Moonv6 Network Project

November Test Set Observations and Results

Executive Summary

Networking equipment and application vendors have heeded calls from commercial service providers and U.S. government agencies to begin manufacturing products compatible with the next generation of the Internet protocol, Internet Protocol version 6 (IPv6). There is an increasing interest in testing IPv6 technology to ensure legacy networks will sustain the market's shift to a new underlying data communications transport protocol. Through service-provider-driven, multi-vendor interoperability tests and demonstrations, the Moonv6 project has globally demonstrated that the underlying infrastructure of IPv6 is stable for small deployments. Network vendors are now refining their IPv6 implementations to improve the interoperability and scalability of their products with legacy equipment and other vendors' products.

While continued testing is needed in areas such as firewalls, security and IPSec, the Moonv6 November Test Set, which ran from Oct. 30 until Nov. 12, 2004, pushed multi-vendor IPv6 testing into new territory. New areas explored inlucded VoIP via Session Initiation Protocol (SIP), wireless LANs and streaming video via multicast. Specific protocols tested included IPsec, DNS, Dynamic Host Configuration Protocol (DHCP), and iSCSI. Routing, tunneling and QoS were also included. Few abnormalities were found in the wide area network; some remote sites experienced configuration issues largely associated with staff inexperience with the protocol.

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 backbone to date. Together with the "IPv6 Ready" logo program, administered by the IPv6 Forum, the Moonv6 project is successfully testing and promoting IPv6 for deployment benefits such as improved multimedia streaming, IP mobility and an alternative to inherently less scalable and secure network address translation (NAT) Internet strategies.

Sixteen vendors participated in the third round of interoperability tests on Moonv6,. The tests revealed that the basic functionality of IPv6 is stable and capable of running key data communications applications such as voice-based services and multicast, which are increasingly of interest to service providers.

Introduction

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. 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.

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

- a functioning interoperability setting designed to reduce time to market and ease deployment;
- 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.

Phase I of the Moonv6 project established the largest next-generation Internet (native IPv6 network) in North America. 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 over a carrier-class architecture involving an international roster of service providers. The results of Moonv6 Phase II testing provided the North American market with strong validation for IPv6 by revealing its functional stability.

The focus of the November Test Set was to move IPv6 technology forward through a new round of advanced deployment and functionality scenarios. The November event combined test plans from multiple network operators, the UNH-IOL, the JITC, MCI and participant equipment vendors. 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 and DNS; Transition Mechanism Comparisons; Dual Stack Routing; Static Tunnel and additional mechanisms (tunnel broker, DSTM); IPv4/IPv6 QoS network level testing and applications testing.

Participants

The latest round of testing involved multiple service providers and networking companies, including Agilent Technologies, AT&T, Check Point Software Technologies, Cisco Systems, Extreme Technologies, Hitachi, Hewlett-Packard, Ixia, Juniper Networks, Lucent Technologies, Native6, Nortel, MCI, Microsoft, Panasonic, Secure Computing, Spirent, Sun Microsystems and Symantec.




Test Scenarios and Results

The November Test Set of Moonv6 used the same basic concept as earlier phases of testing. The core network connected all sites in a static manner. As the final network topology was being constructed, protocol-specific test plans were executed at both the UNH-IOL and the JITC Ft. Huachuca sites. MCI also conducted testing in their facility and contributed test results to this paper. Engineers at each of the other sites executed test activities for network applications and some also tested more advanced functionality.

At the high-level, several routers implemented IPv6 only in software. Most of these software-only implementations could only handle a small, unrealistic number of routes. These devices could not forward at wire rate speeds and suffered performance degradation when the traffic volume and/or packet size increased. The routing processes on these devices could not handle more than a few hundred routes. Evidently, compressed IPv6 address notation differs among vendor implementations. In the event that a configured address was 2001:408:0:0:3:0:0:1/128, different devices display this address in the following two compressed formats: 2001:408::3:0:0:1/128 and 2001:408:0:0:3::1/128

Network Applications Testing

The November Test Set verified basic network operation using common network applications also tested in earlier Phases. These included HTTP, FTP, TFTP, Telnet and SSH. Microsoft Windows Media Server and Player and Panasonic IPv6-controlled Web-enabled video cameras operated smoothly over the native IPv6 network topology.

Additionally, testers ran applications such as VoIP and multicast streaming over the backbone network.

