Moonv6 Network Project

June 2007 Technology Status, Test Observations and Results

The Moonv6 project brings together users and suppliers to answer the tough questions about what pieces of IPv6 are ready for prime time. The testing documented in this paper took place in June 2007 and focused on end-to-end basic office application demonstrations including Printing scenarios, NFS (Network File Share), Web design tools, collaboration tools, and SHIM6 (Site Multihoming by IPv6) in a multi-vendor environment.

With the increase of broadband deployments and mobile wireless connections over the last few years it is apparent that the allocation of the current Internet IPv4 addresses will be exhausted. Emerging electronic and consumer equipment including simple home appliances and cellular phones that include a network interface will require an IP address thus we need to look towards the new Internet. Currently IPv4 has a fixed number of addresses and the below figure would indicate that IANA will allocate all available IPv4 addresses within the next few years.

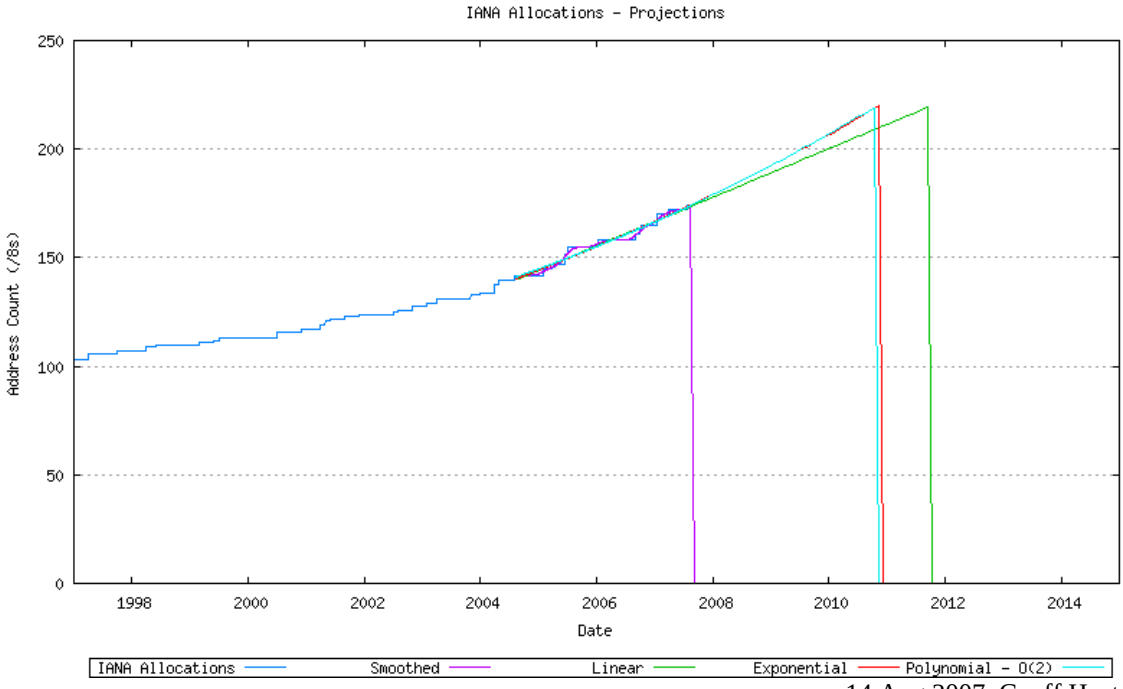

14 Aug 2007, Geoff Huston

Figure 1: IPv4 Address Allocation

Why take interest in IPv6? It's not just a government mandate but new plumbing that will open up possibilities only limited by the imagination. Providers will be deploying NGN (next generation networks), First responders establishing Emergency Managed Networks, Industries including but not limited to Financial and Manufacturing will all gain from the new Internet.

Introduction

Moonv6 is a collaborative project led by the North American IPv6 Task Force (NAv6TF) and includes the University of New Hampshire InterOperability Laboratory (UNH-IOL), U.S. Government agencies and Internet2 (I2). The Moonv6 network, based at the UNH-IOL and the Joint Interoperability Test Command (JITC) Located at Ft. Huachuca, AZ, has rapidly deployed the most aggressive multi-vendor IPv6 network to date. Together with the "IPv6 Ready" logo program, administered by the IPv6 Forum, the Moonv6 project has taken a key role in the IPv6 technology adoption process. Below is a timeline of how different aspects of IPv6 are tested and verified.

