



# Moonv6 Network Project



---

## Phase I Observations and Results

---

### Executive Summary

The North American IPv6 Task Force (NAv6TF) has implemented a plan of action for Internet Protocol version 6 (IPv6) deployment, embodied as the Moonv6 project. Based at the University of New Hampshire InterOperability Laboratory (UNH-IOL) and the Joint Interoperability Test Command (JITC), Moonv6 has established the largest and most diverse next-generation Internet (native IPv6 network) in the world. As deployed in October 2003, the network reached more than 3,000 miles from Durham, N.H. to San Diego, C.A., and involved approximately 80 servers, switches and routers configured in dual stack mode, running IPv4 and IPv6 side by side. A professionally designed architecture employing multi-site connectivity and operational test scenarios, Phase I of Moonv6 demonstrated that current IPv6 networking technology is stable, resilient and ready for integration with today's Internet. More than 30 organizations pooled their products, technologies and engineering resources in an industry showcase that confirmed:

- numerous vendors have developed robust, stable, interoperable implementations of IPv6;
- multiple interests – government, educational and commercial – can act collectively to deploy IPv6; and
- IPv6 is ready for widespread deployment throughout North America and the world.

Moonv6 is a collaboration between the UNH-IOL, the U.S. Department of Defense (DoD), the North American IPv6 Task Force (NAv6TF) and Internet2 (I2). Executed simultaneously in multiple locations nationwide, Phase I required precise planning and execution. Service providers helped design test scenarios, while engineers facilitated testing at each of nine individual sites. An always-on videoconferencing link tunneling IPv4 over IPv6 between Durham and Fort Huachuca, Ariz. aided communications and further assayed the protocol.

### Introduction

Today's Internet's is almost 20 years old – a very long time against the backdrop of the accelerated pace of computer technology's development. The Internet's planners could not have foreseen how massively popular the World Wide Web would become. No one expected that of IPv4's four billion total possible addresses, the allotment in Asia would be nearly exhausted by 2004 and in all of Europe not long after. This growing scarcity of IP addresses has resulted from the popularity and proliferation of IP-addressed mobile phones, PCs, laptops and networked printers. The temporary fixes that engineers have devised to continue the operation of IPv4 risk compromising network performance and end-to-end network security. IPv6 was designed to solve these and other problems symptomatic of the Internet's global wildfire-paced growth. The new version of the protocol expands the IP address space from 32 bits to 128 bits, enabling virtually unlimited IP addresses. In addition to this vast addressing capability, IPv6 enables end-to-end security, improved mobility support and simplified address configuration and management. IPv6 will be the backbone of the next-generation Internet. Moonv6 was created to advance the interoperability and deployment of the IPv6 protocol and to promote it throughout the industry.



An important achievement of launching Moonv6 has been the development of a vital focal point for IPv6 technology information from which every participating organization can benefit.

An ongoing project, Moonv6 will continue to empower service providers and equipment suppliers from every sector to work hand in hand in the design and testing of operative end-to-end network solutions to address large pieces of the interoperability challenge. The benefits include:

- Reduced time to market for deployment of a technically sound, interoperable solution;
- Decreased costs and resources to solve interoperability problems;
- An established framework to cooperatively design end-to-end networking solutions in the future.

Phase I of Moonv6 enlisted commercial service providers, the DoD and commercial equipment vendors to design a set of test plans and deployment scenarios. Phase I brought together expertise from Europe, Asia and North America to further develop and share IPv6 products and knowledge. The Phase I Interoperability event demonstrated the following aspects of IPv6:

- Common Network Applications
- Base Specifications
- Transition Mechanisms
- Routing Protocols
- Security
- Mobility

The test plans were subject to intense review by the DoD and participating commercial service providers. These test plans were based on network operational requirements from the commercial service providers and the JTA RFC specification requirements.

### **Benefits of the Moonv6 Phase I Event**

The Phase I Moonv6 IPv6 test bed network will remain in place as a nationwide proving ground for industry, universities, research labs, Internet service providers, the Department of Defense and other government agencies to help facilitate wide-scale adoption and deployment of IPv6 throughout North America. As such, Moonv6 has established a broad and solid platform for current and future IPv6 network design and testing.