DNS and DHCP Testing

Several of the participating vendors tested Domain Name Server (DNS). The devices ran primary and secondary domain servers in various configurations and tested basic functionality including zone transfers and resolution of names and addresses. Overall, the results of these tests were positive. However, certain implementations lacked comprehensive IPv6 support. Specifically, several name server implementations could not handle AAAA zone transfers and certain types of resource records when using IPv6 addresses. These implementations, however, could resolve names residing in a local database. Some name server implementations also evinced difficulties in properly handling DNS Notifies over an IPv6 transport, which resulted in failed zone transfers. Support for authentication was also not prevalent. Further testing also demonstrated that some popular clients were not able to use a native IPv6 transport when communicating with DNS servers. Successful advanced DNS testing included ENUM related queries (NAPTR resource records) and GSS-TSIG updates.

The use of DNS was found to work with applications like HTTP, WIKI server and Microsoft Media Client. Various web browsers displayed varying degree of IPv6 support.

Browser issues seemed to be mostly related to the IPv6 functionality of resolvers and the relative ability of the application to utilize an IPv6 transport.

Lastly, DHCPv6 server and client testing included both stateful and stateless transactions as well as the processing of prefix delegation requests. Vendors that offered stateless DHCPv6 clients and servers implementations generated positive results. Stateful DHCPv6 tests were not as successful due to a lack of server implementations with comprehensive support for stateful DHCPv6. Similarly, testing prefix delegation requests using DHCPv6 also had limited success as many DHCPv6 server implementations lacked support for the same. Lastly, testing of the integration between DHCPv6 and DNS, or DDNS (RFC2136), showed positive results. Unfortunately, few DHCPv6 server implementations offered the feature.

Routing Testing

Previously, Phase II extensively tested routing protocols. Regression testing in this area assayed the stability of several new code builds. MCI labs performed additional tests in this area. With regards to the IS-IS routing protocols, all participant routers announced IPv4 and IPv6 capabilities. However, not all implementations supported CLNP announcements and ES-Neighbor TLV (Type 3). In BGP testing, IPv4 eBGP sessions were forwarded with an IPv4 NLRI, and IPv6 eBGP sessions were forwarded with an IPv6 NLRI. The participant routers differed in their implementation of the FIB next-hop. Although both should be acceptable, the FIB next-hop can be either a link-local or a global address. Similar to the results from the Phase II testing, none of the routers supported IPv4 NLRI support for the peer over an IPv6 BGP session. This implementation issue will cause issues when BGP routers need to share IPv4 routing information over an IPv6-only link. A temporary solution to this issue would be to create a static IPv4 tunnel across the IPv6 link. This was not tested, however it would be a practical future test item.

Routing protocol scalability testing yielded some interesting results. The number of routes or paths a router can hold depend on its memory. This is true in all routing environments, including those with current IPv4-only networks. Technically, IPv6 route records are larger than IPv4 route records. When routers ran out of memory during the testing, one of two scenarios occurred. In one case, the BGP sessions stopped. The BGP process continued to hold the allocated memory until it was rebooted. In the other case, the router reported that the memory was exhausted and that the prefix was not learned. For scalability testing, 260K IPv6 BGP routes and 520K IPv4 BGP paths were injected into the routers. Testers incrementally increased the IPv6 routes/paths injection until teaching the ceiling of the device under test.. Convergence time upon reboot was 11 minutes. This is listed to give the reader an example of how long it takes for a router to converge when there is a link down. It is, however, important to realize that this is the result for a specific router with a specific amount of memory and does not reflect for every router. Data traffic was not passed in these scenarios, but this is an important future test item.

Firewall Functionality Testing

Testing the functionality of IPv6 firewalls is a high priority of the network operators that participate in Moonv6. To date, the largest issue in this technology area has been the lack of industry specifications that describe firewall behavior. Nevertheless, using input from participating network operators, the November Test Set made new discoveries with regards to firewall capabilities. As in routing, dual stack operation was of primary interest among the participants. As such, firewalls must maintain simultaneous stateful TCP connections across it for both IPv4 and IPv6. The November Test Set verified this capability.

Similar to the Phase II Moonv6 testing, the tested firewalls provided packet filtering capabilities based on the following parameters:
• Source/Destination IPv6 address
• UDP and TCP port numbers
• ICMPv6 packet types

Access Policy was also tested not only in firewall devices but also in routers. Engineers configured simple stateless accept and deny parameters with IPv6 and verified them with different traffic types. This seems to be a feature that will eventually be present in all routers, however not all routers were observed to support access policy functionality during the testing.

In terms of IPSec tunnel operation there were some unexpected observations. In some devices, user interface limitations were observed with some tunnel configurations. One device could not configure a different address for a tunnel and an endpoint. This led to the situation described in the figure below. While the topology of connecting a host to a firewall is correct, a firewall should support both connecting to a host and connecting to another firewall.

