

Engaging with IPv6: Addresses for All

First Author¹ Mark Mitchell

University of Glasgow

University of Glasgow, Glasgow, G12 8QQ, Scotland

E-mail: Mark.Mitchell@glasgow.ac.uk

Another Author David Crooks

University of Glasgow

University of Glasgow, Glasgow, G12 8QQ, Scotland

E-mail: David.Crooks@glasgow.ac.uk

Another Author David Britton

University of Glasgow

University of Glasgow, Glasgow, G12 8QQ, Scotland

E-mail: David.Britton@glasgow.ac.uk

Another Author Colin Cooper

University of Glasgow

University of Glasgow, Glasgow, G12 8QQ, Scotland

E-mail: Colin.Cooper@glasgow.ac.uk

Abstract

This paper discusses the deployment of IPv6 on a dual IPv4/IPv6 network stack running Scientific Linux 5 specifically for Grid computing. The focus of this research is on the deployment of this protocol within a Tier-2 environment embedded within a University Campus network.

This paper describes the setup requirements of the hosts and network and discusses some of the advantages and limitations of the preferred deployment at the Glasgow Scotgrid test cluster as part of the wider HEPIX IPv6 programme [1].

*EGI Community Forum 2012 / EMI Second Technical Conference,
Munich, Germany
26-30 March, 2012*

¹ Speaker

1. Introduction

The current global decline in number of publically routable IPv4 address blocks worldwide has led to the requirement to engage with IPv6 [2]. While IPv6 was proposed in 1998 by the IEEEE and has been shipped with both commercial and open source software platforms for over 5 years, it is still be considered as an optional configuration parameter within the network stack, due to the scale of IPv4 deployments elsewhere within the commercial and academic sectors. The present uptake of IPv6 addressing is being driven by the expansion of networking in parts of Europe, Africa and South East Asia [3]. However, uptake of this newer technology has been slower than expected globally. Within W-LCG, a working group has been set up by HEPIX to investigate the issues around the deployment of this technology.

"The HEPiX IPv6 Working Group has been investigating the many issues, which feed into the decision on the timetable for a transition to the use of IPv6 in HEP Computing. The activities include the analysis and testing of the readiness for IPv6 and performance of many different components, including the applications, middleware, management and monitoring tools that are essential for HEP computing. A distributed IPv6 test bed has been deployed and used for this purpose and we have been working closely with the HEP experiment collaborations and related IPv6 activities in the European Grid Infrastructure (EGI). This working group is also considering other operational issues such as the implications for security arising from a move to IPv6."[4]

As part of this test community, the University of Glasgow's Tier-2, Scotgrid, has begun investigations into how smaller networks, which do not control their IPv4 or IPv6 address allocations, may engage with IPv6 and the programme of work defined by the HEPIX working group.

1.1 Network Configuration for Scotgrid Glasgow

The deployment of the Tier-2 at Glasgow relies upon several core services from the University Campus network. These include:

- Access to the Joint Academic Network (JANET)
- DNS
- Firewalling and Intrusion Detection

The network connectivity to and from the test cluster at Glasgow is based upon a Layer 2 connection. This is physically supplied via a 1 Gigabit multimode fibre optic connection via an OM3 structured network to a core Cisco 6509 aggregation switch within the Campus. The 1 Gigabit switch port connected to the test cluster is set up as an access port only. All relevant IPv6 configuration for this install was configured on this 6509 and associated port. The wider

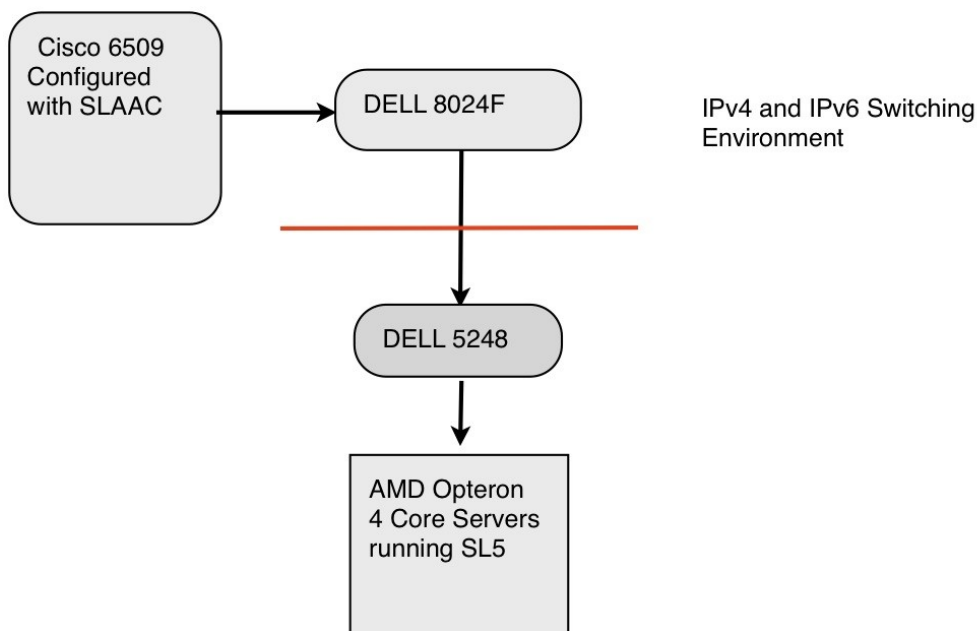
University addressing scheme is deployed across all of these core aggregation devices and core border routers.

