Moonv6 Network Project

Phase II Observations and Results

Executive Summary

Moonv6 is a collaboration between the University of New Hampshire InterOperability Laboratory (UNH-IOL), the U.S. Department of Defense (DoD), the North American IPv6 Task Force (NAv6TF) and Internet2 (I2). The Moonv6 project, 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 backbone to date. Together with the "IPv6 Ready" logo program administered by the IPv6 Forum, the Moonv6 project is designed to test and promote Internet Protocol version 6 (IPv6), the next generation of the Internet Protocol (version 4) that currently runs most of the Internet. The results of Moonv6 phase II testing provided the North American market with strong validation for IPv6 by revealing its effectiveness under operating conditions. The completion of Moonv6 phase II formally launched the network as a native IPv6 backbone available for network peering worldwide.

Phase I of the Moonv6 project established the largest next-generation Internet (native IPv6 network) in the world. Deployed in October 2003, the first phase tested basic applications crucial to commercial rollout of IPv6, including file transfer protocol (FTP), Telnet and videoconferencing applications. Phase II, which ran from February 2nd through April 7th 2004, completed the initial testing phases by successfully demonstrating high speed links, advanced routing functionality, firewalls, quality of service (QoS) and other key features of IPv6.

A carrier-class architecture involving an international roster of service providers, Phase II further demonstrated that current IPv6 networking technology is stable and resilient in some of the scenarios tested, but more testing is necessary before it is ready for integration with today's Internet. More than two-dozen vendors participated in the tests. Service providers AT&T, Chunghwa Telecom, France Telecom and Nippon Telegraph and Telephone Corp. (NTT) helped design the test scenarios, while engineers facilitated testing at nine separate sites.

Introduction

The IPv6 protocol, developed by the IETF, vastly expands the possible number of IP addresses needed by every new device that connects to any IP network, including the Internet. As address space decreases and the stop-gap solutions that make the Internet usable create hurdles to flexibility, the question of an IPv6 transition becomes one of "when" rather than "if."

In addition to providing more flexible internetworking capabilities, IPv6's increased address space will facilitate easier network administration, tighter security, and a hierarchical addressing scheme that is expected to boost scalability. The U.S. Department of Defense has stated it is interested in IPv6 because the current Internet protocol, IPv4, has been in use for almost 30 years and cannot support emerging requirements for address space, mobility and security in peer-to-peer networking.

The Moonv6 Phase II testing demonstrated both mature and preliminary proof-of-concept networks. The following technologies were tested: QoS forwarding, basic firewall functionality, transition techniques, mobile IPv6, Domain Name System (DNS) and IPv4/IPv6 routing protocols including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Intermediate System-Intermediate System (IS-IS) and other applications such as DCTS, JWLI, E-Mail and PKI.

The NAv6TF's future vision for Moonv6 is to create a virtual Internet backbone with the ability to do pre-production IPv6 testing for security, multimedia, roaming devices, and other services as vendors and system integrators begin leveraging the innovative opportunities inherent in IPv6. As such, the core network will remain up and running and available for peering from anywhere in the world. Moonv6 will serve as an ongoing test bed and continue to empower service providers and equipment suppliers from every sector, including industry, universities, research laboratories, Internet service providers and the U.S. military.

It will also offer participants who wish to test IPv6-capable technology:

- an operative interoperability setting designed to reduce time to market;
- compressed research, debugging and development cycles enabling faster and smoother creation of end-to-end networking solutions;
- an ongoing platform for global IPv6 education and knowledge enhancement.

Participants

In addition to the four service providers named above, the Moonv6 Phase II test event included the following organizations and equipment suppliers: Internet2, KDDI Labs USA, Native6, Root Server Test Bed, U.S. Department of Defense, UNH-IOL, Agilent, Ixia, Spirent Communications, 6Wind, Check Point, Cisco Systems, Extreme Networks, Foundry Networks, Fujitsu, Hewlett Packard, Hitachi, Hexago, Lucent Technologies, Microsoft, NEC, Netscreen Technologies, Nokia, Panasonic, Procket Networks, SUN Microsystems and Symantec.

Participants

Agilent

(Netscreen Security)

Scenario Implementation

Moonv6 conducted the Phase II multi-vendor, multi-provider native IPv6 interoperability testing event between February and April, building a final topology that will remain operational indefinitely. Testing took place simultaneously at 11 locations listed below. Future locations are currently connecting to the network include Ixia and Spirent as well as Helsinki University of Technology.

- UNH-IOL in Durham, NH
- JITC in Ft. Huachuca, AZ
- JITC in Indianhead, MD
- AFCA, Scott Air Force Base, IL
- CERDEC, Ft. Monmouth, NJ
- Marine Corps Network Operations and Security Command (MCNOSC) in Quantico, VA
- Technology Integration Center (TIC) in Ft. Huachuca, AZ
- USAFE IEF Ramstein, Germany
- AT&T Labs, NJ
- France Telecom R&D Lab, CA
- Auburn University, AL