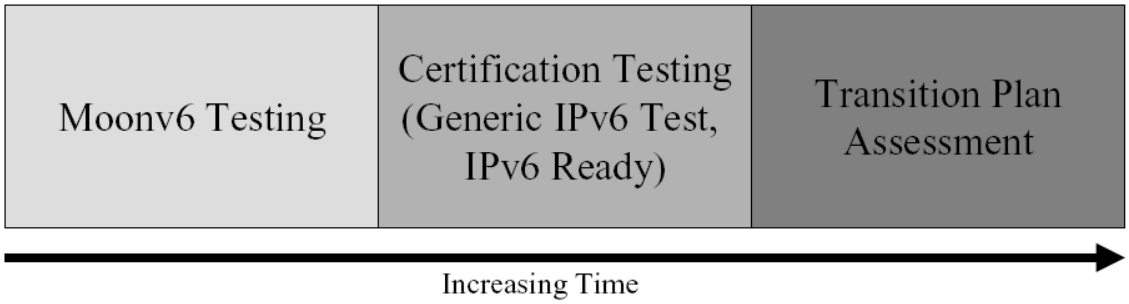

Figure 1: IPv6 Testing Progress

Moonv6 offloads the "bleeding edge" of new areas by providing the users and developers fast feedback in their new implementations. The test plan, results and technical expertise from Moonv6 are used to accelerate the development of new certification processes. Both Moonv6 and the certification processes can reduce the load on enterprises and government agencies when assessing transition plans and deployment strategies.

Phase I of the Moonv6 project established the largest next-generation Internet (native IPv6 network) in North America. Initialized in October 2003, Moonv6 has tested or demonstrated every aspect of IPv6. The first two phases focused on functional stability of routing infrastructure, host connectivity and basic applications.

Further testing in 2004 through 2006 moved IPv6 technology forward through a new round of advanced deployment and functionality scenarios. Test items included Mobile IPv6 (via IEEE 802.11 wireless LANs); Ethernet networks; Applications/Data traffic; Firewalls; Access Policy; Stateful Firewall Functionality; Network-level testing and deployment; IPSec and Applications between Firewalls; DHCP; DNS; VoIP; Transition Mechanisms; Dual Stack Routing; Static Tunnel and additional mechanisms (tunnel broker, DSTM); IPv4/IPv6 QoS network level testing and applications testing.

The NAv6TF's future vision for Moonv6 is to create a virtual Internet backbone with the ability to do pre-production testing for security, multimedia, roaming devices, and other services. Going forward, Moonv6 will serve as a deployment test bed and continue to empower service providers and suppliers from every sector, including industry, universities, research laboratories, Internet service providers and U.S. government agencies.

Participants

The latest round of testing involved service providers, networking companies, including:

Test Scenarios and Results

The June 2007 testing of Moonv6 used similar concepts that were established in the earlier phases of testing. The core network connected all participating sites in a static manner. Using the final network topology established in previous test events this area of testing included basic office applications both over the Local Area Network (LAN) as well as the Wide Area Network (WAN).

Printing Scenarios over IPv6

In order to use IPv6 for practical everyday office activities, applications including printing capabilities will need to be supported. The following demonstration included printing over native IPv6 and dual stack multi-vendor environments using three different printing company vendor devices. The IPv6 enabled printers successfully printed test pages and documents from an IPv6 enabled client using PostScript over a TCP/IP connection. The test procedures included testing both on-link and off-link topologies utilizing the Moonv6 Test Network, refer to Figure 3. This included the verification of using both Link-local and Global addressing syntax as well as the security policy described below.

An IPsec policy was enabled between both printer and host devices in one test scenario as well as between both printer and home gateway in a second test scenario using IKEv1 to establish a secure connection. This security policy was transparent from an applications perspective when printing. Refer to Figure 4.