The immediate benefits of Phase I participation included intensive hands-on training and product development information that could not be replicated in an individual company laboratory. In exchange for the participation fee and the engineering investment, participants drew from the work of many scientists and innovators working for some of the most prestigious laboratories in the world.

The engineers that participated in the event had the rare opportunity to learn about IPv6 in an intense environment. They gained hands-on experience not only in configuring their products, but also in troubleshooting and solving problems in an operative, heterogeneous IPv6 environment. The hard-won insight that such experience yields into the direction and implementation of IPv6 removes a considerable amount of the guesswork involved with technology development and helps prevent individual vendors from “reinventing the wheel” in their own laboratory settings.

Equipment vendors also benefit by pooling their resources with others in working implementations of diverse standard protocols and architectures. These are areas that do not

provide any competitive advantage, as all companies must create products that contain a certain set of baseline features and functionality.

### Participants in Moonv6 Phase I

#### Test Laboratory Sites



#### Participating Service Providers



#### Participating Equipment Vendors

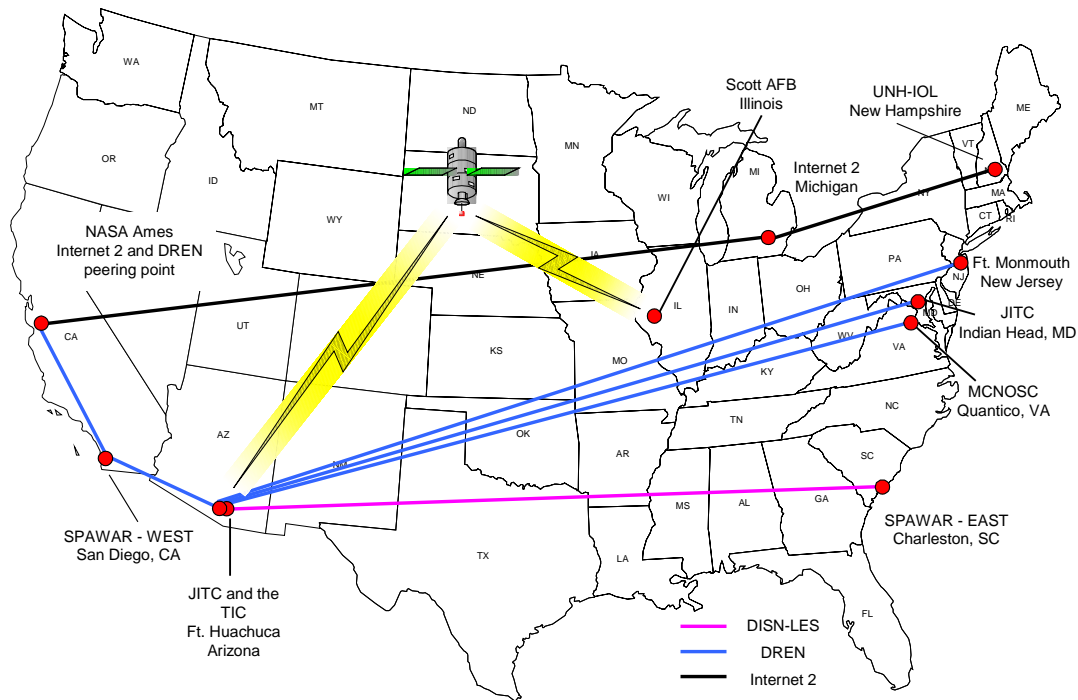




### The Scenario Implementation

Moonv6 conducted the Phase I multi-vendor, multi-provider native IPv6 interoperability testing event between October 6 and October 17, 2003. Testing took place simultaneously at nine locations:

- UNH-IOL in Durham, NH
- JITC in Ft. Huachuca, AZ
- JITC in Indianhead, MD
- AFCA, Scott Air Force Base, Il
- SPAWAR East in Charleston, SC
- Marine Corps Network Operations and Security Center (MCNOSC) in Quantico, VA
- Technology Integration Center (TIC) in Ft. Huachuca, AZ

