network Management (NM) server tools.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

8 (8.1, 8.1.1)

9 (9.1, 9.1.1, 9.1.2, 9.1.3)

Configuration

The DoD goal is for a dual-stack environment; thresholds were set based on this premise. Thus, if a tool uses only IPv4 to handle IPv6 MIB information, that tool meets the threshold. The objective and threshold goals were considered met if the majority of the tools met the goal. Since five tools were reviewed, three out of the five (3/5) would be the majority. Here, only server tool results were included. This test encompassed designing and implementing a dual-stack test bed, and conducting several tests to evaluate NM tools and determine whether effective/equivalent NM could be performed, as is typical for DoD installations. The test explored the advertised IPv6 capabilities of the following network management tools: What’s Up Gold v11.01, Smarts InCharge 6.5.1, NeuralStar 8.0.3, HP Openview’s Network Node Manager 7.5, and CiscoWorks LAN Management 2.5.1. It also explored how the use of the major network management protocol, SNMP, was integrated within the dual stack realm.

Results

Decomposition 9.1.1 - Basic Protocol Support

Of the tools reviewed, 5/5 could support SNMPv1 and SNMPv2 over the IPv4 protocol, while 4/5 could support SNMPv3 over the IPv4 protocol. However, 1/5 of the tools could provide support for SNMPv1, 2, and 3 over the IPv6 protocol.

Decomposition 9.1.2 - Basic Monitoring / Fault Functionality

Of the tools tested, 3/5 of the tools would poll hosts by sending ICMPv6 requests; all tools had some sort of display result (e.g., icon) from the verification received of the nodes presence on the network. Devices with IPv4 and IPv6 addresses were displayed for all of the tools. However, the display of the IPv6 address depended upon the client used. Regarding the GetRequests sent by the tools, 5/5 sent Management Information Base (MIB) queries from the Generic MIB, RFC 4293, Cisco, and Juniper MIBs over IPv4, while only 1/5 sent these queries over IPv6. For MIB queries from RFC 4113, RFC 2466, and RFC 2452, 4/5 were sent over IPv4, while only 1/5 sent these queries over IPv6. For MIB queries from RFC 2465, 3/5 were sent over IPv4, while only 1/5 sent these queries over IPv6.

Decomposition 9.1.4 - Autodiscovery / Discovery Behavior

For automatic discovery of IPv4 devices on the network, 4/5 tools accomplished this, but only 1/5 could do the same for IPv6 devices. With the use of a seed file to discover network devices, 5/5 tools accomplished this when IPv4 devices were being discovered, but only 2/5 could do so for IPv6 devices. Thus, 3/5 tools could perform some form of auto-discovery of IPv4 and IPv6 enabled devices.

Decomposition 9.2.1 - Basic Configuration Capability

In sending SetRequests, 4/5 NM tools could send them using the IPv4 address of the devices being set, but no tools could do the same using the IPv6 address of the nodes.

Polling with IPv6 – An “Out of the Box” Look

One would assume that using IPv6 would take more time for the polling to complete, since the IPv6 header size is 40 bytes to IPv4's default header size of 20 bytes. Even though times were close, IPv6 took less time to discover the devices. An explanation could be from unique tool-centric steps in executing polls that are different between the protocol implementations, or possibly an apparent increase in the amount of traffic on the wire during the IPv4 tests. Two packet captures look virtually identical. The frame size is 20 bytes more for the IPv6. For the IPv4 and IPv6 header breakouts, there are 18 different fields to be processed during the traffic exchange, whereas there are only eight in the IPv6 header.

Table D-17 presents the decomposition, and results of testing.

Table D-17 Network Management Results

Level 3 Decomposition	Level 4 Decomposition		Result
9.1.1 Basic Protocol Support	9.1.1.1	Tool: SNMP Version Support	G
	9.1.1.2	Client: SNMP Version Support	APR 08 Report
9.1.2 Basic Monitoring/Fault Functionality	9.1.2.1	Tool: NM Tool SNMP Get request (three tests: Generic MIB, IPv6 MIB, Private MIB) to dual-stack Server/Router. Displays results. Also ICMPv6 send/rcv/display.	Y
	9.1.2.2	Client: Dual-stack client generates SNMP get Response (Generic MIB, IPv6 MIB, Private MIB). Responds to ICMPv6.	APR 08 Report
	9.1.2.3	Client: Dual stack client generates trap	APR 08 Report
	9.1.2.4	Tool: NM tool receives basic trap and appropriately displays.	G
9.1.3 Help Support	9.1.3.1	Tool: NM Tool – Documentation v6 support, online v6 support, Help Desk v6 support	Y
	9.1.3.2	Client: Server/Router Client – Documentation v6 support, online v6 support, Help Desk v6 support	APR 08 Report
9.1.4 Autodiscovery/Discovery Behavior	9.1.4.1	Tool: Dual stack autodiscovery is done how? (i.e., autodiscovery initiation)	Y
	9.1.4.2	Tool: Compare and contrast how autodiscoveries are performed from v4 vs. dual stack (i.e., autodiscovery execution)	G
	9.1.4.3	Tool: What are the differences in the autodiscovery results of v4-only vs. dual stack (i.e., post-autodiscovery)	Y
9.1.5 Performance Scalability	9.1.5.1	Tool: Test ability of NM tool to manage 1000s of nodes (general management and scaled trap handling)	N/A
9.1.6 Display	9.1.6.1	Tool: Compare and contrast how NM Tool displays MIB information	G
9.1.7 Display Scalability	9.1.7.1	Tool: How well does the NM Tool display a large enterprise comprised of v4 and dual stack clients?	N/A
9.2.1 Basic Configuration Capability	9.2.1.1	Tool: NM Tool SNMP Set (e.g. 1.3.6.1.2.1.1.4)	G
	9.2.1.2	Client: Sets the value as identified in SNMP Set.	APR 08 Report
	9.2.1.3	Client: Capable of being configured as dual stack	APR 08 Report
9.2.2 Advanced Configuration Capability	9.2.2.1	Tool: Capable of recognizing when a client has gone from v4-only to dual stack	Y
9.3.1 Basic Accountability	9.3.1.1	Tool: NM Tool identifies and correctly displays dual stack client’s information as queried under Basic Fault Management	Y
	9.3.1.2	Tool: NM Tool correctly identifies which devices are IPv4 only and which are dual stack; clearly apparent that and IPv4 and IPv6 address are coming from a single physical device.	Y
9.3.2 Out-of-the-Box Performance	9.3.2.1	Tool: (Compare v4 vs. dual stack using out-of-the-box (i.e., default) MIB queries for v4 vs. v6) Handling of various client distributions (v4 vs. dual stack) over increasing number of clients over different IP distributions	APR 08 Report
	9.3.2.2	Tool: (Compare v4 vs. dual stack using out-of-the-box (i.e., default) traps for v4 vs. v6) Trap handling (handling of increasing number of traps over same type of distributions above)	APR 08 Report
9.3.3 “Perfect Performance”	9.3.3.1	Tool: (Compare v4 vs. dual stack using a one-to-one mapping of similar queries for v4 vs. v6) Handling of various client distributions (v4 vs. dual stack) over increasing number of clients over different IP distributions	N/A
	9.3.3.2	Tool: (Compare v4 vs. dual stack using a one-to-one mapping of similar queries for v4 vs. v6) Trap handling (handling of increasing number of traps over same type of distributions above)	N/A
9.3.4 Complex Network Accounting	9.3.4.1	Tool: Complex handling of an environment that includes multiple: dual stack enabled clients/routers, v4 only clients/routers, tunnels, IPsec, DNS instantiations, DNSSEC, other network-layer interactive devices	N/A
Results Column Key			
Met Objective Goal (60% or more of tools met “Objective” criteria)		G	
Met Threshold Goal Only (60% or more of tools met “Threshold” criteria)		Y	
Met Neither Goal		R	
Due in Upcoming Report		APR 08 RPT	
No Test Scheduled		N/A	

Conclusions/Recommendations

A dual stack network can be managed using a combination of IPv4 and IPv6 tools.

D.16 Special Interoperability Test Certification of Ambriel ATX-S Series IPv4/IPv6 Translator device for Internet Protocol Version 6 (IPv6) Capability Combined

Testing Organization and Publication Date

Joint Interoperability Test Command
April 10, 2008

Summary

This report presents the results of two Special Interoperability Test Certifications of the Ambriel ATX-S Series IPv4/IPv6 Translator. These certifications are for Ambriel Devices running Red Hat Enterprise 4 Nahunt update 5 and Linux Kernel 2.6.9-55. This device meets the IPv6 Capable interoperability requirements of a Simple Server and is certified for listing as IPv6 Capable. This special certification is based on IPv6 Capable testing conducted by JITC at Fort Huachuca, Arizona. Testing commenced on January 14 and was completed on January 18, 2007 at JITC's Advanced IP Technology Capability.

This test was conducted by installing the Ambriel ATX-S Series IPv4/IPv6 Translator on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.2, 2.3)
8 (8.1.1, 8.2.1)

Configuration

Table D-18 lists the configuration of the device as it was certified.

Table D-18 Ambriel Configuration

Ambriel ATX-S Series IPv4/IPv6 Translator		
Component	Firmware/Software	Interface
2EA - ATX-S4400	Linux Kernel 2.6.9-55 / Red Hat Enterprise 4 nahunt update 5	RJ45 10/100 Mbps Ethernet
LEGEND:		
Mbps	Megabits Per Second	RJ Registered Jack

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-19 lists the functional category requirements for both certifications, and verifies if those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-19 Ambriel Technologies Interoperability Status Summary

Ambriel Technologies AT-X-S4400 IPv4/IPv6 Translator		
Functional Category	Critical	Verified
Base IPv6	M	Yes
Network Service	M	Yes
IPsec	S+	No
Transition Mechanisms	M	Yes
Quality of Service	S	No
Other Requirement	S	No
LEGEND: IPv6 Internet Protocol Version 6 S Should IPsec Internet Protocol Security S+ Should+ M Must		
NOTE: The terms Must, Should, and Should+ are used to reference specific required Request for Comments from the Internet Engineering Task Force, the DoD Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The Ambriel ATX-S Series IPv4/IPv6 is certified for listing as an IPv6 Capable simple server.

D.17 Special Interoperability Test Certification of TechGuard PoliWall Version 1.21.00 with Ethernet Interface and TechGuard PoliWall Version 1.21.00 with Fiber Interface for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
March 6, 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the TechGuard PoliWall. This device meets the IPv6 Capable interoperability requirements of a network appliance. This special certification is based on IPv6 Capable testing conducted by JITC at Fort Huachuca, Arizona. Testing commenced on November 26 and was completed on November 30, 2007 at JITC's Advanced IP Technology Capability.

This test was conducted by installing the TechGuard PoliWall on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

The device was tested as a network appliance only. Testing of the firewall function of this device was not conducted.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1.1.1, 1.2.1.1, 1.3.1.1)

2 (2.1, 2.3)

8 (8.1.1)

Configuration

Table D-20 lists the configuration of the device as it was certified.

Table D-20 Certification of TechGuard PoliWall Test Configuration

TechGuard PoliWall		
Equipment Name	Firmware/Software	Interfaces
TechGuard PoliWall – Ethernet Interface	1.21.00	RJ45 10/100 Mbps Copper Ethernet
TechGuard PoliWall – Fiber Interface	1.21.00	1 Gbps Ethernet on MMF
LEGEND:		
Gbps	Gigabits Per Second	MMF Multimode Fiber
Mbps	Megabits Per Second	RJ Registered Jack

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-21 lists the functional category requirements and verifies if those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-21 Test Results for Functional Test Category

TechGuard PoliWall with Ethernet Interface and TechGuard PoliWall with Fiber Interface		
Functional Category	Requirement	Verified
IPv6 Base	M	Yes
Network Service	M	Yes
IPsec	S+	Yes
Transition Mechanisms	S	Yes
QoS	S	No
Mobility	S	No
Other Requirements	S	No
LEGEND:		
IPsec	Internet Protocol Security	QoS Quality of Service
IPv6	Internet Protocol Version 6	S Should
M	Must	S+ Should +
NOTE: The terms Must, Should, and Should+ are used to reference specific required Request for Comments from the Internet Engineering Task Force, the DoD Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The TechGuard PoliWall is certified for listing as an IPv6 Capable network appliance.

D.18 Special Interoperability Test Certification of Quantum Autoloader SuperLoader3 backup device Running Build Number v55-0 and InterNiche 3.1 Dual Stack Core and Quantum Scalar i500 Midrange Scalable Tape Library backup device Running Firmware Version 410G.GS007 and Linux Kernel 2.6.11-1 for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
March 4, 2008

Summary

This report presents the results of the Special Interoperability Test Certification, the Quantum Autoloader SuperLoader3 backup device configured running Build number v55-0 and InterNiche 3.1 dual stack core, and the Quantum Scalar i500 Midrange Scalable Tape Library backup device configured running firmware Version 410G.GS007 and Linux Kernel 2.6.11-1. This device meets the IPv6 Capable interoperability requirements of a simple server.

This test was conducted by installing the Quantum Autoloader on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)
8 (8.1.1)

Configuration

Table D-22 lists the configuration of the device as it was during certification.

