Title: **Deliverable D2.1.2** **Report of 2nd Deployment Case Study**	**Document Version:** 1.2

Project Number: 223794	**Project Acronym:** 6DEPLOY	**Project Title:** IPv6 Deployment Support

Contractual Delivery Date: 31/03/2009	**Actual Delivery Date:** 26/08/2009	**Deliverable Type* - Security**:** R – PU

* Type:	P – Prototype, R – Report, D – Demonstrator, O – Other
** Security Class:	PU- Public, PP – Restricted to other programme participants (including the Commission Services), RE – Restricted to a group defined by the consortium (including the Commission Services), CO – Confidential, only for members of the consortium (including the Commission Services)

Responsible and Editor/Author: Tim Chown	**Organization:** University of Southampton	**Contributing WP:** WP 2

Authors (organizations):

James Morse (University of Southampton), with review from Carlos Friacas (FFCN) and Janos Mohacsi (NIIF).

Abstract:

This report describes the IPv6 deployment within the School of Electronics and Computer Science (ECS) at the University of Southampton. It includes a description of the process of IPv6 deployment, including network, systems and applications aspects. The deployment is currently live and spans a network of up to 3,700 hosts and over 2,000 users.

Keywords:

IPv6, enterprise deployment, campus, case study

Disclaimer

The 6DEPLOY project number 223794 is co-funded by the European Commission under Framework Programme 7. This document contains material, which is the copyright of certain 6DEPLOY beneficiaries and the EC, and may not be reproduced or copied without permission. The information herein does not necessarily express the opinion of the EC. The EC is not responsible for any use that might be made of data appearing herein. The 6DEPLOY beneficiaries do not warrant that the information contained herein is capable of use, or that use of the information is free from risk, and accept no liability for loss or damage suffered by any person using this information.

Executive Summary

This report documents the experiences of deploying IPv6 at the School of Electronics and Computer Science (ECS) at the University of Southampton, UK. It covers aspects of deployment for a network spanning roughly 3,700 hosts (of which 2,600 are wireless devices) within a single campus area.

The report focuses on a number of specific areas, including:

- Rationale for IPv6 deployment;
- Steps in planning IPv6 deployment;
- IPv4 integration;
- Technical platforms and solutions;
- New IPv6-enabled services;
- Lessons learned and open issues.

Our basis for IPv6 deployment has been to deploy IPv6 services dual-stack, i.e. where IPv6 capability is available it is enabled on the same existing infrastructure that supports IPv4. We first enabled IPv6 ‘on the wire’ in our network infrastructure, after which we were able to add IPv6 capability to software services as and when that capability became available.

This approach allows our software, systems and networks to use either protocol for internal or external IP communications. The benefit of this approach is to facilitate the use of IPv6 where possible, without the need for IP ‘translation’ devices. The cost of this approach is that both IP protocols need to be supported and managed.

By introducing IPv6 capability to all our core services, we in principle allow IPv6-only devices to operate within our network environment. We did not choose to attempt an IPv6-only deployment within ECS because we already had a significant IPv4 infrastructure, and sufficient global IPv4 addresses available (the campus has a /16 IPv4 prefix, allocated prior to 1994). We felt that operating dual-stack was a better strategy than attempting IPv6-only operation internally and using ‘translation’ tools to communicate with external IPv4 services.

IPv6 has now been running in production within ECS for several years. During this time

the available IPv6 services have been expanded, and the IPv6 capabilities improved. It is worth noting that this deployment has not adversely affected our IPv4 deployment, and that we believe that while some IPv6 capability is still missing from certain applications and services, IPv6 is mature enough to be deployed in production campus environments provided those ‘limitations’ are understood.

Table of Contents

Figure Index

Introduction

IPv6 has, in its most basic form, been defined by IETF standards since 1998 [RFC2460]. At the time of writing, it has now been over 10 years since the initial standard was agreed for the protocol's operation. In those 10 years, an amount of progress towards full deployment has been made, but, as a percentage of total traffic crossing Internet backbone routers today, IPv6 traffic remains low. Support for the protocol in OS and router stacks has improved to the point where all major vendors now support IPv6, many shipping with IPv6 enabled by default.

In academic networks, IPv6 is deployed through the backbone networks of the national research networks (NRENs) and the GÉANT network that interconnects them. Many other NRENs worldwide also support IPv6 natively, including those in the US (Internet 2) and Japan. All these networks have deployed IPv6 dual-stack, i.e. running IPv4 and IPv6 and associated routing protocols on the same infrastructure. For most users on those networks the application performance they will experience - whether their applications use IPv4 or IPv6 - will be very similar.

The main challenge for IPv6 deployment in academic networks lies in deployment at the edge, i.e. in the campuses. While the research backbones can carry IPv6 traffic natively, it is down to the edge sites to deploy the protocol and use it, so generating increased IPv6 traffic. To date, deployment in campuses has been limited, perhaps in part to most long-established universities having 'ample' global IPv4 address space. A university that gained Internet connectivity before the mid 1990's probably has a /16 IPv4 prefix, which offers over 65,000 unique IPv4 global addresses for it to use. It is newer universities, or those who were connected using IPv4 NAT, that are likely to have a shortage of global IPv4 addresses in the (relatively) near future.

Pressure on the IPv4 address space is growing. Currently 28 of the /8 blocks remain, meaning there is only around 10% of the available address space available. The best reference [POTAROO] suggests exhaustion of the IPv4 address space by July 2011, with regional registries running out by April 2012. Once this happens, and no new IPv4 global address space is available, the only way to obtain it will be by 'trading' with other holders. It is thus very important that universities and higher education sites consider IPv6 deployment sooner rather than later. In the next section we expand on the reasons for IPv6 deployment at university campuses.

1. RATIONALE FOR IPV6

There are many reasons for a university or higher education site/campus to consider deploying IPv6.

1.1 Internet and Application Architecture

The original premise of the Internet architecture was ‘end to end’. While the precise meaning of this phrase is often over-discussed, an implicit part of the concept is that hosts should have globally unique addresses (which currently also double up as both ‘locators’ and ‘identifiers’) with which they communicate on the network.

Over time, more and more networks have been deployed that use private IP address space [RFC1918], usually in conjunction with Network Address Translation [NAT]. As a result, the IP address a device uses for communications is often not the address the remote peer sees, and potentially vice versa. Addresses are manipulated by the NAT middleboxes, or by application layer gateways (ALGs).

Use of unique global addresses simplifies network management and support, while also making an application developer’s task much easier. Network managers do not have to contend with potentially ambiguous private IP address space and also have an easier task when accounting for network activity, while applications can be developed without the need for NAT traversal methods to be considered.

If sites wish to have global IP address space available for their IP infrastructure, systems and hosts, the opportunity to secure enough IPv4 address space to do so is fast receding. The predictions cited above indicate that no new IPv4 global address space will be available for allocation from IANA to the Regional Internet Registries after 2011-12, and that assumes consumption continues at the current rate and doesn’t accelerate as supply approaches exhaustion. Once there is no new address space being released to the RIRs, IPv4 address space will need to be secured from existing resources (trading or reclamation), which is highly unpredictable.

While many universities have enough IP address space for the immediate future, the introduction of expanded demands on IT infrastructure may raise the requirement, e.g. increased networking in student halls/dormitories, growth in WLAN-enabled devices (PDAs, iPhones, etc) and increased used of IP devices in campus infrastructure (security, building control, etc). While part of this demand could be met by the use of private IP space, it is prudent to ensure that a site has enough global IP address space for all its

IP-enabled devices for the future.

If a site already has enough IPv4 global address space for its current needs, it can deploy IPv6 dual-stack such that both global IPv4 and IPv6 addresses are in use. Otherwise if a site is already using IPv4 NAT it can use NATed IPv4 dual-stack with global IPv6 addresses. As more applications are able to take advantage of IPv6, the limitations of using NAT can then be avoided by use of IPv6 instead.

While NAT is often ‘marketed’ as offering additional security benefits, there is no reason why these benefits cannot be achieved using IPv6 with global addresses, as described in [RFC4864].

1.2 Supporting Teaching and Research

At Southampton, the initial IPv6 deployment has been in the School of Electronics and Computer Science (ECS). It is quite probable that the ‘computer science’ departments of universities will be the first places that see IPv6 being introduced. The reason for this is that computer science departments tend to be on the leading edge of technology for a campus, researching and deploying new systems and protocols ahead of their wider adoption elsewhere. In that light, it is natural for a computer science department to deploy IPv6 to assist in teaching (both directly in computer network courses but also indirectly via student projects, PhD work, etc) and for the benefit of research projects.

Southampton’s deployment of IPv6 has also led to IPv6 being used in student-developed services and applications that would otherwise not have happened, e.g. an IP TV service using IPv6 Multicast (ECS TV) and in the student-run Southampton Open Wireless Network (SOWN), which provides network outreach to students in the local area and in outdoor areas around campus. These examples are discussed in more detail in a later section.

1.3 Early IPv6 Experience

IPv6 will become a very important part of the Internet in the longer term. Exactly when IPv6 will ‘take off’ is still hard to say, but the pressure on available global IPv4 address space suggests its time is getting much nearer. The alternative is a network world with much more fragmented use of IPv4 address space and multiple layers of IPv4 NAT.

Given that widespread IPv6 deployment will happen, it’s very useful for IT staff to understand its operation and its implications as soon as possible. Delaying that exposure will simply lead to a bigger step up in knowledge being required in the future.

Understanding IPv6 now allows its impact on the overall campus infrastructure to be better understood, so that all aspects of planning and resourcing can be in sync at an early stage. This has an impact on procurements (requiring ‘IPv6 capable’ systems and software) and on management, operations and policy setting. Even if IPv6 is not put into production immediately, being prepared is important. The cost of deploying more rapidly later is likely to be greater, particularly if re-procurements are required.

It is not hard to establish an IPv6 testbed independently of a production IPv4 network, and to source appropriate training for staff. Piloting activities are discussed in more detail in the next Section.

1.4 Outreach

There are emerging networks, especially in Asia, that are deploying IPv6-only environments. While doing this in non-greenfield environments would probably be a little premature, it is important as an educational institution to consider outreach to all parts of the world, from where people may be seeking access to your online material, or information about studying at your institution.

In that light, deploying IPv6 is about enabling additional ways for people to access your Internet presence and material. Of course, at the current time new IPv6-only networks will need to establish their own methods for access existing IPv4-only information, but in the longer term we feel being ‘present’ on the Internet via IPv6 natively is important to us.

In the context of SOWN, we make IPv6 available in student households, such that students can potentially have the capability to deploy more feature-rich networks in their homes.

1.5 Security Implications

IPv6 capability is included in most if not all major OS and router platforms now. It is also usually on/enabled by default. As a result it is quite possible, even if IPv6 is not administratively configured in a network, that applications or services are using it. Administrators should have the tools to determine if IPv6 traffic is present in their network, even where they know that routers and supported services are in principle IPv4-only.

If IPv6 traffic is flowing in a network, this may potentially span filters/ACLs intended to block IPv4-related traffic. However, a site may have IPv6-in-IPv4 tunnels in use, potentially unknown to the site administrators. The ISATAP [RFC4214] protocol allows automatic IPv6-in-IPv4 tunneling within a site; this may also potentially cause IPv4 security measures to be bypassed. The IPv6 tunnel broker [RFC3053] offers IPv6 connectivity to a dual-stack host in an IPv4-only environment via a remote tunnel server. Tunnels should generally not be permitted unless for specific known purposes, particularly when spanning a firewall.

IPv6 may also – perhaps more stealthily – be encapsulated in UDP using the Teredo [RFC4380] protocol (as supported in Microsoft Windows). Teredo serves a similar purpose to a tunnel broker, but has the advantage of working well behind most IPv4 NATs. The use of Teredo might be detected by looking for traffic to the default Teredo server port externally (UDP/3544). While of course other user applications may also tunnel unknown traffic in UDP (e.g. to traverse a stateful firewall or NAT device) it is not ideal to have uncontrolled IPv6 traffic leaving/entering a site.

