IPv6 Enhanced Council

IPv6 for Universities

Ian Hallissy & Colin Donohoe Technological
University of the Shannon TUS Midlands,
Athlone Campus, Ireland

Contents

Executive Summary

Without Higher Education computer networks, the internet may not have developed into the versatile tool it has become today. Sir Tim Berners-Lee is known as the inventor of the World Wide Web while at Cern in the 1980s developing [i.1] a system combining hypertext with the internet that would allow researchers globally share information. To continue and future-proof the benefits of knowledge sharing globally it is vital Higher Education institutes and Universities now embrace migration to IPv6.
Globally steady progress has been made in IPv6 adoption in the last five years. However, there is a clear anomaly with adoption in the Universities where adoption in Ireland's National Research and Educational Network (NREN) HEAnet at 7% is accepted as typical for this key sector. Tools are not available to give overall NREN adoption levels and only available on a per client basis. It is well-accepted that technological innovation is crucial for Higher Education. This was highlighted during the recent Covid pandemic where a pivot to online delivery was required. Still, IPv6 adoption has not been embraced. IPv6 is an essential technology for the future growth and development of the global internet. As an example, the vast IPv6 address space will play a foundational role in the provisioning of billions of IoT devices. The 2022 Google statistics report adoption of IPv6 at 43% [i.2]. While adoption is growing, there is significant variation from country to country and industry to industry.

Telco ISPs, mobile cellular carriers, and cloud and Content Delivery Networks (CDN) have embraced and boosted IPv6 adoption [i.3]. They have provided scale, financial and technical benefits to their businesses. These sectors have seamlessly deployed IPv6 with little or no impact on or input from their customers. Surprisingly, the enterprise sector has seen the slowest adopters and reluctant in deploying IPV6. Enterprise networks tend to be slow-moving deployments with legacy hardware and applications. In addition, there is a perception that a large financial cost in deploying IPv6 exists. However, the current Network hardware is IPv6 compliant. Furthermore, there are misconceptions about IPv6 security which has also delayed deployments.

In terms of the university sector, the EU-funded 6NET project has already enabled deployment in Géant and the EU NRENs back in 2002-2004 [i.4] [i.5]. However, very few EU universities have deployed IPv6. So, the Universities and Higher Educational Institutes were in a privileged position in that their NREN could provide them with their IPv6 prefix and the live IPv6 infrastructure to access the global IPv6 internet while also providing advice and encouragement to deploy. In terms of deployment of IPv6, it demands a fundamentally different approach in terms of architecture and new design e.g., on a Local area network, IPv6 adds new features such as Neighbour Discovery Protocol (NDP) utilising multicast, replacing classic ARP and broadcast [i.6]. This delivers simpler networks with fewer VLANs and more hosts per VLAN. The change requires in-depth network planning and extensive staff training.

This has been potentially one of the reasons towards low adoption rates ranging 5% to 10% in this sector. In this context, this document outlines the TUS: Athlone campus's IPv6 deployment journey towards IPv6 only network.
It is motivated by a desire to encourage other universities to adopt this next-generation protocol. There are several advantages and reasons for the deployment of IPv6 within the university sector. It can inform and ensure valid adoption of IPv6 from a teaching and training perspective, and IPv6 can be a key tool enabling research and innovation in real-life scenarios. Training and Education in IPV6 are central to Networking professionals having the knowledge and confidence to utilise the new protocol.

In terms of migration towards IPv6 in TUS: Athlone, a phased approach was employed, which included the coexistence of both IPv4 and IPv6 in dual-stack deployment. This aligns with best practices. Each of the steps taken and tasks completed in Athlone's five phases of deployment is described here. The result is that TUS: Athlone now boasts an IPV6-only wireless infrastructure. These phases included: address planning; public-facing services; wireless networks; wired networks; and finally, IPv6-only wireless. In the remainder of this document, it provides an overview of the steps taken to achieve each of these phases, and as such this document act as an instructional aid for networking professionals.

Introduction

The Internet protocol (IP) [i.7] provides internet-enabled devices with unique addresses allowing them to communicate with each other. These addresses are used as a crucial aspect towards delivering the future of the global internet. IP version 4 or IPv4 was the first protocol deployed in the early 1980s. The Internet Engineering Task force IETF [i.8] released the IPv6 draft standard in 1998. It was subsequently ratified as an Internet Standard on 14 July 2017 [i.9] as the successor to IPv4. A key feature (one of several changes from IPv4) was the expanded addressing capabilities with IPv6. The IPv4 32-bit address with 4.1 billion addresses was upgraded to the new 128-bit IPV6 with 340 trillion, trillion, trillion IP addresses 3.4 X 10^{38}. The IPv4 address space is depreciated. Other features beyond the address space include multicast replacing broadcast and Stateless address autoconfiguration (SLAAC) providing new ways of deploying IP networks.

At the Irish NREN HEAnet [i.10] annual conference in 2010, they advised and encouraged Educational and Research clients to deploy IPv6. This was motivated by several reasons, but primarily to ensure education and research clients avoided complications and the risks associated with IPv4 depletion. The NREN HEAnet provided each of the University and Higher Educational clients with the opportunity to take advantage of being a fully enabled IPv6 infrastructure. The network team in Athlone embraced this technical challenge regarding how to add this new protocol and to future-proof its network. For the team, this was an exciting challenge with significant research and investigation required to understand the new protocol and philosophy while undertaking a phased approach towards its integration.

In this document, the different phases of integrating IPv6 in Athlone over ten years are presented. These include the initial task of designing the addressing plan right through to the delivery of a unique IPv6-only WI-FI 6 [i.11] wireless infrastructure in 2020. The overarching objective was to provide an improved connectivity experience to university staff and students whether on-site or remotely. IPv6 delivers this by returning to the original internet concept of endpoint-to-endpoint connectivity with globally unique IPv6 addresses.

IPv6 positively affects web performance via its simpler IP header (reduced processing overhead) and improved router performance among other features. By producing this paper, this document aims to share Athlone's experiences and assist enterprises in their decisions towards the adoption of IPv6. IPv6 is acknowledged as the only viable option for the future of the Internet, and it is only a matter of when and not if IPv6 will become the most used protocol. This protocol is complex and different from IPv4, so a phased planned deployment is important to avoid a rushed deployment.

1 Scope

The present document is to provide a use case of a phased deployment of IPv6 in a University enterprise network. It aims to show how a medium-sized University network infrastructure can migrate to IPv6 in a phased dual-stack approach. Explaining how these older legacies IPv4 networks are already segmented by VLAN and how it is possible to identify service types that are ready to utilise a dual-stack approach initially. This provides a way of gaining experience with the new protocol and realising the benefits it can bring in simpler scalable networks. The wireless network service where browser-based access to the internet is the main requirement makes it an ideal place to start. It also shows that IPv4 services (such as voice in Athlone's case) that are not IPv6-ready can continue to function as before until the inevitable IPv6 service becomes available.

The document will outline the five phases undertaken.

- Address Plan
- Enabling Public facing services
- Wireless network
- Wired Network
- IPv6-only Wireless network

The document is written as an instructional aid with each phase broken down into the operational steps required to achieve IPv6 connectivity. It aims to ease the operational process and is not viewed as a teaching aid.

The subject area of IPv6 is vast but a work breakdown structure identifying key tasks at each phase provides all the information to deliver a successful migration to IPv6.

The key takeaway is to get started on the IPv6 journey as a phased planned project as it is preferable to a rushed deployment.

2 References

2.1 Normative References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

The following referenced documents are necessary for the application of the present document.

2.2 Informative References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the referenced document (including any amendments) applies.

The following referenced documents are necessary for the application of the present document.