Table D-22 Quantum Configuration

Quantum Autoloader SuperLoader3 and Quantum Scalar i500		
Component	Firmware/Software	Interface
Quantum Autoloader SuperLoader3	Build Number v55-0/InterNiche 3.1 Dual Stack Core	RJ45 10/100 Mbps Ethernet
Quantum Scalar i500	Firmware Version 410G.GS007/Linux Kernel 2.6.11-1	RJ45 10/100 Mbps Ethernet
LEGEND:		
Mbps	Megabits Per Second	RJ Registered Jack

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jtc.fhu.disa.mil/apl/>.

Table D-23 lists the functional category requirements, and verifies that those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-23 Quantum Test Results for Functional Test Category

Quantum Autoloader SuperLoader3 and Quantum Scalar i500		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
Network Service	M	Yes
IPsec	S+	No
Transition Mechanisms	S	Yes
Quality of Service	S	No
Other Requirement	S	No
LEGEND:		
IPv6	Internet Protocol Version 6	S Should
IPsec	Internet Protocol Security	S+ Should+
M	Must	
NOTE: The terms Must, Should, and Should+ are used to reference specific required Request for Comments from the Internet Engineering Task Force, the DoD Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The Quantum Autoloader is certified for listing as an IPv6 Capable simple server.

D.19 Special Interoperability Test Certification of the IBM Storage System TS3100 Tape Library Express and IBM Storage System TS3200 Tape Library Express Families of Tape Libraries Running Nucleus Net Version 5.4b, Nucleus Net Internet Protocol (IP) Version 6 (IPv6) Version 1.4b, Firmware Version 6.20/2.6EZ, and Nucleus Version 1.15 Operating System (OS) Running a Linux-Based Kernel and the IBM Storage System TS3400 Tape Library Running CENTE Version 1.30, Firmware Version 0001.6000, and uITRON Version 4.0 OS Running a Linux-Based Kernel for IPv6 Capability410G.GS007 and Linux Kernel 2.6.11-1for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
May 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the IBM Storage System TS3200 Tape Library Express backup. This device has met the IPv6 Capable interoperability requirements of a Simple Server. While only the TS3200 was tested, the other server within this family (TS3100) is architecturally equivalent to the IBM Storage System TS3200 Tape Library Express Simple Server and utilizes the same OS; therefore, this certification applies to the family of simple servers. Interoperability testing was conducted from February 18-29, 2008 at JITC's Advanced IP Technology Capability.

This test was conducted by installing the IBM Storage System TS3200 Tape Library on a dual stack IP network, and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)

8 (8.1.1)

Configuration

Table D-24 lists the configuration of the device as it was certified.

Table D-24 IBM Storage System Tape Library Configuration

IBM TS3200 and TS3400		
Component	Firmware/Software	Interface
IBM TS3200 Tape Library	Dual Stack Core; Firmware Version 6.20/2.6 Nucleus (Linux) Kernel Version 1.15	RJ45 10/100 Mbps Ethernet
IBM TS3400 Tape Library	Dual Stack Core; Firmware Version 0001.6000 uITRON (Linux) Kernel Version 4.0	RJ45 10/100 Mbps Ethernet
LEGEND:		
Mbps	Megabits Per Second	RJ Registered Jack

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-25 lists the functional category requirements, and verifies if those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-25 IBM Storage System Tape Library Test Results for Functional Test Category

IBM TS3200 and IBM TS3400		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
IPsec	S+	No
Transition Mechanisms	S	Yes
Quality of Service	O	No
Other Requirement	S	No
LEGEND:		
IPsec	Internet Protocol Security	O Optional
IPv6	Internet Protocol Version 6	S Should
M	Must	S+ Should+
NOTE: The terms Must, Should, Should+, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the DoD Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The IBM Storage System Tape Library is certified for listing as an IPv6 Capable simple server.

D.20 Special Interoperability Test Certification of Cisco Catalyst 4500 Family of Layer 3 Switches with Supervisor Engine 6-E Running Internetworking Operating System Version 12.2(40)SG, for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
April 10, 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Cisco Catalyst 4510R Layer 3 Switch with Supervisor Engine (Sup) 6-E running IOS Version 12.2(40)SG. This device meets the IPv6 Capable interoperability requirements of a Layer 3 Switch. While only the 4510R was tested, the other switches within this family (C4503, C4503-E, C4506, C4506-E, C4507R, C4507R-E, C4510R, C4510R-E) are architecturally equivalent to the Cisco Catalyst 4510R Layer 3 Switch and utilize the same IOS; therefore, this certification applies to the entire Cisco 4500 family of Layer 3 switches with Sup 6-E running IOS Version 12.2(40)SG. The Cisco Catalyst 4510R with Sup 6-E running IOS Version 12.2(40)SG successfully completed the related IPv6 Interoperability portions of the “DoD IPv6 Generic Test Plan (GTP) Version 3,” August 2007. This certification test was conducted from October 11, 2007 through November 21, 2007.

This test was conducted by installing the Cisco Catalyst 4510R Layer 3 on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)

8 (8.1.1)

Configuration

Table D-26 lists the configuration of the device as it was certified.

Table D-26 Cisco Catalyst 4510R Layer 3 Switch Configuration

Cisco Catalyst 4510R Layer 3 Switch		
Component	Firmware/Software	Interface
Cisco Catalyst 4510R	Cisco IOS Version 12.2(40)SG	RJ45 10/100/1000 Mbps Ethernet SFP 1000 Mbps Ethernet
LEGEND:		
IOS	Internetworking Operating System	RJ Registered Jack
Mbps	Megabits Per Second	SFP Small Form Factor Pluggable

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-27 lists the functional category requirements, and verifies if those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-27 Cisco Layer 3 Switch Test Results for Functional Test Category

Cisco Catalyst 4510R Layer 3 Switch		
Functional Category	Critical	Verified
IPv6 Base	M	Partial
IPsec	S	No
Transition Mechanisms	O	Yes
Network Management	CM	Yes
Other Requirements	S	No
LEGEND:		
CM	Conditional Must	M Must
IPsec	Internet Protocol Security	O Optional (May)
IPv6	Internet Protocol Version 6	S Should
NOTE: The terms Must, Conditional Must, Should, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the Department of Defense Information Technology Standards Registry Department of Defense IPv6 Standard Profiles for IPv6 Capable Products Version 2.0, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The Cisco Catalyst 4500 family of Layer 3 switches is certified for listing as IPv6 Capable Layer 3 switches.

D.21 Special Interoperability Test Certification of Cisco Catalyst 6500 Family of Layer 3 Switches with Supervisor Engine 720 Running Internetworking Operating System Version 12.2(33)SXH for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
May 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Cisco Catalyst 6506-E Layer 3 Switch with Supervisor Engine (Sup) 720 running IOS Version 12.2(33)SXH. This device meets the IPv6 Capable interoperability requirements of a Layer 3 Switch. While only the 6506-E was tested, the other switches within this family (C6503-E, C6504-E, C6509-E, C6509-NEB-A, C6513, ME 6524) and supervisor engines (Sup 720-10G, Sup 32, Sup 32-Pisa) are architecturally equivalent to the Cisco Catalyst 6506-E Layer 3 Switch, and utilize the same IOS; therefore, this certification applies to the entire Cisco 6500 family of Layer 3 switches with IOS Version 12.2(33)SXH.

This test was conducted by installing the Cisco Catalyst 6506-E Layer 3 Switch on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)

8 (8.1.1)

Configuration

Table D-28 lists the configuration of the device as it was certified.

Table D-28 Cisco Catalyst 6506-E Layer 3 Switch Configuration

Cisco Catalyst 6506-E Layer 3 Switch		
Component	Firmware/Software	Interface
2 Cisco Catalyst 6506-E	Cisco IOS Version 12.2(33)SXH	RJ45 10/100/1000 Mbps Ethernet
LEGEND:		
IOS	Internetworking Operating System	RJ Registered Jack
Mbps	Megabits Per Second	

D.22 Special Interoperability Test Certification of Cisco 2800 Integrated Services Router (ISR) Family of Routers Running Internetworking Operating System Version 12.4(11)T bundled with the 7600 Family of Routers Running Internetworking Operating System (IOS) Version 12.2(33)SRB1 System and Cisco 3800 ISR Family of Routers Running Internetworking Operating System Version 12.4(11)T bundled with the 7600 Family of Routers Running IOS Version 12.2(33)SRB1 System for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command

May 2008

Summary

In previous testing, the Cisco 2800 and 3800 families of routers were certified as IPv6 capable routers. This test verified those two router's individual firewall capability in two bundled systems with the Cisco 7600 family of routers. The Cisco 2851 ISR Running IOS Version 12.4(11)T bundled with the 7609 Router Running IOS Version 12.2(33)SRB1 (2800/7600 System) met the IPv6 Capable interoperability requirements of an exterior router as described in the DoD Information Technology Standards Registry, "DoD IPv6 Standard Profiles for IPv6 Capable Products Version 2.0," August 1, 2007. The Cisco 3845 ISR Running IOS Version 12.4(11)T bundled with the 7609 Router Running IOS Version 12.2(33)SRB1 (3800/7600 System) also met the IPv6 Capable interoperability requirements of an exterior router. However, there are routers within these families (2801, 2811, 2821, 2851, 3825, 7603-S, 7604, 7606, 7606-S, 7609-S, and 7613) that were not tested, but the routers are architecturally equivalent and utilize the same IOS, and JITC analysis determined them to be functionally identical for certification purposes.

This test was conducted by installing the tested router on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1.1.1, 1.2.1.1, 1.3.1.1)

2 (2.1, 2.3)

8 (8.1.1)

Configuration

Table D-30 lists the configuration of the device as it was certified.

Table D-30 Cisco 2800/7600 and 3800/7600 Integrated Services Router Configuration

Cisco 2800/7600 System		
Component	Firmware/Software	Interface
Cisco 2851	Cisco IOS Version 12.4(11)T	RJ45 10/100 Mbps Ethernet
Cisco 7609	Cisco IOS Version 12.2(33)SRB1	RJ45 10/100/1000 Mbps Ethernet
Cisco 3800/7600 System		
Component	Firmware/Software	Interface
Cisco 3845	Cisco IOS Version 12.4(11)T	RJ45 10/100/1000 Mbps Ethernet
Cisco 7609	Cisco IOS Version 12.2(33)SRB1	RJ45 10/100/1000 Mbps Ethernet
LEGEND:		
IOS	Internetworking Operating System	RJ Registered Jack
Mbps	Megabits Per Second	T New Technology

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-31 lists the functional category requirements, and verifies if those categories were tested and met the criteria identified in accordance with the associated RFCs.

Table D-31 Cisco 2800/7600 and 3800/7600 Test Results for Functional Test Category

Cisco 2800/7600 and 3800/7600 Systems		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
IPsec	M	Yes
Transition Mechanisms	M	Yes
Quality of Service	M	Yes
Network Management	M	Yes
Interior Router	M	Yes
Exterior Router	M	Yes
LEGEND:		
IPv6	Internet Protocol Version 6	M Must
IPsec	Internet Protocol Security	N/A Not Applicable

Conclusions/Recommendations

The Cisco 2800/7600 and 3800/7600 routers are certified for listing as IPv6 Capable routers.

D.23 Special Interoperability Test Certification of Datatek IPv4/IPv6 Translator device for Internet Protocol Version 6 (IPv6) Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
08 May 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Datatek IPv4/IPv6Transformer running software Version 2.1.4. This device meets the IPv6 Capable interoperability requirements of a host. The Datatek IPv4/IPv6 Transformer was granted a waiver by the DoD IPv6 Standards Working Group for IPsec RFC 4301 and IKE Version 2 RFC 4306, therefore it had to meet the IPsec RFCs 2401, 2402, and 2406, and the IKE Version 1 RFCs 2407, 2408, 2409, and 4109. This special certification is based on IPv6 Capable interoperability testing conducted by JITC at Fort Huachuca, Arizona. Interoperability testing was conducted from March 24, 2008 through April 3, 2008 at JITC's Advanced IP Technology Capability

This test was conducted by installing the Datatek IPv4/IPv6 Translator on a dual stack IP network and verifying with a network sniffing device that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.2 2.3)

8 (8.1.1, 8.2.1)

Configuration

Table D-32 lists the configuration of the device as it was certified.

Table D-32 Datatek Configuration

Datatek IPv4/IPv6 Transformer		
Component	Firmware/Software	Interface
Datatek IPv4/IPv6 Transformer	Datatek Software Version 2.1.4	Ethernet 10/100 Mbps
LEGEND:		
IPv4	Internet Protocol Version 4	Mbps
IPv6	Internet Protocol Version 6	Megabits Per Second

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitc.fhu.disa.mil/apl/>.

Table D-33 lists the functional category requirements, and verifies if those categories were tested and met the criteria identified in accordance with the associated RFC.