It is also possible that a host could (accidentally via Windows ICS or deliberately if malicious) send out IPv6 Router Advertisements in a network where other hosts have IPv6 enabled, which would potentially cause those other hosts to route IPv6 traffic to the ‘offending’ host. This may cause unpredictable behaviour in hosts. Using tools to monitor for such traffic (e.g. [RAMOND]) is desirable.

If a site is not using IPv6 locally, it is probably wise to consider disabling or turning off IPv6 in hosts where administrative control to do so exists, and to monitor for use of ‘transition’ aids such as Teredo, ISATAP and Protocol 41 (IP in IP) tunnels.

Obviously the thrust of this document is that campus sites should deploy IPv6, and in doing so be in control of the use of IPv6 in their networks. If IPv6 is deployed natively (dual-stack), the need for Teredo, ISATAP or Protocol 41 tunnels is removed.

2. PLANNING

In this Section we discuss the stages of planning involved in deploying IPv6 in a campus environment. While we only consider the deployment for a School (department) network, since ECS runs all its own infrastructure (routing, DNS, mail, etc) the process is also applicable at a wider campus scale.

The main difference if deploying at a campus scale would be the choice of internal routing protocols (OSPFv3 or IS-IS) in the planning – at a campus scale such choices may have different drivers to those on a smaller scale. This case study does not include any significant discussion of local routing, as our service is not large enough to warrant it. Nevertheless the ECS experience traverses some 2,000 or more users with hosts in approximately 20 IPv4 subnets, using over 50 network edge devices and delivering service to up to 3,700 hosts (of which 2,600 are wireless devices registered in the past three years).

2.1 Phased planning

Although ECS deployed its IPv6 infrastructure over a long period of time in multiple steps (the initial ‘push’ being between 1997 and 2002, with further work in the 6NET project [6NET]), we can look back over the process and describe how we would have achieved the same thing were we starting out today. Our original deployment took a long time because many of the required components were not ready in our initial deployment period – these however are now available as standard, in particular the very solid router and host OS support.

We believe that there are three main phases to the planning, and a post-deployment ongoing ‘monitoring’ phase. There is some flexibility to consider certain aspects in different phases, but we recommend the approach described here.

2.2 Advanced Planning

The advanced planning phase could begin as soon as a campus has made a decision that it wishes to begin the process of moving towards IPv6 deployment. This phase does not commit the campus to actual deployment, but it begins the preparation process.

The question of when to deploy has no exact answer. The scenario will vary from campus to campus, or site to site. All that can be said at the moment is that IPv6 will be needed ‘soon’ and that it’s prudent to be ready for more widespread uptake of IPv6 now rather than leaving it for later and then finding the costs and timescales become significantly more expensive/rushed.

The points to address in this phase are:

- Have key staff take appropriate training to understand basic IPv6 principles, including the ability to understand what is meant by ‘IPv6 capable’ for a given system or service. Management may require a broader, high-level training while administrative staff may require technical training.

- Survey systems, applications and services for IPv6 capability. Identify gaps in the capability and begin planning for bridging those gaps. For example, it may be necessary to upgrade to a newer version of a software package to obtain required functionality, or perhaps switch to a different package or solution altogether.

- Ensure that all future tenders for procurement include IPv6 capability for the system or service in question. This is most important for network infrastructure, service software and all supported applications.

- Talk to your ISP (for a UK campus invariably a NREN or an interconnecting regional network, but for some other countries it may be a commercial provider) to discuss IPv6 connectivity options. This may be arranged natively or via some tunnel service. At this stage also request an IPv6 address/prefix allocation, which at this stage should by default be /48 in size. The allocation should also include information on forward and reverse IPv6 DNS delegation.

- Discuss and document IPv6-related policies, which may include security policies. It should be a goal to ensure IPv6 security is as strong as IPv4 security, and that deploying IPv6 does not compromise IPv4 security.

These initial steps should help ensure that the task of IPv6 deployment is better understood, and that the readiness of various components in the local infrastructure has been assessed. Future procurements should ensure IPv6 support is built into products where required.

Note that a key issue here is in understanding what is meant by ‘IPv6 capable’ systems. This can in part be derived from studying standards, and following available guidance, but getting appropriate training, and talking to sites that have already deployed can be very useful at this stage.

2.3 Testbed/Pilot Operation

In this phase the site undertakes an IPv6 pilot deployment in a limited scale, to gain and improve hands-on experience in the use of IPv6 on a day-to-day basis. This would typically be a restricted deployment, possibly as small as a single subnet testbed.

The action points in this phase include:

- Formulating an initial IPv6 address plan for the site. It is most likely that the IPv4 and IPv6 subnets will at least initially be congruent, i.e. hosts will lie in the same IPv4 and IPv6 subnets. We discuss addressing in more detail later in this document. At this stage, a prefix should be reserved from the /48 allocation for test purposes.

- Acquire an appropriate IPv6-capable router and determine the testbed topology. While the longer-term plan may be production dual-stack deployment with an edge router handling both IPv4 and IPv6, for the initial testbed a tunneled connection upstream should suffice. If the testbed is to be dual-stack, it will also need IPv4 connectivity anyway. One option is to have a single subnet that is dual-stack, with separate routers providing IPv4 and IPv6 connectivity (e.g. by providing IPv6 connectivity into an existing IPv4 DMZ area, which could be a cut-down version of the DMZ architecture of Figure 1).

- Establish IPv6 connectivity upstream and test.

- Decide which services you wish to deploy on the testbed. A minimum deployment might simply be an IPv6 web server (e.g. Apache), but other services that deserve initial attention include DNS (e.g. BIND), mail (e.g. sendmail) and (ssh) login/file transfer services. If the tesbed does not include DNS, then you will need to use your regular DNS (IPv4) servers to carry IPv6 DNS (AAAA) records and assume connecting hosts can query DNS via IPv4 also (i.e. have a dual-stack resolver).

- Ensure IPv6 is enabled/configured appropriately on the testbed systems. At this stage servers will most likely have manually configured IPv6 addresses and any clients would use IPv6 stateless autoconfiguration (i.e. no DHCPv6).

- Configure appropriate filters/ACLs on the access router (or deploy an appropriate filtering device in series) and enable connectivity towards the internal subnet(s).

- Deploy appropriate monitoring systems (e.g. mrtg to track IPv6 traffic on the router) and undertake whatever tests you choose.

This process should boost confidence in a site's ability to deploy IPv6, and help a site understand possible issues to consider and resolve in a fuller deployment. Ideally the testbed should include all the services that are planned for the initial production deployment, which implies you need to think ahead a little to that phase.

Another option for a site to experiment with IPv6 connectivity is to use an IPv6 Tunnel Broker. There are a number of commercially supported and NREN-supported brokers available.

2.4 Production Deployment

Once the advanced planning is completed, and appropriate experience in the use of IPv6 has been gained, a site can consider a production deployment. A general strategy can be to enable IPv6 'on the wire' (in the routing infrastructure) first, and then enable IPv6 for selected services as desired. A 'big bang' approach to enabling application services is not required.

- Plan which parts of your network and which subnets are to be IPv6-enabled, i.e. made dual-stack. This may be only server subnets, or it might be specific elements of a network (e.g. some or all of a DMZ), and/or perhaps the local wireless network.

- Ensure your IPv6 addressing plan is completed.

- Confirm how production IPv6 capability will be delivered. The best approach is a dual-stack edge router handling both IPv4 and IPv6, but it could be separate connection points for IPv4 and IPv6, possibly with separate firewall devices for each protocol. Regardless of the connectivity method, the internal network and subnets can be fully dual-stack (and not delivered on separate infrastructure).

- Ensure your IPv6 connection point is secure, and has appropriate IPv6 filtering capability (integrated to the router or via a separate firewall).

- Ensure your management and monitoring tools can handle/process IPv6 where they need to do so.

- Enable IPv6 on the wire, by configuring local IPv6 routing, and IPv6 Route Advertisements on the desired subnets.

- Enable IPv6 in your chosen production services, e.g. web, mail, DNS, remote logins.

- Add IPv6 addresses (AAAA records) for appropriate systems in your DNS infrastructure, and configure the DNS to respond over IPv4 or IPv6 transport.
- Ensure your IPv6 security model is at least as strong as your IPv4 model. Include IPv6 in all security assessments/tests.

At this point your network should have the capability to support IPv6 operation via dual-stack networking. Each host on a dual-stack part of the network can communicate with other IPv4-only, IPv6-only or dual-stack systems via either protocol. Devices on the local network should be able to operate IPv6-only for basic services, should they need to do so.

2.5 Ongoing Review

Having deployed, there is a natural process to follow of ongoing review to inform future planning and enhancement of the deployment, by identifying remaining ‘gaps’ in the deployment, and opportunities to deploy new IPv6 services to enhance the campus network environment (e.g. exploring services such as MS DirectAccess, Mobile IPv6 or IPv6 multicast).

3. IPv4 INTEGRATION

One of the most difficult aspects of an IPv6 deployment is integrating IPv4 and IPv6 services. Documents often talk of a 'transition' to IPv6, where in reality – at least for the immediate future – the challenge is deploying IPv6 alongside your existing IPv4 infrastructure and providing a consistent user experience via either protocol.

3.1 Approaches and Requirements

The requirements for IPv4 integration can be viewed from a number of perspectives. These will depend on your underlying strategy. The current 'best practice' advice is to deploy IPv4 and IPv6 dual-stack, running both protocols on all systems and networks where possible, leaving the choice of which protocol to use down to the application (which will have that choice based upon the information in the DNS response it gets for a given named host it wishes to communicate with).

One requirement is that all the systems in a site can communicate with any given destination, whether it is IPv4-only, IPv6-only, or dual-stack. If any given system is made dual-stack then, routing/connectivity permitting, it can do so.

This means the challenge is how to integrate new IPv6 capability into the existing infrastructure. While at a strategic level the systems and network operations may be dual-stack there are new features and modes of operation in IPv6 that need consideration – not everything has a direct IPv4 counterpart. Identifying and understanding these can be important.

Also, operating dual-stack effectively means a campus is managing two networks, albeit on the same infrastructure. There will be some extra complexity and cost in doing this. At present, however, that cost is less than trying to deploy IPv6-only, migrating all internal services to IPv6, and establishing 'translation' services at the edge of your network to talk to existing external IPv4 services.

This dual-stack approach of course only works for global addressing on both protocols while there are still available IPv4 addresses to do so. Based on well-respected predictions [POTAROO] it is likely that new IPv4 global address space will be difficult to acquire from 2012 onwards. At this point the only option, should a site wish to run dual-stack, is to use private IPv4 addresses and NAT alongside global IPv6 addresses. So the network is still dual-stack, but only globally addressable via IPv6.

A typical campus network today probably has enough global IPv4 addresses for existing systems, but future network growth may change that position. Thus at some point the site needs to consider migrating services to IPv6, or perhaps more realistically in the medium term making them available via both protocols. This implies some level of application porting may be required, or that at least updated versions of applications are required from vendors.

Another requirement for IPv6 deployment is that deploying IPv6 should not reduce the performance or reliability of the existing IPv4 infrastructure. From a user's perspective there should be no perceivable difference. In the earlier days of IPv6 deployment, where in particular many IPv6-in-IPv4 tunnels were used, routers that implemented encapsulation in software often caused performance problems when non-negligible IPv6 traffic was being processed. For this reason many early deployments used a parallel IPv6 infrastructure – however, current host and router platforms generally no longer have these issues so a dual-stack deployment using common infrastructure is generally considered both the most cost effective and robust solution.

A third important requirement is that IPv6 deployment should not compromise IPv4 security, or rather not introduce new, unmitigated risks. Some early IPv6 deployments had limited external filters/firewalls, allowing possible exposure of systems over IPv6 where IPv4 paths were filtered. Some connectivity tunnels also completely bypassed existing firewalls. A site also has to understand potential new vulnerabilities, as well as the IPv6 equivalent to existing IPv4 issues (e.g. with IPv6 Neighbour Discovery as opposed to IPv4 ARP).