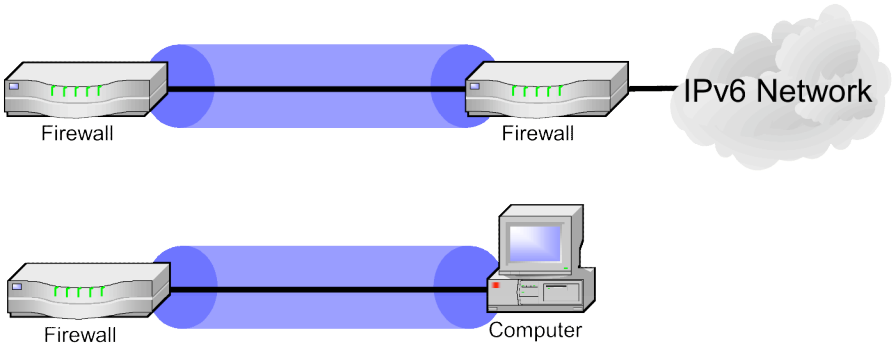

The other observation with tunnels was that one firewall had Neighbor Discovery running on that interface, but not through the tunnel. As a tunnel is a virtual interface, Neighbor Discovery must also run over the tunnel.

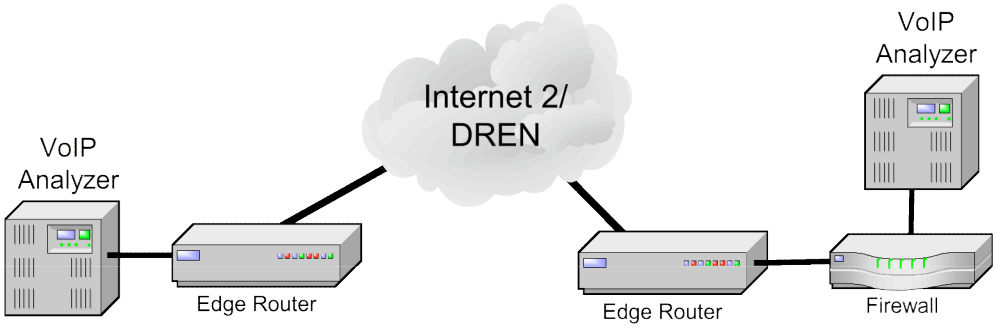

The November Test Set tested stability considerations including a number of stateful TCP connections, a number of HTTP sessions, and a number of VoIP sessions. The VoIP was tested across a firewall and to a remote site. Results demonstrated that some firewall devices could support up to 40,000 stateful TCP connections. Some firewall devices could support more than 1000 sessions of HTTP traffic while improper HTTP traffic was discarded. In this environment, VoIP SIP calls were added. 40 calls were simultaneously supported as active for 20 minutes. This area of multi-protocol testing with voice and data must be further explored to fully understand IPv6 firewall operation.

As a summary of results, the foregoing should not suggest that all participating devices had the same results. These are only preliminary tests, and future testing will need to address many additional firewall functionalities.

iSCSI Demonstration

The Small Computer System Interface (SCSI) is an established set of standards (ANSI T10) for a protocol and a set of I/O buses by which a computer communicates with all types of peripheral devices, such as disks, tapes, printers, etc. Internet SCSI (iSCSI) is a new standard protocol (RFC3720) that removes the distance limitations of traditional SCSI buses by encapsulating SCSI commands and data for transport over a standard TCP/IP network. With iSCSI, a host computer deals with a remote SCSI device as if it were attached via a local SCSI bus.

During the Moonv6 November Test Set, iSCSI was demonstrated to operate over IPv6. The implementations that were tested contained engineering grade software (or "alpha" code) and were not finished products, but this did not inhibit interoperability. An IPv6 connection was properly established between an iSCSI initiator (host computer) from one company and an iSCSI target emulator (storage device controller) provided by UNH-IOL. An iSCSI discovery session properly revealed to the initiator several target devices (logical units) available behind the target controller. Using iSCSI, the initiator successfully formatted, mounted, wrote data to and read data from these remote disks as if they were local SCSI disks.

Further testing needs to be done with iSCSI-IPv6 products from a greater number of vendors to study the effects of various parameter settings which can be negotiated by iSCSI, to investigate problems which may occur when iSCSI is used over longer distances (i.e., timeout settings) and in congested situations (i.e., error recovery), and to

integrate iSCSI with the storage management system on the initiator hosts (including processes for discovery, redirection, authentication and security).

Local and Wide Area Network Topology

The complete testing topology for the multicast and VoIP testing consisted of the below sites. This included local and wide area links for the various participant labs.