Table D-33 Datatek Technologies Test Results for Functional Test Category

Datatek IPv4/IPv6 Transformer		
Functional Category	Critical	Verified
Base IPv6	M	Yes
IPsec	M	Yes
Mobility	CM	No
Bandwidth Limited Networks	O	No
Transition Mechanisms	M	Yes
Quality of Service	O	No
IPv6 Capable Software	M	Yes
Host	M	Yes
LEGEND:		
CM	Conditional Must	IPv6
IPsec	Internet Protocol Security	Internet Protocol Version 6
IPv4	Internet Protocol Version 4	M
		Must
		O
		Optional
NOTE: The terms Must, Conditional Must, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the Department of Defense Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

The Datatek IPv4/IPv6 Transformer was granted a waiver by the DoD IPv6 Standards Working Group for IP Security (IPsec) including:

- Security Architecture for the Internet Protocol (RFC 4301)
- IKEv2 Protocol (RFC 4306).

Therefore, it had to meet the following RFCs:

- IPsec (RFCs 2401, 2402, and 2406)
- IKE Version 1 (RFCs 2407, 2408, 2409, and 4109)

All RFCs are listed in the DoD IPv6 Standard Profiles for IPv6 Capable Products.

Conclusions/Recommendations

The Datatek IPv4/IPv6 is certified for listing as an IPv6 Capable host with waiver.

D.24 Mobile IPv6 Implementation

Testing Organization and Publication Date

Air Force Research Laboratory
April 2008

Summary

The objective of this effort was to provide analysis, design, development, integration, and testing in support of demonstrating the ability of moving network elements to other locations while maintaining connectivity via their original IPv6 addresses using Network Mobility version 6 (NEMOv6) within the Joint Capability for Airborne Networking (JCAN) system. The objective of this report is to capture the differences between JCAN Mobile IP version 4 (MIPv4) with NEMO extensions and NEMOv6, and assess the way forward for integrating NEMOv6 into the JCAN system.

Test and Evaluation Method

Engineer Analysis

Joint Staff Operational Criteria Tested

2 (2.1, 2.2, 2.3)
7 (7.1.1, 7.1.2, 7.1.3)
8 (8.1, 8.1.1, 8.1.2)

Configuration

The JCAN system architecture consists of three major elements: an Airborne Mobile Node (MN), a Ground Node (GN), and one or more Ground Entry Sites (GESs).

An aircraft was modified to integrate the JCAN MN and associated interface hardware. The JCAN MN computers manage data routing, application services, data logging, and the user interface for JCAN system monitoring and control. The MN provides the interface between the LAN and aircraft radios and network. The connection between the MN and the LAN is accomplished via a standard Ethernet connection. The aircraft LAN supports 18 operator workstations and KY-58 crypto units to provide secure data operation.

The JCAN GN is similar to the JCAN MN; it interfaces with the JCAN enabled radios at one or more GESs. The GESs can be collocated with the JCAN GN or geographically separated. Each GES is configured with multiple radios, KY-58 crypto units, antennas and a JCAN Serial Interface to Military Radios (SIMR) shelf. The JCAN SIMR shelf is used to interface to the radio/crypto equipment at the GES. The GESs are connected to the JCAN GN through a satellite interface. The JCAN GN computers manage data routing, application services, data logging, and the user system interface for JCAN system monitoring and control. The JCAN GN can remotely

monitor the status of the radio links and control which radios are available for JCAN use. The data is sent over the air using the legacy radios via an IP tunnel down to the GES where the JCAN Foreign Agent (FA) resides. This same path exists when using IPv6 with NEMO to provide the network mobility to the Joint Surveillance Target Attack Radar Systems (JSTARS) platform.

Results

The additional address space that IPv6 offers will provide more flexibility in defining the airborne network. The capabilities in the JCAN IPv4 system can be carried to the IPv6 solution and alleviate some of the overhead induced by the IPv6 headers. Capabilities such as mobility mode using the JCAN FA can reduce the IP overhead across the wireless links. The performance enhancing proxies also can be incorporated to minimize the data that traverses the wireless links. New capabilities such as robust header compression to further reduce the impact of IPv6 headers in the mobile environment also can be considered. There are capabilities within JCAN, such as concurrent multipath routing, that can be used to further the deployment of NEMOv6. Working groups are investigating the ability to transport IPv4 packets over NEMOv6 to provide a transition mechanism. This makes it more feasible to pursue an IPv6 solution and still support legacy IPv4 over the same infrastructure.

Conclusions/Recommendations

MIPv6 and NEMOv6 provide network mobility similar to the current IPv4 implementation that JCAN uses. There are areas in which MIPv6 and NEMOv6 can still be improved.

D.25 Assessment Report For Evaluating Milestone Objective 2 IPv6 To IPv4 Architecture

Testing Organization and Publication Date

Air Force Communications Agency
April 30, 2008

Summary

Assessment of the IPv6-to-IPv4 tunneling mechanisms for potential use on the Air Force Enterprise Network was performed. The assessment examines the implementation of this technology and the level of assurance the tunneling mechanism provides with respect to the configuration established for each Air Force base. Communications data between router and client were analyzed as well as the routers ability to filter tunneled packets.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.2, 2.3)

8 (8.1.1, 8.1.3)

Configuration

The enclave boundary routers perform the routing of IPv6 prefixes. Static routes were used and routed through the IPv4-only interface connected to the main base infrastructure. A point-to-point tunnel was implemented and a default IPv6 route was used to forward all IPv6 traffic to the tunnel endpoint destination. The firewall rule was modified to allow the IPv6-to-IPv4 relay router prefix. This permitted the IPv6-to-IPv4 prefix, but denied all other IPv6 prefixes. A secondary rule was created to allow the Internet Control Message Protocol version 6 (ICMPv6) messages necessary for stateless auto-configuration. The Teredo and IPv6-to-IPv4 services in windows were enabled.

All tests were performed in an isolated test environment. Besides the tunneling services within the assessment, other applications and services were operational to ensure a simulated Air Force Enterprise Network environment. Application and services within the test environment included those listed in Table D-34.

Table D-34 Assessment Report for Evaluating Milestone Objective 2 IPv6 to IPv4 Architecture Enabled Applications and Services

Enabled Services	
6to4 Client	Microsoft Exchange
6to4 Tunnel Services	Microsoft WINS
IPv6 Helper Services	Monitoring
IPv4/IPv6 Background Traffic	Neighbor Discovery Spoofer
Microsoft Active Directory	Production Gateway
Microsoft DHCP	Statistical Analysis
Microsoft DNS	Traffic Analysis
Legend	
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
DHCP	Dynamic Host Control Protocol
DNS	Domain Name Service Windows Internet Naming Service
WINS	Service

Background traffic consisted of several protocols to include:

- TCP/UDP 53 (DNS)
- TCP/UDP 135 (EPMAP)
- TCP/UDP 137 (NETBIOS-NS)
- TCP/UDP 138 (NETBIOS-DGM)
- TCP/UDP 139 (NETBIOS-SS)
- TCP/UDP 389 (LDAP)
- TCP/UDP 445 (Microsoft-DS)

Results

Client-to-Router Communication

This test presented expected results from the router and clients with no modifications or changes to the network. Traffic was sent in clear text format, but availability of IPsec was only supported in the router and Windows Vista client. Windows XP does not provide IPsec support. A feature to suppress the Router Advertisement (RA) from clients worked well within the router, which then made it necessary for static configuration of IPv6 addresses for each client. When suppressing RA from within the router, it reduced the availability of stateless auto-configuration from within the IPv6-to-IPv4 tunneling mechanism. No additional CPU or memory usage was seen when communicating through the tunnels within the clients.

Bandwidth consumption for Internet Control Message Protocol (ICMP) traffic was around 58.6% of total available bandwidth as opposed to 61.8% for ICMPv6. All rates were based on a FastEthernet (100 Mbps) connection from the client to the router.

Client-to-Client Communication

This test proved the expected performance results with no modifications or changes to the network. Traffic was sent in clear text format because availability of IPsec was only supported in the router and Windows Vista client. Windows XP does not provide IPsec support. No additional CPU or memory usage was seen when communicating through the tunnels within the clients.

Router-to-Router Communication

This test showed responses from each router and their communication transactions. When the tunnel interfaces were first configured, no initial solicitations were made by either relay router. All round trip packets saw each other as residing in the same local link. A traceroute execution verified a single hop designation between the two IPv6-to-IPv4 relay routers. The assessors found the tunnels might be susceptible to a malformed Distribution & Communication Protocol European Telecommunications Standard Institute (DCP-ETSI) packet when using IPv6-to-IPv4 prefixes as the source and destination.

Filtering Protocol 41

The IPv6-to-IPv4 relay router provided the granularity to filter all networks from accessing the tunnel via Protocol 41 and native IPv6 for that matter. Other tunneling protocols are available for filtering. Both egress and ingress filters worked well for all or specific IPv6-to-IPv4 addresses. The Microsoft Windows XP SP2 does not provide the granularity to filter Protocol 41; its only control mechanism is to disable or enable the IPv6-to-IPv4 service via the *Netsh* Command Line Interface (CLI). Its successor, Microsoft Windows Vista, provides the granularity to filter specific hosts or subnets that are servicing the IPv6-to-IPv4 tunnels. By default, Vista blocks all Protocol 41 traffic from leaving the system. An ingress filter is preferred and should only allow the IPv6-to-IPv4 relay router and subsequent enclaves to enter the system. The default egress filter should be disabled.

Conclusions/Recommendations

Client-to-Router Communication

Client-to-router communication using the IPv6-to-IPv4 tunneling mechanism works with some modification to the router. These modifications result in lowered functionality.

Client-to-Client Communication

Client-to-client communication using the IPv6-to-IPv4 tunneling mechanism works with no modifications or changes to the network.

Router-to-Router Communication

Router-to-router communication using the IPv6-to-IPv4 tunneling mechanism works with no modifications or changes to the network.

Filtering Protocol 41

In the tunnel filter test, the router and Vista client provided the necessary mechanisms to filter the IPv6-to-IPv4 traffic.

D.26 2007 Ethernet Switch Comparison Report

Testing Organization and Publication Date

Army
January 2008

Summary

This comparison report presents evaluation results for core building and edge Ethernet switches provided by five vendors (Extreme Network, 3COM, Cisco Systems, Foundry Networks, and Enterasys Networks) for use in the Installation Information Infrastructure Modernization Program (I3MP). The top performing core devices were the Foundry XMR series switches (4000, 8000, and 16000). The 3COM and Cisco core devices that the Technology Integration Center (TIC) evaluated also performed well and were recommended. The only building switch recommended this year was the Cisco 3750-E.

All vendors met throughput and other performance requirements during 2007. However, the devices are lacking when it comes to their IPv6 capabilities and their Multicast performance.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

3 (3.1.1.2, 3.2.1.2, 3.3.1.2)
9 (9.1.1.1, 9.1.2.2)

Configuration

The devices were not tested as part of a network, but rather as stand-alone. The devices were subjected to Layer 3 traffic generated and received from automated test equipment.

Results

Devices are rated on a point system outlined in the test procedures. They were evaluated and awarded points for performance, security, and network management capabilities. Devices with the higher score performed better than those with a lower score. Although no device achieved a perfect score, the highest possible points for each device was 10.

Table D-35 illustrates the point score for the individual devices. Devices are listed by manufacturer and model tested, category each was tested in, and the resulting point given. Devices that did not receive passing scores were not included in Table D-35.

Table D-35 Ethernet Switch Test Results

Core Switches						
Manufacturer	Model	Points		Manufacturer	Model	Points
Foundry Networks	Netiron XMR 8000	8.4		3COM	8814	8.1
Foundry Networks	Netiron XMR 4000	8.4		Cisco Systems	6509E	8.0
Foundry Networks	Netiron XMR 16000	8.4		Cisco Systems	7609S	8.0
3COM	8807	8.3		Cisco Systems	7609S	8.0
3COM	8810	8.2				

Building Switches						
Manufacturer	Model	Points		Manufacturer	Model	Points
Cisco Systems	3750E Stack	8.2				

Edge Switches						
Manufacturer	Model	Points		Manufacturer	Model	Points
Cisco Systems	Catalyst 3750G-48PS	9.0		Enterasys Networks	Matrix N1 -25	8.1
Cisco Systems	Catalyst 3750G-24PS	9.0		Enterasys Networks	Matrix N7	8.0
Cisco Systems	Catalyst 3560G-48PS	9.0		Enterasys Networks	Matrix N3	8.0
Cisco Systems	Catalyst 3560G-24PS	9.0		Enterasys Networks	Matrix N1-49	8.0
Foundry Networks	FGS648P-POE	8.9		Enterasys Networks	Matrix N5	7.8
Foundry Networks	FGS624P-POE	8.9		Enterasys Networks	N-Series Standalone	7.8
Cisco Systems	Catalyst 3750E-48PD	8.8		Enterasys Networks	Matrix N3	7.8
Foundry Networks	Fastiron SX800	8.6		Enterasys Networks	Matrix N1-25	7.8
Foundry Networks	Fastiron SX1600	8.6		Enterasys Networks	Matrix N7	7.7
Cisco Systems	Catalyst 3750E-48PD	8.6		Enterasys Networks	Matrix N1-49	7.7
Cisco Systems	Catalyst 3750E-24PD	8.6		Enterasys Networks	Matrix N5	7.6
3COM	5500G-EI PWR 48	8.4		Enterasys Networks	Summit 450E-24	7.4
3COM	5500G-EI PWR 24	8.4		Enterasys Networks	Summit 450A-48	7.4
3COM	5500G-EI PWR 52	8.2		Enterasys Networks	Summit 450E-48	6.7
3COM	5500G-EI PWR 26	8.2		Cisco Systems	Catalyst 3750E-48	6.4
3COM	5500G-EI 52	8.2		Cisco Systems	Catalyst 3750E-24	6.4
3COM	5500G-EI 28	8.2		Cisco Systems	Catalyst 3750 Stack	6.3
Enterasys Networks	N-Series Standalone	8.1				

Conclusions/Recommendations

Performance and security features were excellent and scored high based on the test rating scheme for all the evaluated switches, however the IPv6 management capabilities were still lacking on all the devices. The 3COM core switches were the only devices this year that supported SNMP over an IPv6 transport, and none of the element managers that were provided by the vendors supported IPv6 management. Multicast improvements also need to be made, especially with virtual routing and forwarding.