[i.1]	Tim Berners-Lee: “The birth of the web”
Note	Available at https://home.cern/science/computing/birth-web
[i.2]	Google.com: “Ipv6 – Google”
NOTE	Available at https://www.google.com/intl/en/ipv6/statistics.html
[i.3]	Akamai.com: “IPv6 – Akamai”
NOTE	Available at https://www.akamai.com/internet-station/cyber-attacks/state-of-the-internet-report/ipv6-adoption-visualization
[i.4]	Bernard Tuy “6net,Empowering the Internet generation” (2001)
Note	Available at https://meetings.ripe.net/ripe-40/presentations/6net.pdf

[i.5]	Geant.org: "IPv6 – Geant"
NOTE	Available at https://geant3plus.archive.geant.net/Pages/Network/IPv6.html
[i.6]	IETF RFC 4861 (2007): "Neighbor Discovery for IP Version 6 (IPv6)".
Note	Available at: https://www.rfc-editor.org/rfc/rfc4861
[i.7]	IETF RFC 791 (1981) "Internet Protcol"
Note	Available at https://www.rfc-editor.org/rfc/rfc791
[i.8]	Ietf.org: "IETF – Internet Engineering Task Force"
Note	Available at: https://www.ietf.org/
[i.9]	IETF RFC 8200 (2017): "Internet Protocol, Version 6 (IPv6) Specification".
Note	Available at https://www.rfc-editor.org/rfc/rfc8200
[i.10]	Heanet.ie: "HEAnet – Higher Educational Authority Network"
Note	Available at: http://www.heanet.ie
[i.11]	Wi-fi.org: "Wi-fi alliance Wi-fi 6"
Note	Available at: https://www.wi-fi.org/news-events/newsroom/wi-fi-alliance-introduces-wi-fi-6
[i.12]	IETF RFC 4862 (2007): "IPv6 Stateless Address Autoconfiguration".
Note	Available at https://www.rfc-editor.org/rfc/rfc4862
[i.13]	IETF RFC 2663 (1999): "IP Network Address Translation (NAT) Terminology and considerations
Note	Available at: https://www.rfc-editor.org/rfc/rfc2663
[i.14]	home.cern: "LHC – Large Hadron Collider"
Note	Available at: https://home.cern/science/accelerators/large-hadron-collider
[i.15]	Wlcg.web.cern.ch: "Wlcg – Worldwide Large Hadron Collider Grid"
Note	Available at: https://wlcg.web.cern.ch/
[i.16]	IETF RFC 1918 (1996): "Address allocation for private networks"
Note	Available at: https://www.rfc-editor.org/rfc/rfc1918
[i.17]	Facebook.com: " Facebook – Legacy support on IPv6-only infra'
Note	Available at: https://engineering.fb.com/2017/01/17/production-engineering/legacy-support-on-ipv6-only-infra/
[i.18]	LinkedIn.com: "LinkedIn -IPv6 inside LinkedIn part 2"
Note	Available at:https://engineering.linkedin.com/blog/2016/08/ipv6-inside-linkedin-part-ii--back-to-the-future

[i.19] IETF RFC 922 (1984): 'Broadcasting internet datagrams in the presence of subnets"

Note Available at: https://www.rfc-editor.org/rfc/rfc922.html

[i.20] IETF RFC 4443 (2006): "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification".

Note Available at: https://www.rfc-editor.org/rfc/rfc4443

[i.21] Novell.com: "Novell – IPX protocol"

Note Available at: https://www.novell.com/documentation/nw6p/pdfdoc/ipx_enu/ipx_enu.pdf

[i.22] Apple.com: "Apple – AppleTalk protocol"

Note Available at: https://developer.apple.com/library/archive/documentation/mac/pdf/Networking/Introduction_to_AppleTalk.pdf

[i.23] IETF RFC 8421 (2018): "Guidelines for multihomed and IPv4 AND IPV6 Dual-stack Interactive Connectivity Establishment"

Note Available at: https://www.rfc-editor.org/rfc/rfc8421.html

[i.24] IETF RFC 4380 (2006): "Tunnelling IPv6 over UDP through Network Address translations (NATS)"

Note Available at: https://www.rfc-editor.org/rfc/rfc4380

[i.25] IETF RFC 6147 (2011): "DNS64"

Note Available at: https://www.rfc-editor.org/rfc/rfc6147

[i.26] IETF RFC 6146 (2011): "Stateful NAT64"

Note Available at: https://www.rfc-editor.org/rfc/rfc6146

[i.27] IETF RFC 1035 (1987): "Domain Names Implementation and Specifications"

Note Available at: https://www.ietf.org/rfc/rfc1035.txt

[i.28] IETF RFC 2674 (1999): "Definition of Managed objects for bridged classes, Multicast filtering and Virtual LAN extensions"

Note Available at: https://www.ietf.org/rfc/rfc2674.txt

[i.29] Iso.org: "Open System Interconnection -OSI"

Note Available at: https://www.iso.org/obp/ui/#iso:std:iso-iec:7498:-1:ed-1:v2:en

[i.30] IETF RFC 8519 (2019) "YANG Data model for Access Control List -ACL"

Note Available at: https://datatracker.ietf.org/doc/rfc8519/

[i.31] IETF RFC 3584 (2003): "Coexistence between Version1, Version2 and Version 3 of Internet standard Network Management Protocol"

Note Available at: https://datatracker.ietf.org/doc/rfc3584/

[i.32]	IETF RFC 4251 (2006): “The Secure Shell [SSH] Protocol Architecture”
Note	Available at: https://datatracker.ietf.org/doc/rfc4251/
[i.33]	Ripe.net: “Preparing an IPv6 address plan”
Note	Available at: https://www.ripe.net/support/training/material/IPv6-for-LIRs-Training-Course/Preparing-an-IPv6-Addressing-Plan.pdf/view
[i.34]	Isc.org: “ BIND – Berkely Internet Name Domain”
Note	Available at: https://www.isc.org/about/
[i.35]	IETF RFC 2888 (2000): “Secure Remote Access with L2TP”
Note	Available at: https://www.ietf.org/rfc/rfc2888.txt
[i.36]	IETF RFC 8415 (2018): "Dynamic Configuration Protocol for IPv6 (DHCPv6)".
Note	Available at: https://www.rfc-editor.org/rfc/rfc8415
[i.37]	Eduroam.org: “ Eduroam – Educational Roaming”
Note	Available at: https://eduroam.org/
[i.38]	IETF RFC 6105 (2011) “ Router Address Guard”
Note	Available at: https://datatracker.ietf.org/doc/html/rfc6105
[i.39]	Microsoft.com “ Microsoft Windows – Dynamic DNS”
Note	Available at: https://learn.microsoft.com/en-us/windows/win32/ad/active-directory-servers-and-dynamic-dns
[i.40]	IETF RFC 2865 (2000) “ Remote Authentication Dial In User Services (RADIUS)”
Note	Available at: https://www.ietf.org/rfc/rfc2865.txt
[i.41]	IETF RFC 4944 (2016) “ IPv6 over Low-Power Wireless Personal Area Network (6LoWPAN)
Note	Available at: https://datatracker.ietf.org/doc/html/rfc8025
[I.42]	IETF RFC 9030 (2021): “An architecture for IPv6 over the Time-Slotted Channel Hopping Mode of IEEE 802.15.4 (6Tisch)”
Note	Available at: https://datatracker.ietf.org/doc/rfc9030/

3 Abbreviations

6LOWPAN	IPv6 over Low-Power Wireless Personal Area Networks
6BONE	IPv6 testbed network established in 1996
802.11AX	Licensed by Wi-Fi Alliance as WI-FI 6
A	IPv4 DNS address record
AAAA	IPv6 DNS address record
ACL	Access control list
AD	Active Directory
AP	Access Point
AR	Augmented reality
ARP	Address resolution protocol
BIND	Berkley Internet Name Daemon, Domain name internet software system
CAPEX	Capital Expenditure
CDN	Content delivery network
DHCP	Dynamic host configuration protocol
DHCPV4	Dynamic host configuration protocol version 4
DHCPV6	Dynamic host configuration protocol version 6
DMZ	Demilitarised zone
DNS	Domain name system
DNS64	Domain name system IPv6 to IPv4
DUID	Dynamic unique identifier
EDUROAM	Educational roaming wireless service
EU	European union
EUI-64	Extended unique identifier – 64bit
GEANT	European network for Research and Education
HTTP	Hypertext transfer protocol
HTTPS	Secure Hypertext transfer protocol
HYPERV	Hyper Visor
IETF	Internet Engineering Task Force

IoT	Internet of things
IIoT	Industrial Internet of things
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IPX	Internetwork Packet Exchange
IRF	Intelligent Resilient Framework
ISP	Internet service provider
LHC	Large Hydron Collider
MIS	Management Information systems
MPLS	Multi-Protocol
NAT	Network address translation
NAT64	Network address translation IPv6 to IPv4
NTP	Network time protocol
NMS	Network Management system
NREN	National Research and Educational Network
OSPFv3	Open Shortest Path First version 3
PBX	Private Branch Exchange
POE	Power Over Ethernet
POP	Point Of Presence
PTR	Pointer record
P2P	Peer to peer
RAGUARD	Router Advertisement Guard
RADIUS	Remote Authentication Dial-in User Services
RDNSS	Recursive Domain Name System Service
RIPE	Regional Internet
SLAAC	Stateless Automatic address configuration
SNMP	Simple Network Management protocol
SNMPV3	Simple Network Management protocol version 3
SSID	Service Set Identifier

TELCO	Telecommunication company
VDOM	Virtual domain
VLAN	Virtual Local Area Network
VLE	Virtual Learning Environment
VR	Virtual reality
VRF	Virtual Routing and Forwarding
WAN	Wide Area Network
WI-FI 6	WI-FI version 6

4 Motivation to Deploy IPv6 in Academia and LAN Design Implications

In this section, the motivation for the deployment of IPv6 on academic networks is presented. Furthermore, the design considerations in the context of both coexistence and transitioning from IPv4 to IPv6 is presented. While IPv6 was originally developed by the IETF in response to IPv4 depletion the technology is now viewed as a driver for innovative business models to transform how enterprises operate.