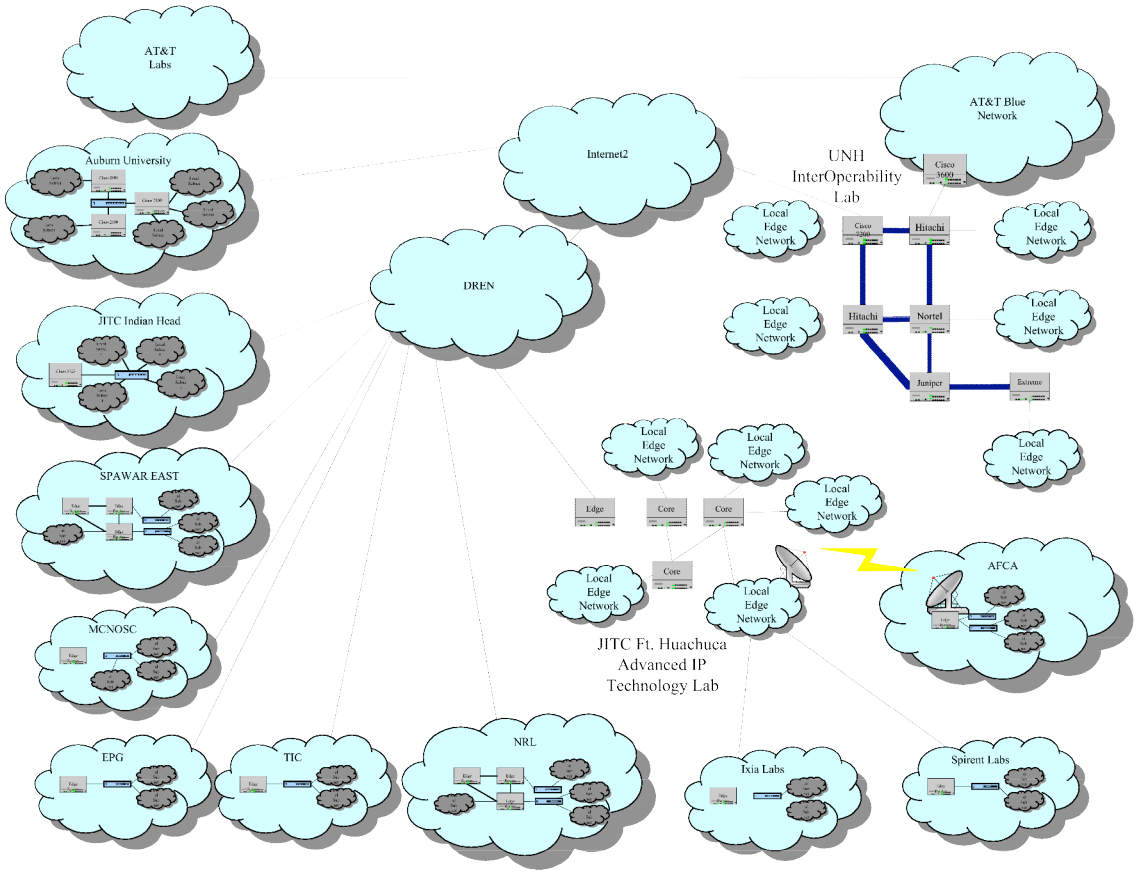

Conclusion

The largest hurdles to IPv6 deployment and adoption that Moonv6 has identified have been either specific device implementation or user configuration issues. Naturally, the transition to IPv6 will involve a learning curve for system and network administrators. That said, there is still a great need for future testing in firewalls, security, IPsec, SIP, multicast, streaming video, mobility and other applications and routing protocols.

Distinct advantages to service providers regarding IPv6 include the enhanced addressing space that will be needed for new applications and overseas customers. IPv6 also plays a key role in restoring the Internet's network address organization and enabling secure reachability across disparate networks. Service providers will 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 ongoing 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.

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.
BRAS	Broadband Remote Access Server.
DoD	United States Department of Defense.
DNS	Domain Name Server.
DSCP	Diff-Serv Code Point. Used to differentiate different types of traffic. Uses the ToS bits in a packet header.
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.
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
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 Translation. This concept is used to solve the problem of lack of IP addresses within an AS.
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. Used to 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.
SPT	Shortest Path Tree.
SYN	Synchronize bit in a TCP handshake.
TCP	Transmission Control Protocol. A connection-oriented Layer 4 protocol.
TSP	Tunnel Server Protocol.
UDP	User Datagram Protocol. A connectionless Layer 4 protocol.
VLAN	Virtual Local Area Network.

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:

Jim Bound, Hewlett Packard and North American IPv6 Task Force Chair
Alan Bavosa, Juniper Networks
John Brzozowski, Lucent Technologies
Ankur Chadda, Spirent Communications
Leigh Huang, Microsoft
Yurie Rich, Native6
Robert D. Russell, UNH-IOL
Ben Schultz, UNH-IOL
Greg Stilwell, MCI
Chris Volpe, UNH-IOL
Beth Vu, MCI
Erica Williamsen, UNH-IOL
Tim Winters, UNH-IOL