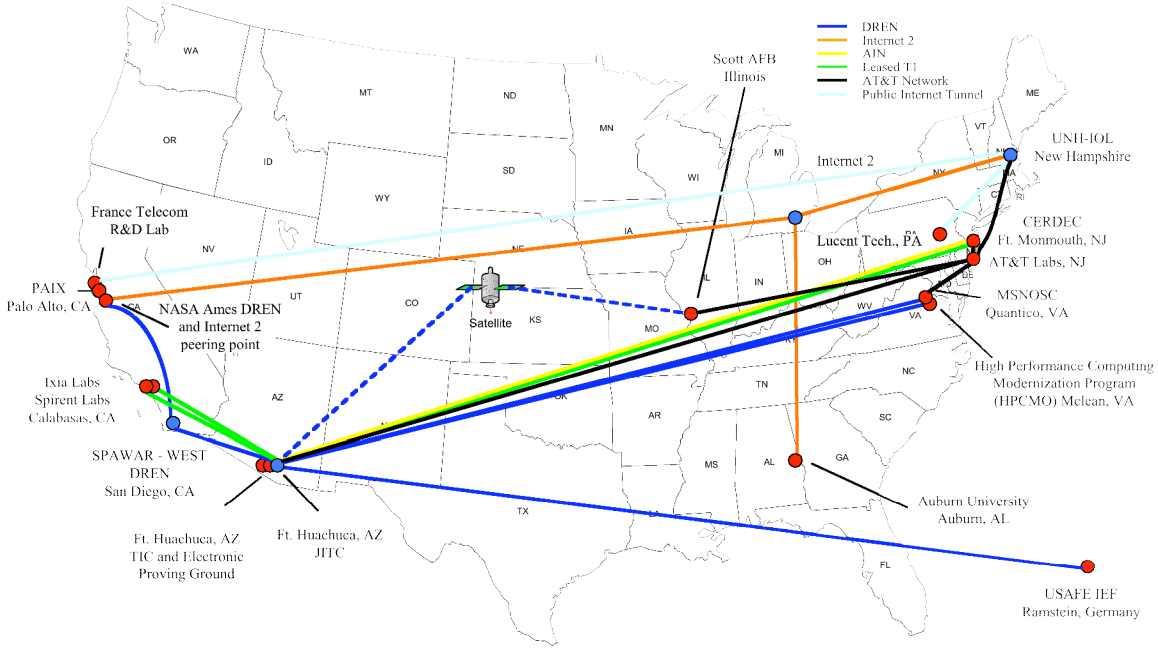

Figure 1: Geographical Map of Moonv6 Connectivity

UNH-IOL and JITC Ft. Huachuca hosted the primary interoperability test locations. Equipment vendors invested a significant number of engineering and equipment resources at these locations. Similar to the architecture of the Phase I testing, the empirical and scenario-based interoperability testing primarily took place at these

locations. The additional sites hosted application servers and routers and were also connected to the backbone.

The Defense Department provided the majority of the network's sites with each branch of the military present for Phase II. IPv6 testing equipment was temporarily loaned to distant sites throughout the network to facilitate testing and deployed services for IPv6. Each branch of the military has a problem analogous to that of a commercial service provider; building any network necessitates the purchase of equipment and deployment of reliable network services. Commercial service provider AT&T provided T1 connections between several of the participating sites.

In the private sector, NAv6TF is the leading industry task force behind the propulsion of the IPv6 protocol into widespread North American deployment. The NAv6TF membership provided logistical support in the form of assistance in developing the initial concept of Moonv6, participating in planning sessions and test scenario creation, marketing the Moonv6 concept, providing educational assistance to various participants from the DoD community prior to Phase I, and lending technical expertise and engineering support where needed. Moonv6 is the implementation of the goals of the NAv6TF.

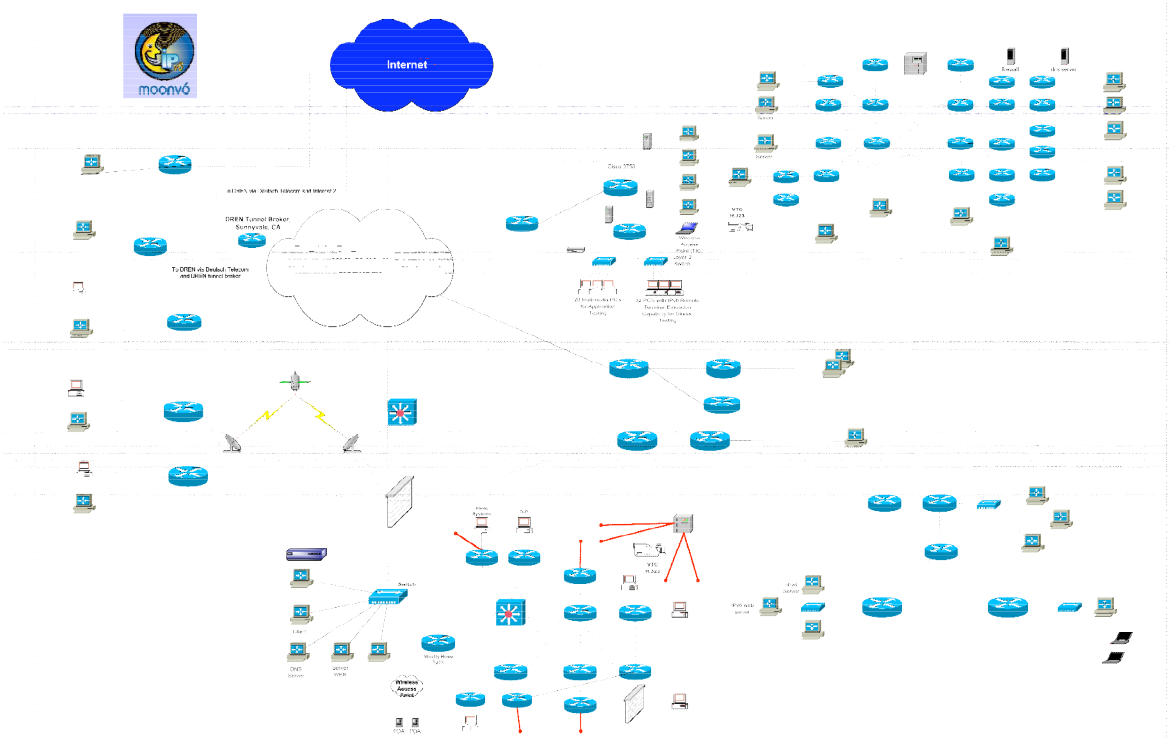

Figure 2: Logical Map of Moonv6 Connectivity

Another crucial component of Moonv6, Internet 2 is a consortium comprising 205 universities working in partnership with industry and government. Internet 2 operates an advanced backbone that supports the development and deployment of advanced network applications and technologies among university and research member organizations. This process accelerates the creation of tomorrow's Internet. As seen in Figure 1, Internet 2 provided the link between UNH-IOL and the Defense Research and Engineering Network (DREN), the Defense Department's recognized research and engineering network.