4.1 Why IPv6 now in Academia

IPv6 is the next-generation version of the Internet protocol which is the worldwide system of computer networks allowing information sharing. It was developed by the IETF to deal with the address exhaustion of the original protocol IPv4. This is achieved by increasing the address space from 2^{32} (4.1 billion) to 2^{128}(340 trillion trillion). It provides other technical benefits including:

- Auto-configuration – Stateless Address Auto-configuration (SLAAC) [i.12]
- Simpler header format
- Simplified more efficient routing.
- True Quality of Service
- Better Multicast routing
- Built-in authentication and privacy support IPsec
- No more NAT (Network address translation) [i.13]

4.1.1 Replace NAT with Original Endpoint to Endpoint Connectivity

A primary motivator for migrating to IPv6 is to remove Network Address Translation (NAT) and return to the original Internet concept of endpoint-to-endpoint connectivity using Global unique IPv6 addresses. This key principle of each endpoint possessing a globally unique address allows for simple & scalable networks with the obvious benefit of designing secure solutions based on the IP address identity.

4.1.2 Athlone's VLE Performance Enhanced by IPv6

An example of improved performance for Athlone is the Virtual Learning Environment (VLE) Moodle. This business mission-critical service is hosted in Amazon Web services and dual-stacked with IPv4 and IPv6. This provides that all connections from staff and students from the Athlone campus are native IPv6 connections end to end avoiding Network address translation (NAT). This delivers improved network performance and monitoring capabilities as well as enhanced analytics.

Remote connectivity for staff and students particularly during the recent pandemic also benefited from this capability where a number of Ireland's ISP's provide IPv6 services.

4.1.3 Global Research and High-Performance Computing

A Global example of IPv6 in High-Performance Computing in Educational Research space is the Large Hadron Collider (LHC) [i.14] incorporating the Worldwide LHC Grid (WLCG) [i.15]. This iconic project is a global collaboration of around 170 computing centres in 40 countries linking grid infrastructures mainly over IPv6 to analyse 200 Petabytes of data expected each year from the Large Hadron Collider at Cern. The large address space and routing enhancements are the features of IPv6 utilised for this impressive project.

4.1.4 CAPEX Versus OPEX and Training Savings

Capital Expenditure (CAPEX) investment in IPv6 based network assets delivers investment protection which removes the need and reduces the capital and operating costs (OPEX) of deploying and maintaining NAT devices. Deploying IPv6 in a phased basis allows for incremental training steps in gaining a full appreciation and understanding of this new protocol. This is preferable to expensive commercial training when a rushed deployment is required.

4.2 Design Considerations for the Coexistence of IPv4 and IPv6

Athlone's wired infrastructure of over 4000 nodes is a dual-stack system. IPv4 and IPv6 coexist and operate alongside each other. This document will outline the reasons for this and the ultimate journey to IPV6 only and the objective of turning off the legacy IPv4 protocol. Network address translation (NAT) was the patch solution to allow a single public IPv4 address to masquerade for large private address spaces and to preserve the limited IPv4 addresses [i.16]. However, the NAT layer can make response times slower and where multiple NATs occur network services are unworkable.

IPv6 solves these problems by returning to the original internet concept of unique endpoint-to-endpoint connectivity due to the extended address space.

Turning off the legacy IPv4 protocol is the goal as it will reduce complexity and cost but as several legacy systems are only compatible with IPv4 this is an ongoing process. Big tech such as Google, Amazon, Facebook, Microsoft, Apple, and Content Delivery networks like Akamai and Cloudflare are all accessible over IPv6 and some like Facebook [i.17] and LinkedIn [i.18] are IPv6 only in their internal networks. All current modern desktop operating systems including Windows 10, Apple iOS, and Linux support IPv6 and in fact will connect over IPv6 first in a dual-stack environment. However, it is in the mobile cellular 5G area that IPv6 adoption is growing IPv6, and this is a reason to have your Technology enhanced learning services available natively over IPv6.

At the start of this process in 2010, a dual-stack design was the only viable option available, and the design requirement was how to add IPv6 capability to a live network. Athlone utilised a segmented design which comprised three layers, external services layer, firewall layer and internal layer.

- At external services, NREN HEAnet used MPLS to provide multiple services over a single trunk.
- At the firewall layer, virtual firewall instances were utilised to segment services required to be isolated for security or compliance purposes.
- At the internal Layer for both Wired and Wireless Cores the use of VRFs (Layer 3 routed instances) further segmented services based on design requirements.

4.2.1 IPv6 Design Change: Multicast Replaces Broadcast

In IPv6, there is no longer any broadcast. This allows for fundamental design changes in how to layer two broadcast domains are utilised [i.19]. Instead, network discovery functions are built upon multicast groups and ICMPv6 [i.20]. This potentially allows for a larger number of IPv6 address clients per VLAN. The convention is to break IPV6 subnets or prefixes into \64. Therefore 2^{64} equals 274,877,906,944 clients. With larger IPv6 address scopes allowing more clients per VLAN this would potentially reduce the number of VLANs required. However, as initially utilising a dual-stack design there is a need to be cognizant of the number of IPv4 clients per VLAN to avoid the potential risk of broadcast storms. Thus, adding a /64 IPv6 subnet to match the current IPv4 one is implemented. The current IPv4 private address class c subnets have 254 clients as per convention. While this may seem like an inefficient use of IPv6 addresses at this phase of transition and coexistence it is required. In an IPv6-only design, smaller VLAN numbers with larger address space will be the norm and deliver simpler scalable IP networks.

4.2.2 Consideration of Current IPv4 Network

The transition between IPv4 and IPv6 which started with a draft protocol in 1998 is taking a long time as they are two separate protocols. IPv6 is not backwards compatible with IPv4 due to the fundamental differences in the IPv4 and IPv6 headers. IPv4 hosts, and network devices cannot communicate directly with IPv6. One cannot send an IPv4 packet to an IPv6-only device or send an IPV6 packet to an IPv4-only device. The transition to IPv6 requires the integration and co-existence of interaction between IPv4 and IPv6 networks. While the IPv6 transition means migration from one protocol to another the deployment is more of a coexistence as the IPv4 internet is going to be around for some time yet. This concept of multi-protocol is not new and is not a concern. As IP made its way onto Local area networks in the early 1990s to join the then incumbent protocols of Novell Netware IPX [i.21] and Apple's AppleTalk [i.22] among others. It is most likely as all those protocols disappeared from network card configurations over time without much notice as everyone settled on TCP/IP IPv6 will eventually replace IPv4.

There are three strategies to be followed for the transition from IPv4 to IPv6

- Dual Stack [i.23] – Running both protocols IPv4 and IPv6 on all hosts and devices and communicating with each other using either protocol. Where a dual-stack host communicates with a dual-stack destination whether IPV6 or IPv4 is used is down to the application.
- Tunnelling – Transporting IPv6 through an IPv4 network transparently with encapsulation techniques like Teredo [i.24]. These are not used very much now and as IPv6-only networking becomes more available in ISP IPv4 as a service in IPv6 tunnels is more common.
- Translation Technique at a boundary router between an IPv4 and IPv6 network a translation process maps an IPv4 address to an IPv6 address or vice versa. DNS64 [i.25] and NAT64 [i.26] were the translation technique utilised by Athlone on IPv6 only wireless infrastructure.

The selection of the mechanisms utilised to transition is dependent on the status and design of the Enterprise LAN. Ten years ago in Academic Institutions, it was also common practice to have separate wired and wireless physical infrastructure air-gapped due to security concerns with wireless infrastructure only providing internet access. This also helped the process of a phased migration by being able to isolate components of infrastructure and upgrade while lowering the risk of issues.

4.2.3 Domain Name System

DNS [i.27] is a critical service in IPv6 and is fundamental to the dual-stack approach. Dual-stack transition is based on DNS where the dual stack client device queries the name of a destination node and DNS returns an IPv4 address (a DNS A record) it forwards IPv4 packets. If DNS responds with an IPv6 address (a DNS AAAA record) it forwards an IPv6 packet. This provides the option to connect to IPv6 nodes where available but will have the fallback of the IPv4 option. The client device operating system will prefer IPv6 over IPv4.