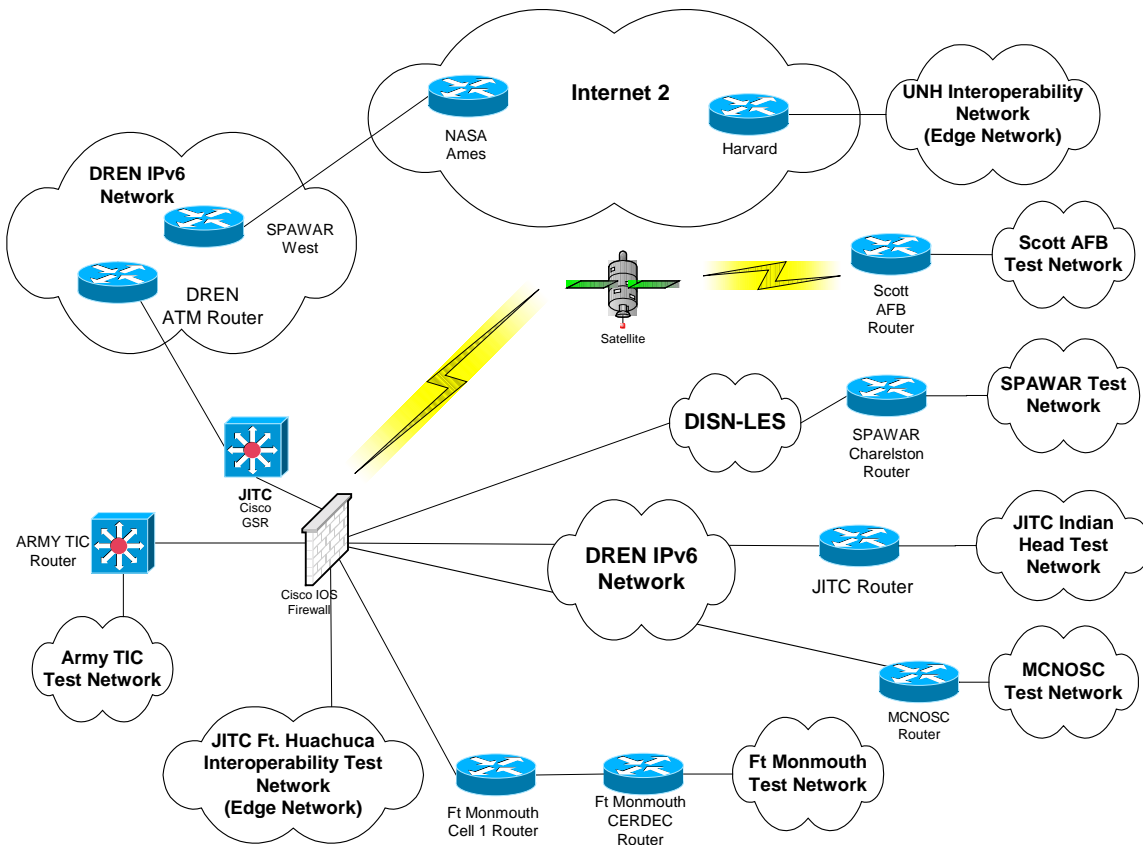


**Figure 1: Network Map**

The two primary interoperability test locations were UNH-IOL and JITC at Ft. Huachuca. These two laboratories hosted multiple equipment vendors and performed empirical and scenario-based interoperability testing. The results of these initial tests informed the design and implementation of complex real-world topologies. The resulting networks were then connected to the Moonv6 backbone. The additional seven sites hosted application servers and routers and were also connected to the backbone. Engineers at multiple locations executed applications across the entire network.

The Defense Department sites provided the majority of the network’s points of presence. In a sense, each branch of the military is also a “service provider” to an extent. Each faces similar problems of purchasing equipment and deploying services to its forces in the field and in the office.

However, in the private sector, NAV6TF is the leading industry task force behind the propulsion of the IPv6 protocol into widespread North American deployment. NAV6TF contributed in several critical areas including organization, event promotion and technical direction. As a NAV6TF project, Moonv6 is the implementation of the goals of the NAV6TF. Another crucial component of Moonv6, Internet2 is a consortium comprising 205 universities working in partnership with industry and government. Internet 2 operates an advanced backbone network 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.



**Figure 2: Logical Network Design**

The DREN is a robust, high-capacity, low-latency nation-wide network. The DREN provides connectivity between and among 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.