3.2 Use of Transition Tools

Despite 'transition' not being the best name for the introduction of IPv6, a number of so-called 'transition tools' exist to assist sites, users or applications to gain IPv6 connectivity when required.

The fundamental assumption when deploying dual-stack is that there is no requirement to use such tools internally, because all systems can communicate via either protocol between any two points on the network. However there may be a small number of scenarios where such a requirement does arise, e.g.:

1. *A site is only partially dual-stack, and a host with IPv6 enabled is in an IPv4-only part of the network and wants to access an IPv6-only service (perhaps some new application).* In such cases the site can seek to expand the dual-stack network coverage to where the host resides, it can seek to retrospectively add

IPv4 capability to the IPv6-only service, or it could deploy some form of automatic intra-site tunneling protocol to enable connectivity, e.g. in this scenario ISATAP [RFC4214] might be well-suited (or perhaps 6rd [6RD] may be appropriate, though as yet it has not been widely tested). The preference here should be expanding the dual-stack coverage.

2. *An IPv6-only device in the dual-stack network wants to connect to an IPv4-only service, e.g. to print to an IPv4-only printer.* In this case the site can seek to add IPv4 capability to the device, attempt to add IPv6 capability to the service, or deploy some form of 'translation' service to facilitate communication (e.g. at the application layer a dual-stack print server, or at the network layer a more crude translation method (e.g. [NAT64], which is replacing [NAT-PT]). The preference here should be to acquire a dual-stack printer (e.g. entry-level network printers from at least one major manufacturer now have IPv6 support).

It can be argued that, in general, a strategy to avoid the use of new IPv6-only devices may make network operation 'simpler' in the short term, but is not helpful in a site progressing to the end game of predominantly IPv6 operation. Where possible IPv6-capability should be added to existing systems, which should minimize the requirement to use specific 'transition tools'.

Where dual-stack nodes exist in an IPv4-only part of the site infrastructure, ISATAP may provide a solution, offering automatic IPv6-in-IPv4 tunneling from the host to a configured ISATAP router. While this has its attractions, it should be noted that creating virtual IPv6 subnets that potentially span many IPv4 subnets will add to management complexity. The preferred approach is to enable IPv6 on the wire where required.

3.3 Supporting Remote IPv6 Access

A campus deploying IPv6 services may wish to consider whether and - if so - how it gives assistance to its users to access services via IPv6 when those users are off-site and connected to IPv4-only networks.

In general, the campus can continue to provide IPv4-only versions of the services that users require for the foreseeable future. However there may be some users (e.g. students or researchers) that specifically need IPv6 access for academic reasons. In the longer term some applications or services will emerge that are IPv6-only (though the timeline for that cannot be predicted as yet).

There are two general approaches to this access problem:

1. *Leave it in the hands of the user and their ISP*. The user may be able to utilize Teredo [RFC4380] or 6to4 [RFC3056] to gain 'automatic' IPv6 connectivity via the general IPv4 Internet. How well this may or may not work is not predictable. Or they may choose to use a 'public' IPv6 tunnel broker [RFC3053] like xs4all, which may have performance issues if the first hop is remote (e.g. a popular tunnel broker is based in Canada).

2. *Provide IPv6 access services at the campus to assist remote users*. Here the site may choose to run an appropriate method for its users, e.g. an IPv6 tunnel broker or possibly 6to4 (where your users' 6to4 routers are configured to use a 6to4 relay at the edge of the ECS network). If it does so, it should probably restrict usage to just its own users. In the ECS case, we deployed a production IPv6 tunnel broker both for our own users and those of the JANET community [BROKER]. Similar steps have been taken by other NRENs including Renater.

3. *Provide VPN service at the campus with dual-stack IPv4 and IPv6 support*. A site may choose to support VPN access to restricted services, and enable a dual-stack service via the protected channels. This type of solution can be deployed using PPTP, L2TP, IPsec or OpenVPN (in ECS we tested the latter approach).

In the longer run, we can expect ISPs to deploy IPv6 to their customers, and so this concern will be reduced and eventually disappear. In the medium term, sites should consider whether their remote users specifically require IPv6, and if they do how they can best give some level of assured access.

3.4 IPv6-only Deployment?

The dual-stack approach is considered best practice for the short to medium term. However at some point IPv6 deployment will grow and, eventually, IPv6 will become the dominant IP protocol on the Internet. It's impossible to predict exactly when IPv6 will 'take off' in a significant way (traffic-wise), but it is reasonable to assume that it's inevitable in the long term.

Similarly, at some point campus sites may consider deploying IPv6-only networks without supporting IPv4 (with or without NAT), because they consider that to be the most cost-effective way to run their network. There are some important considerations for when this strategy becomes practical, e.g.:

- Are all the systems and services in your network capable of running IPv6-only?

- Do you have legacy equipment/applications that are IPv4-only, that cannot be replaced or ported to support IPv6?

- Are there appropriate tools to use within or at the edge of your network to facilitate communication between IPv6-only and IPv4-only systems or services?

The last point is perhaps of greatest concern. The implication is that a site will need to run a pretty significant ‘translation’ service to realise this – some combination of [NAT64] with perhaps some application-specific ALGs. The question is whether the cost of doing so is more attractive than supporting both versions of IP within the local network.

As yet, experiments in NAT64 and its variants are in their early stages – how these mature is as yet unknown. Hence, currently the dual-stack approach is favoured, unless perhaps your site can rely purely on specific ALGs for external IPv4 Internet access.

4. TECHNICAL PLATFORMS AND SOLUTIONS

Having discussed the reasons for IPv6 deployment in the ECS context, and the considerations on how a campus might approach the task of deploying, we now look at the technical specifics of an IPv6 deployment.

4.1 The ECS Scenario

The ECS network includes up to 3,700 systems or active IP devices and has a user base of over 2,000 users. The network spans four buildings on campus, and includes approximately 20 IPv4 subnets (of various sizes). The network infrastructure is predominantly Cisco, with a 6509 for core routing and edge stacks of Cisco 3750's (totaling around 50 such devices).

ECS runs all its own network services, including DNS and DHCP for IPv4, supports various wireless access methods, and has an IPv4 firewall (Checkpoint platform) between it and the rest of the campus network. ECS has its own mail and web servers, and also runs some corporate applications (almost exclusively web-based) and a variety of commercial and open source services. Thus the ECS network is in essence a 'mini campus' network in its own right, and forms an excellent sandbox for developing and deploying new systems and services, including IPv6.

ECS's upstream connectivity is to the Southampton campus and from there via the LeNSE regional academic network to JANET (the UK academic network) and thence the wider Internet. Both LeNSE and JANET support IPv6 natively, dual-stack.

With the upstreams supporting dual-stack, and an established IPv4 network, the natural approach in ECS was to introduce IPv6 in dual-stack mode alongside IPv4. This approach maintains compatibility with systems using either protocol, while also enabling new IPv6 applications and services (as discussed in a later section).

The ECS deployment evolved over a longer period of time (starting in 1997 with an IPv6 64kbit/s X.21 link to UUNet in London). As the deployment has grown, a number of studies have been undertaken, with resulting documentation, e.g. the JANET Bermuda project, 6NET [6NET], the JANET IPv6 technical guide [JANETv6] as well as various training materials [TRAIN, 6DISS, 6DEPLOY].

Were the deployment one from a fresh start, we would follow the phased process as described previously in this document, i.e. preparation, deployment on the wire with associated security provision, and finally enabling services as and when ready.

4.2 Address Planning

There are three main aspects to IPv6 address planning; firstly obtaining an IPv6 prefix to use on your network, secondly determining how to allocate your address space within your network, and finally deciding how hosts and end systems will obtain their address configuration.

4.2.1 Allocation

In ECS' case, obtaining a network prefix for IPv6 was easy. We contacted JANET Customer Services (as it is now) and made a formal request. We were then allocated a /48 for University of Southampton use, specifically 2001:630:d0::/48.

In some NRENs, the regional networks offer address space to universities, but in the UK all sites go directly to JANET. There may be some advantages to considering using regional network assigned prefixes if universities are multihomed via those networks, but in the UK JANET is the single IP provider. At the time of writing JANET is approaching 100 IPv6 /48 allocations.

Having obtained a /48 for University use, we discussed the allocation with our campus computing service organization (now iSolutions) and agreed that ECS would initially use a /52 prefix, specifically 2001:630:d0:f000::/52.

In terms of planning IPv6 address allocations, we chose to assign IPv6 subnet prefixes, which are always /64 in size, to be congruent with our IPv4 subnets. Thus a host will always be in the same 'administrative' IPv4 and IPv6 subnets.

IPv6 has the advantage that you do not need to resize subnets depending on host utilization, so whether the IPv4 subnet is /23, /24 or even /28, the IPv6 subnet used for the same hosts will be a /64, and never need to change.

For point-to-point links we have used, over time, /126's, /112's and /64's. Issues with point-to-point link addressing are discussed in [RFC3627]. We currently use /64's for point-to-point links. This topic has certainly caused some debate in the community. It has been noted that anything shorter than a /127 can in principle allow a ping-pong packet amplification attack, but we have not experienced any such attack as yet, and will review policy and mitigation techniques as and when such an attack is detected.

We have not yet found a need to use IPv6 ULAs [RFC4193] for local unicast addressing; these are loosely equivalent to IPv4 private addresses, except they are generally intended to be used in parallel with global IPv6 addresses purely for intra-site communication. Given our allocated IPv6 prefix is stable (we do not expect to renumber away from JANET), we do not see a need to run with ULAs at this time. There is also not yet much reported deployment experience on the use of ULAs in this way.

4.2.2 Management

There are three choices for assigning addresses to end hosts:

- Manual configuration;
- IPv6 Stateless Address Autoconfiguration [RFC4862];
- DHCPv6 [RFC3315].

For servers, we have chosen to manually configure IPv6 addresses on interfaces. For servers running specific services, we often use the port number as the host address, e.g. a DNS server running on <prefix>::53. We recommend not using autoconfigured addresses for servers/services, since these may change should hardware (and thus MAC address) change.

We use DHCP for IPv4, in part so that administrators feel ‘in control’ of the address assignments, and because there is then some increased accountability between devices and IP addresses that are assigned to them. We would like to use DHCPv6, but as yet implementations are in their relative infancy. We had hoped to achieve a DHCPv6 deployment in time for this report (which caused some of the delay in its production) but use of DHCPv6 is still considered ‘experimental’ in ECS.

The ISC DHCP implementation includes IPv6 support, as does Windows Vista and Windows 7, and we expect to make further progress soon. We are also testing DHCPv6 relay support in our network equipment. There is also the issue of DHCPv6’s use of DUID’s in client requests – in DHCP for IPv4 the servers can tie a request to a specific MAC source, but in IPv6 the DHCP request carries a unique, but host-generated, ID that is not known apriori by the DHCPv6 server.

For these reasons our non-server hosts currently use IPv6 stateless address autoconfiguration, and they rely on IPv4 DHCP for other configuration information (DNS resolvers, domain suffix, etc). Because our hosts are dual-stack, those hosts can operate ‘normally’ in the environment without DHCPv6 for that additional configuration information. IPv6-only hosts can manually configure an IPv6 DNS resolver if required.

A final consideration for address assignment is the use (or non use) of IPv6 Privacy

Addresses [RFC4941]. By default, some host operating systems generate new Privacy Addresses on a regular (e.g. daily) basis for use by the host when initiating outbound communications. While RFC 4941 is designed to assist in the avoidance of hosts (and thus users) being tracked by the appearance of a fixed MAC address-based host part of their autoconfigured address when they attach to different networks, it also provides a notable problem for network management whether the device is static or not. It may be difficult to identify which addresses are associated to which hosts, or the same host may appear as many different hosts in network monitoring software that cannot correlate the changing IP addresses. It's thus recommended to disable by default the privacy extensions in campus environments – as suggested in RFC 4941 – at the very least in the fixed desktops. Of course, in many such environments laptops are not under administrative control, so it may not be possible to disable the use of Privacy Addresses on such systems, and thus the onus falls on the network administrator to use improved tools for system monitoring (e.g. that can tie changing IP sources addresses to observed fixed MAC addresses).