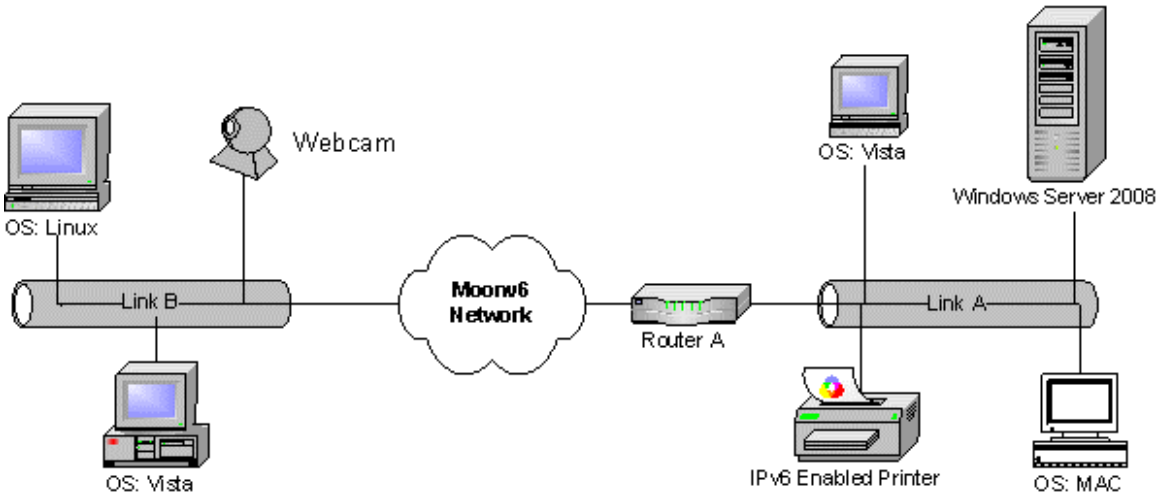

Figure 3: Printing over IPv6 Common Topology

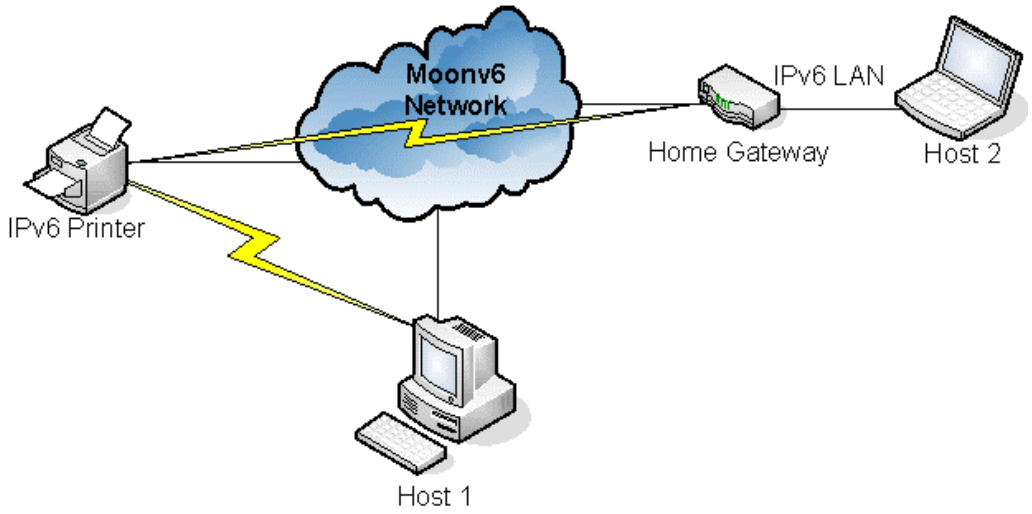

Figure 4: IPsec Printing Topology

NFS over IPv6

Network File System (NFS) allows clients to access files over a network as if they were local. Having the data saved in one location allows for easy relocation and less space consumption. The June 2007 Moonv6 NFS testing consisted of NFS mounting directories through IPv6. Both Link-Local and Global addresses were used to show the ability of export functions between client and server. Both on and off-link directories were created in order to verify that the clients are able to manage directories over NFS.