**Test Scenarios and Results**

The core network formed a stable backbone for Phase I of the Moonv6 test. Based on pre-defined network topologies defined in the Scenario Implementation section, devices from each vendor were rotated through the ephemeral edge networks at JITC at Ft. Huachuca and UNH-IOL. Engineers at these sites performed specific protocol interface tests in six categories for each rotation of vendor equipment. All other sites executed test activities for the common network applications across the network. After the completion of the vendor equipment rotations, an end-to-end network evaluation was performed. The measurement criteria for this test included passing a mix of IPv4 and IPv6 traffic over the network using participating test vendors and/or a subset of the tests found in the common network applications tests.

Overall the majority of the issues encountered were small configuration or implementation problems that were quickly fixed. Some of these simple issues that, if caught early, could save

time in future IPv6 deployment or operation, are noted below. Others, such as configuration typos, were quickly identified, and in any case were not unique to IPv6.

### ***Common Network Applications***

Common network applications included software applications that run natively over an IPv6 network connection. These applications could use peer-to-peer or client-server models for communication.

Items in this area that were tested and proven included:

- HTTP and HTTPS
- FTP and TFTP
- Telnet and SSH
- DNS
- DHCP

The above tested and proven items worked well over both point-to-point connections and across the network infrastructure.

Multiple participant vendors did not yet support LDAP, SNTP and NTP. Although a native version of SIP for IPv6 was not present, SIP was nonetheless demonstrated to work in an IPv6 in IPv4 static tunnel. And while no participant brought mail software that ran natively on IPv6, SMTP was demonstrated to work by the server vendors by using telnet to port 25 and manually sending mail over that interface.

Issues encountered here included an occasional HTTP connection refusal when accessing Web pages, the cause of which is currently under investigation.

When exchanging files with FTP, there were version issues with EPRT and LPRT. The EPRT command allows for the specification of an extended address for the data connection and is defined in RFC 2428. The extended address MUST consist of the network protocol as well as the network and transport addresses. The LPRT command allows users to specify a “long” address for the transport connection over which data are transferred and is defined in RFC 1545. As EPRT and LPRT have the same function, some devices only supported LPRT and others only supported EPRT. Although EPRT is the most recently defined command structure, this situation did not allow FTP to establish a connection.