4.2.4 VLAN's Design and Subnetting

This section will outline layer two VLAN [i.28] design at the three service layers and how adding IPv6 prefixes accomplishes dual stacking with IPv4 subnets at layer three of the OSI model [i.29].

4.2.4.1 External Services

At the External services layer, Athlone's ISP and NREN HEAnet provides a 10Gbps interface at both diverse Athlone POP locations. Each 10Gbps interface is subdivided into sub-interfaces. Each service is tagged with a VLAN to segregate the traffic over the physical interface and breakout as follows.

1. General internet.
2. Management Information Systems (MIS) applications.
3. Externally hosted Voice solution.
4. Commercial internet.

At the firewall layer, three virtual firewall instances are created to be the bridge between Internal and External networks and provide the necessary security mechanisms.

1. The primary root instance has external access to the general internet and DMZ and internal access to the wired LAN and wireless infrastructure.
2. A second virtual instance delivers an externally hosted Voice PBX system over IPv4.

Finally, commercial internet connectivity is provided for non-educational/research clients also IPv4 only.

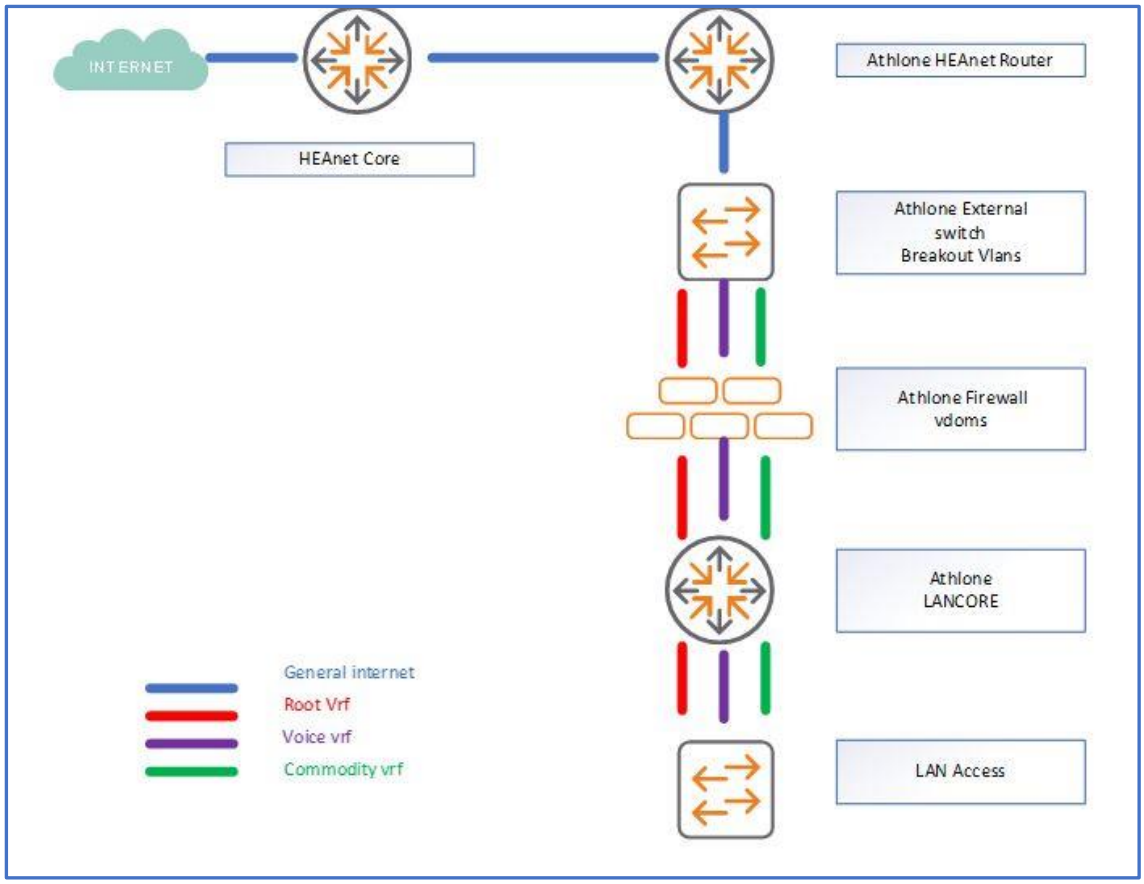

Figure 1 External Services

4.2.4.2 Internal Services

This section outlines the local area network VLAN design and Access control posture for Cyber Security. VLANs were logically divided between staff and students at the VLAN Access layer. At the server farm layer - Staff and Student Server VLANs were configured. Voice access VLANs were deployed campus-wide at each wiring closet. Student access VLANs are permitted only access to Student server VLANs utilising an Access control policy (ACL) [i.30] Staff access VLANs are permitted access to both Staff and student server VLANs for shared applications. Access between staff and student access VLANs was not permitted and enforced with Access control lists on core switches. This delivered a secure micro-segmented network on IPv4. This structure allowed a simple task of mapping IPv6 addressing to VLAN and also applying IPv6 access control.

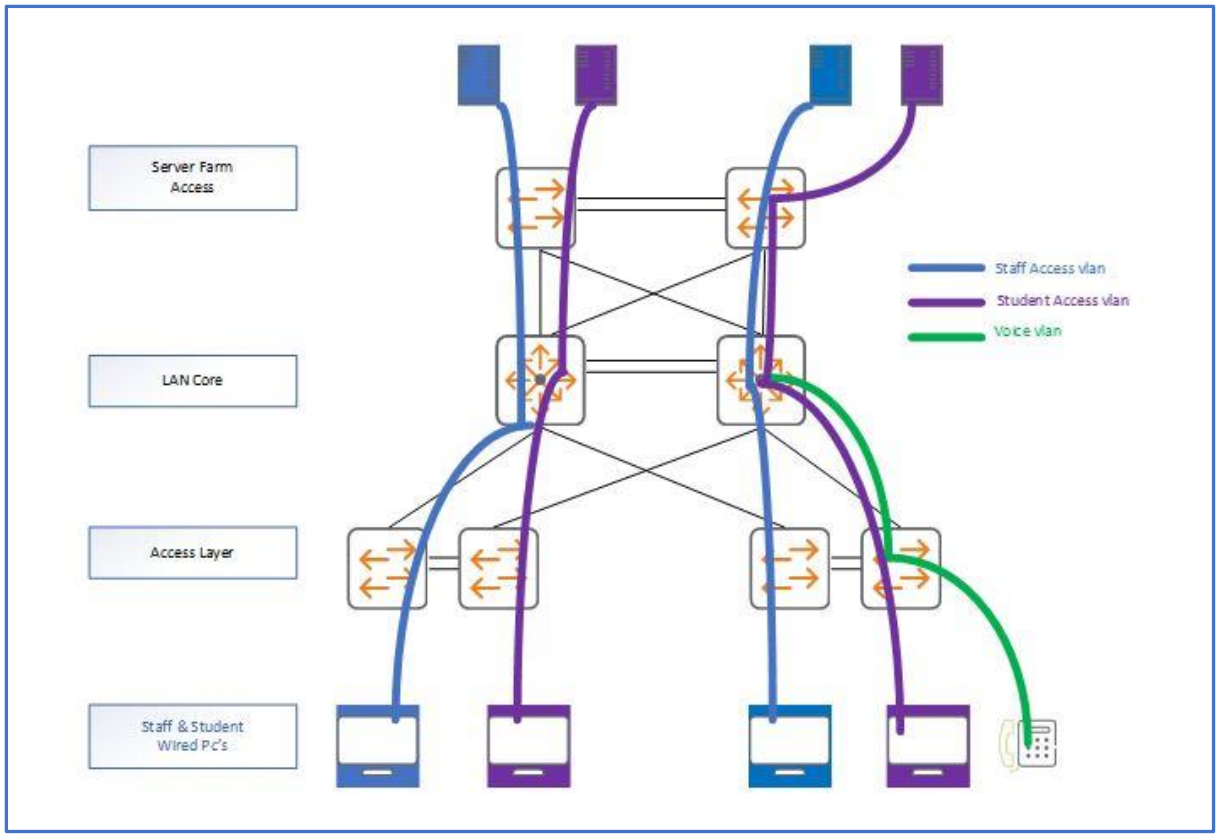

Figure 2: Internal Services VLAN Design

4.3 Network Infrastructure Hardware Readiness

A hardware and more crucially a software audit of all network equipment is the first step in the deployment of IPv6. Ten years ago, when Athlone started its journey IPv6 features were not as available as are currently the situation. IPv6 is a protocol so compatibility with Ipv6 networking is mainly a software or firmware issue. Switch functionality is performed at the mac address any switch model will be able to switch IPv6 packets in hardware. Crucially Athlone's Firewall and LAN Core and access switches had IPv6 routing capabilities. When IPv6 functionality is enabled on the firewall there is a performance hit due to extra processing and memory requirements of two protocols. The layer three switches also must be capable of dealing with the new method of dealing with the mac layer and neighbour discovery features Routing functionality is where IPv6 support comes into play. Like many University networks, Athlone's ageing switch hardware consisted of three different hardware vendors, but as open standards are utilised this did not present a major problem. However, there were limitations in IPv6 functionality on the older hardware.