Multiple clients were configured to mount a directory on an NFS server through Link-local or Global addresses. Clients were connected within the server's local network and outside the network. Clients were successful in creating directories within the shared directory; this was done in order to verify the managing ability of the clients. Both the local network and the off-link networks obtained mount status. Refer to Figure 5.

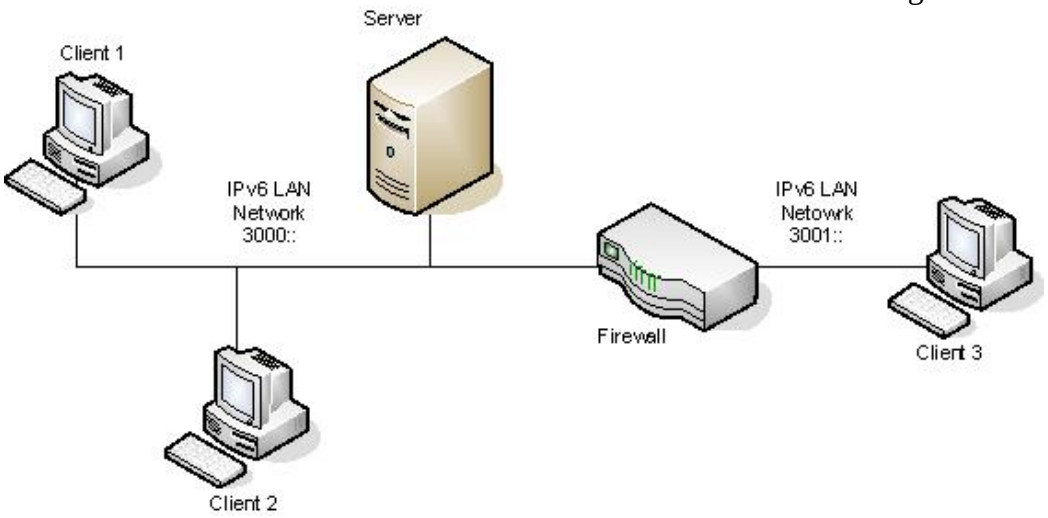

Figure 5: NFS IPv6 Topology

SHIM6

SHIM6 [SHIM6] is a multi-homing solution in IPv6. It is a network layer approach for providing the split of locator/identifier of an IP address [LOCSPLIT], so that multi-homing can be provided for IPv6 with transport-layer survivability.

In essence it specifies a layer 3 shim approach and protocol for providing locator agility below the transport protocols, so that multi-homing can be provided for IPv6, with failover and load spreading properties. This is without assuming that a multi-homed site will have a provider independent IPv6 address prefix that is announced in the global IPv6 routing table. The hosts in a site which have multiple provider allocated IPv6 address prefixes, can use the shim6 protocol to setup state with peer hosts, so that the state can later be used to failover to a different locator pair, should the original one stop working.

A SHIM6 endpoint can use a constant IP address as an Upper Layer Identifier (ULID) for an association. In any case SHIM6 uses multiple IP addresses as locators (L) for routing packets. For each Upper Layer Protocol (ULP) connection, SHIM6 establishes a context state by using four signaling messages: I1, R1, I2 and R2, so the SHIM6 context, associating a ULID pair with a set of locators for endpoints, performs as a per-host header address mapping function. This functionality is indicated in Figure 6.