While Scotgrid deploys its own DNS service internally for the Cluster it is fully dependent upon the University Campus DNS for external name resolution. New AAAA records for the IPv6 enabled elements of the test cluster were generated for this test. These records were moved to a test DNS environment specific to this deployment to ensure that no campus production services were affected by this initial installation.

Also new Access Lists for IPV6 were generated on the border devices within the Campus to block all ICMP traffic into the test cluster. Trace route functionality was also blocked to this IPV6 network to enhance the security of the test environment. These services are generally denied as standard procedure within the university network environment for experimental network services. These services are blocked utilising ACLs within the Campus network environment. Due to the security implications of publishing these configurations they have been omitted from this report.

1.2 Network Overview

The basic network topology for the IPv6 test cluster is shown below:



This test deployment was designed to emulate the hardware deployed within the GU Scotgrid production cluster in 2011. The module used to connect to the Campus network was a 1 Gigabit Long Reach SFP connected via multimode OM3 fibre into the 8024F.

The primary uplink on the Dell 8024F to the Cisco 6509 was configured with a 1 Gigabit interface and access VLAN which transported all external traffic into and out of the test environment. This external VLAN on the Dell 8024F was not configured with any IPv4 or v6 addressing and both broadcast and multicast Layer 2 traffic was allowed. All other physical connections other than the external link were based upon Cat 5e connections within the test cluster itself.

The external VLAN was also emulated on the Dell 5248 switch which acted as the primary aggregation point for the test cluster hosts. Additionally an internal VLAN for RFC 1918 [5] traffic was configured to emulate the configuration of servers within the production environment.

Again this VLAN was also configured without any static IP addressing and both Broadcast and multicast Layer 2 traffic was allowed on the links.

1.3 Test Cluster Overview

The systems installed within the test cluster for the initial connectivity test included five X AMD Opteron 4 core servers which were configured with Scientific Linux 5. The Network Interface Cards (NIC) were 2 X 1 Gigabit Ethernet on board interfaces.

The current Glasgow test cluster consists of approximately 30 hosts from various manufacturers, which support both Intel and AMD chipsets. The devices utilised in this primary test were 6 years old and were chosen due to the age of the network hardware in an attempt to identify any inherent issues with the Stateless Address Auto Configuration (SLAAC) [6] deployment being considered.

The configuration requirements for an IPv4/IPv6 with a Grid environment were as follows:

Two DNS entries per host examples are shown below:

IPv4

dev001-v4.gla.scotgrid.ac.uk 130.209.239.130

dev002-v4.gla.scotgrid.ac.uk 130.209.239.131

IPv6

dev001-v6.gla.scotgrid.ac.uk 2001:630:40:ef0:230:48ff:fe5a:4b7

dev002-v6.gla.scotgrid.ac.uk 2001:630:40:ef0:230:48ff:fe5a:2f

Two Grid certificates per host each specific to the IP instance being run. This decision was undertaken as part of the HEPIX IPv6 test deployment framework, which is partially standardised across multiple WLCG sites. For a full list of sites currently involved in this test framework please refer to w3.hepix.org/ipv6-bis. This deployment allowed for the

simplification of trialing IPv6 services be allowing both the IPv4 and IPv6 network stack instances to be accessed independently of one another. The Standard SL5 deployment for Worker Nodes based upon the Scotgrid Glasgow build model was deployed within the initial testing phase.