The DREN is a robust, high-capacity, low-latency nation-wide network. The DREN provides connectivity throughout the High Performance Computing Modernization Program (HPCMP)'s geographically dispersed High Performance Computing (HPC) user sites, HPC Centers, and other networks. The DREN Wide Area Networking (WAN) capability is provided under a commercial contract. The DREN WAN service provider has built the DREN as a virtual private network based on a commercial infrastructure. As seen in Figure 1, the DREN is the primary connector between the participating DoD laboratory sites. The DREN enables over 4,300 scientists and engineers at DoD and other government laboratories, test centers, universities, and industrial locations to use HPCMP computing resources. Since its inception, the DREN has been very active in transferring leading edge network and security technologies across DoD and other federal agencies. Since users and resources are scattered throughout the United States, strong interconnectivity with other major networks and high performance test beds at key interconnect points are critical for optimal use of DoD HPC resources.

Test Scenarios and Results

Moonv6 Phase II involved a test setup similar to that used in the Phase I testing. The core network formed a stable backbone. Engineers ran protocol-specific test plans at the UNH-IOL and the JITC Ft. Huachuca sites before building the final topology for each of the edge networks. JITC Ft. Huachuca designed and built a final topology before the end-to-end testing began. Engineers at each of the other sites executed test activities for network applications and some also tested more advanced functionality.

Initial issues were either small configuration or implementation problems that were quickly fixed. Such smaller problems can easily add up to larger ones, and finding them early is essential to any technology's stability. There were some significant issues encountered. Some devices randomly rebooted when sent fragmented packets and a significant number of devices do not support mobile IPv6 functionality. While most applications run in a dual stacked or tunneled environment, few applications beyond HTTP, FTP, and TELNET/SSH support native IPv6 environments.

Note: All testing below, with the exception of applications testing, was tested at UNH-IOL or JITC Ft. Huachuca, and not tested in an end-to-end manner.

Beyond Common Network Applications

Phase II verified basic network operation using common network applications also tested in Phase I. These included HTTP, FTP, TFTP, Telnet and SSH. Microsoft Windows Media Player and Panasonic IPv6-controlled Web-enabled video cameras operated smoothly over the native IPv6 network topology. Additionally, these tests also demonstrated the efficacy of several commercially available media conferencing software applications, including France Telecom's eConf and KDDI's QualityMeeting. These applications transform PDAs equipped with a miniature camera into mobile videophone devices. They are designed to operate via a wireless link to receive the video stream over IPv6. These applications were tested with Hexago's Migration Broker in order to demonstrate IPv6 connectivity over existing IPv4 wireless networks.

Domain name server testing took place between several of the participating host vendors. The participants ran primary and secondary authoritative domain servers and tested the resolution of names across these systems.

Routing Protocol Testing

Routing protocols were extensively tested during Phase II. Among these, OSPF and BGP convergence in a dual stacked network were tested in an empirical manner.

OSPF Configuration (OSPFv2 and OSPFv3 dual-stack testing)
After building the physical topology and configuring the routers, Phase II engineers implemented two scenarios: pulling the link or configuring the link with a higher metric value. When the link was pulled, the network converged in a reliable manner. When the metric was changed, the network converged without dropping any IPv4 or IPv6 traffic.

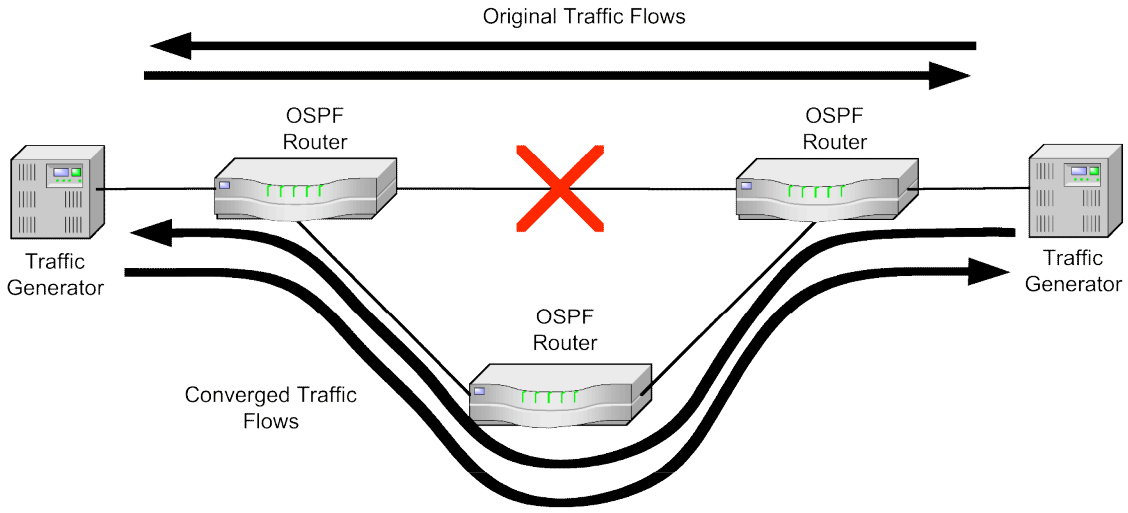

Figure 3: OSPF Configuration

BGP-4 Configuration

This test had two scenarios, first pulling the link or configuring an AS Path Prepend to prefer a different path. When the link was pulled, the network converged in a reliable manner. When the AS Path Prepend was configured, the network converged without dropping any IPv4 or IPv6 traffic. Some implementations of BGP-4 showed difficulty in running the same process on two different ASes simultaneously, one for IPv4 and one for IPv6.