D.27 Transition Mechanisms Study AFATDS over IPv6

Testing Organization and Publication Date

Army
January 9, 2008

Summary

This test was conducted in participation with JUICE 2007. An IPv4 only baseline measurement was taken on a simple network that employs Advanced Field Artillery Tactical Data System (AFATDS) and Simulator-Stimulator (SISTIM) systems, and used for comparison against a dual stack configuration and an address translation configuration. Both of the later two tests were run in two modes, one on a network with no background traffic and the next with background traffic. The command line operation of “ping” was run from end to end to provide simple IPv6 background traffic onto the network. The intent was to see if there was any change in the system performance with background traffic, versus the network without this additional traffic.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

8 (8.1.1.1, 8.1.2.1)

Configuration

Traffic was sent across a Cisco Networks 3825 router configured with a mirrored port from a laptop to laptops with the specific program software installed. Wireshark, a network packet sniffing tool, recorded this traffic.

Results

Dual Stack

A review of the data showed that the dual-stack transition mechanism had little effect on the overall transmission of the AFATDS data. For both scenario runs (with and without background traffic), the test traffic followed an almost identical pattern compared with its specific baseline. Since the IPv4 and IPv6 stacks were available, AFATDS (an IPv4 only application using the IPv4 stack) should perform the same as the baselines.

Address Translation

The data shows that using the Datatek Transformer transition mechanism had an effect on the overall transmission of the AFATDS data. For the scenario run without IPv6 background traffic, the traffic followed a similar pattern when compared with its specific baseline. Some minor variance was expected, due to the added steps the messages went through during the translation.

The scenario run with IPv6 background traffic showed a significant variance from the baseline. This variance was generated by operator error and the test data was discarded.

Conclusions/Recommendations

Dual Stack

This transition mechanism allowed the IPv4 only AFATDS system to work seamlessly. Metrics gathered and compared to the baseline determined that dual-stacking the network had minimal effect on the AFATDS system traffic.

Address Translation

Using the Datatek Transformer, the IPv4 only AFATDS software could communicate across the IPv6 only backbone to another IPv4 only AFATDS system. The impact on the critical AFATDS messaging with the additional steps of translation was not visibly perceptible. Operators running the test scenarios noted no noticeable change in system operation or performance. The metrics gathered and compared to the baseline determined that translation had minimal effect on the AFATDS system traffic.

D.28 Network Management IPv6 Initiative (NM12) (Client Analysis)

Testing Organization and Publication Date

Air Force Research Laboratory
April 30, 2008

Summary

NMI2 is in support of accomplishing high priority IPv6 transition planning and coordination for the DoD. The test involved looking at characteristics in the areas of monitoring, configuring, and accounting for IPv6 devices by network management tools.

Threshold, objectives, and goals were identified for each test plan category. If a tool uses only IPv4 to handle IPv6 MIB information, that tool meets the threshold. Since seven types of clients were examined, four out of the seven (4/7) were considered a majority. A final recommendation of a “snapshot” of the state of NM and IPv6 in meeting Joint Staff Operational Criterion 9 will be identified in a follow-on report to include results from all testing and analysis performed to date. Additional results show that network management polling performed over the IPv6 protocol will take less time than equivalent polling over the IPv4 protocol. In addition, “Out-of-the-Box Performance” results for one tool showed network usage requirements for network management is higher for IPv6 vs. IPv4.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

9 (9.1.1, 9.1.2, 9.1.4, 9.1.6, 9.2, 9.3.1, 9.3.2)

Configuration

Each product has its own method of implementing IPv6. The goal was to create a test plan that could review capabilities important for a successful and seamless IPv6 network management transition implementation, and allow the flexibility to take these differences into account. Additionally, a comparison of the network management products was not a focus. This is true for these reasons: different DoD organizations use different network management tools; finding the “best” tool/client would not aid in the seamless IPv6 transition of DoD networks, since the current methods involve the use of various tools/clients; and not all of the organizations using network management tools make use of these tools in the same way. To perform an accurate and fair analysis, it was necessary to implement a test plan that would have generic testing characteristics and would offer the best general conclusion to the readiness and condition of a sampling of network.

Results

The threshold and objective goals were considered met if the majority of the clients met the goal.

- Within the *monitoring* category, objective goals were met for clients that could support SNMPv1, v2, and v3 over either protocol and clients that could generate SNMP. Threshold goals were met for NM tool help support. All sections tested in this category met either the objective or the threshold goals.
- Within the *configuration* category, objective goals were met for clients that could send SetRequests over either IPv4 or the IPv6 protocol and clients that could be configured as dual stack. All sections tested in this category met the objective goals.
- Within the *accounting* category (which, in this document, is a follow up to previous testing work and thus not client-focused), no goal was met for comparing IPv4 vs. IPv6 out-of-the-box MIB queries and comparing IPv4 vs. IPv6 out-of-the-box trap queries.
- Regarding areas in the accounting category, IPv6 requires more network use to send equivalent data (since it has a bigger header). Therefore, it requires more bandwidth/more time on the wire. The performance categories did not meet their goals. Of the categories reviewed for *client* testing, 83% resulted in objective goals being met, while the remaining 17% resulted in only the threshold goals being met.

Conclusions/Recommendations

All of the categories tested produced acceptable results sufficient to execute equivalent network management capabilities during an IPv6 transition as seen in a purely IPv4-only environment.

D.29 Assessment Report for Evaluation Milestone Objective 2 Virtual Local Area Network Architecture

Testing Organization and Publication Date

Air Force Communications Agency
May 16, 2008

Summary

This report provides the product and process assessments for the migration of the Air Force to IPv6. The purpose of the Virtual Local Area Network (VLAN) architecture assessment was to assess the implementation of this technology and IPv6 with regard to the current security posture the Air Force bases provide today. The assessment examined the level of IA the VLANs provides with respect to the configuration established for each Air Force base. An assessment of known vulnerabilities associated with spoofing Neighbor Discovery (ND) and RA was performed. The testing used passive and penetration type methods to include known vulnerability testing.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

- 1** (1.4.1)
- 2** (2.2, 2.3)
- 8** (8.1.1, 8.1.3)

Configuration

The VLAN technology used in this assessment was configured in a dual stacked environment. Tunnels required to exchange traffic with neighboring enclaves were established in accordance with the *DITO IA Guidance for MO2*. The VLAN tag is 16 bits and normally follows the 48-bit site prefix. VLANs are constructed using a 64-bit Extended Unique Identifier (EUI-64) format. Equivalent application of security policy was provided to the IPv6 path, similar to the IPv4 path. To segregate the authorized IPv6 hosts on the VLANs, IPv6 ACL was applied to the interfaces. These filters allow and deny the specific IPv6 subnets or hosts to the designated VLAN configured. Managing IPv6 enclaves is a key component to the transition of the Air Force Enterprise networks to IPv6.

Results

Client-to-Router Communication

The client-to-router test presented expected results with no modifications or changes to the network. A feature to suppress the RA from clients worked well within the VLAN sub-interfaces, which made it necessary for static configuration of IPv6 addresses for each client. When suppressing RA within the router, it reduced the availability of stateless auto-configuration within the IPv6 architecture. No additional CPU or memory usage was seen when communicating to the VLANs from the clients.

Client-to-Client Communication

The client-to-client test provided expected performance results with no modifications or changes to the network. VLANs using IPv6 did not diminish the use of other VLANs in the IPv4 infrastructure. When using the Gigabit (1000 Mbps) throughput and the FastEthernet (100 Mbps) connectivity for clients, there was no potentially degrading performance seen by the routers. A *Low* risk rating was given to CPU utilization of less than 1% for routers. Anything above eight megabytes of throughput was given a *High* risk rating because of the increased CPU utilization experienced by the routers, above 60% when using tunnels.

Filter Egress/Ingress IPv6 Subnets

In the filter egress/ingress IPv6 test, the router and Microsoft Windows Vista client provided the necessary mechanisms to filter the IPv6 traffic. The VLAN sub-interfaces provided the granularity to filter all networks from accessing the particular VLAN Identifier (ID). Both egress and ingress filters worked well for all specific VLAN traffic or IPv6 subnets. The Microsoft Windows XP SP2 SDC v1.3 did not provide the granularity to filter IPv6 traffic; its only control mechanism was to disable or enable the IPv6 service. Its successor, Microsoft Windows Vista SDC v2.0.3, provided the granularity to filter specific hosts or subnets that were using IPv6.

Filter Neighbor Discovery Advertisements from Surrounding VLANs

While testing filter neighbor discovery advertisements from surrounding VLANs, it was found that the VLAN sub-interfaces adequately filtered specific IPv6 address blocks for ND. VLAN sub-interfaces could distinguish between specific IPv6 addresses for ingress and egress filtering. This allowed systems only in a specific VLAN to formulate IPv6 addresses using the router. Other VLANs being routed through the network did not have access to those specific VLAN sub-interfaces.

Mitigate Neighbor Discovery Attacks

To test the preventative protocols that should mitigate neighbor discovery attacks, the assessors used an open source tool to initiate a spoofing attack. The client and router were subjected to the ND attack. The router was susceptible to ND attacks, which prevented clients from accessing the VLAN sub-interface. ND attacks are predominant on VLAN sub-interfaces and could prevent systems from communicating with clients or the router.

Mitigate Router Advertisement Attacks

To test the ability to mitigate RA attacks, the assessors used an open source tool to initiate a spoofing attack. The router was subjected to the RA attack, which was only mitigated once the attack stopped. Only access to the configured VLAN sub-interface was denied when the router was under an RA attack. All other interfaces, including any that may have been dual-stacked, were operational and accessible by the network. VLAN sub-interfaces could mitigate the RA attacks by suppressing the discovery phase with “*ipv6 nd suppress-ra.*” Manual or static configuration of clients’ IPv6 addresses was required when invoking this feature.

Conclusions/Recommendations

The assessors set forth recommendations to ensure only trusted source addresses were used to establish IPv6 connectivity across the core network using VLANs to establish enclaves. In addition, the assessors recommend using ESP with Null encryption for tunnels that would be used if the infrastructure did not allow enclaves to share an end building node. The IPv6 Information Assurance Group (IIAG) assessors do not recommend expanding the use of the VLAN technology for IPv6 outside of an Air Force base or permitting accessibility from all segments within the base. One or more designated VLANs should be utilized to isolate IPv6 clients and thwart a potential attack against other segments of the base. The VLAN technology has adequate security practices and support of access control mechanisms for distribution of enclaves. ACLs could be used to ensure specific IPv4 blocks or hosts of addresses are filtered through each VLAN sub-interface. Only use trusted IPv4 addresses for the stateless auto-configuration of the IPv4 hosts. Vulnerabilities for RA are mitigated when static configuration from a server is used.

D.30 Assessment Report For Evaluating Milestone Objective 2 Intra-Site Automatic Tunnel Addressing Protocol Architecture

Testing Organization and Publication Date

Air Force Information Operations Center/Information Operational Assessment Division
May 16, 2008

Summary

The purpose of the ISATAP Assessment was to evaluate the implementation of this technology with regards to the current security posture that Air Force bases provide today. The assessment examined the level of IA the tunneling mechanism could provide with respect to the configuration established for each Air Force base.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.3)

8 (8.1.1.2, 8.1.3.2)

Configuration

Testing recorded the communication between the ISATAP tunneling mechanism relay and the clients being serviced. The communication was recorded using a network protocol analyzer for analysis of the handshake between the ISATAP relay and the clients using stateless address auto-configuration. Testing also ensured the protocols and ports associated with the ISATAP tunnels were secure. Then the traffic negotiated between enclaves and their clients was assessed to address severe increases in bandwidth utilization. Finally, an assessment of known vulnerabilities associated with spoofing ND and RA was accomplished. The testing used passive and penetration type methods to include known vulnerability testing.

Results

Client-to-Router Communication

Assessment Objective: The primary focus of this assessment was to find any irregularities in the communication handshake between a client and router using the ISATAP services.

Results: Test showed expected results from the router and clients with no modifications or changes to the network. Traffic was sent in clear text format, but availability of IPsec was only supported in the router and Windows Vista client. A feature to suppress the RA from clients worked well within the router, which then made it necessary for static

configuration of IPv6 addresses for each client. No additional CPU or memory usage was seen when communicating through the tunnels within the clients. Bandwidth consumption for ICMP traffic was around 57.4% of total available bandwidth compared to 59.8% for ICMPv6. Tested protocols/services include:

- TCP/UDP 53 (DNS)
- TCP/UDP 135 (EPMAP)
- TCP/UDP 137 (NETBIOS-NS)
- TCP/UDP 138 (NETBIOS-DGM)
- TCP/UDP 139 (NETBIOS-SS)
- TCP/UDP 389 (LDAP)
- TCP/UDP 445 (Microsoft-DS).