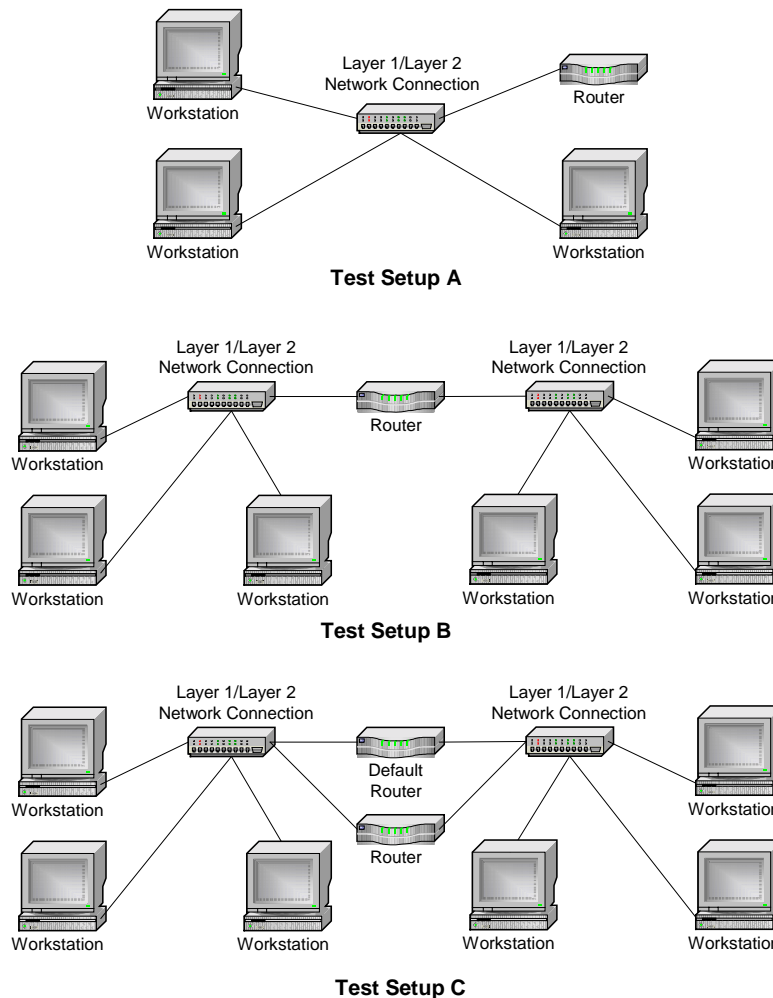
### ***Base Specifications***

The IPv6 Base Specifications include the following RFCs: 2460, 2461, 2462, 2463, and 1981. Unlike IPv4, the functionality of layer 2 address resolution of neighboring hosts and default routers is built into the IPv6 base specifications. The tests in this section included a verification of the following basic operations:

- ICMP Echo Requests and Replies (both on-link and off-link)
- ICMP hop limit exceeded
- Neighbor Unreachability Detection (both off-link and with the loss of the default router)
- The proper transmission and reception of ICMP Redirect messages
- Path MTU Detection and Fragmentation/Reassembly

- TCP/UDP interoperability
- Address Autoconfiguration and Duplicate Address Detection
- Multiple Prefixes and Network Renumbering

Successfully passing these tests, along with additional protocol operations testing, would indicate that a node is reasonably capable of supporting IPv6 requirements. The topologies used to implement these scenarios included simple hosts (and/or routers) is demonstrated in Figure 2 below.



**Figure 3: Test Setup for Base Specifications Scenarios**

The Base Specifications are the most stable, longest existing standards in IPv6. Extensive interoperability testing has been performed on these implementations over the last several years. The interoperability of hosts and routers in this area is extremely stable. As expected, few issues were encountered here.

One implementation issue that was discovered was address selection. This could result in ICMP packet transmission to a location across the network that contained a global destination address with a source address of a link-local address. Because the source address was only link-local, no device that was beyond that link could reply to the ICMP echo requests. Vendors should be careful in address selection implementation.




## Routing Protocols


Routing Protocols allow routers to learn the network addresses for those networks to which they are not attached, and to find the optimal interface to forward traffic to the appropriate destination. There are currently three interior gateway protocols (RIP, OSPF and IS-IS) that support IPv6 and one exterior gateway protocol (BGP-4) that supports IPv6.

Time constraints allowed for only OSPF and BGP-4 to be investigated during Phase I. The project did include some small scale informal RIP testing, but no official results were collected. Much of the testing took place in environments where IPv4 and IPv6 were running simultaneously. This included scenarios with OSPFv2 (OSPF for IPv4) and OSPFv3 (OSPF for IPv6) running at the same time like “ships in the night.” No direct interference between the processes was observed, but further experimentation with routing protocols is necessary, as it is impossible to test every case across all possible topologies.

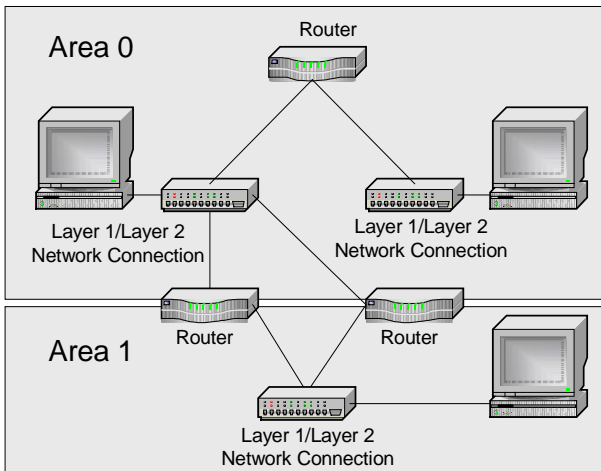
The routing protocol tests included a verification of basic functionality and more advanced rerouting scenarios. The following topologies and test items were tested and verified:




**Figure 4: OSPF Hello Protocol, DR Election and, DR/BDR Failure**




**Figure 5: Intra-Area Route Verification**



**Figure 6: Inter-Area/Intra-AS Route Verification**



**Figure 7: OSPF Multi-Protocol ASes**



**Figure 8: BGP Functionality**

Overall the above test scenarios had a nearly perfect rate of success. However, some of the participating OSPF implementations had yet to be finalized with multi-area routing and/or virtual link capabilities. Configuration and physical connection took a good amount of the testing time for these scenarios.