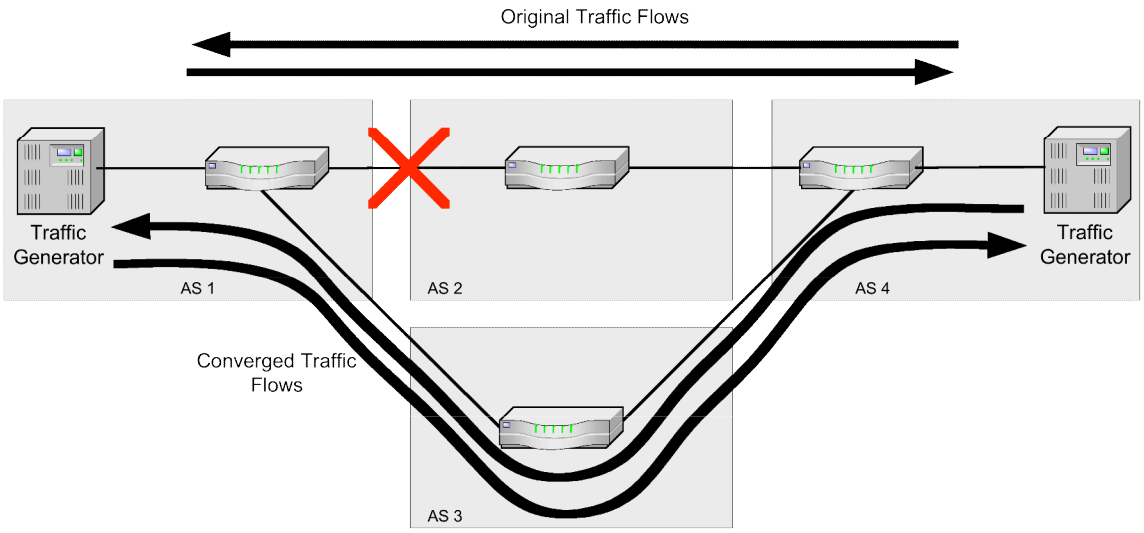

Figure 4: BGP Configuration

IS-IS Configuration

IS-IS testing involved the same topology as that of OSPF. The network worked properly with the same link costs for IPv4 and IPv6. Fewer routers, however, supported IS-IS than OSPF. The test found little support for multi-topology IS-IS among the participating routers. When multi-topology was unsupported, the configured metric value for IPv6 was not able to change. Although this result is not a surprise, it could limit the deployment flexibility of IPv6 in IS-IS networks. As with IPv6 deployments, legacy routers that do not support IPv6 will likely remain in operation for some time. Integrating IPv6 onto these networks without multi-topology IS-IS would be impossible.

Combination Configuration

Two ASes were setup with IGPs. This topology was operational when the equivalent policies were running on all routers. AS1 was running IS-IS and AS2 was running OSPF internally.

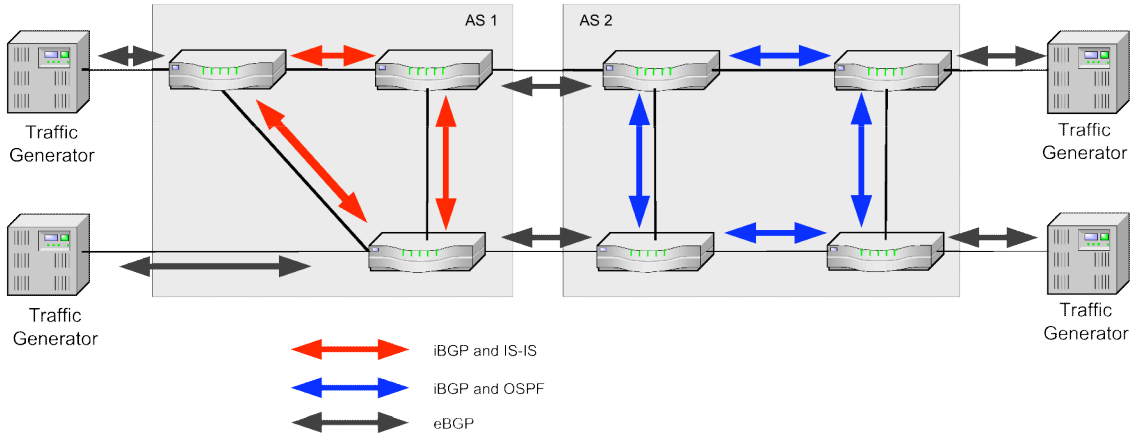

Figure 5: Combination Configuration

When this topology was running and stable, BGP-4 routes for IPv4 were being exchanged over IPv4 transport and BGP-4+ routes for IPv6 were being exchanged over IPv6 transport. The topology placed each traffic generator into a different AS. Each traffic generator generated 5500 IPv4 routes and 5500 IPv6 routes and a 30% traffic load. The routers in AS 1 supported the feature to exchange all BGP routes over IPv4. However, when the test engineers configured this feature, the network ceased forwarding traffic through AS 1, and thus it was concluded that this feature does not operate properly. None of the participating router vendors supported transmitting IPv4 and IPv6 routes over IPv6 transport.

Firewall Functionality Testing

Security is a high priority in today's networks; not surprisingly, there has been significant interest in testing the functionality of IPv6 firewall technologies for IPv6. The Phase II firewall testing areas included the multiple functionality tests described below.