4.3 Networking

The ECS network has a Cisco 6509 at its core, with stacks of Cisco 3750 edge devices providing service to users. The 6509 handles the main routing function for the internal IPv4 and IPv6 subnets, with the switch stacks generally providing the switched VLAN service to end users and other network services (wireless access points, printers, alarm systems, etc).

4.3.1 Routing

IPv4 and IPv6 are supported fully dual-stack across the 6509 and 3750 stacks. However the paths from the core router to the campus uplink are still separate for IPv4 and IPv6. This is in part because the existing IPv4 edge firewall device does not have all the IPv6 functions we require, but also because, as a computer science research School, we wish to be able to perform some level of experimentation on our IPv6 firewall and also with IPv6 multicast. As a result, we use a separate IPv6 firewall (currently iptables on Linux) and a separate router as our IPv6 Multicast PIM Rendezvous Point (RP) (a Cisco 7206, though this function could also be run on our 6509).

This 'split' connectivity is illustrated in Figure 1 where the paths that are IPv4-only, IPv6-only and dual-stack are shown. Hosts in the DMZ have separate default IPv4 and IPv6 routers.

Routing where used between our internal IPv6 routing systems uses RIPng [RFC2080].

If our routing were campus scale, we would almost certainly makes use of IS-IS for IPv6 (and for IPv4 too).

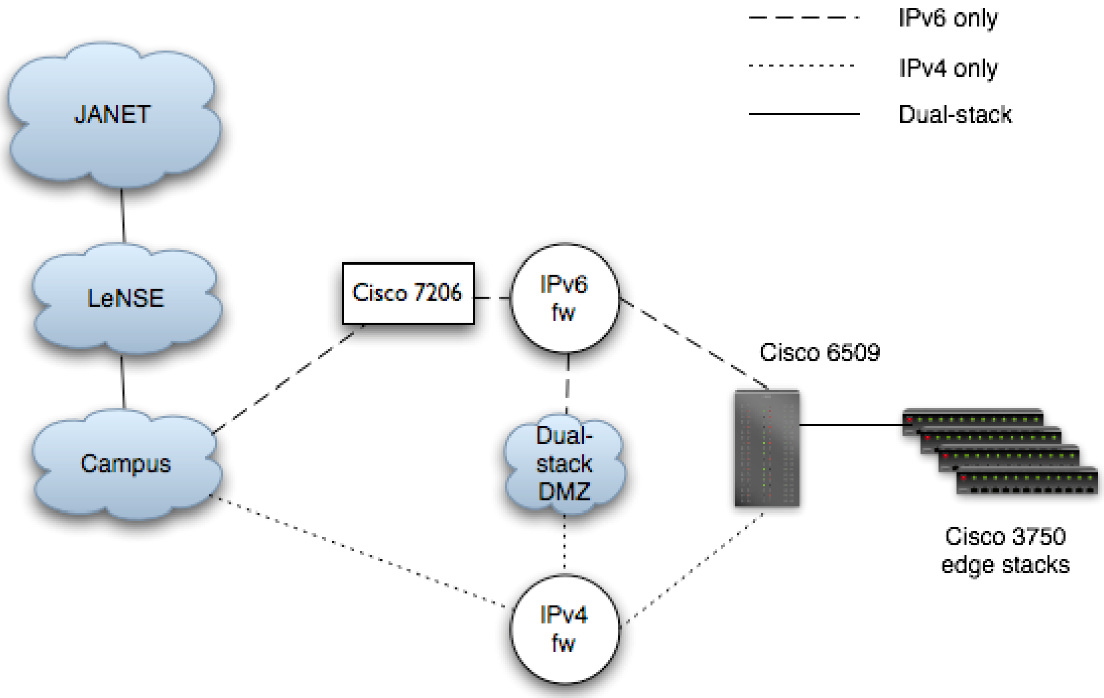

Figure 1: Overview of ECS campus connectivity

In due course, the IPv4 and IPv6 firewalling will be united into a single device/topology. At present the ability to 'experiment' with the IPv6 filtering, e.g. for specifics of Mobile IPv6 headers, is useful.

4.3.2 Switches

The edge switch stacks are deployed in server rooms within each of our four buildings. The switches carry IPv4 and IPv6 traffic in the same VLANs. There is no specific different configuration for layer 2 for IPv6.

There is one exception though; the network, and hosts on it, can benefit from IPv6 multicast flooding control just as they do with IPv4. Where IGMP snooping is used for IPv4 multicast, MLD snooping can be used for IPv6. MLDv1 and MLDv2 snooping for IPv6 are parallels to IGMPv2 and IGMPv3 snooping for IPv4.

Our core Cisco Catalyst router does not support MLDv1 snooping in hardware; the 3B (and 3BXL) variant of the Supervisor 720 only recognizes MLDv2. To handle both protocols, a 3C (or 3CXL) variant is required.

4.4 Hosts

Support for IPv6 in host operating systems has improved significantly over recent years. Windows 7 has excellent support, improving further on Vista and XP before it. Support in current releases of Linux, Solaris and BSD is also very good.

While MacOSX has many IPv6 features, and ships with IPv6 enabled by default, it lacks a couple of important elements, namely MLDv2 support (source specific multicast), has no support for RFC 3484 IPv6 address selection, and has no DHCPv6 client (which is a significant omission for a large enterprise environment).

Older operating systems (e.g. Windows 2000) are unlikely to ever see good IPv6 support. For a dual-stack deployment these should be phased out.

4.5 Services

In this section we discuss the deployment and use of core network services including DNS, mail relay (MX), web and multicast.

4.5.1 DNS

Support for IPv6 in DNS comes in two elements, one being the ability to add an IPv6 record for hosts in the DNS, and the other being the ability to communicate with the DNS server using IPv6 transport.

An IPv6 record is expressed in the DNS just like an IPv4 record, except that an IPv6 record uses AAAA where IPv4 uses A. For example a forward record might look like this:

```
ipv6lab.ecs.soton.ac.uk    IN    A           152.78.63.249
ipv6lab.ecs.soton.ac.uk    IN    AAAA        2001:630:d0:7000::9:2
```

Delegation for DNS also happens the same as IPv4, with glue records being used. Reverse DNS is also delegated similar to IPv4, e.g.:

```
0.63.78.152.in-addr.arpa            IN NS ns0.ecs.soton.ac.uk.
7.0.d.0.0.0.3.6.0.1.0.0.2.ip6.arpa        IN NS ns0.ecs.soton.ac.uk.
```

Reverse, nibble-based delegations can be used, e.g.

```
2.0.0.0.9.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.7.0.d.0.0.0.3.6.0.1.0.0.2.ip6.arp
a PTR \ ipv6lab.ecs.soton.ac.uk
```

Note that use of the .int TLD for reverse IPv6 DNS entries has been depreciated, and .ip6.arpa should now be used.

In ECS we use the ISC BIND [BIND] software for DNS, currently at the BIND9 version. There are no real ‘surprises’ in configuring IPv6 in BIND. The IPv6 syntax is as you might expect, e.g.:

- listen-on-v6 { any; };
- “transfer-source-v6 *” to specify IPv6 source address for transfers;
- “query-source-v6 address * port *” to specify IPv6 source for queries.

IPv6 addresses can of course also be used in ACLs, e.g. to restrict zone transfers.

In ECS all our DNS resolvers are dual-stack, and can thus be used by IPv4-only or IPv6-only clients if required. Until DHCPv6 is deployed, clients using DNS will get DNS resolver configuration information via DHCPv4. While an Experimental RFC defines DNS configuration via Router Advertisements [RFC5006], there are as yet no known implementations, and for ECS we will need DHCPv6 anyway to define other configuration data for hosts.

4.5.2 Mail relays

Enabling mail relays (MX servers) to support IPv6 is also not particularly tricky. You need to add AAAA records for your MXes to the DNS, and configure your relays to listen on IPv6. As per IPv4, you may also wish to configure filters and limitations on which hosts/prefixes can use the relays.

Suggested best practice for dual-stack relays is described in [RFC3974], which recommends using both A and AAAA records for each MX. We have followed this advice in ECS, where our records appear as follows:

```
mx.ecs.soton.ac.uk.     3600    IN      AAAA    2001:630:d0:f110::25c

mx.ecs.soton.ac.uk.     3600    IN      AAAA    2001:630:d0:f102::25b

mx.ecs.soton.ac.uk.     3600    IN      AAAA    2001:630:d0:f102::25c

mx.ecs.soton.ac.uk.     3600    IN      AAAA    2001:630:d0:f110::25b

mx.ecs.soton.ac.uk.     3600    IN      A       152.78.68.132

mx.ecs.soton.ac.uk.     3600    IN      A       152.78.68.137
mx.ecs.soton.ac.uk.     3600    IN      A       152.78.71.14

mx.ecs.soton.ac.uk.     3600    IN      A       152.78.71.210
```

ECS uses sendmail, which has supported IPv6 out of the box since version 8.10. The only specific configuration that may need adding is enabling the inet6 Family in sendmail.mc.

ECS monitors its mail traffic volumes using MRTG. In the following Figures we show the volumes of IPv6 email as viewed during June 2009, showing data for IPv4, IPv6 and tagged ‘spam’ messages.

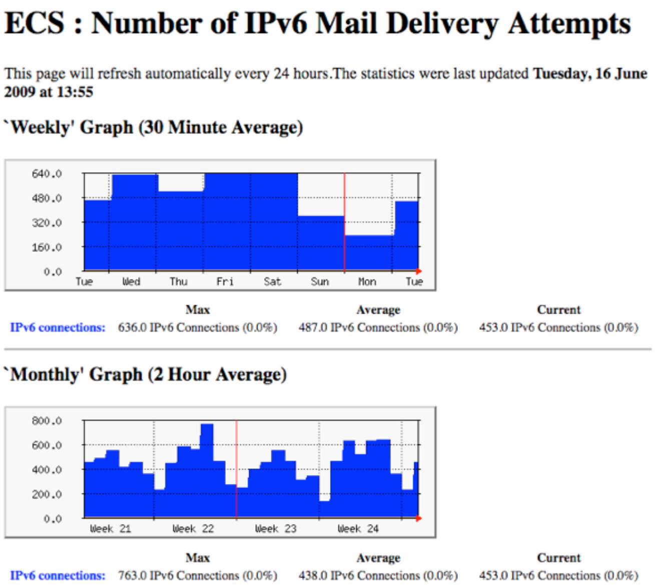

ECS : Number of IPv6 Mail Delivery Attempts

This page will refresh automatically every 24 hours.The statistics were last updated **Tuesday, 16 June 2009 at 13:55**

`Weekly' Graph (30 Minute Average)

	Max	Average	Current
IPv6 connections:	636.0 IPv6 Connections (0.0%)	487.0 IPv6 Connections (0.0%)	453.0 IPv6 Connections (0.0%)

`Monthly' Graph (2 Hour Average)

	Max	Average	Current
IPv6 connections:	763.0 IPv6 Connections (0.0%)	438.0 IPv6 Connections (0.0%)	453.0 IPv6 Connections (0.0%)

Figure 2: E-mail entering ECS over IPv6

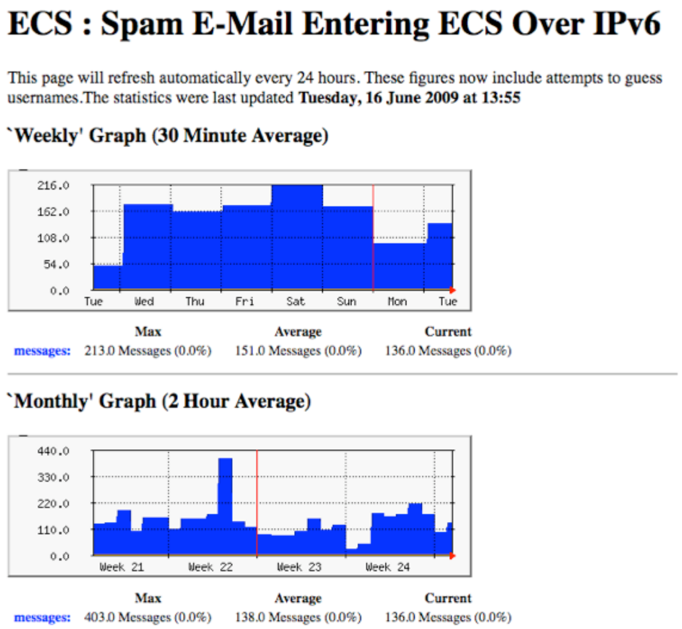

ECS : Spam E-Mail Entering ECS Over IPv6

This page will refresh automatically every 24 hours. These figures now include attempts to guess usernames.The statistics were last updated **Tuesday, 16 June 2009 at 13:55**