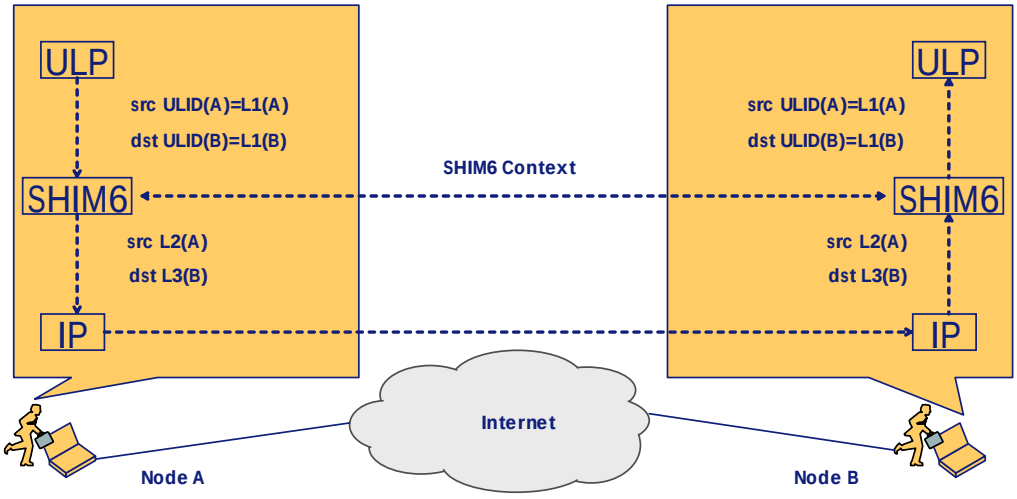

Figure 6: SHIM6 Overview

From the figure, we can see that, the SHIM6 protocol of communicating endpoints (e.g. node A and node B), the ULP selects the initial locator pair (e.g. L1(A) and L1(B)) being the ULID pair [IPV6MHS], which avoids introducing a new identifier name space as well as the modification of ULP. The SHIM6 context provides a set of associations between endpoint identifier pairs (e.g. L1(A) and L1(B)) and locator sets (e.g. L2(A) and L3(B)).

When packets are passed from the ULP to the IP, the endpoint identifiers of ULP are mapped to a current pair of locators. The reverse mapping is applied to incoming packets, where the incoming locator pair is stripped off the packet, and the packet header is rewritten with the mapped endpoint identifier pair. Packets are then passed to the ULP.

SHIM6 Testing

There was a very simple objective for the SHIM6 testing between the TSSG [TSSG] and Moonv6 [MOONv6] test beds and that was to "Verify that SHIM6 can work across the open Internet".

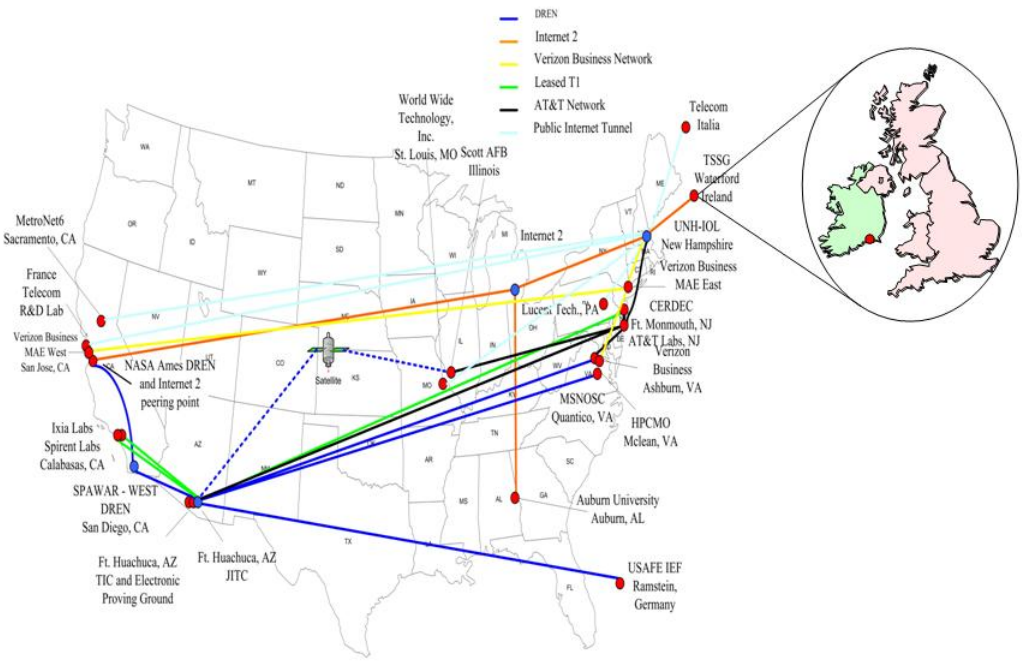

Figure 7: Connectivity between Moonv6 & TSSG Test beds

The first step towards this objective was to find a software implementation of SHIM6, and to this end the one made available by the Universite Catholique de Louvain [UCL] which can be used on GNU Linux, was utilized. The TSSG created a Ubuntu LiveCD [SHIM6 LiveCD] so that both test beds could load identical software implementations of SHIM6.

