SLAC-PUB-8902

IPv6 in ESnet

Warren Matthews, Principal Network Specialist, Stanford Linear Accelerator Center (SLAC)
Bob Fink, Energy Sciences Network (ESnet) Research Staff, Co-chair NGtrans WG, IETF and Project Lead for the 6bone Project, IETF.
Susan Hicks, Oak Ridge National Laboratory (ORNL)
Vyto Grigaliunas, Network Analyst, Fermi National Accelerator Laboratory (FNAL)

Abstract

The importance of the Internet to modern High Energy Physics collaborators is clearly immense, and understanding how new developments in network technology impact networks is critical to the future design of experiments.

The next generation Internet Protocol (IPv6) is being deployed on testbeds and production networks throughout the world. The protocol has been designed to solve todays internet problems, and many of the features will be core Internet services in the future.

In this talk the features of the protocol will be described. Details will be given on the deployment at sites important to High Energy Physics Research and the network services operating at these sites. In particular IPv6 deployment on the U.S. Energy Sciences Network (ESnet) will be reviewed. The connectivity and performance between High Energy Physics Laboratories, Universities and Institutes will be discussed.

Keywords: Network, Technology, Internet, Protocol, IP, IPv6

1 Introduction

High Energy and Nuclear Physics (HENP) experiments around the world generate huge amounts of data, much of which is transferred across networks to collaborators for analysis. Network technology has become critical to the success of experiments. The requirements of HENP researchers often means the networks they use become early implementors of new network technology. In this paper the next generation of the Internet Protocol (IP) is reviewed. In particular the experiences of HENP Laboratories connected to the Energy Sciences Network (ESnet) will be discussed.

2 What is IPv6 ?

IPv6 is simply, version 6 of the Internet Protocol (IP). It is the next generation Internet Protocol developed by the Internet Engineering Task Force (IETF) [1]. Similarly to IPv4 today, IPv6 is designed to enable computers all around the world to communicate across the Internet.

3 Benefits of IPv6

One of the clearest benefits of IPv6 is the increased address space. IPv4 was defined before it was understood just how explosive the Internet would be. Consequently available addresses for IPv4 hosts are running out. Many analysts predict hundreds of millions of mobile phones will be in use and require network connectivity. There will be literally trillions of network devices within a decade.

IPv6 also benefits from the experience of the protocol developers. Many new ideas and features that have been added-on to the IPv4 specification can now be a fundamental part of the new protocol. This certainly helps the various features work as intended with each other.

The specification also enables an IPv6 host to be automatically configured by the upstream router. As the host boots up it will solicit a router for its address prefix and other

Work supported in part by US Department of Energy contract DE-AC02-76SF00515.

information that is usually manually configured. This scheme also makes it much easier to renumber the machines, for example if the site changes service provider. In the case of multi-homed sites there will be multiple addresses per system.

To some, the most important feature is restoring the end-to-end nature of the Internet. This ideal has been eroded by private networks behind Network Address Translator (NAT) boxes. Improved Security is possible with end-to-end connections and in fact, some versions of IP security (IPsec) which is built-in to the IPv6 specification require it. Some games too cannot operate through a NAT.

4 IPv6 Deployment

Most of the large players in the Internet world have released IPv6 versions of their products. Many router vendors including Cisco have IPv6-enabled products. Sun and Microsoft have IPv6-aware operating systems and TCP/IP stacks. IPv6 is also available for Linux, FreeBSD and most other operating systems. Many networks are deploying IPv6. Also, sites are using tunneling techniques to transport IPv6 over their existing IPv4 network. This allows sites unable to install native IPv6 connections to participate early in the effort and gain useful operational experience.

The third generation (3G) mobile phone developers have also settled on using IPv6 as the core of their system. It is clear the future of mobile communication is not limited to voice and the ability to transfer data, whether that is a web page or video conferencing will greatly enhance their product.

5 IPv6 in HENP

In section 1, it was noted that HENP provides a tremendous network challenge. IPv6 is not a solution for these challenges, but it may possibly provide advantages for Physics in the developing countries that do not have adequate address space. As advanced users of the network HENP stands to be amongst the first to benefit.

6 ESnet IPv6 Network

The Energy Sciences network (ESnet) is funded by the U.S. Department of Energy (DoE) to provide wide-area network connectivity for many of the DoE funded basic energy research laboratories and facilities throughout the United States.

ESnet is almost entirely based on ATM technology. This means native IPv6 services can be easily deployed through a permenant virtual circuit (PVC).

6.1 History of IPv6 in ESnet

ESnet staff has been highly active from the beginning of the IETF IPv6 standards development process, and were key participants in organizing the 6bone (See section 8). ESnet was one of the first half dozen transit network providers supporting IPv6 transport in the 6bone in late 1996, which it did by use of IPv6 over IPv4 tunnelling to take advantage of existing IPv4 networks. Subsequently ESnet evolved to a native IPv6 network operating over its ATM virtual circuits. Today, ESnet still provide project management oversight for the worldwide 6bone network.

In 1999 ESnet started the first US-based native IPv6 network peering point, the 6TAP, located at StarTAP facilities at the Chicago NAP (courtesy of NSF). It continues to operate this important peering facility today, and is expanding the services to include New York and Palo Alto (near San Francisco in California) to further encourage early deployment and interconnection of worldwide IPv6 networks.