`Weekly' Graph (30 Minute Average)

	Max	Average	Current
messages:	213.0 Messages (0.0%)	151.0 Messages (0.0%)	136.0 Messages (0.0%)

`Monthly' Graph (2 Hour Average)

	Max	Average	Current
messages:	403.0 Messages (0.0%)	138.0 Messages (0.0%)	136.0 Messages (0.0%)

Figure 3: Spam E-mail arriving over IPv6

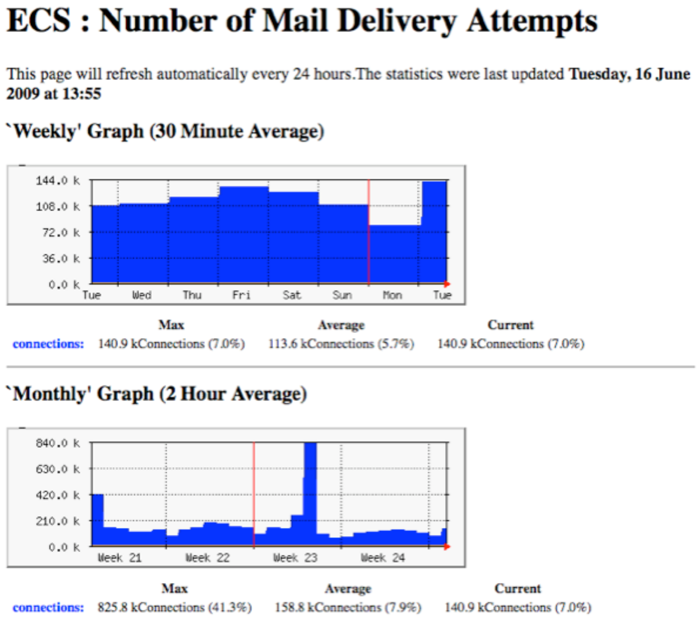

ECS : Number of Mail Delivery Attempts

This page will refresh automatically every 24 hours.The statistics were last updated **Tuesday, 16 June 2009 at 13:55**

`Weekly' Graph (30 Minute Average)

	Max	Average	Current
connections:	140.9 kConnections (7.0%)	113.6 kConnections (5.7%)	140.9 kConnections (7.0%)

`Monthly' Graph (2 Hour Average)

	Max	Average	Current
connections:	825.8 kConnections (41.3%)	158.8 kConnections (7.9%)	140.9 kConnections (7.0%)

Figure 4: Emails entering ECS over IPv4

We use the open source MailScanner package (developed in ECS) to detect spam and viruses in e-mails. This is IP version agnostic, though as yet there are no IPv6 DNS-based spam 'blacklists' (a major supplier of such lists has told us this may still be two years away, but presumably as the need arises support will be forthcoming).

In June 2009 we received around 158,000 messages per day over IPv4 of which 81%

were identified as spam while we only received 438 per day over IPv6, of which 32% were spam. Almost all the IPv6 spam comes from mail lists run on dual-stack mail list servers.

4.5.3 WWW

Enabling IPv6 access to web servers is also quite straightforward in current web server applications. In ECS we almost exclusively use Apache, where you can choose to use ‘Listen 80’ to listen on a single socket (receiving IPv6 and IPv4 mapped connections), or ‘Listen [::]:80’ will listen for IPv6 only and ‘Listen 0.0.0.0:80’ will listen for IPv4 only.

There are no special considerations except updating the configured web server access controls to support IPv6 also. Sites should check that their web log analysis software supports IPv6 address formats.

4.5.4 Separating IPv6 namespace?

When deploying IPv6 support in services the natural approach is to use the same namespace for both IPv4, and IPv6, e.g. you add IPv4 and IPv6 records to the DNS for www.ecs.soton.ac.uk.

However, some sites prefer to initially make, in particular, IPv6 web services available via a separate namespace, e.g. rather than using www.domainname.org for both, the site uses ipv6.domainname.org for IPv6 access. The thinking behind such caution is that if a client/host receives an IPv6 AAAA record when an application looks up a domain, if the host has IPv6 enabled but connectivity is poor or non-existent, the service may be adversely affected. This may be a prudent first step during a pilot phase – the decision is one for each site to make.

As yet in ECS we have not had significant issues arise from using a common namespace. We feel that if problems do arise this helps us understand better how to improve the connectivity for a given user or remote host. In principle, as IPv6 becomes more widely deployed the connectivity concerns should reduce.

4.5.5 Multicast

We discuss IPv6 Multicast in a later section. In ECS we use IPv4 multicast for many applications, including videoconferencing (H.323 as well as AccessGrid), viewing streamed video content, and workstation configuration (Ghost).

We chose to deploy IPv6 Multicast support in the network to encourage new services and applications to be developed and deployed. The principle requirements to achieve this were to enable multicast routing (PIM), to ensure scope boundaries were configured appropriately, to deploy a site PIM Rendezvous Point (RP) and to enable MLD snooping in switch-router devices.

4.6 Network Management and Monitoring

One of the most important aspects of introducing IPv6 using a dual-stack strategy is to be able to monitor and manage both protocols effectively.

It is thus important that your network management and monitoring tools support IPv6, but also that where tools are used to manage both protocols they do so in an efficient and consistent manner.

4.6.1 Netflow

By using routers that support Netflow v9, network flow data for both IPv4 and IPv6 can be sent to a collector for subsequent processing and analysis.

In ECS we send data from our core Cisco 6509 router to a collector that runs the nfsen [NFSEN] visualization package, which we have found to be very flexible and rich in features for traffic analysis.

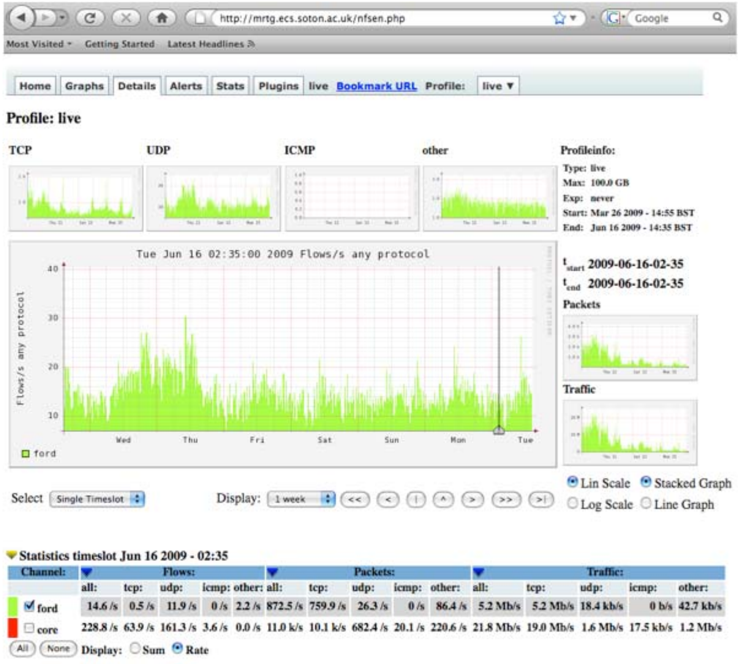

Profile: live

Profileinfo:

Type: live
Max: 100.0 GB
Exp: never
Start: Mar 26 2009 - 14:55 BST
End: Jun 16 2009 - 14:35 BST

t_{start} 2009-06-16-02-35
t_{end} 2009-06-16-02-35

Packets

Traffic

Statistics timeslot Jun 16 2009 - 02:35

Channel:	Flows:					Packets:					Traffic:				
	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:	all:	tcp:	udp:	icmp:	other:
🟩 ☑ ford	14.6 /s	0.5 /s	11.9 /s	0 /s	2.2 /s	872.5 /s	759.9 /s	26.3 /s	0 /s	86.4 /s	5.2 Mb/s	5.2 Mb/s	18.4 kb/s	0 b/s	42.7 kb/s
🟥 □ core	228.8 /s	63.9 /s	161.3 /s	3.6 /s	0.0 /s	11.0 k/s	10.1 k/s	682.4 /s	20.1 /s	220.6 /s	21.8 Mb/s	19.0 Mb/s	1.6 Mb/s	17.5 kb/s	1.2 Mb/s

Display: ○ Sum ● Rate

Figure 5: nfsen view of IPv6 flow data

The nfsen interface allows various views of the flow data, but also allows specific queries to be passed to the underlying nfdump tool to search for specific traffic instances.