The second step was to implement a usage test scenario, and for the purpose of this SHIM6 test there was the following:-

Pre-conditions:

a) Node A (Moonv6 Test bed) with 2 network interface cards (NIC), with each NIC connecting to a separate VLAN. Node A must also have an IPv6 enabled web browser.
b) Node B (TSSG Test bed) with 2 network interface cards with each NIC connecting to the same VLAN. Node B must also host a non-descript but large file (300MB) for download.
c) Verified IPv6 connectivity between Node A and Node B.

Test Steps:

1. Node A (Moonv6 Test bed) to use the web browser to locate and then initiate a download of a file (300MB) from Node B (TSSG Test bed).
2. While the download is active, the Node A side tester disables the active NIC in this transportation flow.
3. Observe SHIM6 behavior and the download transport-layer survivability.

Shim6 Lessons learned

While in theory it was unnecessary for both Node A and Node B to have 2 NICs each, for the purpose of this particular test it was necessary so that the basic transport connection between the Moonv6 and TSSG test beds could be verified.

Once a basic connection between the nodes was verified the TSSG had some issue in getting the SHIM6 layer enabled. The problem was eventually identified with the IPv6 access control lists for the TSSG test bed in particular that IP Protocol 61 [Protocol 61] (which SHIM6 uses for control and data transfer) was disallowed. Once enabled, the SHIM6 layer between Nodes A and B became active.

Now the test steps were carried out and it was found that the SHIM6 layer would stop working at Step 1. exactly at the point where the file transfer would be initiated. The problem was identified as a bug in the SHIM6 software implementation used, in that the MTU had to be set to any value less than 1492. SHIM6 adds 8 bytes to the packet. This meant that due to a bug in the implementation, without user intervention, the packet size was increasing to 1508 bytes and thus Ethernet Switches in the route of the test were refusing to pass packets. Reducing the MTU size on both hosts easily rectified this.

Other than these items there were no other issues found while observing the SHIM6 behavior. It successfully used the protocols' locator update messages (i.e. Update Request message and Update Acknowledgement message), the Locator Preference option and a Locator List option for changing the set of locators dynamically.

And once SHIM6 determined that there was a failure in that a locator could no longer be used, the alternative locator was used to preserve the established communications (file download).

Application Testing (Adobe Dreamweaver, MS Meeting Space)

Adobe Dreamweaver allows users to create content and upload files to a central server using FTP, SFTP and WebDAV. To demonstrate this functionality a web page was created and then transferred to a web server using the Adobe Dreamweaver software. The transfer was performed using IPv6 in both IPv6 only and dual-stack environments. IPv6 addresses were entered into the application in a variety of ways including suppressed and unsuppressed forms, using uppercase and lowercase letters, and the domain name of the web server. This was done to demonstrate the ability of the application to handle the different forms of IPv6 addresses.

The testing was performed with the application installed on multiple operating systems including, Microsoft Vista, and Apple OS X. A variety of FTP, SFTP, and WebDAV servers were also used. The application was able to connect using the above file transfer protocols and successfully transferred an uncorrupted copy of the file. This was verified by using a web browser to view the uploaded web page over IPv6.

Microsoft Meeting Space allows users to collaborate by enabling users to share files, send notes and share desktop applications for remote users to view and control. Users were able to successfully connect to a meeting using IPv6 over the Moonv6 Test Network. Once in the meeting, users were able to send messages to each other and allow other users to control their desktop in real time.

During the June 2007 Moonv6 test event an IPv6 only public safety application was demonstrated over a Wide Area Network. This application integrates remote IPv6 based sensors, personnel tracking, and video streaming for use in an IPv6 based public safety network. This demonstration included an engineer at the UNH-IOL participant site to administer a command center for tracking first responders using the public safety software to a mobile network in Herndon, VA. The application utilized the sensor systems in order to monitor and obtain data including location, health and environment statistics to the command center. This application may be used over an Emergency Management Network (EMN) to support first responders in an emergency event. See below for more details on EMN.