4.4 Application Readiness

Like network hardware, Enterprise application software needs to support IPv6. In this Higher Education environment where a wide array of software applications is used this is a crucial assessment that needs to be carried out. Microsoft Office 365 email and collaboration tools are cloud-delivered browser-based applications. Faculty software applications are generally locally installed with network-based licencing. Athlone did not encounter any legacy application that did not continue to function in a dual stack environment.

4.5 Network Management Software

Athlone utilises an open standards Network Management Software (NMS) system that integrates fault management, element configuration and network monitoring from a central vantage point. Devices are managed securely with authentication and encryption using SNMPv3 [i.31] and SSH [i.32] access. The wired network and wireless are both managed by the same NMS. A Windows Server hosting NMS is dual stacked allowing network switches to be managed over IPv6 and over IPv4. Currently over 90% of all switches are managed over IPv6 only.

4.6 HEAnet Context

HEAnet is the Irish Higher Educational Authority network NREN (National Research and educational network. It connects over one million users to the internet, linking them to each other and the global internet via its secure high-capacity educational network. The national backbone consists of over 2,500 Km of fibre and interconnects all Higher Education institutes in Ireland as well as providing backhaul for over 4,000 primary and post-primary schools nationwide. The ‘HEAnet Network’ is made up of three distinct layers: Optical, IP and MPLS (Multi-Protocol-Label-Switching) [i.34]

5 Address Planning

IPv6 address planning is the critical function in the initial phase of the overall process of designing and implementing an IPv6 deployment project. With IPv4 address planning the focus due to the smaller address space is around utilising every address sparingly but this is not the case with IPv6. It is important to avoid IPv4 thinking so as not to shoehorn IPv4 limitations into the vast IPv6 space.

5.1 IPv4 & IPv6 Addressing Plan

VLANs and Subnets perform a similar function in the segmentation of IP networks with VLans at layer 2 and Subnets at layer 3 of the OSI model and it is common to create a one-to-one relationship. Segmentation improves security by preventing attacks from spreading across a network and reducing the size of the broadcast domain. Athlone's design was to segment managed student access clients from accessing staff resources. Each campus location network stack had a Staff, Student VLAN deployed with a class C private IPv4 address 192.168.x.y supplied via the DHCP server of Windows Active directory infrastructure. Layer three ACLs restrict student access VLAN to student server VLAN at the Server farm switch infrastructure. Staff Access VLANs have ACL-controlled access to both Staff and Student server VLans. This results in many IPv4 subnets, and the numbering convention has evolved with Staff identified with an even number and Students with an odd number. The one-to-one relationship between the subnet and VLAN is maintained with the VLAN number matching the subnet number of the network i.e., Vlan22 corresponds to network 192.168.22.x. In contrast, IPv6 addresses are 128 bits long, so unlike when creating an IPv4 addressing plan, with IPv6 there is no real concern about the number of IP addresses available within a subnet. There are just so many addresses. The IETF recommendation is to use /64 prefixes when creating IPv6 subnets. This is the smallest subnet that you can use if auto-configuration is required. The number of IP addresses you get with a /64 is 2^{64} which is 18,446,744,073,709,551,616 IP addresses.

5.2 NREN ISP

HEAnet is Ireland’s National Education and Research Network provides high-speed internet connectivity and ICT shared services to all levels of the Irish education sector and provided Athlone with an IPv6 address allocation which was a /48 subnet i.e. 2001:770:50::/48. This meant that 65536 were available subnets. The task in IPv6 address planning is to determine the best way to break this up and utilise it on the network.

There were also several aspects to consider when creating the IPv6 addressing plan for Athlone:

1. As dual stacking was being used, had to be able to align the IPv6 subnets with the IPv4 subnets that were already in place.
2. The IPv6 address was going to be new to all technical staff in Athlone so there had to be a way of making the IPv6 addressing as easy to become familiar with as possible.
3. The addressing plan had to provide scope for scalability with the possibility of needing new subnets in the future.

After carrying out research and looking for best practices and recommendations on how best to allocate the IPv6 subnets; a document that provided very good guidance and the one that the Athlone IPv6 addressing plan is based on is a paper published by RIPE [i.33] called "Preparing an IPv6 Address Plan". This document explains some of the different approaches that can be used when creating the basic structure for your IPv6 address plan.

Use Type First

One of the approaches outlined is to assign the addresses first by use type and then by location. When the user type comes first and is the primary subnet, it makes it much easier to implement a security policy as most firewall policies are done based on the type of use and not the location.

For Example:

2001:770:50: [Primary Use Type] [Sub Use Type] [Primary Location] [Sub Location]::/64

In this example, there are 4 bits assigned to the Primary Use Type, 4 bits assigned to the Sub Use Type and then 4 bits for the Primary Location followed by 4 bits for the Sub Location. This means that we can have 2^4 or 16 Primary Use Types with each of those having up to 2^4 or 16 Sub Use Types that can then be allocated in up to 16 Primary Locations with up to 16 or 2^4 Sub locations.

An example of these would be:

Primary use types - Student, Staff, Research, Wireless etc.

Sub-use types – Servers, Desktops, Access points etc.

Primary Locations – Engineering Building, IT Centre, East Campus etc.

Sub Locations – locations of comms rooms or comms cabinets

An example of an IPv6 subnet for student desktop pcs in the Nursing Department located on the East Campus might be:

2001:770:50: **[Primary Use Type] [Sub Use Type] [Primary Location] [Sub Location]**::/64

2001:770:50:**1234**::/64

- Primary Use Type is 1 – this primary use type is Student
- Sub Use Type is 2 – this sub use type is Desktops
- Primary Location 3 – this primary location is for East Campus
- Sub Location 4 - this sub-location is for Nursing

5.3 External DNS for IPv6

As IPv6 adoption grows and with the complexity of IPv6 addresses, DNS is becoming even more important as a mechanism to help users to reach the most appropriate IP addresses to reach their intended destination. An IP address forward lookup refers to the process of retrieving an IP address for an Internet domain name or in the case of a reverse lookup, retrieving an Internet domain name for an IP address. For example, www.ait.ie resolves to the IPv4 address 193.1.30.24. To enable IPv6 the requirement was to set up an external address to get our web server's addresses to resolve to both an IPv4 and an IPv6 addresses. In AIT there is a split DNS design like many other organisations. This design means that you do not need to share the DNS information for the internal network with the Internet, both for security reasons and as private IP addresses, are not routable on the Internet. In Athlone, there is an Internal DNS for all internal DNS requirements and an external DNS configured using Berkley Internet Name Domain (BIND) [i.34] where to publish DNS information for the DMZ-hosted nodes or public-facing web servers.

There are several ways of configuring BIND to enable IPv6. One of the options is to have a separate zone file for IPv4 and IPv6. But in Athlone, the same AIT zone file is used for both IPv4 and IPv6 DNS records. This means that once a web server was Dual Stacked, all that was needed was to add an AAAA record with the ipv6 address of the server when adding the A record for the IPv4 address.

To get the External DNS resolving to both IPv4 and IPv6 addresses there were several BIND configuration changes required.

These included the following:

- the named.conf configuration file to include ipv6 reverse lookup zones and ipv6 loopback.
- the ait.ie. zone file was updated to include an AAAA record for each of the web servers that already existed with an IPv4 A record.
- The addition of an IPv6 reverse lookup zone file with PTR records for each of the web servers.

There was also a requirement for HEAnet to delegate IPv6 lookup to Athlone nameservers.

6 Public Facing Service

Making the public face of the institutes network accessible over IPv6 is a key priority and a major milestone when achieved s it provides global internet connectivity over IPv6. Providing access to users who are on IPv6 only networks is becoming increasingly important for the University sector.