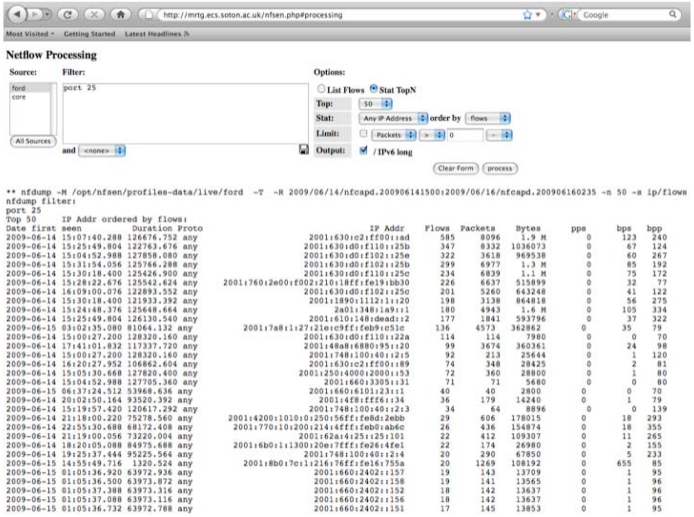

Netflow Processing

Source: ford, core

Filter: port 25

and <none>

Options: ○ List Flows ● Stat TopN

Top: 50

Stat: Any IP Address order by flows

Limit: □ Packets > 0 -

Output: ☑ / IPv6 long

```
** nfdump -M /opt/nfsen/profiles-data/live/ford  -T  -R 2009/06/14/nfcapd.200906141500:2009/06/16/nfcapd.200906160235 -n 50 -s ip/flows
nfdump filter:
port 25
Top 50    IP Addr ordered by flows:
Date first seen          Duration Proto                                   IP Addr   Flows  Packets     Bytes       pps      bps   bpp
2009-06-14 15:07:40.288 126676.752 any                      2001:630:c2:ff00::ad     585     8096     1.9 M         0      123   240
2009-06-14 15:25:49.804 122763.676 any                     2001:630:d0:f110::25b     347     8332   1036073         0       67   124
2009-06-14 15:04:52.988 127858.080 any                     2001:630:d0:f102::25e     322     3618    969538         0       60   267
2009-06-14 15:31:54.056 125766.288 any                     2001:630:d0:f102::25b     299     6977     1.3 M         0       85   192
2009-06-14 15:30:18.400 125426.900 any                     2001:630:d0:f110::25c     234     6839     1.1 M         0       75   172
2009-06-14 15:28:22.676 125542.624 any     2001:760:2e00:f002:210:18ff:fe19:bb30     226     6637    515899         0       32    77
2009-06-14 16:09:00.076 122893.552 any                     2001:630:d0:f102::25c     201     5260    643248         0       41   122
2009-06-14 15:30:18.400 121933.392 any                      2001:1890:1112:1::20     198     3138    864818         0       56   275
2009-06-14 15:24:48.376 125648.664 any                           2a01:348:1a9::1     180     4943     1.6 M         0      105   334
2009-06-14 15:25:49.804 126130.540 any                      2001:610:148:dead::2     177     1841    593796         0       37   322
2009-06-15 03:02:35.080 81064.132 any          2001:7a8:1:27:21e:c9ff:feb9:c51c     136     4573    362862         0       35    79
2009-06-14 15:00:27.200 128320.160 any                     2001:630:d0:f110::22a     114      114      7980         0        0    70
2009-06-14 17:41:01.832 117337.720 any                     2001:48a8:6880:95::20      99     3674    360361         0       24    98
2009-06-14 15:00:27.200 128320.160 any                      2001:748:100:40::2:5      92      213     25644         0        1   120
2009-06-14 16:20:27.952 106862.604 any                      2001:630:c2:ff00::89      74      348     28425         0        2    81
2009-06-14 15:05:30.668 127820.400 any                    2001:250:4000:2000::53      72      360     28800         0        1    80
2009-06-14 15:04:52.988 127705.360 any                         2001:660:3305::31      71       71      5680         0        0    80
2009-06-15 06:37:24.512 53968.636 any                        2001:660:6101:23::1      40       40      2800         0        0    70
2009-06-14 20:02:50.164 93520.392 any                          2001:4f8:fff6::34      36      179     14240         0        1    79
2009-06-14 15:19:57.420 120617.292 any                      2001:748:100:40::2:3      34       64      8896         0        0   139
2009-06-14 21:18:00.220 75278.560 any        2001:4200:1010:0:250:56ff:fe8d:2ebb      29      606    178015         0       18   293
2009-06-14 22:55:30.688 68172.408 any         2001:770:10:200:214:4fff:feb0:ab6c      26      436    154874         0       18   355
2009-06-14 21:19:00.056 73220.004 any                      2001:62a:4:25::25:101      22      412    109307         0       11   265
2009-06-14 18:20:05.088 84975.688 any         2001:6b0:1:1300:20e:7fff:fe26:4fe1      22      174     26980         0        2   155
2009-06-14 19:25:37.444 95225.564 any                       2001:748:100:40::2:4      20      290     67850         0        5   233
2009-06-15 14:55:49.716  1320.524 any           2001:8b0:7c:1:216:76ff:fe16:755a      20     1269    108192         0      655    85
2009-06-15 01:05:36.920 63972.936 any                         2001:660:2402::157      19      143     13709         0        1    95
2009-06-15 01:05:36.500 63973.872 any                         2001:660:2402::158      19      141     13565         0        1    96
2009-06-15 01:05:37.388 63973.316 any                         2001:660:2402::152      18      142     13637         0        1    96
2009-06-15 01:05:37.088 63973.116 any                         2001:660:2402::156      18      142     13637         0        1    96
2009-06-15 01:05:36.732 63972.788 any                         2001:660:2402::151      17      145     13853         0        1    95
```

Figure 6: Using nfsen to search for IPv6 netflows on port 25 (SMTP)

The collector can draw traffic from multiple sources, e.g. we also send IPv6 flow data from our 7206 router to the same collector.

4.6.2 NAV

To monitor our switch-router stacks we use an open source package called NAV [NAV] that was developed by UNINETT, the Norwegian NREN.

This tool has proven to be very useful and offers an excellent range of configurable views of our Cisco network equipment, e.g. you can:

- Look at a selection of active prefixes/VLANs;
- Look at active IPv6 addresses in a subnet;
- Look at active IP addresses on a MAC address.

These are illustrated in the following Figures.

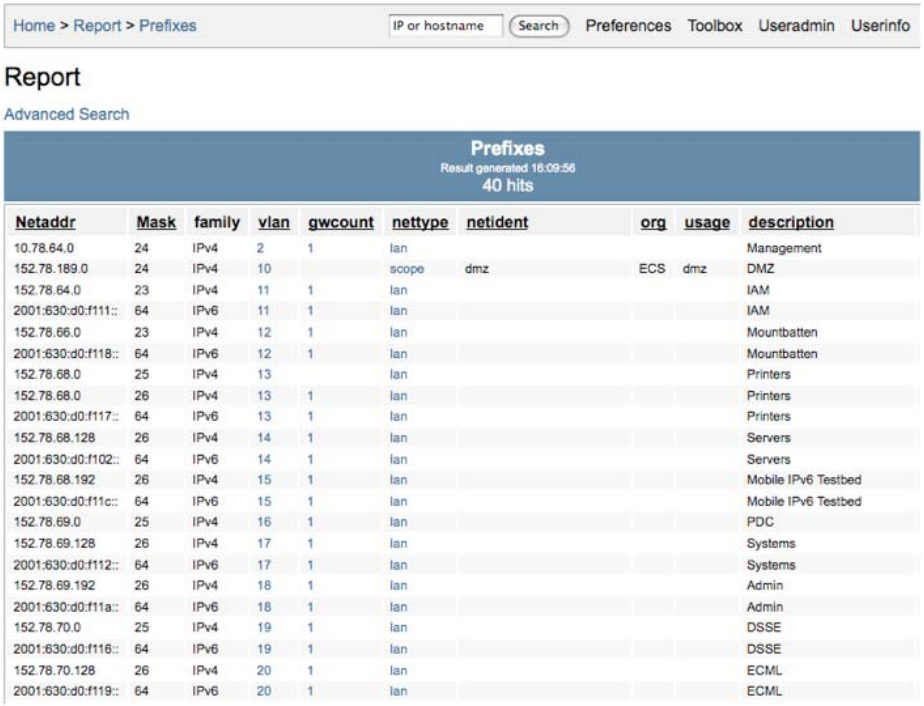

Report

Advanced Search

Prefixes
Result generated 16:09:56
40 hits

Netaddr	Mask	family	vlan	gwcount	nettype	netident	org	usage	description
10.78.64.0	24	IPv4	2	1	lan				Management
152.78.189.0	24	IPv4	10		scope	dmz	ECS	dmz	DMZ
152.78.64.0	23	IPv4	11	1	lan				IAM
2001:630:d0:f111::	64	IPv6	11	1	lan				IAM
152.78.66.0	23	IPv4	12	1	lan				Mountbatten
2001:630:d0:f118::	64	IPv6	12	1	lan				Mountbatten
152.78.68.0	25	IPv4	13		lan				Printers
152.78.68.0	26	IPv4	13	1	lan				Printers
2001:630:d0:f117::	64	IPv6	13	1	lan				Printers
152.78.68.128	26	IPv4	14	1	lan				Servers
2001:630:d0:f102::	64	IPv6	14	1	lan				Servers
152.78.68.192	26	IPv4	15	1	lan				Mobile IPv6 Testbed
2001:630:d0:f11c::	64	IPv6	15	1	lan				Mobile IPv6 Testbed
152.78.69.0	25	IPv4	16	1	lan				PDC
152.78.69.128	26	IPv4	17	1	lan				Systems
2001:630:d0:f112::	64	IPv6	17	1	lan				Systems
152.78.69.192	26	IPv4	18	1	lan				Admin
2001:630:d0:f11a::	64	IPv6	18	1	lan				Admin
152.78.70.0	25	IPv4	19	1	lan				DSSE
2001:630:d0:f116::	64	IPv6	19	1	lan				DSSE
152.78.70.128	26	IPv4	20	1	lan				ECML
2001:630:d0:f119::	64	IPv6	20	1	lan				ECML

Figure 7: NAV view of subnet / VLAN data

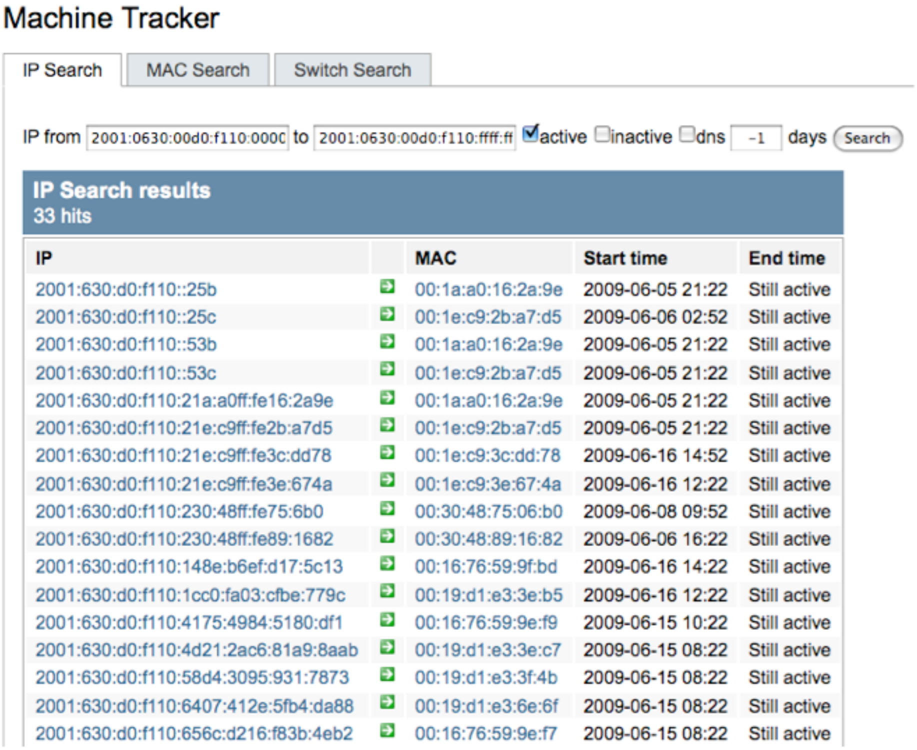

Machine Tracker

IP from 2001:0630:00d0:f110:000(to 2001:0630:00d0:f110:ffff:ff ☑active □inactive □dns -1 days Search

IP Search results
33 hits

IP		MAC	Start time	End time
2001:630:d0:f110::25b	➔	00:1a:a0:16:2a:9e	2009-06-05 21:22	Still active
2001:630:d0:f110::25c	➔	00:1e:c9:2b:a7:d5	2009-06-06 02:52	Still active
2001:630:d0:f110::53b	➔	00:1a:a0:16:2a:9e	2009-06-05 21:22	Still active
2001:630:d0:f110::53c	➔	00:1e:c9:2b:a7:d5	2009-06-05 21:22	Still active
2001:630:d0:f110:21a:a0ff:fe16:2a9e	➔	00:1a:a0:16:2a:9e	2009-06-05 21:22	Still active
2001:630:d0:f110:21e:c9ff:fe2b:a7d5	➔	00:1e:c9:2b:a7:d5	2009-06-05 21:22	Still active
2001:630:d0:f110:21e:c9ff:fe3c:dd78	➔	00:1e:c9:3c:dd:78	2009-06-16 14:52	Still active
2001:630:d0:f110:21e:c9ff:fe3e:674a	➔	00:1e:c9:3e:67:4a	2009-06-16 12:22	Still active
2001:630:d0:f110:230:48ff:fe75:6b0	➔	00:30:48:75:06:b0	2009-06-08 09:52	Still active
2001:630:d0:f110:230:48ff:fe89:1682	➔	00:30:48:89:16:82	2009-06-06 16:22	Still active
2001:630:d0:f110:148e:b6ef:d17:5c13	➔	00:16:76:59:9f:bd	2009-06-16 14:22	Still active
2001:630:d0:f110:1cc0:fa03:cfbe:779c	➔	00:19:d1:e3:3e:b5	2009-06-16 12:22	Still active
2001:630:d0:f110:4175:4984:5180:df1	➔	00:16:76:59:9e:f9	2009-06-15 10:22	Still active
2001:630:d0:f110:4d21:2ac6:81a9:8aab	➔	00:19:d1:e3:3e:c7	2009-06-15 08:22	Still active
2001:630:d0:f110:58d4:3095:931:7873	➔	00:19:d1:e3:3f:4b	2009-06-15 08:22	Still active
2001:630:d0:f110:6407:412e:5fb4:da88	➔	00:19:d1:e3:6e:6f	2009-06-15 08:22	Still active
2001:630:d0:f110:656c:d216:f83b:4eb2	➔	00:16:76:59:9e:f7	2009-06-15 08:22	Still active