Simple functionality of firewalls demonstrated packet filtering capabilities in which specific packets were denied access or accepted and forwarded. These decisions were made based on the following parameters:

- Source/Destination IPv6 address
- UDP and TCP port numbers
- ICMPv6 packet types

Moonv6 engineers also tested combinations of the above parameters and time-based authorization, as well as the ability to block the following packet floods:

- SYN Flood
- UDP Flood
- ICMPv6 Flood

Testing verified all of the above on the participating vendors' device functionality. The

testing also set the devices to log and/or notify the administrator of the network activity. One host-based firewall was also tested with a subset of the above features. It should be noted that this does not test all possible firewall functionality. These are only preliminary test items and many more functionality tests are needed in the future.

Quality of Service Testing

QoS is an increasingly important aspect of network functionality. As voice, video and data networks converge to share the same network paths, IP routers must differentiate between real-time and best-effort applications. Moonv6 Phase II built several test scenarios to determine whether multiple queues would be supported for IPv6 traffic.

The majority of the QoS testing focused on ingress functionality. QoS uses a 6-bit field to differentiate traffic flows. Mapping a traffic flow based on source or destination IPv6 address or from one DSCP field to another is essential for basic operation. Testing found that these features were widely supported among the participants. Two routers were set up to test this operation. The ingress router did the marking and the egress router forwarded the test traffic.

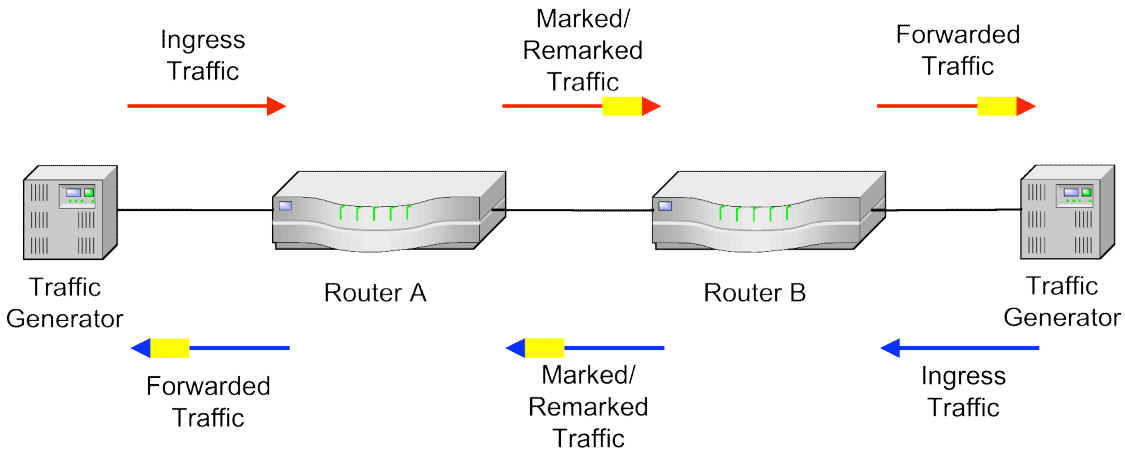

Figure 6: Simple QoS Topology

None of the participants were capable of mapping a layer 2 VLAN tag onto a DSCP field. While lack of this functionality will not hurt QoS for all deployments, some networks may require it.

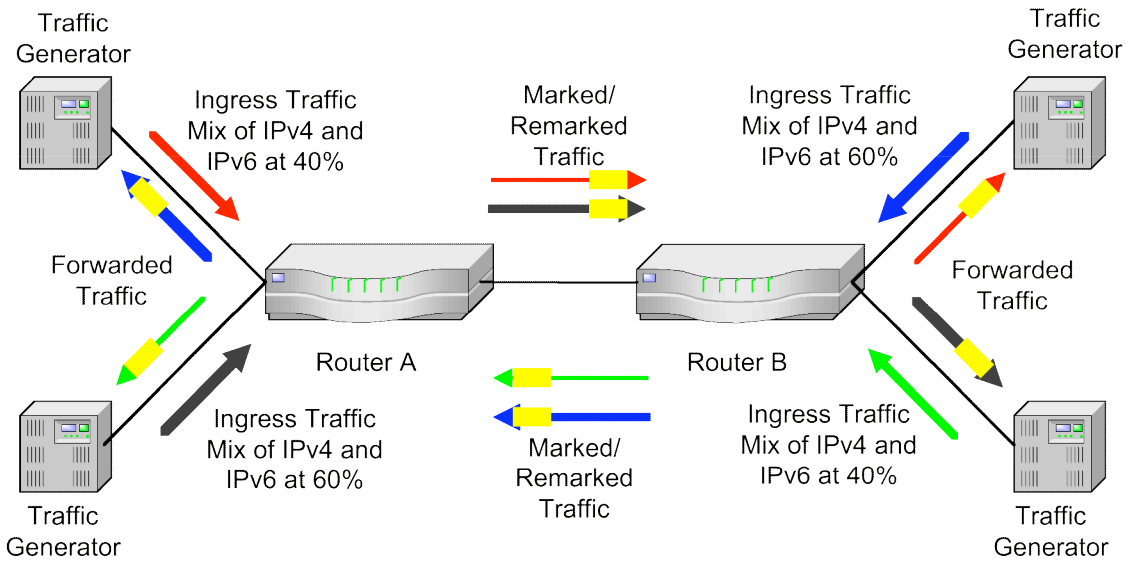

Figure 7: Interface Overload QoS Topology

QoS testing also focused on verifying the proper function of the queuing mechanism on a per-port basis. To test this, the network was set up as in the figure above. Overloading the interface receiving the ingress traffic caused congestion. This also tested the queues for QoS by forcing the forwarding mechanism in the device under test to choose between a priority stream and a non-priority stream. The QoS testing successfully executed this using Gigabit Ethernet interfaces.