Issues noted with OSPF testing:

- Network LSA and Router LSA inconsistencies caused certain routes not to be installed in the routing table. These were manually fixed and then worked properly.
- In a stable network, if a link went down, one device set DR and BDR to the previous DR holder. This created problems.
- One device did not dynamically update changes in the network, and its process had to be continuously restarted.
- Inter-Area-Prefix-LSA only had the prefix of the link to which the router was connected and did not contain any area information.
- Intra-Area-Prefix LSAs with zero prefixes in it. The LSAs were rejected as invalid. This resulted in a routing “loop.” The solution was to modify the code to accept the LSAs that do not contain any advertised prefixes.

### OSPF Rerouting

The early success of the small topologies allowed for two larger OSPF networks to be built. With dual OSPFv2 and OSPFv3 operation enabled, an OSPFv2-only router was placed in the middle of the network. Rerouting testing was then performed with both link down and metric change scenarios.

As shown in figure 9, OSPFv2 and OSPFv3 successfully operated simultaneously in all participating routers except for the central IPv4 only router. The topology change took the two traffic flows (the original flows were light blue and red) and rerouted them over two different links (the rerouted flows were orange and dark blue). This result demonstrated that it is important for service providers to carefully design the IPv4 and IPv6 route metrics so that rerouted traffic flows can be predetermined.




**Figure 9: OSPF Rerouting**


The BGP reroute tests revealed similar findings; however there were some important implementation issues noted. Among the behaviors observed were the following:

- Learning BGP routes, but not allowing them to expire. This creates convergence issues.
- No hardware detection for link down, thus a software timeout must detect link-down. This delays convergence time.
- Loss of BGP peers when a link is lost. This creates convergence issues.
- Route reflectors learning and propagating the wrong e-BGP routes.
- The ASBR delivered the ORIGINATOR\_ID Type code transparently to an exterior AS. This is illegal behavior and will cause network problems.

Figures 10 and 11 below note the BGP topologies used in the BGP reroute tests. Although there are only two reroute scenarios displayed, four reroute tests were run in total. The preliminary results collected would suggest that the IPv6 networks running BGP-4+ are resilient, as in each case they were successfully able to converge.



**Figure 10: BGP Reroute Topology 1**



**Figure 11: BGP Reroute Topology 1**

### *Mobility*

Mobile IPv6 is a defined protocol subset of IPv6 that enables a mobile node to attach at multiple points in the network without requiring a change of address on that node. Packets may be routed to the mobile node using a pre-designated address regardless of that node's current point of attachment to the network. The mobile node may also continue to communicate transparently with other nodes (stationary or mobile) after establishing a new point of attachment. This allows higher-level (TCP, UDP and SCTP) applications to continue offering application transport services regardless of where the node is attached.

There are three important entities in the mobile IPv6 exchange. The mobile node is the device that is attaching to different points on the network at different times. At each attachment point the mobile node receives a new IPv6 address. The home agent is a device (usually a router) on the home network of the mobile node. The home agent receives a new registration each time the mobile node attaches and receives a new IP address. The home agent takes all IPv6 packets destined to the home address of the mobile node and forwards them to the newly registered care-of address. The third entity is the correspondent node that is exchanging IPv6 traffic with the mobile node. The correspondent node could be a fixed node or a mobile node.

Phase I of Moonv6 tested and proved several key areas of mobility. Basic Mobile Node to Correspondent Node and Mobile Node to Mobile Node Communication tests worked without an issue. Various scenarios of Home Network Renumbering were also successfully tested. Dynamic Home Agent Address Discovery testing revealed that from remote locations, mobile nodes could properly detect the proper home agent when multiple home agents existed on their home networks from remote locations. Thus, the home agent was verified to defend the mobile node against a station on its home network with a duplicate address. One thing that was found was that overall configuration and setup time were intensive for Mobile IPv6. Configuration of Home Agents was significantly more complex than normal router operation. Time constraints and a finite number of implementations limited the number of scenarios tested.

## ***Security***

IPSec for IPv6 was proven to work with ICMP and TCP in the direct host-to-host scenario. The most significant issue emerged in the user-unfriendliness of the key exchange. An important feature of the protocol is to allow for the display of encrypted keys to an application when an administrator is configuring the device for its preliminary connection. This could save time and energy troubleshooting later. However, it is necessary to prevent access to this feature for normal operation, as it is a security hole.