The lack of multicast support prevents the use of automatic router discovery. ISATAP hosts must resolve the ISATAP router through DNS to be assigned an address. There is a possible spoofing attack in which spurious IP Protocol 41 packets are injected into an ISATAP link from outside the enclave. Using encryption between the clients and relays provides an adequate solution for most of these spoofing attacks. An IPv4 ingress filter can be used to filter or block all inbound traffic using Protocol 41.

Client-to-Client Communication

Assessment Objective: The primary focus of this analysis was to find irregularities in the client-to-client communication handshake using the ISATAP services in two different enclaves.

Results: Tests expected performance results with no modifications or changes to the network. No additional CPU or memory usage was seen when communicating through the tunnels within the clients. The same security implications from the previous test were applicable to this assessment. When DHCP was utilized, ISATAP clients would continually expire or renew their address based on the expiration policy of the IPv4 addresses.

Filtering Protocol 41

Assessment Objective: The objective of this analysis was to ensure that all filters on the clients and relay routers could provide the level of IA in which no type of tunneled (Protocol 41) traffic traversed the Air Force base without explicitly granting that service.

Results: The tested router and Vista client provided the necessary mechanisms to filter the ISATAP traffic. The ISATAP relay router provided the granularity to filter all networks from accessing the tunnel via Protocol 41 and native IPv6 for that matter. Both egress and ingress filters worked well for all or specific ISATAP addresses. The Microsoft Windows XP SP2 SDC v1.3 did not provide the granularity to filter Protocol 41, its only control mechanism was to disable or enable the ISATAP service via the CLI. Its successor, Microsoft Windows Vista SDC v2.0.3, provided the granularity to filter specific hosts or subnets that were servicing the ISATAP tunnels. By default, the Vista

SDC v2.0.3, blocked all Protocol 41 traffic from leaving the system. The ISATAP relay routers did not have the ability to identify whether other relays were authoritative.

Filter ISATAP Stateless Auto-configuration

Assessment Objective: In the analysis of the communication transactions, a client and relay router exchange was examined.

Results: The client and relay router exchanged information in order for the clients to generate their IPv6 addresses. Routers could distinguish between specific ISATAP addresses for ingress and egress filtering. A client on a separate segment of the network could obtain access to the ISATAP relay router if a DNS query was performed. Outside clients did not have to reside within the same VLAN as the ISATAP clients. The security posture of the network could be protected if the authentication and/or confidentiality of data were invoked.

Mitigate Neighbor Discovery Attacks

Assessment Objective: An analysis of the preventive controls that may mitigate ND attacks was examined.

Results: The client and relay router were subjected to the ND attack. ISATAP did not support multicast and acts as a Non-Broadcast Multi-Access (NBMA) link. NBMA links did not support multicast or broadcast traffic. The ISATAP tunnels were not susceptible to the ND. Existing countermeasures for tunneling mechanisms should be used accordingly.

Mitigate Router Advertisement Attacks

Assessment Objective: An analysis of the preventive controls that may mitigate RA attacks was examined.

Results: The relay router was subjected to the RA attack, which was only mitigated once the attack stopped. When the relay router was subjected to the attack, only access to the configured ISATAP tunnel was denied. All other interfaces, including any that may have been dual-stacked, were operational and accessible by the network. Routers could mitigate the RA attacks by suppressing the discovery phase. If hosts used static configuration, the attacks based on RA were mitigated. The SEND protocol was also applicable for mitigating attacks based on RA vulnerabilities.

Conclusions/Recommendations

IPv6 enclaves can be deployed throughout the base utilizing the ISATAP tunneling mechanism to allow development and testing of applications that require or include IPv6 support.

D.31 Assessment Report for Evaluating MO2 Microsoft Windows IPv6 to IPv4 Architecture

Testing Organization and Publication Date

Air Force Communication Agency
May 22, 2008

Summary

This assessment reports on the evaluation of the IPv6 transition mechanism IPv6-to-IPv4. Tests evaluated communications between hosts utilizing an IPv6-to-IPv4 tunnel. Tests were chosen to verify that a set of common applications (e.g., web browsers, FTP, Telnet) would function properly using IPv6-to-IPv4.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

1 (1.1.1.1, 1.2.1.1)

2 (2.3)

8 (8.1.1, 8.1.3)

Configuration

The first three tests were designed to show communication between client-to-server, server-to-server, and client-to-client. The clients and servers were located on two different subnets. Once the IPv6 protocol was installed, the IPv6-to-IPv4 tunnel interface was created.

Results

IPv6-to-IPv4 support was provided by the Microsoft Windows IP helper service. When a host had an IPv4 address assigned but no IPv6-to-IPv4 RA was received, the IP helper service automatically configured an IPv6-to-IPv4 address to its tunneling pseudo-interface.

One application test case included FTP. Accessing an FTP site involves using an IP address within the address box of an Internet browser. Using this IPv6-to-IPv4 address created no new security concerns.

Another application test involved Internet Explorer. When using an IPv6 literal address to browse a website, brackets were needed to enclose the address. Accessing web pages using an IPv6 address, this did not open new security holes.

A third application test involved Telnet. Telnet services can communicate whether an individual uses an IPv4 address or its IPv6-to-IPv4 address to connect to another client.

A security test verified that a properly configured Windows Vista firewall denied transit of IPv6-to-IPv4 packets from unauthorized computers. Computers were explicitly authorized to communicate with Windows Vista Standard Desktop Configuration (SDC) clients using Protocol 41, while unauthorized computers were blocked by the firewall.

Another security test involved DNS and undesired AAAA record propagation beyond the enclave. The configuration under test was unable to prevent AAAA records created by the IPv6-to-IPv4 tunnels from propagating beyond the enclave's end point. This creates a security issue.

Conclusions/Recommendations

Implementation of preventive measures to block the FTP server's DNS records from getting past the routers at the service delivery point is important. In addition, administrators should block web addresses beginning with the 2002::/ prefix. This was the default prefix for IPv6-to-IPv4 interface addresses. IPv6-to-IPv4 tunnels need AAAA records to communicate within an enclave. Disabling the IP helper service to stop the AAAA records was not feasible, as it terminated any communication using IPv6-to-IPv4 tunnels. Solutions for DNS records propagating past the service delivery point are:

- Creating two split DNS servers, one internal to the network and one external
- Not allowing zone transfers between those servers.

IPv6-to-IPv4 tunneling is a viable option on the Air Force Enterprise Network; however, network administrators must address security concerns. Solutions for IPv6-to-IPv4 security concerns are to allow only authorized endpoints to establish tunnels. This could be accomplished using static routes or virtual private network connections.

D.32 Special Interoperability Test Certification of the Novell SuSE Linux Enterprise Server 10, Service Pack 2 Running on an IBM P-Series High Volume Open Power Personal Computer Server, IBM X-Series 226 x86 Server, Dell Precision M6300 32 and 64-bit x86 Laptop, and Dell Precision T5400 32 and 64-bit x86 Desktop for IPv6 Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
July 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Novell SuSE Linux Enterprise Server 10, Service Pack 2. This device meets the IPv6 Capable interoperability requirements of an Advanced Server and Host. Interoperability testing was conducted from May 15-20, 2008 at JITC's Advanced IP Technology Capability.

This test was conducted by installing laptop and servers loaded with Novell SuSE Linux Enterprise Server 10, Service Pack 2 on a dual stack IP network. A network sniffing device verified that the proper sequence of packets was passed back and forth across the network during communications required by the chosen RFCs. When the proper sequences of packets were recorded, the tested RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

- 1** (1.1.1.1, 1.2.1.1)
- 2** (2.3)
- 8** (8.1.1, 8.1.3)

Configuration

An IBM P-Series High Volume Open Power Personal Computer Server, IBM X-Series 226 x86 Server, Dell Precision M6300 32-bit and 64-bit x86 Laptop, and Dell Precision T5400 32-bit and 64-bit x86 Desktop were used to test the Novell SuSE Linux Enterprise Server 10, Service Pack 2. A test network was constructed to send and receive test packets across the network. A router in the network was configured with a mirrored port to allow a packet sniffing device to record the packets as the traversed the network from server to client. Proper packet conversation between server and client was recorded in accordance with the applicable.

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-36 lists the category of testing outlined for host and advanced servers in the DISR, and whether the devices met the requirements.

Table D-36 SuSE Test Results for Functional Test Category

Novell SuSE Linux Enterprise Server10, Service Pack 2		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
IPsec	M	Yes
Transition Mechanisms	M	Yes
Quality of Service	O	No
Mobility	CM	No
Bandwidth Limited Networks	O	No
Server	M	Yes
Host	M	Yes
LEGEND:		
CM	Conditional Must	M Must
IPsec	Internet Protocol Security	N/A Not Applicable
IPv6	Internet Protocol Version 6	O Optional
NOTE: The terms Must, Conditional Must, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the Department of Defense Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The Novell SuSE Linux Enterprise Server 10, Service Pack 2 is certified for listing as an IPv6 Capable host and advanced server.

D.33 Special Interoperability Test Certification of the Red Hat Enterprise Linux 5.2 Server and Client running on an IBM P-Series High Volume Open Power Personal Computer Server, IBM X-Series 226 x86 Server, Dell Precision M6300 32 and 64-bit x86 Laptop, and Dell Precision T5400 32 and 64-bit x86 Desktop for IPv6 Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
June 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Red Hat Enterprise Linux (RHEL) 5.2 Server and Client running on an IBM P-Series High Volume Open Power Personal Computer Server, IBM X-Series 226 x86 Server, Dell Precision M6300 32-bit and 64-bit x86 Laptop, and Dell Precision T5400 32-bit and 64-bit x86 Desktop. This device meets the IPv6 Capable interoperability requirements of a host and advanced server. This special certification is based on IPv6 Capable Interoperability testing conducted by JITC at Fort Huachuca, Arizona. Interoperability testing was conducted from June 9-18, 2008 at JITC's Advanced IP Technology Capability.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

- 1** (1.1.1.1, 1.2.1.1)
- 2** (2.3)
- 8** (8.1.1, 8.1.3)

Configuration

The DUTs were divided into two categories all ran RHEL 5.2. The host category DUTs were the Dell Precision M6300 32-bit and 64-bit x86 laptops, and the Dell T5400 32-bit and 64-bit x86 desktops. The advanced server DUTs were the IBM P-Series HVO Power PC Server and IBM X-Series 226 Server. Each device can act as a host (workstation running client-side applications) and advanced server (server running server-side applications).

A test network was constructed to send and receive test packets across the network. A router in the network was configured with a mirrored port to allow a packet sniffing device to record the packets as they traversed the network from server to client. Proper packet conversation between server and client was recorded in accordance with the applicable RFC.

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jitic.fhu.disa.mil/apl/>.

Table D-37 lists the category of testing outlined for host and advanced servers in the DISR and whether the devices met the requirements.

Table D-37 Red Hat Test Results for Functional Test Category

Red Hat Enterprise Linux 5.2 Server and Client		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
IPsec	M	Yes
Transition Mechanisms	M	Yes
Quality of Service	O	No
Mobility	CM	No
Bandwidth Limited Networks	O	No
Server	M	Yes
Host	M	Yes
LEGEND:		
CM	Conditional Must	M Must
IPsec	Internet Protocol Security	N/A Not Applicable
IPv6	Internet Protocol Version 6	O Optional
NOTE: The terms Must, Conditional Must, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the Department of Defense Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

Conclusions/Recommendations

The RHEL 5.2 running on the Dell workstations and IBM servers is certified as an IPv6 Capable host and advanced server.

D.34 LOSSKNOT Section IV, Test Plan and Results

Testing Organization and Publication Date

National Security Agency
August 2007

Summary

The testing focused on two key functional areas, components and system testing of unique LOSSKNOT solutions. Component testing demonstrated how the individual components support the technical requirements as described in the Minimum Requirements Document (MRD). The system tests demonstrated how the system supports the overall objectives as described in the MRD.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1.1.1, 2.2.1.1, 2.3.1.1)

3 (3.1.1.2, 3.2.1.2, 3.3.1.2)

8 (8.1.1, 8.1.2)

9 (9.1.1.2, 9.1.2.2)

Configuration

Hardware tests focused on Foundry network products that had interfaces with capacities as high as 10 gigabytes. These routers consist of the NetIron MLX core router and the FastIron SX and FES X424-POE-PREM edge routers. Testing consisted of unit, integration, and system testing. The majority of the testing was functional in nature; however, there were some performance tests. Table D-38 lists and describes the tests that were performed and are required by the MRD.