MetroNet6 Network Infrastructure Testing

Emergency Management Networks are designed to easily disseminate information between the Local, State and Federal Agencies in an emergency event. The June Moonv6 testing completed the first steps towards the MetroNet6 vision of providing an EMN test bed that public safety agencies can rely on to test interoperability between IPv6 enabled devices and applications. Moonv6 has been established as a temporary, redundant command and control center in which MetroNet6 and Moonv6 will remain up and running permanently. Communication between UNH-IOL and the MetroNet6 point of presence at iStreet Solutions has been confirmed. The Moonv6 network provides an IPv6 native backbone peering network to help solve the communication issues Metronet6 is addressing. These communication issues include; communication between first

responders in the same network, first responders on different networks, city to city and state to state. Testing will be ongoing and incremental.

Conclusion

Basic office applications were successful for supported features and functions in a dual-stack and IPv6-only environments. The end user transparently operated in either environment without discrepancy.

The largest hurdles to IPv6 deployment and adoption that Moonv6 has identified have been specific device implementation, user configuration issues or lack of support. Naturally, the transition to IPv6 will involve a learning curve for system and network administrators. The first crucial step to utilizing the protocol is training and testing. It is important to educate administrators on IPv6 enabled devices that properly work together. The Moonv6 network is permanently maintained to support plug and play requests from equipment vendors, agencies or system administrators. Interoperability is vital to an emerging technology and Moonv6 helps identify any issues before deployment. That said, there is still a great need for future testing in firewalls, security, PKI, IKE, IPsec, SIP, IMS, multicast, streaming video, mobility and other proprietary applications.

Terminology

DoD	United States Department of Defense.
DISR	DoD Information Technology Standards Registry. Formerly the Joint Tactical Architecture. The list of standards that the U.S. DoD uses as requirements in its networks.
DNS	Domain Name Server.
DSCP	Diff-Serv Code Point. Used to differentiate different types of traffic. Uses the ToS bits in a packet header.
DHCP	Dynamic Host Configuration Protocol
ICMP	Internet Control Message Protocol. ICMP Echo Requests and Replies facilitate troubleshooting at Layer 3 for both IPv4 and IPv6. IPv6 has built extra features into ICMP.
IGP	Interior Gateway Protocol.
IPSec	Internet Protocol Security
IPv4	Internet Protocol Version 4. The first widely deployed Layer 3 data networking protocol. The 32 bit address is creating an address limitation on the growth and development of the modern internet and creating an interest in IPv6.
IPv6	Internet Protocol Version 6. A next generation Layer 3 data networking protocol. The 128 bit address space and additional features in the design creates a flexible alternative to IPv4.
JITC	Joint Interoperability Test Command
JUICE	Joint Users Interoperability Communications Exercise
NAT	Network Address Translation. This concept is used to temporarily solve the problem of lack of IP addresses.
NAv6TF	North American IPv6 Task Force. The NAv6TF supports and drives the IPv6 US Summits in North America, promotes IPv6 with industry and government, provides a technical and business center of expertise for the

	deployment of IPv6, provides white papers, briefings, and presentations for public consumption, and works with the IT sector to understand the effects of IPv6 transition on the enterprise. The NAv6TF is implementing a plan of action for IPv6 deployment through Moonv6.
OSPF	Open Shortest Path First. An Internal Gateway Protocol (IGP) for IP Routing primarily used in large enterprise and service provider networks.
SHIM6	Layer 3 multihoming shim protocol
SIP	Session Initialization Protocol. Primarily used to setup and facilitate Voice over IP (VoIP).
SMTP	Simple Mail Transfer Protocol. A protocol designed to transfer e-mail reliably and efficiently between servers.
SYN	Synchronize bit in a TCP handshake.
TCP	Transmission Control Protocol. A connection-oriented Layer 4 protocol.
UDP	User Datagram Protocol. A connectionless Layer 4 protocol.

References

[DoD IPv6] DoD IPv6 Standards Profiles for IPv6 Capable Products, Version 1, dated 1 June 2006.

[IPV6MHS] M. Bagnulo, “Updating RFC 3484 for multihoming support,” draftbagnulo-ipv6-rfc3484-update-00.txt (work in progress), Dec. 2005.