6.1 IPv6 Transition and Coexistence with IPv4

IPv6 transition includes the integration, co-existence, and interoperation between IPv4 and IPv6 networks and devices. As outlined in section 4 Athlone chose the Dual stack transition method as being the only viable option at this time in 2012. There were a lot of gaps in IPv6 support in most infrastructure vendors which allowed ability to progress with IPv6 deployment with limitations Having audited the components Athlone were confident all public-facing infrastructure could support a dual-stack design. This was the first experience of adding a 128-bit v6 address to interfaces of public Internet-facing services. The initial requirement was to confirm IPv6 layer three connectivity with HEAnet's External routers and Athlone Institute of Technologies WAN (Wide Area Network) infrastructure. A cautious approach was adopted due to a lack of experience with IPv6 and to avoid the risk of disruption to IPv4 services.

6.1.1 IPv6 only Virtual Firewall Instance

Athlone decided to test IPv6 connectivity to the internet isolated from the production network. This was achieved by creating a new virtual firewall instance or vdom (virtual Domain) on the Firewall and using layer 2 VLans on either side of the firewall to route the traffic. The diagram shows the IPv6 only vdom in blue and root vdom in red. Addressing was assigned to deliver connectivity as per the following diagram and connectivity over IPv6 was verified at each zone from the External router to Internal clients. Client devices with static IPv6 addresses assigned made successful access to IPV6-only internet delivering such services as Google search, YouTube, Facebook, and other Content Delivery networks over IPv6.

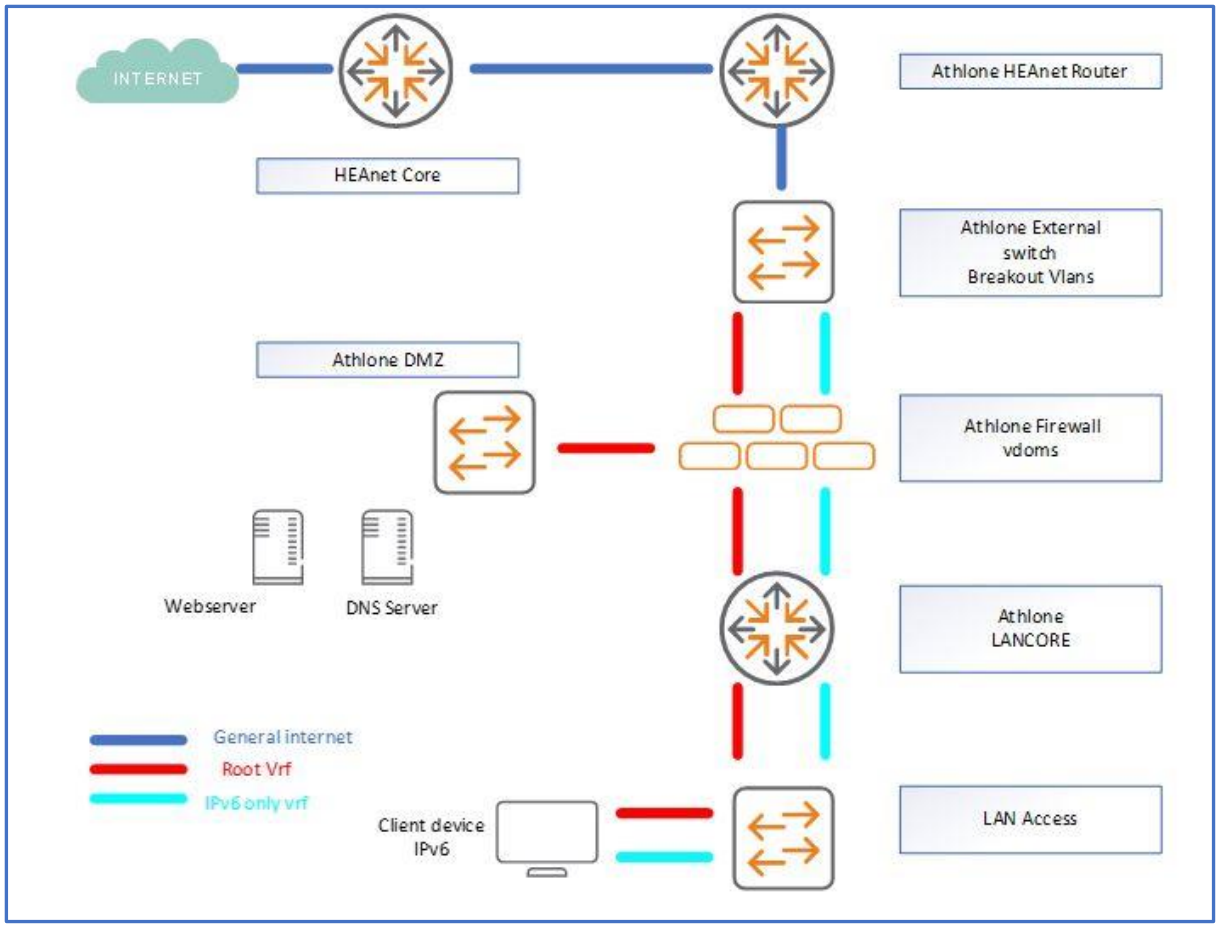

Figure 3: IPv6 only and Root VRF

6.1.2 HEAnet Router

HEAnet manages and configures the router infrastructure for Athlone. The router with dual diverse 10 Gig routed links to the HEAnet core network. HEAnet added a new VLAN for our IPv6-only vdom with its own /64 network. A static route provided connectivity to the physical interface on the firewall. The current general Internet service VLAN had the new IPv6 gateway applied. This allowed the first ping6 connectivity test from the firewall to the Athlone HEAnet router and onto the IPv6 HEAnet core.

6.1.3 Firewall

Institute firewall is the critical infrastructure component of IPv6 deployment both from a routing and security point of view and crucially had IPv6 feature support in 2010. The addition of IPv6 does add to resource utilisation increases of memory and CPU on the firewall. Firewall must support Neighbor discovery ICMPv6 message processing. These are steps

- enabling IPv6 support globally on the device.
- assigning IPv6 addresses to relevant interfaces on IPV6 and root vdoms.
- IPv6 only – Internal and external interfaces
- Root vdom Internal external and DMZ interfaces
- configuring routing between interfaces
- applying security/filtering policy for servers in DMZ

6.1.4 Demilitarised Zone

A DMZ [i.35] network is a perimeter network that protects and adds an extra layer of security to Athlone's internal local network from untrusted traffic. The end goal is to allow access to untrusted networks, such as the internet while its LAN remains secure. Organisations typically store external facing services and resources, as well as servers for Domain Names systems (DNS), File Transfer Protocol (FTP), mail and web servers in the DMZ. Previously Athlone implemented a static NAT to a public IPV4 address for each device in the DMZ. A DMZ provides a buffer between the internet and private internal LAN and so is an appropriate network to start IPv6 deployment. Assigning the IPv6 address from the addressing plan to the DMZ interface was completed.

6.1.5 DNS Servers and Webservers

Athlone utilises a split DNS design. Like many organisations, it doesn't want to share the complete DNS information for the internal network with the internet for security reasons as private IP addresses are not routable on the internet. However, it needs to publish DNS information for its DMZ-hosted nodes and does so using BIND which it manages in its DMZ. To enable IPv6 DNS features it was decided to upgrade to the latest version of Bind and to make the locally hosted institute webserver www.ait.ie in DMZ accessible.

- Configuration changes included an update named.
- conf with ipv6 reverse lookup and ipv6 loopback.
- Update the ait zone file with AAAA records.
- Update reverse lookup zone file with PTR record.
- Ask HEAnet to delegate IPv6 lookup to Athlone.

Athlone also manages and hosts its web servers. Each web server had to be upgraded with an IPv6 address to be reachable over IPv6. Each web server DNS record had to have an AAAA record to resolve the new IPv6 address.

6.1.6 Internal LAN Core

The LAN Core consists of three-layer 3 Switches implemented as a single IRF logical layer 3 switch. The split design is spread over two comms rooms with diverse access switch stacks trunked to location cores. For this initial phase, a new VLAN was added with IPv6 layer 3 routing capability to the internal firewall interface.

6.1.7 Security

Network infrastructure security is the process of securing the integrity of networking systems and software assets such as end-user devices against cyber threats.

IPv6 networks must be deployed securely and leverage the lessons learned with IPv4. IPv6 is not dissimilar to IPv4, and the security issues and mitigation techniques are mostly identical with some exceptions. A dual-stack network doubles the attack exposure surface, and this phase begins the practical exposure to IPv6 security. Critical in the early phase of deployment is filtering and monitoring. Filtering is achieved using firewall Access control lists with the same security policies applied for both protocols. E.g., the addition of DMZ policy rules to allow HTTP and HTTPS access over IPv6 from the internet permits this access requirement. Athlone utilises a powerful log management, analytics and reporting platform with a single console to monitor all new IPv6 conversations.