Time constraints and a limited number of implementations prevented security from being tested on a larger scale. Future security testing could include strategic firewall integration into the network architecture design, red team network attacks and worms/viruses that attack IPv6 applications.

## ***Transition Mechanisms***

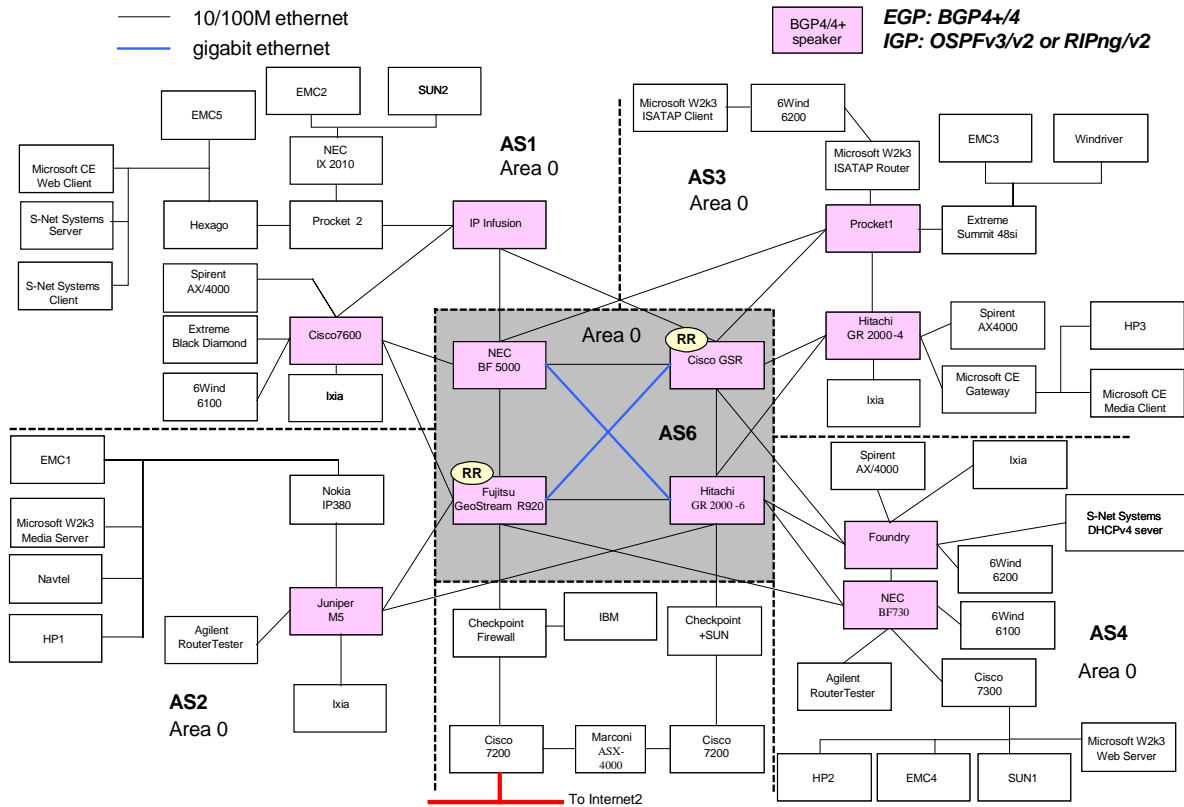
IPv4 is currently the dominant network layer protocol in data communications. To successfully integrate and deploy IPv6, compatibility and seamless integration with existing infrastructure is essential. Numerous transition mechanisms exist in IPv6. Phase I of Moonv6 deployed some of these implementations to demonstrate that they operate properly.

The most extensive testing was executed using IPv6 in IPv4 static tunnels as defined in RFC 2893. Router to Router, Router to Host and Host to Host were tested in this configuration, and ICMP echo requests/replies were successfully exchanged over the tunnels. Three other methods were also tested in the multi-vendor environment:

- Router to Router Tunnel as defined in RFC 3056, also known as 6to4 and verified with an ICMP echo requests/reply exchange.
- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), currently an Internet draft, was verified with an ICMP echo requests/reply exchange.
- Negotiated tunnels using both a Tunnel Broker (RFC3053) and the Tunnel Setup Protocol (TSP, currently an Internet draft) were verified using ICMP exchanges, DNS and HTTP communication.