Table D-38 Microprocessor Library Definition Required Tests

Test	Description
Hardware/Platform Testing	
Cold and Warm Start Test	This test will verify that the DUT performs a proper boot from cold start and warm start and record times.
Hotswap test	This test will demonstrate the DUT's ability to replace components under power.
Hitless L2/L3 Failover with Graceful OSPF/BGP Restart	This test will demonstrate the DUT's High-Availability management features for stateful failover of the Management cards
Hitless Software Upgrades	This test will demonstrate the DUT's operational impacts while undergoing software upgrade capabilities in a real-time operational environment
Software Upload/Upgrade Test	This test will demonstrate the DUT's ability to upload system software and upgrade. TFTP file upload as well as IronView will be utilized.
POE Conformity	This test will verify that the PSE device under test classifies a powered device correctly
Port Aggregation (Trunking) Testing	This test will verify that the DUT/SUT can aggregate multiple 100Mbps, 1Gbps, and 10Gbps Ports using the standard LACP (802.3ad).
Security Testing	
User Accounts Testing – Internal Database	This test will verify the DUT's ability to allow user accounts to be created and different privilege levels assigned.
User Accounts Testing – RADIUS	This test will verify the DUT's ability to accept user accounts and different privilege levels assigned from an external RADIUS server.
Authentication, Authorization, Access (AAA)	This test will verify the DUT's ability to allow for user authentication, authorization, and access levels to be defined.
MAC port Security Testing	This test will demonstrate the DUT's ability to provide network access security via MAC address.
802.1x port Security Testing	This test will demonstrate the DUT's ability to provide network access security via the 802.1x protocol.
Logging Conformity	This test will verify that the DUT/SUT is capable of logging various levels of events to both an internal database and external SYSLOG server
Layer 2 Protocol Testing	
Virtual LAN (VLAN)/802.1q Tagging Conformance Testing	This test will demonstrate the DUT's ability to support in excess of 500 VLANs per Switch and tag VLANs frames between using the 802.1q protocol
Spanning Tree Protocol (STP) Conformance Testing	This test will demonstrate the DUT's ability to support Spanning Tree Protocol (STP - 802.1d), Rapid Spanning Tree Protocol (RSTP - 801.1w), and Multiple Spanning Tree Protocol (MSTP – 802.1s).
Internet Protocol Testing	
RIPv2 Conformity	This test will demonstrate the enabling of RIPv2 on all L3 devices and test to ensure network convergence
OSPF Conformity	This test verify the Device Under Test's (DUT's) compliance with the following capabilities defined in various OSPF RFCs:
	OSPFv2 . RFC 1583, RFC 2328
	OSPF Opaque LSA . RFC 2370
	OSPF NSSA . RFC 1587
	OSPF Database Overflow . RFC 1765
OSPFv3 (OSPF for IPv6) . RFC 2740	
BGPv4(+) Conformity	This test will verify the DUT's compliance with capabilities defined in various BGP specifications: RFC 1771, RFC 1772, draftietf-idr-bgp4-12, draft-ietf-idr-gp4-17.
IPv6 Conformity	This test will verify the DUT's compliance with the following features defined in various RFCs:
	IPv6 (RFC 2460).

	Transmission of IPv6 Packets over Ethernet Networks (RFC 2464).
	IPv6 over PPP (RFC 2474).
	ICMPv6 (RFC 2463).
	Stateless Address Autoconfiguration (RFC 2462).
	Path MTU Discovery (RFC 1981).
	Neighbor Discovery Protocol (RFC 2461).
	Tunneling (RFC 2529, RFC 2893, and RFC 3056).
Multicast Functionality	
Multicast Functionality	Suggested tests are as follows
	IGMP Join and Leave Latency
	MLD Join and Leave Latency
	IGMP Scalability
	MLD Scalability
	Mixed Class Throughput
	Reverse Path Forwarding
	First Hop Router Latency
	Last Hop Router Throughput
	Last Hop Router Latency
	Rendezvous Point Scalability
	Rendezvous Point Throughput
	PIM Join Latency
PIM Prune Latency	
Virtual Router Redundancy Protocol (VRRP) Testing	This test verifies that the DUT/SUT was capable of running the VRRP for gateway High-availability
System Performance Testing	
Layer 2 System Performance Testing (RFC2889)	This test will demonstrate the DUT's ability to perform line rate forwarding of Ethernet frames at all interface speeds and duplex.
Layer 3 System Performance Testing (RFC2544)	This test will demonstrate the DUT's ability to perform line rate forwarding of Layer 3 packets at all interface speeds.
IPv4/IPv6 Dual-Stack Performance Testing	Verify Line Rates at IPv6 using a dual-stack model
Access Control Lists (ACLs) Performance	Verifies that the DUT/SUT can handle a maximum of 500 ACLs per interface at L2/L3 without significant performance degradation
QoS	Measure the baseline performance of the DUT with and without QoS when stateless traffic is injected into the network. The 1st step is to take measurements and collect statistics when QoS is disabled on the DUT. The 2nd step is to take measurement and collect statistics when QoS with Diffserv classifying and DSCP marking is enabled on the DUT.
System Management	
INM Functionality	Configure INM server to auto-discover the network. Verify management capabilities of the INM server on the CAN network devices
SSH/Telnet/HTTP Functionality	
SNMP/MIB/RMON Management Functionality	
	Test connectivity to the Switch CLI using the SSH, HTTP and Telnet applications.
	To enable SNMPv1, v2, and v3 on all network devices with both read and write community strings. Utilize both standard and vendor MIBs in order to gather information and send configuration information to the switch. Capture and analyze RMON information

Results

All devices passed all required tests (Hitless software upgrades were required for the MLX platform, therefore Hitless testing applied only to the NetIron MLX).

Conclusions/Recommendations

This test demonstrated that the LOSSKNOT system complies with the MLD. The test also illustrated that IPv6 performance, interoperability and security met MLD requirements.

D.35 Special Interoperability Test Certification of the Sun Microsystems SPARC T2000 and X86 V40z 32-bit and 64-bit Platforms Running Solaris 10 for IPv6 Capability

Testing Organization and Publication Date

Joint Interoperability Test Command
July 2008

Summary

This report presents the results of the Special Interoperability Test Certification of the Sun Microsystems SPARC T2000 and X86 V40z 32-bit and 64-bit platforms running Solaris 10. This device meets the IPv6 Capable interoperability requirements of a host and advanced server. The Sun Microsystems SPARC T2000 and X86 V40z were granted a waiver by the DoD IPv6 Standards Working Group for IPsec RFC 4301 and IKEv2 RFC 4306. Therefore, the devices had to meet the IPsec RFCs 2401, 2402, and 2406, and the IKEv1 RFCs 2407, 2408, 2409, and 4109.

This test was conducted by installing the Sun Microsystems SPARC T2000 and X86 V40z 32-bit and 64-bit platforms running Solaris 10 loaded on laptops and servers on a dual stack IP network. The network sniffing device was used to verify that the proper sequence of packets was passed back and forth across the network during communications required by the DISR chosen RFCs. When the proper sequences of packets were recorded, the tested DISR RFC requirement was marked as met.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

- 1** (1.4.1)
- 2** (2.2, 2.3)
- 8** (8.1.1, 8.1.3)

Configuration

Table D-39 list the hardware and software configuration of the devices used in the certification testing.

Table D-39 Test Configuration Hardware and Software

Sun Microsystems SPARC T2000 and X86 V40z		
Component	Firmware/Software	Interface
SPARC 64-bit T2000 Server	SunOS 5.10 Generic_120011-14 Solaris 10	Ethernet 10/100Mbps
X86 64-bit V40z Server	SunOS 5.10 Generic_120012-14 Solaris 10	Ethernet 10/100Mbps
LEGEND:		
Mbps	Megabits per second	OS Operating System

Results

JITC distributes interoperability information via the JITC ERD system, which uses NIPRNet email. More comprehensive interoperability status information is available via the JITC STP. The STP is accessible by .mil/gov users on the NIPRNet at <https://stp.fhu.disa.mil>. Test reports, lessons learned, and related testing documents and references are on the JITC JIT at <http://jit.fhu.disa.mil> (NIPRNet), or <http://199.208.204.125> (SIPRNet). Information related to IPv6 Capable testing is at <http://jtc.fhu.disa.mil/apl/>.

Table D-40 lists the category of testing outlined for host and advanced servers in the DISR and whether the devices met the requirements.

Table D-40 Sun Microsystems Test Results for Functional Test Category

Sun Microsystems SPARC T2000 and X86 V40z		
Functional Category	Requirement	Verified
Base IPv6	M	Yes
IPsec	M	Yes
Transition Mechanisms	M	Yes
Quality of Service	O	No
Mobility	CM	No
Bandwidth Limited Networks	O	No
Server	M	Yes
Host	M	Yes
LEGEND:		
CM	Conditional Must	M Must
IPsec	Internet Protocol Security	O Optional
IPv6	Internet Protocol Version 6	
NOTE: The terms Must, Conditional Must, and Optional are used to reference specific required Request for Comments from the Internet Engineering Task Force, the Department of Defense Information Technology Standards Registry, and the Department of Defense Internet Protocol Version 6 Generic Test Plan.		

The Sun Microsystems SPARC T2000 and X86 V40z were granted a waiver by the DoD IPv6 Standards Working Group for the following RFCs:

- Security Architecture for the Internet Protocol (RFC 4301)
- IKEv2 Protocol (RFC 4306).

Therefore, it had to meet the following RFCs:

- IPsec (RFCs 2401, 2402, and 2406)
- IKE Version 1 (RFCs 2407, 2408, 2409, and 4109).

All are listed in the DoD IPv6 Standard Profiles for IPv6 Capable Products.

Conclusions/Recommendations

The Sun Microsystems SPARC T2000 and X86 V40z 32-bit and 64-bit platforms running Solaris 10 are certified for listing as IPv6 Capable host and advanced server.

D.36 IPv6 Transition Mechanism Test Report

Testing Organization and Publication Date

Army, Information Technology Agency (ITA)
September 28, 2007

Summary

The goal of this effort was to test different tunneling and translation mechanisms, ensuring that the use of these tools will be effective solutions in cases where dual-stack is not an option. Initially, these solutions can be used to provide IPv6 access to IPv4-only hosts (translation) or dual-stack hosts on IPv4-only networks (tunneling). Toward the end of the transition phase, when ITA decides to no longer provide an IPv4 service, these transition techniques can be used to allow customers to maintain IPv4 connectivity.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)

3 (3.1)

8 (8.1.1, 8.1.2)

Configuration

Functional and performance testing were conducted on the various transition tools using simple pings, trace routes, HTTP, and SSH sessions, as well as the Spirent Smartbits performance tester. Functional testing involved successful IPv6 communication across an IPv4 network, and/or successful translation between IPv4 and IPv6 networks. Performance testing focused on determining the overhead of an additional header when tunneling. In the case of translation using the Netscreen Firewall, tests focused on the effect of table lookups and header processing. Tests utilized the network devices in the ITA lab, which were configured to mimic an operational environment. Table D-41 lists the devices that were tested and their software version number.

Table D-41 IPv6 Transition Mechanism Test Report Equipment List

Device	Software Version
Extreme 6804	Extremeware 7.6.3.3
Extreme 5i	Extremeware 7.6.3.3
Cisco 2691	IOS 12.3(21)
Cisco 3550	IOS 12.1(22) EA8a
Cisco 6503	IOS 12.2(18)SXF7
Cisco 6506	IOS 12.2(18)SXD7b
Juniper M10i	JunOS 8.2R2.4
Juniper M20 (W1)	JunOS 7.5-20060511.0
Juniper M20 (W4)	JunOS 8.0 R2.8
Netscreen 5200	ScreenOS 6.0.0b3.0
HP Laptop	Windows XP Service Pack 2
Dell Optiplex GX270	Ubuntu Linux 6.10
Hexago Gateway6	HexagoOS 5.0
Hexago Dongle6	Linux Kernel 2.6.18
Datatek Transformer	OS version 2.0.4
Spirent Smartbits	SmartFlow 4.70.022.1
Spirent Smartbits	TeraRouter Tester 5.00.150
Legend	
IOS	Internetwork Operating System
OS	Operating System
Jun	JUNOS

Results

Hexago Gateway6 Functional Testing

The test results verified that an IPv6-in-IPv4 tunnel could be successfully created and connectivity between the Host PC and the IPv6 server could be established. Successful use of an encryption-enabled application SSH through the IPv6-in-IPv4 tunnel was shown.

The results of the performance testing showed that the average response times to retrieve a URL from the IPv6 server began at 156 ms for 200 users and ended at 353 ms for 1000 users when using native IPv6. When accessing the web server and using the tunnel broker in combination with the Dongle6 device, the response times reached approximately 1.2 seconds for 1000 users. When using native IPv6, the client averaged about 930 successful transactions per second. When using the Gateway6 Tunnel broker in combination with the Dongle6 device, the number of successful transactions per second dropped to between 250 and 260.

The performance results also illustrated that when accessing the web server using translation, the number of successful transactions per second was reduced. Using native IPv6, the number of transactions per second averaged approximately 930 transactions per second for user loads up to 1000. When translation was used, the number of successful transactions per second dropped to between 685 and 701 transactions per second.

While performing the baseline tests, the CPU utilization remained near its idle utilization levels of approximately 3% to 7% for all user loads. However, when performing translation, the CPU load spiked to 60% for 200 users and to 89% for 400 users and above.