6.1.8 Observations

This phase confirmed IPv6 connectivity between Athlone's External infrastructure and the Internet and followed best practices to deploy first at the network and infrastructure level before being rolled out to the end user. Static unicast addressing utilising static and dynamic IPv6 routing provided Athlone’s web servers connected to the IPv6 internet extending its internet presence to IPv6 users.

7 Dual-Stack WIFI

7.1 Introduction

In 2012 with the addition of a new Engineering building capital funding was made available to deliver a campus enterprise wireless solution and to embrace the adoption of mobile technology in higher education learning and teaching. The convention at the time in higher education was to airgap wireless client access from wired infrastructure and deploy on separate physical infrastructure. Eduroam the secure worldwide roaming access service was deployed. Staff and students were provided wireless 802.11 browser-based connectivity to cloud-hosted VLE and Microsoft Office 365 email storage and collaboration tools.

The original plan was to deploy IPv6 only to Wireless client nodes. Trialling and testing the DNS64 and NAT64 transition mechanism resulted in poor client performance and issues with peer-to-peer (P2P) applications and was not a viable solution. The network was deployed using IPv4 only. In 2013 Public facing services were dual-stacked successfully. This provided an impetus to add IPv6 capability to the wireless network in a dual-stack configuration.

7.2 Wireless Hardware Design and IPv6 Limitations

The physical network design consisted of two core switches and two controllers in a split design over two diverse comms locations. POE switches populated with access points were trunked to either split core. This physical wireless network infrastructure was isolated from the wired LAN by a firewall interface but SMMPv3 traffic for Network was permitted to NMS server. Furthermore, the diagram shows the logical separation of switch & Access point management from wireless client access by use of separate VRF on the wireless wired core. Access point deployment over IPv6 was not supported by the wireless controller. Each POE Access switch per physical building location had a VLAN assigned and DHCP providing an IPv4 scope per VLAN.

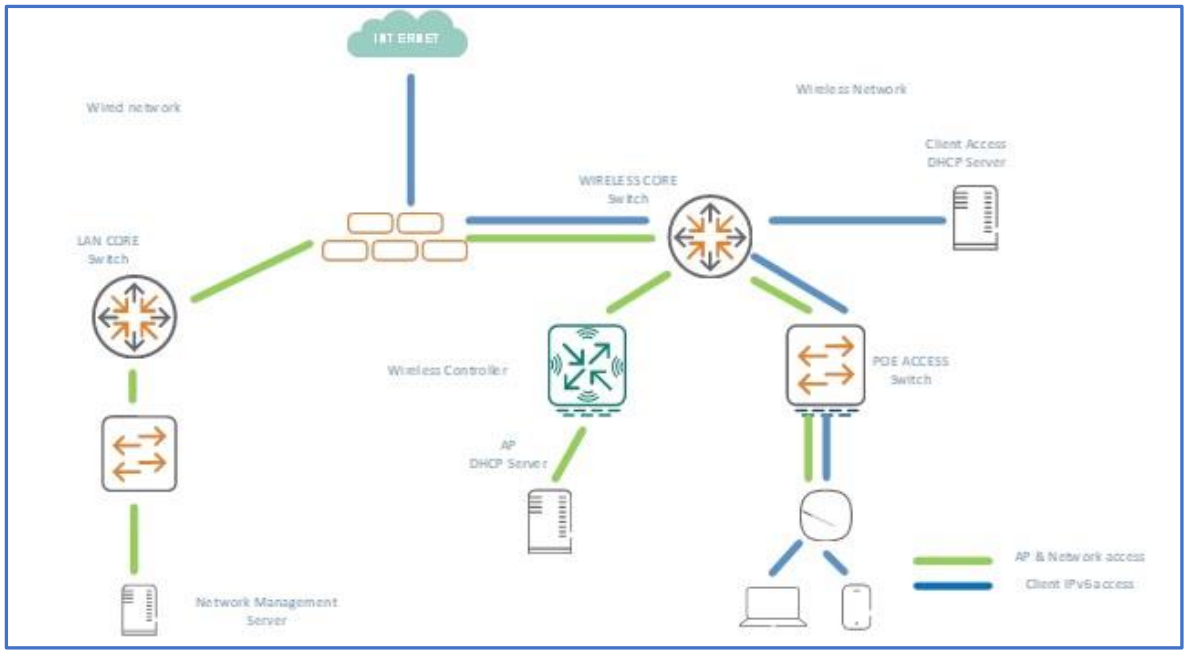

Figure 4: Wireless VRF segmentation

7.3 IPv6 Site Address Allocation

In Dual stack deployment IPv4 addressing remains unchanged with DHCP providing dynamic addressing. In IPv6 there are three ways to assign addresses statically, stateful using DHCPv6 [i.36] and stateless using SLAAC (State Less Auto Address Configuration). The initial plan was to utilise DHCPv6 as it was functionally like DHCP, but Android devices require SLAAC for auto-addressing so both methods were used.

DHPCPv6 server features were enabled on the server by mirroring the IPv4 subnet allocation already in place per VLAN with the new IPv6 prefix. The prefixes were determined using the Address plan in the previous section.

New in IPv6 address assignment is that the layer three router interface interacting with the client node is how it is determined what assignment method is used. Using ICMPv6 Router Advertisement RA goes to all nodes' multicast and intimates to the client what method of assignment is used.

With SLAAC the client generates its own globally unique address using the /64 network prefix advertised by the local gateway in a Router Advertisement RA.

DHCPv6 stateful Through the exchange of DHCPv6 messages between the client and DHCPv6 server a global unique /128 address is assigned to the client by the server. Additional options such as DNS & NTP are also provided to the client.

DHCPv6 stateless was also utilised as this was required to provide DNS server settings to the client.

7.4 Router Address Assignment Configuration

As outlined above the router interface plays a major role in the mode and method of address allocation using RA flags. ON layer 3 interface DHCP relay is used to assign IPv6 DHCPv6 server address. Use managed flag to let the host know to use DHCPv6 and other Flags for the host to use DHCP for other services such as DNS. For SLAAC each host receives a 64-bit prefix, and the last 64 bits are generated using EUI-64. OSPv3 routing was enabled to share new client IPV6 subnet routes with the firewall.

7.5 Firewall

The following configuration additions were added to the institute firewall which also added a considerable increase in resource usage of the device in CPU and memory.

- Internal Wi-Fi interface of firewall IPv6 gateway address added.
- IPv6 subnets were added to the firewall for internet routing.
- IPv6 firewall policy for IPV6 subnets added including application and content filtering rules to mirror IPv4 policy.
- External firewall interface had IPv6 configured from the Public facing services phase and this was utilised to forward traffic to IPv6 internet

7.6 Client Devices

In higher education when utilising the Eduroam [i.37] Wireless service client devices such as laptops and smartphones are unmanaged by the home institute. The prerequisite is that the IPv6 protocol which is enabled by default on devices is enabled to utilise the IPv6 features. This was the first experience of IPv6 automatic address assignment and the while longer Globally unique IPv6 address is new there are also differences in attributes, types, structures and how they are used. It is normal for IPv6 interfaces to have multiple addresses including link-local addresses, stable global addresses and privacy global addresses. Gaining experience with the new IPv6 addresses and types is a key learning milestone in the overall Ipv6 adoption project. Happy Eyeball is an algorithm published by the IETF which can make dual-stack applications more responsive to users by attempting to connect using both IPv4 and IPv6 at the same time (preferring IPv6). This solution helps in avoiding the usual problems faced by users with imperfect IPv6 connections or setups.

7.7 Security

IPv6 security features were limited in Network firmware of the time and features such as RA guard were not available [i.38]. The network was monitored for Rogue Router Advertisements by use of Linux IPv6 security toolkits.

7.8 Observations

A key step on the IPv6 transition journey was completed with the first experience of deployment of IPv6 to client devices using the new modes and means of SLAAC and DHCPv6. Staff and student client devices could connect natively over IPv6 to Institutes VLE and the global internet for IPv6-enabled services. IPv6 accounted for 30% of overall internet traffic from Institute. This phase proved IPv6 as a viable stable protocol on a campus wireless infrastructure.