## **Final Topology**

Testing of the Base Specifications, Routing Protocols and Transition Mechanisms progressed from point-to-point connections to more complex topologies. Applications, mobility and security were tested either point-to-point or over the various test configurations. This approach was effective but could be improved upon. The recommendation is to first test the routing protocols and then build the final topology. Only after, when the network is stable, should the remaining test items be executed. Problems that were discovered and fixed in early testing were naturally less likely to manifest themselves later in the event. Documented problems that were not immediately corrected provided guidance when future issues were encountered.



**Figure 12: Internet Exchange Model Final Topology**

When the addressing architecture was designed for the final topology, a disagreement occurred as to how the point-to-point links were masked. One philosophy was to address the links with 64 bit masks. The other was to address the links with much larger masks, as somewhere between 120 bit and 126 bit masks. There was no agreement on what was correct according to the specifications and operating procedures by the participating service providers. It was agreed that the network would use both systems. AS 2 and AS 3 were addressed with 64 bit masks. AS 1, AS 4 and AS 6 were addressed with much larger masks.

## Conclusion

In a process driven by the North American IPv6 Task Force, Moonv6 created the first large-scale collaborative effort to include Defense Department agencies from every branch of the U.S. military – the Army, Air Force, Navy and Marines – as well as significant commercial service providers and an unprecedented number of IPv6 network equipment vendors.

The Phase I testing had the effect of spurring the DoD's effort to raise awareness of IPv6 in the North American market (the DoD has set the goal to only purchase networking IPv6-capable devices by 2008). It rallied the networking industry to think in a collective manner about IPv6 and assured networking vendors that their devices can interoperate under realistic operative conditions.

Faced with the challenge of selling capabilities for their products to an inexperienced market that currently contains a small but growing interest, equipment vendors are still defining their business

strategies. Vendors are still defining what their customers need and want, which aspects of IPv6 will differentiate their products, and what parts of IPv6 do and do not have a relevant market segment for their product lines. It would seem that, for these vendors trying to gauge a developing but still undefined market, the most effective strategy is to be as active as possible in the IPv6 industry.

Involvement and participation in industry association seminars such as the North American IPv6 Summit meetings can provide input to company decision-makers. Formal and informal discussions illustrate the current state of the industry.

The individual lessons learned from participation in the Moonv6 project run much deeper than those described in this white paper. Individual companies tried their products in a unique environment marked by a strongly diverse network deployment. The innovative ideas that the participants generated will benefit their individual organizations above and beyond the solution of individual problems.

Going forward, Moonv6 will explore IPv6's strengths and limitations in carrier-class (99.999% uptime) networks. This includes the potential for IPv6 network services including MPLS. There is strong interest from participants to perform more extensive DNS, security and mobility scenarios. Adding tests for additional routing protocols (RIP and IS-IS), multicast, anycast and QoS is also under discussion.

Phase II will address a set of the above service challenges and create additional operational test scenarios. Service providers in addition to Internet 2 will peer with the UNH-IOL and connect the Moonv6 network to additional wide-area networks to test service provider-to-service provider interoperability scenarios across the current topology.

The overarching, ongoing goal of Moonv6 is to fight short-term economic gain and promote long-term growth and development in the Internet.

## Terminology

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.
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.
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	Open Shortest Path First. 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.
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.
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.



## 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 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-bgp4-20 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.

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

**Special Thanks To:**

Marc Blanchet, Hexago

Jim Bound, North American IPv6 Task Force Chair and Hewlett Packard Fellow

Major Roswell Dixon, IPv6 Action Officer, JITC

Yasuyuki Matsuoka, NTT Corporation

Steve Pollock, Cisco Systems

Cathy Rhoades, UNH-IOL

Yurie Rich, Native 6

Lt. Col. Jerry Schlabach, JITC

Ben Schultz, UNH-IOL

Shawn Smith, JITC

Chris Volpe, UNH-IOL

Cisco and Marconi (IPv4 only) provided Backbone Routers to connect Moonv6 to Internet 2