Routing protocol issues were revealed in the QoS testing. When an interface was overloaded, OSPF would go down in some cases. This was due to the drop of OSPF hello packets under heavy load. It is recommended that routers should have a mechanism by which protocol control packets, such as OSPF hello or BGP KeepAlive, are handled with higher priority than other packets even if under heavy load.

Transition Mechanisms

Transition mechanisms are essential for IPv4 and IPv6 co-existence, as legacy systems and applications may or may not be ported to IPv6 for some time. There are several techniques to accomplish this task. Moonv6 Phase II focused on two aspects of transition: dual-stack routing and coexistence mechanisms. Dual-stack routing proved to be an efficient way to provide IPv6 when the existing routers support running both IPv4 and IPv6 at the same time.

There will be some networks that only run IPv4 because they do not have the required resources or their administrators simply have not implemented it. Two coexistence mechanisms based on tunneling are considered for these scenarios: tunnel brokers (RFC3053) and Teredo. In the case of the Teredo, due to configuration issues this mechanism could not be evaluated. Conversely, two Tunnel Setup Protocol (TSP) servers were successfully set up at UNH and PAIX allowing IPv6 peer-to-peer communications between IPv4 hosts even when connected to different TSP tunnel brokers. Also the TSP

capability to maintain the same IPv6 address allocated when the end IPv4 host address changes, was demonstrated.

The Hexago implementation of a tunnel broker was demonstrated in a wide variety of topologies. These topologies included networks with simple NAT configuration as well as double NAT configurations. Each of these tests was run both in authenticated mode, in which the IPv6 address is permanent, and in anonymous mode. Automatic DNS name assignment was also successfully tested, as well as IPv6 delegation (with a /64 prefix).

Integration with existing applications was tested with two videoconferencing software packages: KDDI's Quality Meeting and eConf from France Telecom. These tests demonstrated that tunneling techniques can be used effectively with high-bandwidth applications. However some significant packet loss was observed in several tests using UDP encapsulation over wide area network connections.

IPv6 prefix delegation was tested using a laptop as a mobile wireless router and a Panasonic IPv6 camera as a remote device. The camera successfully auto-configured itself and was reachable through the laptop. The camera remained available with short disruptions as the laptop was moving from one IPv4 network to another.

High Speed Proof of Concept

Ad-hoc testing of 10 Gigabit Ethernet links yielded very high throughput. The tests were proof-of-concept that IPv6 is capable of near line-rate throughput in the high-speed, multi-vendor, multi-protocol environment.

Multicast Testing

A small topology was configured to test MLD and PIM-SM. It was first discovered that not all equipment supported MLD. When PIM-SM was enabled, one router supported an optional feature in its hello message. The neighboring router did not support this option and discarded the hello messages, and thus no neighbor adjacencies were formed.

When the topology was rearranged, the PIM exchange was correct up to the register-stop and the shortest path tree (SPT) failed to form. This test was aborted due to a shortage of time.

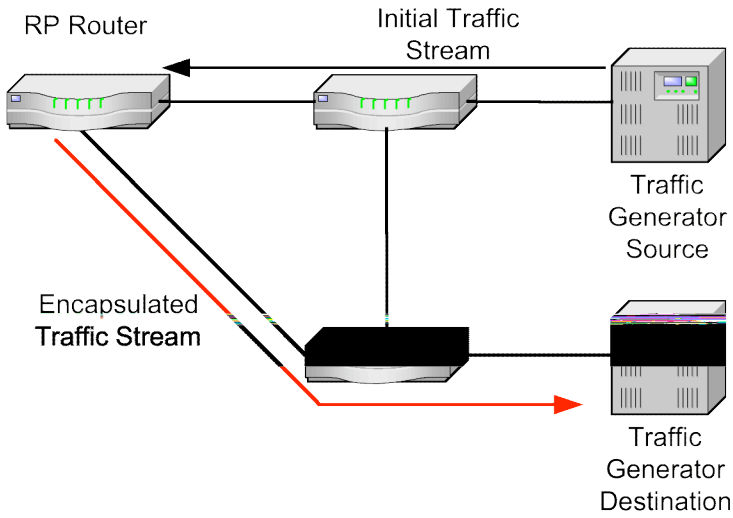

Figure 8: PIM-SM Topology

Mobile IPv6 Testing

Several vendors presented their Mobile IPv6 solutions, including the different Home Agent, Mobile Node and Correspondent Node functionalities. This illustrated the advantages IPv6 enables in the mobility domain. However, some incompatibility issues appeared when trying to interoperate the implementations from different vendors together such as a Home Agent from one company with the Mobile Node of another company. Also, due to lack of time during Moonv6 Phase II event, it was impossible to complete a thorough set of interoperability tests in this area but it is expected that testing will continue in the future.

PPPoE Testing

The PPPoE exchange occurs across a DSL link. The Broadband Remote Access Server (BRAS) performs remote dial-up authentication, and then allows the customer router on the network. Once the customer router registers with the BRAS, the host transmits an ICMPv6 echo request to the BRAS.