1.4 Deployment of IPv6

While multiple methodologies, such as static routing or OSPFv3 were considered for connecting the test cluster environment to the Campus, (SLAAC) was chosen for the following reasons:

The default IPv6 settings within Scientific Linux allow for Neighbour Discovery Protocol and Multicast

This allows for the propagation of IPv6 addressing and routing information to be supplied by the core device to the host.

The SLAAC protocol also maintains the IPv6 address lifetime and delivers Duplicate Address Detection (DAD) to the network which simplifies large scale deployments of addressing. The default lifetime expiry for the addressing under the default configuration is “never” and DAD is designed to identify and block duplicate IPv6 address clashes.

There is no administrative overhead to the System Administrator and the likelihood of manual mis-configuration of logical 128 bit IPv6 addressing is removed. However, the DNS entries for this deployment were entered manually.

Security issues have been reported under DHCPv6 and SLAAC, though the behaviours reported under SLAAC are considered to be more predictable.

This deployment (SLAAC) is considered to be presently the easiest deployment methodology available within the University Campus network architecture.

The main advantages of this deployment scenario will be discussed in the next section.

1.5 Potential Advantages of this deployment methodology

The main advantages of the SLAAC deployment at Scotgrid Glasgow was that the entire IPv6 network configuration was effectively handed off to the Campus Network and the deployment itself allowed for near instantaneous access into JANET's IPv6 mesh. An example traceroute taken from devices in the United States of America 30 minutes after the test deployment is shown below:

traceroute to 2001:630:40:ef0:230:48ff:fe5a:4b7 (2001:630:40:ef0:230:48ff:fe5a:4b7), 30 hops max, 40 byte packets

```

1 2001:1af8:4200:b000::1 (2001:1af8:4200:b000::1) 1.600 ms 1.813 ms 1.882 ms
2 2001:1af8:4100::5 (2001:1af8:4100::5) 1.320 ms 1.392 ms 1.465 ms
3 be11.crs.evo.leaseweb.net (2001:1af8::9) 2.587 ms 2.631 ms 2.619 ms

```

```

4 linx-gw1.ja.net (2001:7f8:4::312:1) 8.475 ms 8.466 ms 8.453 ms
5 ae1.lond-sbr4.ja.net (2001:630:0:10::151) 78.338 ms 78.388 ms 78.376 ms
6 2001:630:0:10::109 (2001:630:0:10::109) 9.900 ms 9.479 ms 9.446 ms
7 so-5-0-0.warr-sbr1.ja.net (2001:630:0:10::36) 13.320 ms 13.196 ms 13.317 ms
8 2001:630:0:10::296 (2001:630:0:10::296) 18.705 ms 18.542 ms 18.793 ms
9 clydenet.glas-sbr1.ja.net (2001:630:0:8044::206) 18.947 ms 18.931 ms 18.948 ms
10 2001:630:42:0:3e::9a (2001:630:42:0:3e::9a) 19.434 ms !X 18.214 ms !X 17.682 ms !X

```

This allowed for testing of a similar network to that used in the Scotgrid production environment. The lack of configuration changes to the switching fabric allows for the removal of potentially complex routing and switching commands. This also has the advantage of handing higher processor intensive tasks to larger core devices than the test cluster switches deployed.

A sample configuration for SLAAC on the Cisco 6509 is shown below:

```

IPv6 dhcp pool IPV6_DHCPPOOL
    dns-server 2001:630:40::1
    domain-name v6-scotgrid.gla.ac.uk
!
interface Ethernet0/0
    IPv6 address 2001:630:40::1/64
    IPv6 enable
    IPv6 nd other-config-flag
    IPv6 dhcp server IPV6_DHCPPOOL

```

A sample configuration of the access switching configuration:

```

vlan database
vlan 2-6,1234
exit
snmp-server location "Room243d,Kelvin_Building,Glasgow"
snmp-server contact "uki-scotgrid-glasgow@physics.gla.ac.uk"
hostname "Scotgrid_10_GigTest_Switch"
snmp unicast client enable
ip address 172.20.239.11 255.255.255.128
ip address vlan 2
no logging console
logging file debug
logging buffered debug
interface vlan 2
name "external"
exit

!
interface ethernet 1/xg23
description 'Link_to_Kelvin_1_Port_320_1_Gig'
switchport access vlan 2
exit
!

```

```
interface ethernet 1/xg24
description 'Trunk_Link_to_Dell_5224_Cthulu_Rack'
switchport mode trunk
switchport trunk allowed vlan add 1-5
exit
!
```

```
Scotgrid_10_GigTest_Switch#
```

An additional benefit of this technology is that the CPU load on the switches and hosts tested remained constant with those seen in the initial IPv4 deployment on the test cluster. While this IP addressing deployment has not been tested with production traffic, presently there appears to be no issue with the IPv6 stack in SL5 causing adverse CPU activity.