In mid-1999 the Internet address allocation registries started to allocate IPv6 production addresses (actually routing prefix blocks). ESnet received the very first production IPv6 address allocation, and today there are more than 80 production IPv6 routing prefix allocations worldwide. In addition, ESnet initiated the 6REN (IPv6 in Research & Education Networks) to further promote early deployment. This effort helped stimulate the creation of the IPv6 Forum activity worldwide, which ESnet is a participant in.

6.2 IPv6 in ESnet

IPv6 is a production service in ESnet and support is provided by the engineers and administrators in a similar fashion to support for IPv4. Many ESnet connected sites are participating in early IPv6 deployment at various levels.

6.2.1 IPv6 at SLAC

The Stanford Linear Accelerator Center (SLAC) is a particle physics laboratory operated by Stanford University on behalf of the U.S. Department of Energy located near San Francisco in California. IPv6 hardware has been at SLAC since summer 1999 and a tunnel was configured in the third quarter of 1999. A 2Mbps native connection was brought up shortly afterwards carved from the SLAC production connection. Within the SLAC local area network (LAN) a seperate VLAN provides connectivity into the SLAC switched network. Currently there are FreeBSD, Linux, Windows2000 and Solaris8 machines operating on the IPv6 network at SLAC.

In particular, SLAC is evaluating several firewall products that understand IPv6. The ip6fw program often run on FreeBSD and iptables for Linux are available. However there seems to be an absence of commercial products. Cisco currently only had standard ACLs, it will be some time before Extended ACLs are IPv6 aware.

6.2.2 IPv6 at ORNL

The Oak Ridge National Laboratory (ORNL) conducts research and development in key areas of science. ORNL is managed by UT-Battelle, LLC for the U.S. Department of Energy. An IPv6 testbed was configured at ORNL in the Spring 2000. This network is a physically separate network from the production ORNL network and connects to the ESnet native IPv6 network via a 100M fastethernet connection. Currently, effort is underway to deploy IPv6 enabled ports of essential internet services such as domain name service (DNS), email and web services using SUN Solaris8 systems. DNS is particularly important considering the nature of an IPv6 address.

6.2.3 IPv6 at FNAL

The Fermi National Accelerator Laboratory (FNAL) is a particle physics laboratory near Chicago, Illinois. As with SLAC, IPv6 is enabled on a seperate subnet. FNAL uses a tunnel to connect to ESnet across the regular IPv4 connection. FNAL also has a connection to the local Metropolitan Research and Education Network (MREN) and it is hoped to bring up a multi-homed IPv6 network.

7 Other Networks

Many other networks have implmented IPv6. In particular the Abilene network is of interest to the HENP community as the backbone for the Internet2 project. Abilene is based on packet-over-sonet (POS) technology and unlike ATM cannot easily deploy IPv6 across the production network so a set of Tunnels between the gigapops have been constructed.

8 The 6Bone

The 6bone is an IPv6 Testbed. It is currently a world wide informal collaborative project, informally operated with oversight from the "NGtrans" (IPv6 Transition) Working Group of the IETF.

The 6bone started as a virtual network (using IPv6 over IPv4 tunneling/encapsulation) operating over the IPv4-based Internet to support IPv6 transport, and is slowly migrating to native links for IPv6 transport.

The initial 6bone focus was on testing of standards and implementations, while the current focus is more on testing of transition and operational procedures.

The 6bone operates under the IPv6 Testing Address Allocation [2], and has formalized rules for participation and becoming a transit provider [3]. Currently the 6bone is deployed in over 50 countries distributed around the world, is fast approaching 1000 networks and end-user sites, and has almost 90 networks providing 6bone backbone transit services. Every implementation of IPv6 available is running somewhere in the 6bone.

Inter-network peering is handled by BGP4 in the 6bone, as it is in the IPv4 Internet. There are also numerous public peering points now in use worldwide, most notably in Amsterdam, Chicago (the 6TAP), London, Munich, New York, Palo Alto, Tokyo. Others are rapidly coming online.

9 IPv6 Performance and Reliability

The succesful PingER [4] project has been ported to work on IPv6 and many connections were monitored. Native connections perform very similarly to IPv4. Tunnels can work well but do not scale.

10 The Contribution from HENP

The HENP community, with its special needs and experiences, is in a position to contribute to the success of IPv6.

It is expected ESnet will implement services as they become available and the user community demands them. Sites connected to ESnet will continue to deploy services and replicate their LAN environments mostly for operational experience. Research into multihoming, multicasting and porting of physics tools is targeted, but this work needs funding !

11 Conclusions

Many developers and researchers are become convinced IPv6 will become the standard Internet Protocol. Of course others are not convinced but the momentum is building up and many vendors are implmenting IPv6. Many applications still need porting, and network administrators need operational experience.

References

[1] The IPng Working Group, IETF, See http://www.ietf.org/html.charters/ipngwg-charter.html

[2] RFC2471 ftp://ftp.isi.edu/in-notes/rfc2471.txt

[3] RFC2772 ftp://ftp.isi.edu/in-notes/rfc2772.txt

[4] Warren Matthews *IPv6 Performance and Reliability*, IPv62000 Conference, Washington D.C., October 2000.