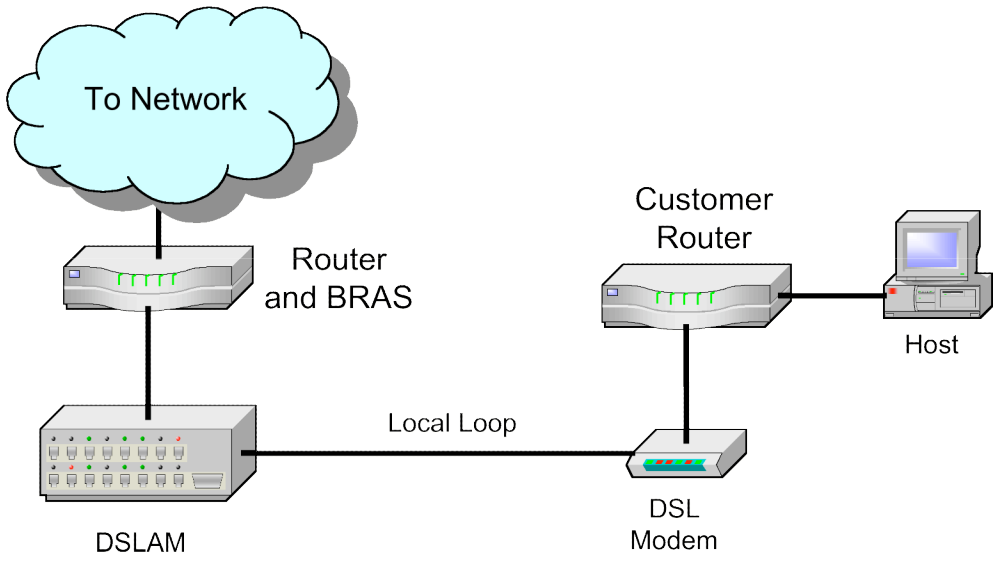

Figure 9: PPPoE Test Setup

PPP negotiation worked properly and the customer router correctly received the EUI64 address from the BRAS. It is important to manually configure a default route on the BRAS. For some customer routers, it is also necessary to add a default route.

Final Topology

Once the testing was completed, a final topology was built for permanent operation. The core routers in this infrastructure will provide the backbone across which long-term applications testing will occur.

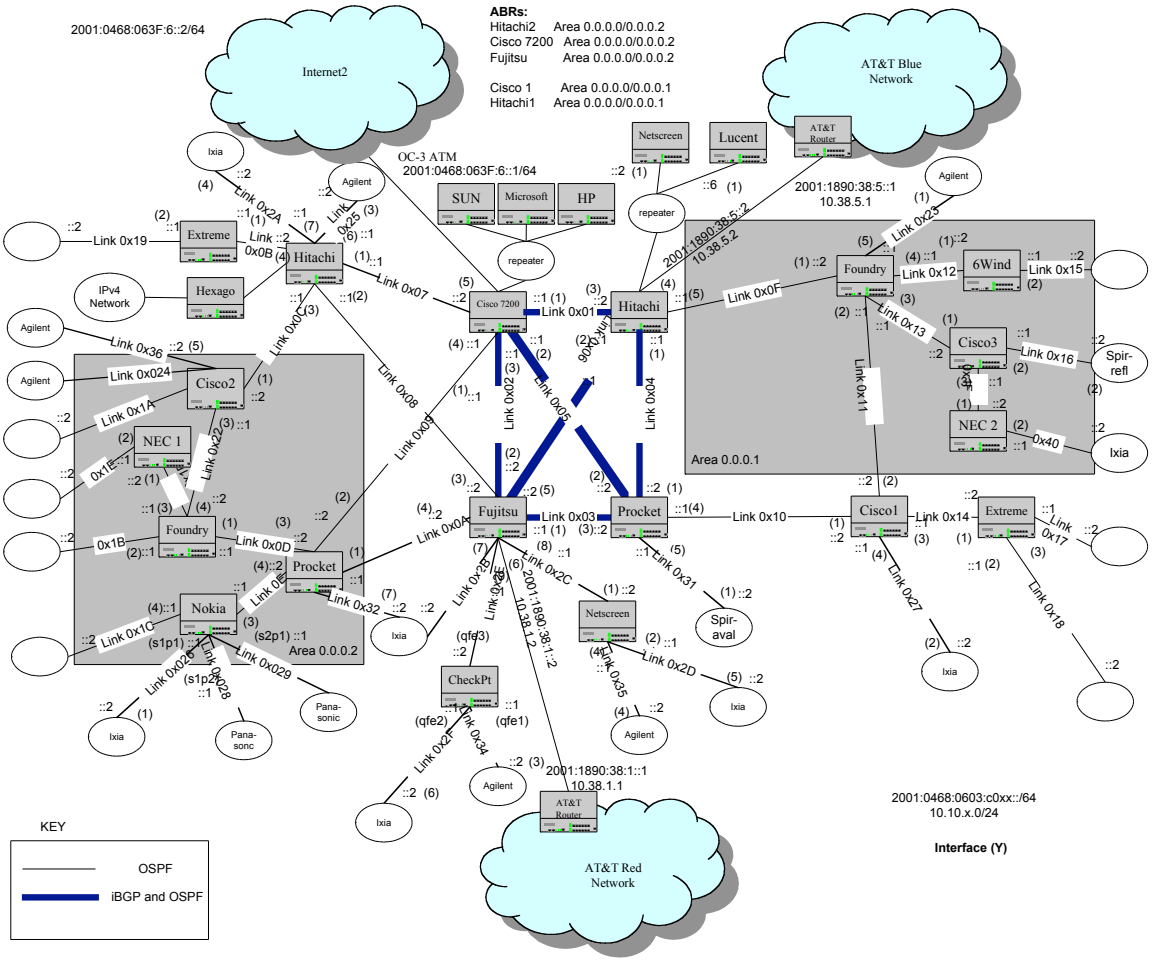

Figure 10: Final UNH Topology for Phase II

Conclusion

International service providers seeing the advantages of IPv6, including the enhanced addressing space that will be needed for new applications and overseas customers, will likely continue to require accepted metrics of interoperability from their equipment vendors. The U.S. DoD, if it continues on its stated course of transitioning completely to IPv6 by 2008, will continue to drive interest in IPv6 in the North American market.