Netscreen Translation Functional Testing

The performance results showed that users accessing the web server using translation had similar response times per URL as compared to the baseline with 200 simultaneous users. As user loads increased, the response times between the tests began to deviate. At the maximum load of 1000 users, the response times when using native IPv6 averaged 353 ms. However, when using translation, the response time at 1000 users averaged 761 ms. The performance results showed that when using native IPv6, the number of transactions per second averaged approximately 930. When translation was used, the number of successful transactions per second dropped to between 685 and 701 transactions per second for the same set of user loads.

While performing the baseline tests, the CPU utilization remained near its idle utilization levels of approximately 3% to 7% for all user loads. However, when performing translation, the CPU load spiked to 60% for 200 users and to 89% for 400 users and above.

Datatek IPv4-IPv6 Transformer

The performance results showed that users accessing the web server with the Datatek Transformer performing translations experienced higher response times relative to the native IPv6 baseline. The response times for 200 users using native IPv6 averaged 152 ms; when using the Datatek Transformer, the response times increased to an average of 282 ms. At 1000 users, the highest number of users tested, the response times averaged 315 ms using native IPv6 and 1.3 seconds using the Datatek Transformer.

Test results indicated that when the Datatek Transformer was used, the average number of transactions per second was reduced in comparison to the native IPv6 baseline. When using native IPv6 to access the web server, the number of transactions per second averaged 930 for user loads up to 1000. When accessing the web server using the Datatek Transformer, the number of transactions per second dropped to an average of 378.

Conclusions/Recommendations

Hexago Gateway6 Functional Testing

Hosts using the Hexago Gateway6 can successfully tunnel across the infrastructure, but also illustrate the negative effects of encapsulation on latency and transaction throughput. It should be noted that the Dongle6 device used for testing was a prototype still under development by Hexago. As stated in the test setup, performance-testing using the client software was not possible, since it could not be installed on the Smartbits testing device. Therefore, the performance of a host personal computer using the client software may differ from the Smartbits results collected during this test.

Netscreen Translation Functional Testing

Overall, the test results showed that the Netscreen firewall could successfully translate IPv4 packets to IPv6 and vice versa. However, latency and transaction throughput degradation caused by protocol translation was noted. Additionally, the spike in CPU processing of the Netscreen 5200 as the number of users increased was of concern, especially since the firewall was configured with a minimal number of rule sets. If this solution is required, a separate Netscreen for translation may be necessary.

Datatek IPv4-IPv6 Transformer

The Datatek Transformer could successfully translate between the IPv4 and IPv6 protocols. As seen with the other translation devices, there was a performance impact associated with using the Transformer. Testing also showed that the device could only support the use of standard FTP in passive mode since it does not currently support application layer translation.

D.37 NIPRNet IPv6 Compliance Demonstration

Testing Organization and Publication Date

Defense Information Systems Agency
June 18, 2008

Summary

DISA, who manages the NIPRNet, determined and assessed the backbone configuration changes required to make the infrastructure IPv6 Capable. IP devices that make up the operational NIPRNet backbone core were configured dual-stack, enabling them to route both IPv4 and IPv6 packets through the network. Tests demonstrated the ability to route IPv6 packets through the NIPRNet core backbone infrastructure and to/from an external network. The results showed that IPv6 connectivity and transport through the NIPRNet was consistent with that of IPv4. The demonstration successfully met the conditions outlined in the Office of Management and Budget (OMB) memorandum M-05-22.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

3 (3.2.1.2)

8 (8.1.1, 8.1.2)

Configuration

This demonstration was designed in accordance with the Federal Chief Information Officer (CIO) Council IPv6 demonstration plan. IPv6 configured laptops were set up at various node locations and were used for initiating and receiving transmitted IPv6 traffic. Utilizing the dual-stacked laptops, a series of *ping* and *traceroute* commands were performed, initiating IPv6 packets within the demonstration. The *ping* test was devised to assess the backbone core's IPv6 connectivity, while the *traceroute* test was used to assess the backbone core's ability to transport IPv6 traffic. One test scenario attempted to transmit IPv6 traffic from a NIPRNet node to an external node, residing on an external network. The scenario was then reversed, passing traffic from the external node back to the NIPRNet node; a second scenario tested IPv6 routability within the NIPRNet, passing traffic from a NIPRNet node to another; and a third scenario tested for completeness, attempting to transmit IPv6 traffic from a backbone core router to its neighboring core routers (core to core). In each of the scenarios, the Continental United States (CONUS) and/or Pacific (PAC) backbone core routers were analyzed for their ability to route IPv6 traffic.

Results

Core Network IPv6 connectivity demonstration

To demonstrate the core backbone's ability to route IPv6 traffic between two laptops, a series of tests were administered. A set of 10 *ping* tests were executed via the *ping* script between hosts. The executed *ping* commands produced a number of responses that were displayed and their statistics recorded. The results indicated that for each initiated IPv6 *ping* command, a 32 byte data packet was sent within one millisecond (ms). Following 10 *ping* attempts, the generated packets were analyzed for packet loss, which was zero. The test was performed several times with no errors. A continuous *ping* test was run for about one hour between the two CONUS hosts calibrating the responses. The results from the tests indicated that IPv6 connectivity was successful and that the operation of the base TCP/IPv6 stack was working correctly across the CONUS backbone core.

Core Network IPv6 Connectivity with External Network Demonstration

A set of 10 *ping* tests were executed via the *ping* script. As opposed to the earlier test, which was confined to one network theatre, this test involved hosts connecting to one another across separate networks. The executed *ping* commands that were initiated between hosts produced 10 responses. The data that was shown indicated that each *ping* initiated a data packet of 32 bytes and had a connectivity time of 123ms with zero packet loss. The test was performed several times with no errors. A continuous *ping* test was set up for about an hour between hosts with responses showing no errors. The reported response time of 123ms was expected and attributed to the distance between CONUS and PAC. The results from tests indicated that the operation of the base TCP/IPv6 stack was working correctly between the CONUS backbone core and external network.

Core Network IPv6 transport demonstration

To demonstrate the core backbone's ability to transport IPv6 traffic between two laptops, a series of tests were administered. A *traceroute* test with a maximum limit of 10 hops was executed. The results displayed the routes the packets took from the source host through to the CONUS backbone core, ending at the chosen destination host. The results also showed that the response times for each hop within CONUS were 1ms. It should be noted that there were routing issues in the backbone core routers, due to security configurations that did not permit the CONUS backbone core routers from responding to *traceroute* requests. Although the routers were fully functional, they purposely did not respond to the requests. The results indicate that IPv6 was successfully transported through the core network.

Core Network IPv6 transport with external network demonstration

Multiple *traceroute* tests over a maximum of 10 hops were executed via the *traceroute* script between hosts. The procedures involved hosts attempting to *traceroute* IPv6 packets to hosts on a different network. The executed *traceroute* command initiated an IPv6 packet that was sent from the source host in search of a destination host, while limiting the number of hops to 10. The results showed that the route taken from the source through the CONUS backbone core to the chosen external destination host was successful. The results indicated that for each hop within CONUS, *traceroute* response times were 1ms while hop times recorded within the external network were 123ms. The observed latency was expected since packets would have to travel across two separate backbone cores before their times were calibrated. As with the CONUS transport test, there were test pings that did not return due to security settings on distant routers. Results from this test appeared normal, indicating that IPv6 traffic was transported across separate networks.

Conclusions/Recommendations

The tested networks were fully capable of routing IPv6 traffic.

D.38 IPv6 Tunnel Broker Transition Test Report

Testing Organization and Publication Date

Air Force Systems Networking
April 11, 2008

Summary

One transition technique required to be demonstrated in support of the MTP v2.0 was the Hexago Tunnel Broker device. Tunnel broker devices are one of the possible solutions for Air Force transition mechanisms that may be used in the future. This evaluation applied varying loads of IPv4 and IPv6 traffic to a simulated Air Force and DISA network in the testing facility.

As many DoD applications and legacy systems continue to use IPv4 (as well as IPv6) well into the future for various logistical and technological reasons, running tunnels over network backbones may provide a solution for legacy network requirements. Network equipment is already heavily taxed providing security with ACLs and Air Force VPN tunnels, as well as other filtering and processor intensive functions that are in use. The addition of running tunnels of IPv4 and IPv6 traffic, as well as the increase in networks and addresses provided by the capability of IPv6, raises some performance concerns for existing and future network equipment. The use of a separate piece of equipment to provide tunneling mechanisms might be advantageous to DoD networks.

Test and Evaluation Method

Experiment

Joint Staff Operational Criteria Tested

2 (2.1, 2.3)

3 (3.1)

8 (8.1.1, 8.1.2)

Configuration

Testing included configurations for IPv4 traffic over IPv6 networks and IPv6 traffic over IPv4 networks. Functionality and performance was evaluated by attempting to pass traffic over the test network. Traffic was generated using IPv6 and IPv4 addressing with the Spirent test device. Network equipment was evaluated for processor utilization, throughput, frame loss, and latency, as well as functionality and other performance issues as pertinent to each respective type of equipment. Testing of TCP sessions was completed using Avalanche software (on the Spirent test device). For client-to-client testing, Chariot test software was used to generate 10 traffic pairs; the results were documented for throughput, transaction rate, and response time. Figure D-1 shows the test network configuration.

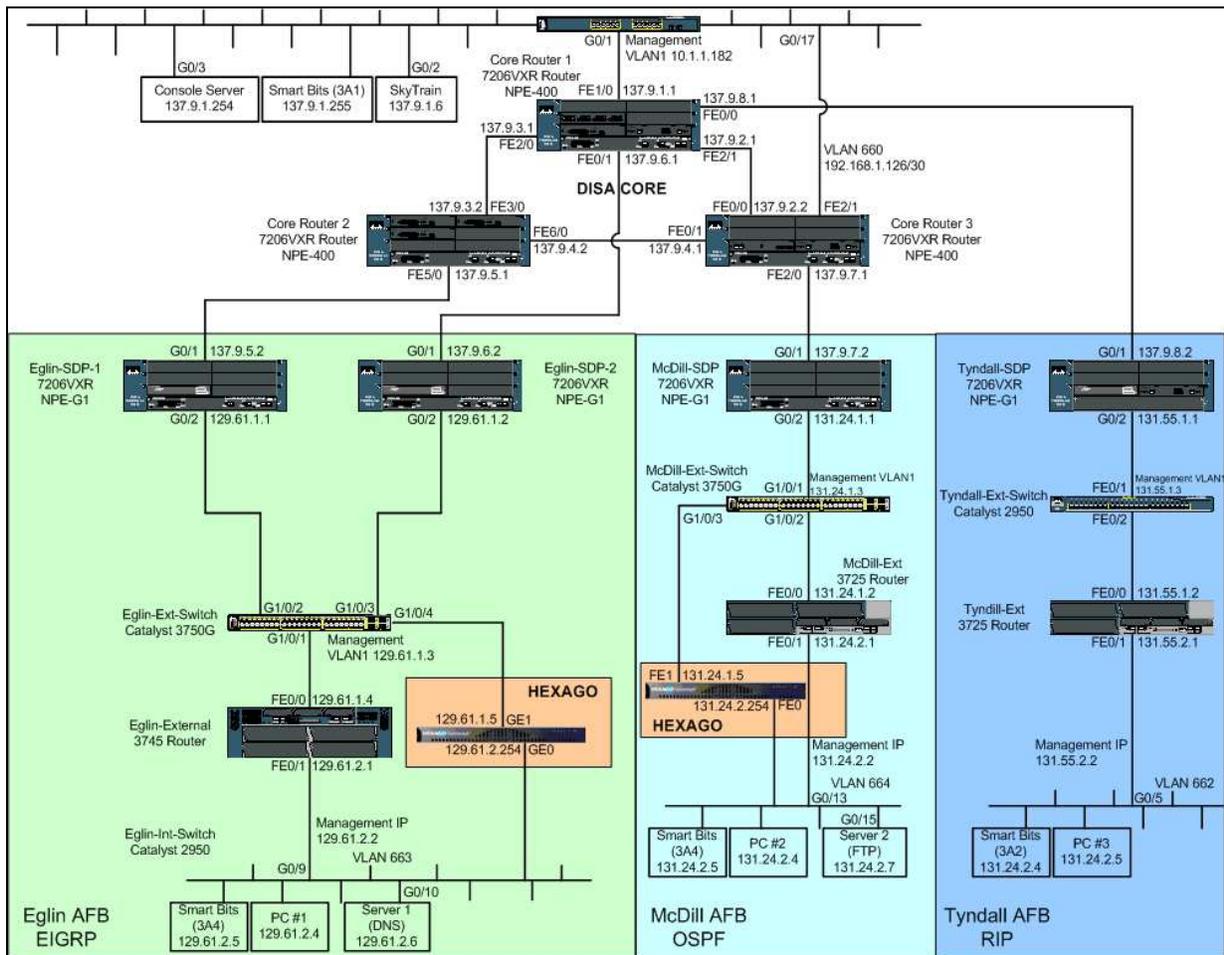


Figure D-1 IPv6 Tunnel Broker Transition Test Diagram

The test network or SUT represents a simulated DISA WAN and three base networks (Eglin, Tyndall and MacDill Air Force Base). The Eglin network setup resembled a future Block 30 or dual diversity/path configuration. The other two bases had architectures that closely resemble the current NIPRNet Air Force architecture.