[LOCSPLIT] B. Aboba, “IAB Considerations for the Split of Identifiers and Locators,” draft-iab-id-locsplit-00.txt (work in progress), IAB, Mar. 2004

[MOONv6] www.moonv6.org

[Protocol 61] Any host internal protocol (IANA)

[RFC 959] J. Postel, J. Reynolds, File Transfer Protocol (FTP), October 1985.

[RFC 1350] K. Sollins, The TFTP Protocol (Revision 2), July 1992.

[RFC 2328] J. Moy, OSPF, Version 2, April, 1998.

[RFC 2408] D. Maughan, D., Schertler, M., Schneider, M., and J. Turner, "Internet Security Association and Key Management Protocol (ISAKMP)", November 1998.

[RFC 2409] Harkins, D., and D. Carrel, D., "The Internet Key Exchange (IKE)", November 1998.

[RFC 2460] Hinden, R., S. Deering, Internet Protocol, Version 6 (IPv6) Specification, December 1998.

[RFC 2461] Narten, T., Nordmark, E., and W. Simpson, Neighbor Discovery for IP Version 6 (IPv6), December 1998.

[RFC 2462] Thomson, S., T. Narten, IPv6 Stateless Address Autoconfiguration, December 1998.

[RFC 2616] R. Fielding, J. Gettys, J. Mogul, H. Frystyk, L. Masinter, P. Leach, T. Berners-Lee. Hypertext Transfer Protocol -- HTTP/1.1. June 1999.

[RFC 2740] Coltun, R., Ferguson, D., Moy, J. OSPF for IPv6, December, 1999.

[RFC 2821] J. Klensin. Simple Mail Transfer Protocol, April 2001.

[RFC 2858] T. Bates, Y. Rekhter, R. Chandra, D. Katz, Multiprotocol Extensions for BGP-4, June 2000.

[RFC 2874] M. Crawford, C.Huitema. DNS Extensions to support IPv6 Address Aggregation and Renumbering, July 2000.

[RFC 2893] R. Gilligan, E. Nordmark, Transition Mechanisms for IPv6 Hosts and Routers, August 2000.

[RFC 3530] S. Shepler, B. Callaghan, D. Robinson, R. Thurlow, C. Beame, M. Eisler, D. Noveck. Network File System version 4 Protocol, April 2003.

[RFC 4301] Kent, S., "Security Architecture for the Internet Protocol", December 2005.

[RFC 4303] Kent, S., "IP Encapsulating Security Payload (ESP)", December 2005.

[RFC 4443] Conta, A., S. Deering, M. Gupta, Ed, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification, March 2006.

[SHIM6 LiveCD] http://enable.tssg.org/ubuntu-shim6v4.3.iso

[SHIM6] E. Nordmark and M. Bagnulo, “Level 3 multihoming shim protocol,” Internet Draft (work in progress), IETF, May 2006. [Online]. Available: http://www.ietf.org/internet-drafts/draft-ietf-shim6-proto-05.txt

[TSSG] www.tssg.org

[UCL] http://gforge.info.ucl.ac.be/projects/shim6/

Special Thanks To:

Jim Bound, Hewlett Packard and North American IPv6 Task Force Chair
Ernie Brown, Ixia
Dave Green, Command Information
Junichiro Hamaguchi, Konica Minolta Business Technologies
Erica Johnson, UNH-IOL
Geof Lambert, MetroNet6
Tim LeMaster, Juniper Networks
Adam Lowe, UNH-IOL
Thomas Peterson, UNH-IOL
Miguel Ponce de Leon, TSSG
Yanick Pouffary, Hewlett Packard and North American IPv6 Task Force Technology Director
John Ronan, TSSG
Chris Sabato, UNH-IOL
Ben Schultz, Microsoft Corporation
George Usi, MetroNet6
Timothy Winters, UNH-IOL
Serena Zhao, Adobe Systems

Figure in Introduction by Geoff Huston, IPv4 Address Report
http://www.potaroo.net/tools/ipv4/index.html

This is to acknowledge that the writing of the SHIM6 Testing section has been partially supported by the European Commission's Information Society Technologies Sixth Framework Programme ENABLE project.