As IPv6 continues to progress in areas that drive new service creation and cost reduction, a key determinant of emerging protocol standardization and commercial adoption is validation in operative networks. To further these efforts, the Moonv6 test events at the UNH-IOL have provided and will continue to provide an aggressive test scenario built around service providers' requirements and real-world deployment characteristics.

Recommendations for further investigation

As a result of the Moonv6 test methodologies and findings described above, several facets of IPv6 technology emerged as compelling candidates for further testing. Among these are the following:

- Performance and security comparison between different transition mechanisms
- Additional Firewall capabilities tests
- Mobile IPv6 testing across core network topologies
- Additional Multicast testing
- MPLS Services testing
- Advanced services and proof-of-concept testing

Additionally, Moonv6 has generated participation interest worldwide. The next 12 months will see the topology of Moonv6 extend beyond the borders of North America and interconnect with dozens of other IPv6 networks from Europe and Asia, as well as numerous universities, vendor networks, and service providers. The Moonv6 committee will need to develop a methodology for connecting with these entities and establish common routing, security, and connectivity policies and processes.

Terminology

AS	Autonomous System. A set of routers under a single technical administration that has a coherent interior routing plan and presents a consistent picture of what destinations are reachable through it.
BGP	Border Gateway Protocol. BGP version 4 is currently the most popular External Gateway Protocol (EGP) for IP Routing.
DoD	United States Department of Defense.
DNS	Domain Name Service. The service that maps names to IP addresses.
DSCP	Differentiated Services Code Point. The field in the IPv4 of IPv6 header that defines the per-hop behavior, or priority of a forwarded packet.
DSL	Digital Subscriber Line.
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. A routing protocol that routes within an Autonomous System such as RIP or OSPF.
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.
IS-IS	Intermediate-System to Intermediate-System. An Internal Gateway Protocol (IGP) for IP Routing primarily used in service provider networks as an alternative to OSPF.
JTA	Joint Tactical Architecture. The list of standards that the U.S. DoD uses as requirements in its networks.

LDAP	Lightweight Directory Access Protocol. A standards based method of remotely accessing information directories based on the X.500 model.
MLD	Multicast Listener Discovery. An IPv6 registration method for hosts to receive multicast data destined to a certain multicast address. Replaces Internet Group Management Protocol (IGMP) for IPv4.
MPLS	Multi-Protocol Label Switching.
NAT	Network Address Translator. A device that temporarily solves the IPv4 address problem through mapping local addresses of stub networks to (usually fewer) globally unique 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.
NTP	Network Time Protocol. A protocol designed to synchronize the clocks of network nodes from a central server or set of servers.
OSPF	Open Shortest Path First. An Internal Gateway Protocol (IGP) for IP Routing primarily used in large enterprise and service provider networks.
PIM-SM	Protocol Independent Multicast, Sparse Mode. A protocol for efficiently routing multicast traffic groups that may span wide-area networks.
PPP	Point-to-Point Protocol. A standard encapsulation method for transporting IP traffic over point-to-point links.
PPPoE	PPP over Ethernet.
RIP	Routing Information Protocol. Currently an Internal Gateway Protocol (IGP) for IP Routing primarily used small home and office networks.

SIP	Session Initialization Protocol. Primarily used to setup and facilitate Voice over IP (VoIP).
SNTP	Simple Network Time Protocol. A lightweight version of NTP.
SMTP	Simple Mail Transfer Protocol. A protocol designed to transfer e-mail reliably and efficiently between servers.
TCP	Transmission Control Protocol. A connection-oriented Layer 4 protocol.
UDP	User Datagram Protocol. A connectionless Layer 4 protocol.
VLAN	Virtual Local Area Network. Defined in the IEEE 802.1Q Bridging specification. A tagging mechanism that allows several bridged Ethernet LANs to be differentiated through an extra tag field at layer 2.

References

RFC 854 J. Postel, J. Reynolds, TELNET Protocol Specification, May 1983.

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

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

RFC 1981 McCann, J., S. Deering, and J. Mogul, Path MTU Discovery for IPv6, August 1996.

RFC 2030 D. Mills. Simple Network Time Protocol Version 4 for IPv4, IPv6 and OSI, October 1996.

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

RFC 2401 S. Kent, R. Atkinson, Security Architecture for the Internet Protocol, November 1998.

RFC 2406 S. Kent, R. Atkinson, IP Encapsulating Security Payload, 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 2463 Conta, A., S. Deering, Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification, December 1998.

RFC 2516 L. Mamakos, et. al, A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999.

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 2710 Deering, S., Fenner, W., Haberman, B., Multicast Listener Discovery (MLD) for IPv6, October 1999.

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

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

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.

draft-ietf-idr-bgp4-23 Y. Rekhter, T. Li, S. Hares, A Border Gateway Protocol 4 (BGP-4).

draft-ietf-mobileip-ipv6-24.txt D. Johnson, C. Perkins and J. Arkko, Mobility Support in IPv6.

draft-ietf-pim-sm-v2-new-09.txt Bill Fenner, Mark Handley, Hugh Holbrook, and Isidor Kouvelas, Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised), February 2004.

Joint Technical Architecture (JTA) List of Mandated and Emerging Standards (LMES) Version 5.1 (Draft) dated 21 July 2003.

Special Thanks To:

Yann Adam, France Telecom
Marc Blanchet, Hexago
Jim Bound, Hewlett Packard and North American IPv6 Task Force Chair
Ron Broersma, Defense Research Engineering Network (DREN)
Major Roswell Dixon, JITC
Yasuyuki Matsuoka, NTT Corp.
Captain Dan Millane, JITC
Steve Pollack, Cisco Systems
Cathy Rhoades, UNH-IOL
Yurie Rich, Native 6.
Ben Schultz, UNH-IOL
Joesp Sole i Tresserres, France Telecom
Shawn Smith, JITC
Jean François Tremblay, Hexago
Chris Volpe, UNH-IOL