1.6 Potential Disadvantages of this deployment methodology

One issue with the Glasgow Scotgrid is that the deployment of SLAAC is in a very controlled and regulated network and multiple security methods are used across both the Campus network and the test cluster, including Campus switching ACLs, line rate firewalling and intrusion detection and OSSEC on the test cluster. While SLAAC is a useful delivery mechanism for IPv6 into a Tier-2 network and cluster imbedded in a University Campus, there are several reported exploits [7] on multiple platforms which may make it unsuitable for larger shared networks. While this methodology allowed for the deployment of IPv6 in a controlled manner to the Scotgrid test cluster it may not deliver a scalable solution for other Tier-2's which are located within University network environments.

Also, SLAAC's ease of use could potentially lead to a scenario where all hosts in a Cluster are automatically configured as IPv6 hosts under the standard SL5 install. Again this may cause software interoperability and stability issues and network failures depending on the deployment configuration within the cluster [8].

1.7 Security implications

The implementation of the IPv6 specific ACLs for this deployment were not developed by GU Scotgrid staff and were developed to meet the security policies deployed by the University of Glasgow. Presently there has been no investigation of how the security software such as ossec interoperates within an IPv6 environment.

Due to the limited global deployments of IPv6 there has been limited work conducted in the field. Currently, it is difficult to predict the exact threat vectors which will be used to compromise IPv6 environments. This does not mean that a strong security model based upon logging, reporting and continuous monitoring of an IPv6 environment should not be implemented as standard.

1.8 Implications for Grid services

Presently the deployment of IPv6 for Grid Services is at best experimental in nature. The HEPIX IPv6 group has been leading these investigations for over 12 months presently and have uncovered several issues with multiple Grid Services including the configurations for Gridftp and Glite 3.1 UIs. The resolution to these issues and more detailed information surrounding the requirements for service deployment for Grid services can be found at :

<https://w3.hepik.org/ipv6-bis/doku.php?id=ipv6:testbed>

Due to the dynamic environment developed around Grid Middleware to meet the requirements of various LHC experiments and non High Energy Physics User communities the majority of TCP/IP stack development work has focussed upon IPv4 implementations as these are the most common global network deployments. The issues encountered by the HEPIX group indicates that more indepth research into all Grid services is required to ascertain how many globally deployed services are compliant with the IPv6.

1.9 Discussion

The deployment of IPv6 is still in its infancy worldwide and many issues surrounding security, backwards compatibility with IPv4 and the adoption of the protocol have still to be tackled. The limited tests at Glasgow indicate that the utilisation of SLAAC is potentially suitable for LHC based clusters such as Scotgrid Glasgow. As an addressing mechanism SLAAC does allow Cluster administrators an easy migration path to IPv6 without major configuration changes to the underlying systems. However, there is a requirement on the Campus networks to develop and deploy IPv6.

The next phase of the Glasgow testing will be to engage in a limited pilot with standard grid services configured specifically for testing WLCG applications and also a further investigation in the actual impact on switching and routing equipment under production loads running IPv6. [8]

The potential for IPv6 in collaborations such as the WLCG is vast as the protocol delivers greater functionality than its progenitor including built in security mechanisms and better control of multicast and broadcast traffic than is presently possible. It will also allow for an underlying addressing framework for greater collaboration between established research networks and programmes and emerging ones.

References

- [1] D Kelsey, *From IPv4 to eternity - the High Energy Physics transition to IPv6*, EGI User Community Conference, Munich 2012.
- [2] S Deering, *Internet Protocol Version 6 Specification*, <http://tools.ietf.org/html/rfc2460>.
- [3] <http://www.ripe.net/internet-coordination/news/industry-developments/global-IPv6-deployment-monitoring-survey-2011>
- [4] D Kelsey, *From IPv4 to eternity - the High Energy Physics transition to IPv6*, EGI User Community Conference, Munich 2012.
- [5] Rekhter et al, *Address Allocations for Private Internets*, <http://tools.ietf.org/html/rfc1918>
- [6] S Thomson, *IPv6 Stateless Autoconfiguration*, <http://tools.ietf.org/html/rfc4862>
- [7] E Carter, *ICMP and Security in IPv6*, <http://blogs.cisco.com/security/icmp-and-security-in-IPv6/>
- [8] J Arkko et al, *Experiences from and IPv6 only Network*, <http://tools.ietf.org/html/rfc6586>