Frame sizes varied from a minimum of 128 bytes to a maximum of 1408 bytes. SmartBits ports used for connectivity during the test were 100 Mbps ports set. For the Avalanche or layer 7 testing, the Smartbits was set up to test five protocols: HTTP, HTTPS, FTP, SMTP, and DNS. Specifically, the ratios were: HTTP 50%, HTTPS 20%, FTP 15%, SMTP 10%, and DNS 5%. Each test was conducted with 1000 simultaneous users, and then repeated with 2000 simultaneous users.

Testing of the client-to-client (or host-to-host) feature of the tunnel broker was accomplished using Chariot test software. Traffic was generated to simulate 10 traffic pairs. Furthermore, to attempt to simulate as the testing environment provided by the Spirent Avalanche tests noted above, the 10 pairs were given the same protocols and similar ratios.

Results

Manual (protocol 41) IPv6 over IPv4 Tunnel Test

Test results showed decreasing losses as frame sizes increased. There was minimal loss during the throughput test for 128 byte frame sizes up to about 16% load. For loads above 16%, there was increasing loss. For loads above 51% (still, for the 128 byte frame size), losses were near 100%. When the frame size was increased to 256 bytes, losses decreased, with minimal loss through 21% load. Then losses increased as the load was increased, reaching near 100% loss at 81% load. With increasing frame sizes, losses continued to decrease, never reaching 100% loss anymore. The largest frame sizes showed minimal loss, with no notable loss until loads exceeded 71%. Note that this pattern is typical, as larger and larger frame sizes are more efficient (with less overhead) and typically produce decreasing losses with increased frame size. This was true in almost all test results observed for any environment. The Avalanche results for the baseline test showed similar results, with increased losses when the simultaneous users were increased from 1000 to 2000.

Generic IPv4 over IPv6 Tunnel Test

Tests were conducted using the SmartBits Smartflow and Avalanche software. Test results showed more loss and performance issues than the Manual or protocol 41 testing. As expected, using IPv4 traffic over an IPv6 network resulted in decreased throughput, and frame loss numbers increased significantly over the results that were seen during other testing. At 128 byte frames with 1% loading, some loss was already noted with nearly 6% loss in one traffic direction and approximately 42% loss in the other direction of traffic flow. This phenomenon (with different path losses in different directions) was observed on almost every result during this test. In addition, as loads increased, losses increased to the point that traffic directions had almost identical losses until 51% load, and then the loss went to near almost 100%. As frame sizes increased, loss decreased little (compared to other testing scenarios). In fact, near 100% loss was recorded for every frame size with loads exceeding 51%; losses jumped substantially when the load increased from 1% to just 6%, and again when the load increased from 6% to 11%. Even the largest frame sizes performed poorly in this test scenario.

6-to-4 Tunnel Test

Tests were conducted using the SmartBits Smartflow and Avalanche software. Test results indicated this tunnel mode was similar in performance to the Manual or protocol 41 testing, though increasing (larger) frame sizes did not reach the same throughput performance as protocol 41 results. The results of this test were very similar to the protocol 41 results, though the phenomenon of different loss in the opposite direction was observed again (just like in the Generic IPv4 over IPv6 testing noted above). In addition, the difference in path loss direction “flipped.” Furthermore, increasing frame sizes did not result in path losses decreasing as much as the case for protocol 41 test results. There were decreases in losses with increasing frame sizes, but not as noticeable or significant as the protocol 41 testing.

Static Tunnels (broker-to-broker) Test

The static (broker-to-broker) test was the first scenario using the Hexago tunnel broker capabilities over an IPv4 network. IPv6 traffic was passed over the network utilizing static tunnels that were created in the broker devices. Tests were conducted using the SmartBits Smartflow and Avalanche software. No Chariot testing was done for this portion of the test. Test results indicated this tunnel mode had more loss than Manual or 6-to-4 tunnels.

Several interesting observations were made during this testing. For the 128 byte frame size, results were slightly better for the Hexago device than earlier tests done for reference. While there were losses at 1% load (8% one direction and 14% the other traffic direction, as opposed to almost no loss on earlier Manual and 6-to-4 testing), as the load incrementally increased for this frame size, losses were noticeably smaller than those in previous test scenarios. Again, the difference in loss for different traffic directions was noted. But when frame size increased to 256 bytes, the results were slightly worse than earlier router-to-router test results. With each increasing frame size, unlike the router-to-router results, performance did not increase with increasing sizes. Instead, performance slightly degraded until the 1280 byte frames were used. This was the last useable frame size. Performance at this size, when compared again to router-to-router results, was worse. No traffic passed in either direction with the last test utilizing the 1408 frame size. It was later discovered the Hexago box has a Maximum Transmission Unit (MTU) of 1280.

Chariot with Dynamic Tunnels (host to host) Test

The testing was first conducted with IPv4 traffic running over the IPv4 network with no tunnels and using Chariot test software (as a baseline to compare for reference). The test was repeated using the Hexago client software and running IPv6 dynamic tunnels over the IPv4 network with Chariot test software. Results indicated the maximum MTU for the device was 1280; there also appeared to be bandwidth limitations of under 2 Mb.

In addition to the bandwidth limitation and MTU limitations noted above, it was noted that FTP files in excess of 2 kb would not pass. Early attempts with FTP file sizes of 1 Mb failed to pass any traffic. Subsequent attempts with 1 kb files were successful. The dynamic tunnels had significant loss of throughput compared to the IPv4 traffic baseline data. The tests that ran with

the dynamic tunnels took two to four times longer to complete for the same test setup. Furthermore, the dynamic tunnel tests failed every second or third attempt. The IPv4 traffic baseline tests ran to completion with no failures

Conclusions/Recommendations

ISATAP tunneling could route IPv6 traffic, although predictable latency degradation was encountered and the effects of low bandwidth links prevented some file transfers.

D.39 Assessment Report for Evaluating Milestone Objective 2 Microsoft Windows Intra-Site Automatic Tunnel Addressing Protocol

Testing Organization and Publication Date

Air Force Information Operations Center/Information Operations Assessment Division
May 19, 2008

Summary

This assessment allowed the evaluation of the Windows operating systems used within the Air Force Enterprise Network: Microsoft Windows Vista Standard Desktop Configuration (SDC) v2.0.3, Windows XP SDC v1.3, and Windows Server 2003 with Service Pack (SP) 2. All OSs were dual stacked. The results obtained by this assessment helped determine which system and hardware configuration settings need to be addressed to implement an ISATAP-based IPv6 network. The assessment did not consider other network security capabilities or preventive measures found within the networks today. Every effort was taken to ensure the assessment recreated real-world scenarios within the confines of MO2.

Test and Evaluation Method

Demonstration

Joint Staff Operational Criteria Tested

2 (2.3)

8 (8.1.1.2, 8.1.2.2, 8.1.3.2)

Configuration

It is important to note that all clients and servers must be configured to enable ISATAP tunneling. The configuration guidelines can be found in the *Security Configuration Guidance for Milestone Objective 2 (MO2) Microsoft Windows Intra-Site Automatic Tunnel Addressing Protocol*. These steps were followed and implemented before testing was conducted.

Table D-42 lists each hardware device with its associated operating system or platform and the version of the software.

Table D-42 Hardware Software Configuration for Microsoft Windows ISATAP Test

Service	OS/ Platform	Hardware	Software Version
Microsoft Active Directory	Microsoft Windows 2003 Server	Dell PowerEdge 1850 Server	Enterprise Edition Service Pack (SP)
Microsoft DNS	Microsoft Windows 2003 Server	Dell PowerEdge 1850 Server	Enterprise Edition SP 2
ISATAP Client	Microsoft Windows XP SDC v1.3	Dell Optiplex GX520	Professional SP 2
ISATAP Client	Microsoft Windows Vista SDC v2.0.3	Dell Optiplex GX520	Enterprise Edition
IPv6 Helper Services	Microsoft Windows XP SDC v1.3	Dell Optiplex GX520	Professional SP 2
IP Helper Services	Microsoft Windows Vista SDC v2.0.3	Dell Optiplex GX520	Enterprise Edition
Monitoring	Wireshark	Dell Precision 670	Version 0.99.6a
Legend: IP Internet Protocol SDC Standard Desktop Configuration ISATAP INTRA-SITE AUTOMATIC TUNNEL ADDRESSING PROTOCOL SP Service Pack			

Results

Client-to-Server ISATAP Tunnel and Server-to-Server ISATAP Tunnel

Assessment Objective: Clients established an ISATAP interface address with a server located within the same subnet

Results: Once the client and server established an ISATAP interface address, they began to exchange ICMPv6 messages establishing communication. Network ping testing confirmed successful connectivity using IPv4 and ISATAP tunnel interface addresses.

The automatic configuration that makes ISATAP easy to implement also makes it more susceptible to potential exploitation. Unauthorized ISATAP tunnels have the potential to bypass firewall rules blocking protocol 41. This has implications on network discovery, which may allow man-in-the-middle attacks because no authentication is required when ISATAP is installed.

ISATAP is designed for intra-site communication not global communication. The site's border router should block incoming and outgoing Protocol 41 (IPv4 encapsulated IPv6 traffic). If this encapsulated IPv6 traffic is blocked on the network firewall, this will add an additional layer of security as ISATAP is designed specifically as an intra-site transition mechanism.

For a client-to-server ISATAP tunnel, source and destination servers must distinguish between authorized servers and unauthorized servers. Servers need to implement ingress and egress filtering. Windows Server 2003 does not support this; however, Windows Server 2008 should. Server firewall settings or ACL on servers must be enforced. In addition to these recommended server configurations, the site's border router and network firewall should block incoming and outgoing encapsulated IPv6 traffic, adding another layer of security as ISATAP is designed specifically as an intra-site transition mechanism.

Client-to-Client ISATAP Tunnel

Assessment Objective: Clients are tested on different enclaves establishing an ISATAP tunnel. Each client established a connection with a server on their subnet, and those servers acted as relays within the enclaves.

Results: Using standard network protocol analysis, the IPv4 address and its ISATAP tunnel interface address connectivity was verified as fully functional.

Just as in the previous two cases, client-to-client ISATAP tunneling with automatic configuration make this scenario equally vulnerable to external. Similar precaution should be used in network security configurations. Windows Vista SDC provides an advanced firewall feature allowing clients to communicate with authorized clients while allowing it to filter IPv6 encapsulated packets from unauthorized clients.

File Transfer Protocol

Assessment Objective: This portion of the assessment involved creation of an FTP server and ensured clients could access the FTP server.

Results: Success in transferring files located on FTP server to clients using ISATAP tunnel interface address. Clients supporting Windows XP SDC and clients supporting Windows Vista SDC accessed the FTP server and transferred files utilizing an ISATAP link local address.

Internet Explorer

Assessment Objective: Ensure Internet Explorer (IE) 7.0 can access a web page using the web page's IPv6 literal address. When using an IPv6 literal address to browse a website, brackets are needed to enclose the address, as the IE 7.0 browser typically treats anything after a colon as a port number.

Results: Following the standards set by RFC 2732, it was confirmed the IPv6 enclave web page was accessed by its IPv6 address using IE 7.0. It was possible to move around and explore the web page, and access links that were on the main page. Accessing web pages using an IPv6 address does not open new security holes; however, it may allow users to by-pass firewall or proxy server settings that would normally prevent access to previously blocked websites.

Administrators should block all web addresses containing or beginning with the 5EFE:: prefix. Although ISATAP is designed for intra-site use only, this would provide a layer of security should outside clients attempt to connect to Intranet websites.

Telecommunications Network

Assessment Objective: Telnet is for the interactive communication of data and commands between clients with the concept of a session. Telnet is a connection-oriented service that uses port 23 with TCP. When a client wants to access a particular server, it initiates a TCP

connection to the appropriate server, which responds to set up a TCP connection using the standard TCP three-way handshake.

Results: There was success in creating a Telnet connection between the client and server utilizing an ISATAP link local address. No IPv6 security implications were determined during this assessment. Telnet services are rarely used within the Air Force Enterprise Network; in fact, Windows Vista does not have Telnet services installed by default. Administrators will need to enforce all existing Ports, Protocols, and Services (PPS) policies for IPv6 and IPv4.

Filtering Protocol 41 on Vista SDC Firewall

Assessment Objective: Ensured Windows Vista firewall denied transit of IPv6-to- IPv4 packets from unauthorized computers. By default, the Windows Vista SDC v2.0.3 has Protocol 41 blocked.

Results: The Windows Vista firewall on Windows Vista SDC v2.0.3 client denied ISATAP communication with an unauthorized source. The Windows Vista SDC v2.0.3 client was able to communicate with an authorized server and clients using its ISATAP link-local address by explicitly identifying authorized host addresses within the firewall settings. The ability of Windows Vista SDC v2.0.3 and its unique firewall prevents unauthorized hosts from accessing the tunnel or end-host. Current filtering on the firewall was enabled on the outbound filter. Rules creating the same filtering on the inbound firewall should help in preventing unauthorized communication between clients.

Conclusions/Recommendations

ISATAP is a reasonably secure and low maintenance mechanism. It can provide isolated dual-stack hosts with IPv6 connectivity to other IPv6 hosts.