Figure 8: NAV showing active IP addresses on a given IPv6 prefix

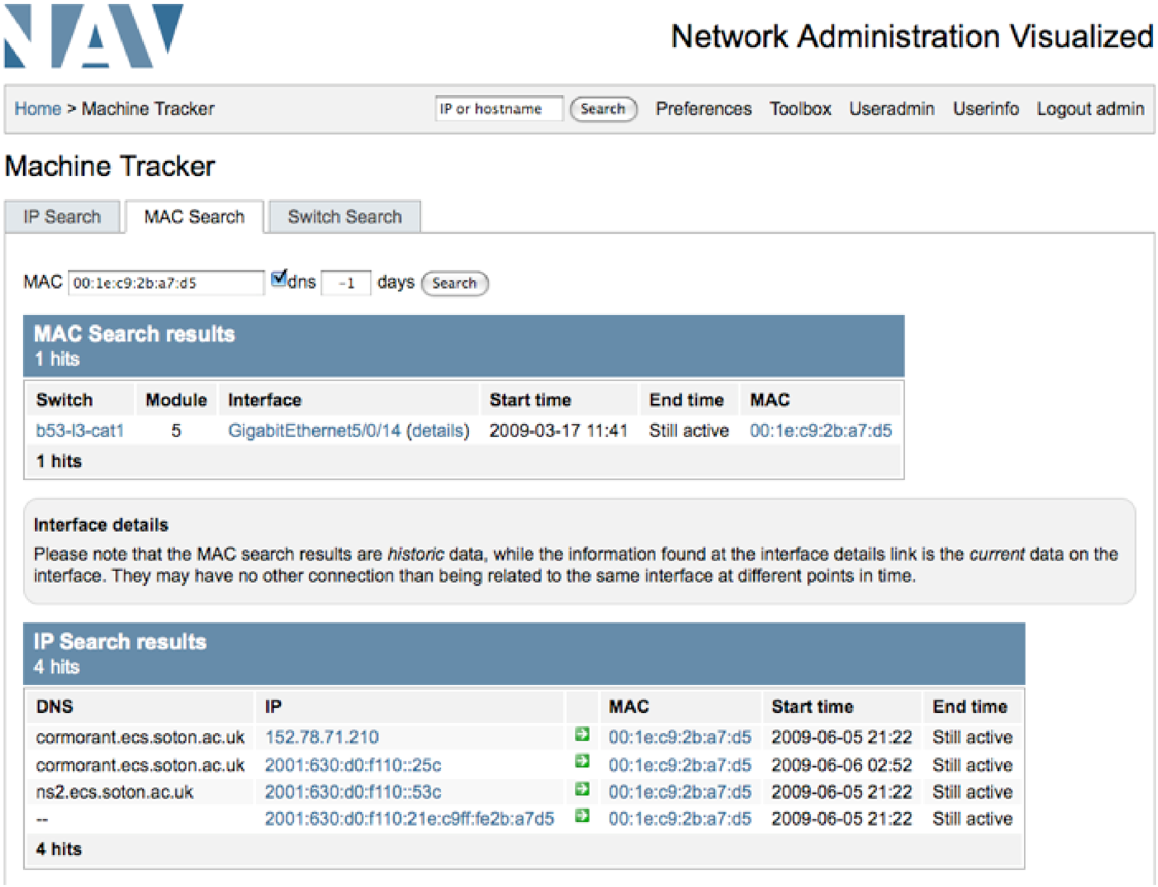

Machine Tracker

MAC 00:1e:c9:2b:a7:d5 ☑dns -1 days Search

MAC Search results
1 hits

Switch	Module	Interface	Start time	End time	MAC
b53-l3-cat1	5	GigabitEthernet5/0/14 (details)	2009-03-17 11:41	Still active	00:1e:c9:2b:a7:d5
1 hits					

Interface details
Please note that the MAC search results are *historic* data, while the information found at the interface details link is the *current* data on the interface. They may have no other connection than being related to the same interface at different points in time.

IP Search results
4 hits

DNS	IP		MAC	Start time	End time
cormorant.ecs.soton.ac.uk	152.78.71.210		00:1e:c9:2b:a7:d5	2009-06-05 21:22	Still active
cormorant.ecs.soton.ac.uk	2001:630:d0:f110::25c		00:1e:c9:2b:a7:d5	2009-06-06 02:52	Still active
ns2.ecs.soton.ac.uk	2001:630:d0:f110::53c		00:1e:c9:2b:a7:d5	2009-06-05 21:22	Still active
--	2001:630:d0:f110:21e:c9ff:fe2b:a7d5		00:1e:c9:2b:a7:d5	2009-06-05 21:22	Still active
4 hits					

Figure 9: NAV can identify all IPv4 and IPv6 addresses used by a host

The last Figure is particularly interesting because it shows that NAV understands, from MAC address correlation, that a single host may have multiple IPv6 addresses, as well as IPv4 address(es).

4.6.3 Nagios

We use the Nagios package to monitor service and system availability in ECS. IPv6 support in Nagios was improved through the work of [6NET]. Currently, you generally need to define separate IPv4 and IPv6 tests. Hopefully, in due course, Nagios will allow more simplified testing of the same service via a single configured test for both IP versions.

4.7 Security

When enabling IPv6, it is important both that security principles applied to IPv4 are applied with equal thoroughness to IPv6, and that introducing IPv6 does not adversely affect existing IPv4 security. Thus where an IPv4 perimeter firewall is used, the same (effective) ruleset should be applied 'service for service' for both protocols, whether implemented in the same device or not. And where IPv6 is enabled, IPv4 security should not be subverted, e.g. by inappropriate use of tunneling.

There are some specific considerations for IPv6 ICMP filtering as discussed in [RFC4890]. Note that IPv6 requires the use of certain ICMPv6 messages for proper PMTU discovery.

4.7.1 Firewalls

In ECS we have an IPv4 firewall as shown in Figure 1 which we have chosen to keep for only IPv4 use at this time. The Checkpoint product does not yet have all the features we would like to see (including some IPv6 Multicast capability) and we also wish to be able to do some research/experimentation on our IPv6 firewall, e.g. to handle specific existing or new IPv6 headers.

We currently use a Linux iptables firewall for IPv6 filtering, and we also make use of ACLs on our Cisco equipment.

Cisco IOS has access lists that use the same principle as the IPv4 filters, e.g.

```
ipv6 access-list-name permit tcp 2001:0db8:0300:0201::/64 eq 22
```

The traffic-filter command can also be used to then apply named rules inbound or outbound, e.g.:

```
interface ethernet 0
ipv6 traffic-filter access-list-name in
```

Support for host-based IPv6 filtering is also good now in most operating systems.

4.7.2 IDS

ECS uses Snort for IPv4 intrusion detection. Since v3.0, Snort has also supported IPv6 inspection for application-oriented intrusion pattern detection, e.g. attempts to exploit a web server via an IPv6 connection. ECS runs a separate instance of Snort on its IPv6 external link to detect possible suspect traffic on that link.

As yet Snort does not include specific tools to detect issues with IPv6 headers or header contents.

4.7.3 IPv6-specific issues

One of the potential dangers from a security perspective in introducing IPv6 is that new classes of security risks may arise.

One such class lies in the array of transition tools that may be used, with out without the approval of the site administrators. Users may inadvertently or deliberately use such tools to introduce new connectivity within the site or to external sites. Such concerns can be reduced by deploying IPv6 pervasively in the site, as is the case in ECS, such that transition tools are not required.

The most commonly observed security-related issue we have seen in ECS lies in rogue RAs being observed on the local network, most commonly 6to4-based ones. The problem is discussed in [ROGUE-RA] and one current mitigation is to run a detection tool such as [RAMOND] – in the future we expect to see RA-based filters in Ethernet switch devices much as exist for DHCPv4 snooping/filtering today (see [ROGUE-RA for details).

4.8 Applications

The applications used in ECS are a mixture of open source and commercial ones. In general, the infrastructure software that involves network access is either common open source software or software from Microsoft (Outlook, IIS, Exchange, etc).

Support for IPv6 in most open source packages is now very good. There are no significant concerns about IPv6 support in any such packages used in ECS. Further, while we have yet to fully deploy IPv6 in some of our Microsoft applications, support has been introduced in recent versions of their packages, e.g. Exchange 2007 running on Windows Server 2008.

Some support may be less extensive if IPv6-only operation is expected, but in ECS we currently run all services dual-stack. We expect support for IPv6-only operation to improve in the near future.

Complexity in introducing IPv6 has been mitigated a little due to the fact that most of our ‘corporate’ applications are web-based. Support for IPv6 in Apache and ISS is very well established.

Our experience in porting applications to be IP version-independent has generally been positive. Most of the issues are described in [RFC4038]. The ease of porting largely depends how well the applications have been written, in particular how abstracted the networking components are, and how the data structures are defined and used. Long standing support in C and Java has been complemented by support in scripting languages such as Perl. We have not encountered a piece of software in ECS that we have been unable to port to support IPv6.

5. New IPv6-enabled Services

One of the reasons for deploying IPv6 in ECS was to facilitate and encourage new application development. It allows both undergraduate and postgraduate students to use the protocol for their studies and research respectively.

5.1 Student-developed Applications

In ECS we have seen various uses of IPv6 by students, e.g.:

- In coursework assignments, particularly for networking courses;
- In final year projects, e.g. file sharing applications, and network monitoring and management applications;
- In the student-run Southampton Open Wireless Network (SOWN) where IPv6 connectivity is provided via wireless to student residences and outdoor areas in and around the main campus;
- In the ECS IP TV system, ECS-TV, which has been developed by a number of students. This uses IPv6 Multicast, as described in the next section, to relay a variety of free-to-air channels.

Overall the experience has been very positive. While many of these applications could have been developed as IPv4-only ones, the students have shown great enthusiasm for developing applications with IPv6 capability included and, in the case of ECS-TV, using only IPv6.

5.2 IPv6 Multicast

In ECS we have introduced IPv6 Multicast alongside IPv4 Multicast. We continue to run IPv4 applications such as Ghost and AccessGrid while we have introduced an IPv6 IP TV system as described below.

There are some IPv6 advantages with respect to Multicast, including:

- It is easier to obtain a global multicast group address – these can be generated based upon your unicast IPv6 prefix.
- The new Embedded-RP [RFC3956] protocol means that for ASM there is no need

for the MSDP protocol between PIM domains. Embedded-RP, as the name suggests, embeds the RP address in the group address, so a router knows implicitly where the group RP is.

- Scoping is explicit in the IPv6 group address, making scope boundaries easier to manage. The commonly used scopes are 0x2 (link), 0x5 (site), 0x8 (organization) and 0xf (global). In ECS we use site scope on our boundaries with 0x8 scope reserved for our campus boundary.

The IPv6 Multicast support in the Cisco systems we use in ECS is generally good. The specifics to look for are support for MLD snooping and for Embedded-RP (which relies on all routers on a multicast path understanding the protocol).

There is a good selection of multicast debugging tools available, including most usefully ssmping [SSMPING] and dbeacon [DBEACON].

5.3 ECS-TV

The ECS-TV system is an IPv6-only Multicast IP TV system developed by students in ECS. It uses Any Source Multicast (ASM) with Embedded-RP to provide Multicast service for over 100 FreeView and third party channels.

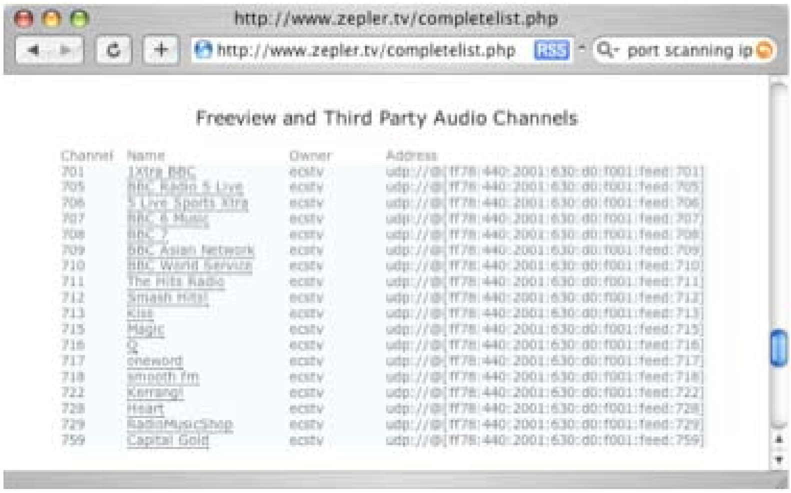
http://www.zepler.tv/completelist.php
http://www.zepler.tv/completelist.php
RSS
port scanning ip

Freeview and Third Party Audio Channels

Channel	Name	Owner	Address
701	1Xtra BBC	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:701]
705	BBC Radio 5 Live	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:705]
706	5 Live Sports Xtra	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:706]
707	BBC 6 Music	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:707]
708	BBC 7	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:708]
709	BBC Asian Network	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:709]
710	BBC World Service	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:710]
711	The Hits Radio	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:711]
712	Smash Hits!	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:712]
713	Kiss	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:713]
715	Magic	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:715]
716	Q	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:716]
717	oneword	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:717]
718	smooth fm	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:718]
722	Kerrang!	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:722]
728	Heart	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:728]
729	RadioMusicShop	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:729]
759	Capital Gold	ecstv	udp://@[ff78:440:2001:630:d0:f001:feed:759]

Figure 10: ECS-TV channels

The VideoLAN software is used for content distribution and viewing [VIDEOLAN]. The client software runs on all common operating systems.

The RP for the ECS-TV groups runs on our Cisco 7206. All the groups are limited to organisation scope for licensing reasons. MLD snooping is used on our Cisco 3750 switch stacks to limit multicast flooding to hosts.

5.4 Mobile IPv6

We have undertaken some initial Mobile IPv6 deployment in ECS. The initial experiments saw a virtual MIPv6 subnet on the Cisco 7206 used as a home link and running the Home Agent (HA) function. We have since used a Linux HA on a physical (wireless) link, and run MIPv6 between the ECS wireless network and SOWN.

A driver for MIPv6 deployment between ECS and SOWN is to allow students to roam between ECS and SOWN wireless coverage while maintaining open network-based applications (e.g. streaming applications, or ssh-based ones).

By supporting IPv6 on its wireless LANs, ECS is a tier 3 JANET Roaming Service [JRS] site and is eduroam-compliant [EDUROAM]. The use of 802.1x for authentication allows access by IPv4-only or IPv6-only devices, where web-redirect portals might otherwise limit access to IPv4-only devices currently.

The main challenge in such a deployment is MIPv6 operation through deployed firewalls – currently the best practice resides only in IETF Internet Drafts [FWVEN], [FWADMIN] and is not available in commercial products. Unless IPv6 traffic is not firewalled/filtered between the Mobile Node, Home Agent and Correspondent Node, some quite intelligent stateful processing is required.

6. REFLECTIONS ON DEPLOYMENT

In this section we reflect on what we found easy and what we found trickier in the deployment process, and on open issues remaining.

6.1 What was easy

The following aspects of deployment were the relatively straightforward ones:

- Obtaining IPv6 connectivity and address space. This was thanks to JANET's forward thinking in acquiring its address space and deploying IPv6 dual-stack, and due to the support from our regional network (LeNSE).

- IPv6 support in host/router platforms. This has taken time to harden, but is now very good all round.

- Enabling IPv6 support in core network services, including DNS, mail relays and our web presence.

- Porting applications/tools to support IPv6. We found no software that could not be ported, and in general adding support was not complex or time consuming. Open source packages now generally have very good support.

- Host autoconfiguration. This works well for hosts to autoconfigure an IPv6 address and default router. However for an enterprise deployment, richer configuration information is required (see next section).

- IPv6 wireless via eduroam. We are able to authenticate devices at Layer 2 using 802.1x/eduroam, rather than relying on finding IPv6 support in web-authentication gateways.

This is not an exhaustive list, but is intended to give a flavour of the 'easier' deployment tasks.

6.2 What was not easy

There are some aspects that certainly are proving trickier to deploy, including:

- DHCPv6 support has been slow to evolve, and is still missing in some key operating systems (in particular MacOSX). DHCPv6 also uses host DUIDs rather than MAC addresses, which has implications for administrative management.

- The Microsoft applications have taken time to become available, but that situation has improved since the release of Windows server 2008. This affects services such as Exchange and VPN (which ECS are in the process of testing).

- Handling multi-addressed hosts including those for Privacy Addresses. In IPv4 hosts tend to have single addresses. Allowing for multi-addressed hosts requires some careful consideration.

- Managing a dual-stack environment is more complex than one just running IPv4. There needs to be better consistency in tools to manage firewall policy, integrated DHCP, etc.

- There have been some new LAN security issues to deal with, most notably rogue RAs.

- Providing equivalent IPv6 services for nomadic users, e.g. enabling users to use IPv6 at home.

An important point though is that at this stage we have taken the approach to enable IPv6 fully on the wire (dual-stack) and for key network services. We can then add application support as this becomes ready. An ‘all in one go’ approach is not required, so application ‘gaps’ are not critical.

6.3 To do

There is still work to do in our ECS IPv6 deployment. Immediate next steps include:

- Further testing of IPv6 support in MS applications, including Exchange for Windows Server 2008 and the MS VPN service. These are expected to be positive given investigation to date.

- We need to develop better tools to ensure consistent firewall policy being applied on our separate IPv4 and IPv6 firewalls. Commercial products also need to allow better definition of rules for dual-stack nodes.

- Further DHCPv6 tests need to be run, initially in our student laboratory spaces, with a view to an ECS-wide DHCPv6 service being available.

- Further work on building tools to support joint DHCPv4/v6 and DNS management; again, this is an example where having integrated tools to perform mission-critical tasks is so important.

- We need to do further work on tools to efficiently manage and monitor our dual-protocol environment; Nagios is the next target.

7. SUMMARY AND CONCLUSIONS

Overall, the IPv6 deployment experience has been very positive for ECS. We have deployed IPv6 on all our existing IPv4 network links, enabling both unicast and multicast IPv6 traffic. We have enabled many network and application services dual-stack, including our externally facing web, mail and DNS services. The dual-stack approach to deployment has worked well for us.

IPv6 is robust in operation, and we have not observed any degrading of our long-established IPv4 services. There have been occasional issues that have arisen, but these have been addressed quite quickly, and none were significant. Some new IPv6 services have been introduced successfully, many from student initiatives.

The main challenge in running a dual-stack enterprise lies in monitoring and managing both protocols, and to ensure consistent operation between the protocols. Some open source and commercial products have room for improvement to allow such consistent operation to be realised.

There are some new IPv6 'ways of thinking' for our systems and networks administrators. These have been taken on board quite quickly and effectively.

The main tasks we have still to complete are to establish a DHCPv6 service for our client systems, and to enable IPv6 support in the remaining Microsoft applications.

8. REFERENCES

[6DISS] The 6DISS project, http://www.6diss.org

[6DEPLOY] The 6DEPLOY project, http://www.6deploy.org

[6NET] The 6NET project, http://www.6net.org

[6RD] IPv6 Rapid Deployment on IPv4 Infrastructures, IETF Internet Draft, draft-despres-6rd-03, R. Despres, April 2009.

[BIND] ISC BIND, http://www.isc.org/products/BIND

[BROKER] JANET IPv6 Tunnel Broker, http://www.broker.ipv6.ac.uk

[DBEACON] dbeacon, a Multicast beacon, http://fivebits.net/proj/dbeacon

[EDUROAM] eduroam, http://www.eduroam.org

[FWADMIN] Guidelines for Firewall Administrators Regarding MIPv6 Traffic, IETF Internet Draft, draft-ietf-mext-firewall-admin-01, S. Krishnan et al, May 2009.

[FWVEN] Guidelines for Firewall Vendors Regarding MIPv6 Traffic, IETF Internet Draft, draft-ietf-mext-firewall-vendor-01, S. Krishnan et al, May 2009.

[GEANT] The GÉANT network, http://www.geant.net

[JANETv6] JANET IPv6 Technical Guide, GD/JANET/TECH/012, T. Chown, July 2006 : http://www.webarchive.ja.net/services/publications/technical-guides/ipv6-tech-guide-for-web.pdf

[JANETv6MC] IPv6 Multicast on JANET, GD/JANET/TECH/013, S. Venaas, T. Chown, October 2006: http://www.webarchive.ja.net/services/publications/technical-guides/ipv6-multicast-web.pdf

[JRS] JANET Roaming Service, http://www.ja.net/services/authentication-and-authorisation/janet-roaming.html

[NAT] Network Address Translation, IETF RFC 1631, P. Francis, K. Egevang, May 1994.

[NAT-PT] Network Address Translation – Protocol Translation, IETF RFC 2766, G. Tsirtsis, P. Srisuresh, February 2000.

[NAT64] NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers, IETF Internet Draft, draft-ietf-behave-v6v4-xlate-stateful-01, M. Bagnulo et al, July 2009.

[NFSEN] nfsen, http://nfsen.sourceforge.net

[POTAROO] The IPv4 Address Report, http://www.potaroo.net/tools/ipv4/index.html

[RAMOND] RAMOND, http://sourceforge.net/projects/ramond

[RFC1918] Address Allocation for Private Internets, IETF RFC 1918, Y. Rekhter et al, February 1996.

[RFC2080] RIPng for IPv6, IETF RFC 2080, G. Malkin, R. Minnear, January 1997.

[RFC2460] Internet Protocol version 6 (IPv6) Specification, IETF RFC 2460, S. Deering, R. Hinden, December 1998.

[RFC3053] IPv6 Tunnel Broker, IETF RFC 3053, A. Durand et al, January 2001.

[RFC3056] Connection of IPv6 Domains via IPv4 Clouds, IETF RFC 3056, B. Carpenter, K. Moore, February 2001.

[RFC3315] Dynamic Host Configuration Protocol for IPv6 (DHCPv6), IETF RFC 3315, R. Droms, Ed, July 2003.

[RFC3627] Use of /127 Prefix Length Between Routers Considered Harmful, IETF RFC 3627, P. Savola, September 2003.

[RFC3956] Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address, IETF RFC 3956, P. Savola, B. Haberman, November 2004.

[RFC3974] SMTP Operational Experience in Mixed IPv4/IPv6 Environments, IETF RFC 3974, M. Nakamura, J. Hagino, January 2005.

[RFC4038] Application Aspects of IPv6 Transition, IETF RFC 4038, M-K. Shin, Ed, March 2005.

[RFC4193] Unique Local IPv6 Unicast Addresses, IETF RFC 4193, R. Hinden, B. Haberman, October 2005.

[RFC4214] Intra-Site Automatic Tunnel Addressing Protocol (ISATAP), IETF RFC 4214, F. Templin et al, October 2005.

[RFC4380] Teredo: Tunneling IPv6 over UDP through NATs, IETF RFC 4380, C. Huitema, February 2006.

[RFC4862] IPv6 Stateless Address Autoconfiguration, IETF RFC 4862, T. Narten et al, September 2007.

[RFC4864] Local Network Protection for IPv6, IETF RFC 4864, G. Van de Velde et al, May 2007.

[RFC4890] Recommendations for Filtering ICMPv6 Messages in Firewalls, IETF RFC 4890, E. Davies, J. Mohacsi, May 2007.

[RFC4941] Privacy Extensions for Stateless Address Autoconfiguartion in IPv6, IETF RFC 4941, T. Narten et al, September 2007.

[RFC5006] IPv6 Router Advertisement Option for DNS Configuration, IETF RFC 5006, J. Jeong, Ed, September 2007.

[ROGUE-RA] draft-chown-v6ops-rogue-ra-03

[SSMPING] ssmping, http://www.venaas.no/multicast/ssmping

[TRAIN] UKERNA IPv6 Training, http://www.ipv6.org.uk/workshop

[VIDEOLAN] VideoLAN, http://www.videolan.org/