8 Dual-Stack LAN

8.1 Information Systems Infrastructure

Following the success of deploying IPv6 to the wireless clients without it was decided to do similar on the campus wired network again adopting a dual-stack approach. A key difference was this was an enterprise production network of over 6000 staff and students with business-critical applications used in a Higher Education Institute of Technology. Athlone's information systems are based on the Microsoft windows client-server model which provides file storage, authentication directory services, messaging collaboration tools and business applications. Computer laboratories provide students with specialist software applications for various disciplines. Applications are predominantly web applications and require a web browser for client access. The requirement was to enable the client devices to communicate with the server infrastructure using the IPv6 protocol. The infrastructure consisted of over 3000 clients and over 100 Physical and Virtual servers.

8.2 VLAN Design & Network Core Configuration

Athlone's internal service 4.2.3.4 details how the VLANs are segmented between Staff and Students and the Windows infrastructure mirrors this design. The key task is adding the IPv6 layer 3 interfaces for the Staff and Student VLANs at the LAN core. No changes were required at layer two VLAN layer.

8.3 IPv6 Site Address Plan and Configuration DHCPv6

The coexistence of two protocols is central to dual-stack deployments and in this phase addition of v6 subnets mirrors the v4 structure. Utilising the institute addressing plan and the convention of use type and location type the IPv6 prefixes for the staff and student subnets were worked out. The original Windows server 2012 DHCPv4 Server had IPv6 addressing added and DHCPV6 features installed and IPv6 scopes added.

With IPv4 design, the VLAN number matched the subnet network octet i.e., vlan22 was easily identified with network 192.168.22.0/24. The IPv6 network prefix 2001:770:50: ABC had no link to the VLAN number and while initially confusing it was more important that the new addressing plan convention was what identified the networks. It is important in this phase not to shoehorn IPv4 thinking into the future IPv6 architecture. The DHCP server would now provide an IPv6 and IPv4 address per client device with the same VLAN numbering as before along with DNS and other configured DHCP options.

Stateful addressing using Windows server DHCPv6 to address enterprise clients were primarily chosen for ease of enterprise client identification purposes logging and monitoring. DHCP unique identifier (DUID) allows the DHCPv6 server to identify the lease and track its lease. DHCPv6 is functionally like DHCPv4 but DHCPv6 does not provide a gateway address to the client.

8.4 Windows DNS

DNS is critical to the operation of IPv6 and on an Enterprise Lan. It is not practical to remember or write the 128-bit IPv6 address, so name resolution avoids this requirement. Like most Enterprise campus networks Athlone's internal DNS is provided by Windows Active Directory (AD) DNS services & is integrated with Windows DHCP services. The key task is to enable these services for windows machine names or NetBIOS names for internal name lookups for internal client-to-client communication. Recursive DNS services are required for internet name lookup for enterprise clients for internet access. Athlone's Windows Active directory AD structure consists of a root DC ait. i.e., and two subdomains of staff.ait.ie and student.ait.ie. For resilience, there are two servers for each domain/subdomain combination of physical and virtual servers requiring configuration changes to six servers. As these services are originally deployed over IPv4 requirement is to enable these capabilities over IPv6.

- Applying IPv6 static address to each DNS server network interface card (Nic).
- Update DNS forwarders from subdomain to root server over IPv6.
- Confirm DNS root hints for internet lookups resolve over IPv6.
- Windows Active Directory (AD) dynamic DNS (DDNS) features manage dynamic mappings of Windows machine names or NetBIOS names to IPv6 addressing [i.39].

At the end of this process, each client device will have its original IPv4 address with an A name DNS record and the addition of a globally unique IPv6 address with an AAAA DNS record. In this dual stack design, this will facilitate communication over both protocols and critically IPv6 first where all configuration is in place.

8.5 Windows Server Static Addressing

All Windows Active Directory (AD) staff and student Servers both physical and virtual require the addition of a static IPv6 address on Network Interface Cards (Nic). IPv6 DNS server setting addresses are added. Each server was tested to resolve DNS over IPv6 and IPv4 as part of the process.

8.6 Windows Client Configuration

All Windows Enterprise clients 8 and later have IPv6 enabled by default since 2013.

Athlone implements a standard windows desktop environment with controls and restrictions enforced by Active Directory (AD) group policy. In a dual-stack deployment where IPv6 fails connectivity, it will then try over IPv4. This may hide issues with IPv6 but assures connectivity.

The most common issues on the Local Area Network were desktop clients to local server infrastructure where the key combination of IPv6 addressing and DNS resolution over IPv6 was the most common misconfiguration. Athlone's Cloud-hosted VLE Moodle showcased v6 endpoint-to-endpoint benefits delivering excellent reliability and performance to this key asset.

8.7 IPv6 Training and Troubleshooting

Training had been provided on basic fault-finding techniques in identifying and resolving issues. The use of ping on both IPv6 and IPv4 for connectivity status. Performing DNS lookups for both protocols would identify most issues which tended to be DNS related. Unfortunately, initially, the most common approach was to disable IPv6 on the client or server when an issue arose. The Key was to continue to outline the need to fault-find both protocols to determine the issue and this did improve with time and experience.

8.8 Observations

The technical tasks in this phase were straightforward and successful for staff and student clients using IPv6 when the endpoint service was available over IPv6. The biggest challenge was getting all technical staff and management on board to work with the new protocol.

IPv6 versus IPv4 usage was 50% each on internet-facing interfaces where did was monitored using sflow both on firewall and external routers. Athlone's most critical educational service is hosted by VLE Moodle and was accessible natively over IPv6 to the Amazon datacentre.

Interestingly it was also visible when students were on campus during term time as Ipv6 usage was higher by 10% due to increased usage of CDN such as google Facebook Apple WhatsApp Instagram

9 IPv6 only WI-FI 6

9.1 Context for Project

The purpose of this project was to deliver a wireless data network that provides complete campus coverage. The wireless network design would need to provide a high-performance secure robust wireless network with 802.11ax (Wifi6) capability. A complete refresh of all wireless-related wired & POE infrastructure was sought. Wi-Fi network will utilise 10 Gbps internet access provided by NREN HEAnet network. A full Rf survey of campus for 802.11ax coverage was carried out utilising survey software and AP on a stick survey methodology. This delivered the Access point deployment requirement for the campus. Critical to our approach from the outset was to where possible deliver an IPv6-only wireless infrastructure. Initially, two wireless services were deployed. For Higher Educational institutes Eduroam is the primary wireless service provided to staff students and researchers. Eduroam authentication is provided using 802.1x and RADIUS [i.40] over IPv6 to the institute's Microsoft Windows Active Directory for user-based authentication. A guest Wi-Fi service with a captive portal and utilising the sponsored guest access feature using email was a new way of securing guest access to infrastructure.

9.2 Motivation for IPv6 and Foundations in Place

Athlone IT was keen to progress work already done with IPv6 and was ready for the challenge that IPv6-only issues that would bring. We wished to manage one protocol instead of two and wish to promote IPv6 first approach.

We were keen to utilise the foundations already in place and apply the same addressing plan and the DNS setup for IPv6 already in use. We continued with Dual stack on original firewall interfaces but all traffic on new wireless interfaces was IPv6 only and all traffic was native IPv6 for Radius authentication and network management features.

9.3 Tender Objective

Due to the financial value of this project a competitive tender was required while also securing the latest wireless networking technologies. Due to our previous experience with IPv6, we wanted to pursue the goal of an IPv6-only infrastructure. A key approach to our tender was to reward IPv6 capability in all our markings and ask questions that could be achieved over IPv6 only. We conducted exhaustive research on multiple Network Vendors' IPv6 capabilities. We also based our requirements on Open standards.

9.4 Physical Design LAN Core and Access

A split core is spread over two diverse comms rooms which provide resilient single-mode fibre interconnectivity for all campus building locations and are also points where diverse internet connectivity to the Internet via HEAnet is provided. Each building POE Access switch has diverse fibre connectivity to each split core. 330 Access points provide full 802.1AX WIF I6 to 64,000 metre campus.

A fully configured and routed wired network is a prerequisite for successful wireless network deployment and is delivered here with a collapsed core/aggregation design consisting of two core switches in a high availability cluster deployed over IPv6 only. The aggregation layer provides connectivity to the core with two VLAN’s for management and Access points.

Dual physical wireless mobility controllers in an IPv6 cluster terminate the Access points.

9.5 Wireless Logical Design

The logical diagram shows the Access, Wireless, Core Aggregation, Wan Edge and physical and Virtual server layers and how they integrate to deliver this secure wireless solution.

This simple scale-able design deploys link and multi-chassis link aggregation between the Core/aggregation and Access layer devices. Routed links are utilised at the core with layer 3 IPv6 path redundancy. Athlone deploys a controller-based design where the controller authenticates the user and bridges the user to the core for internet access.

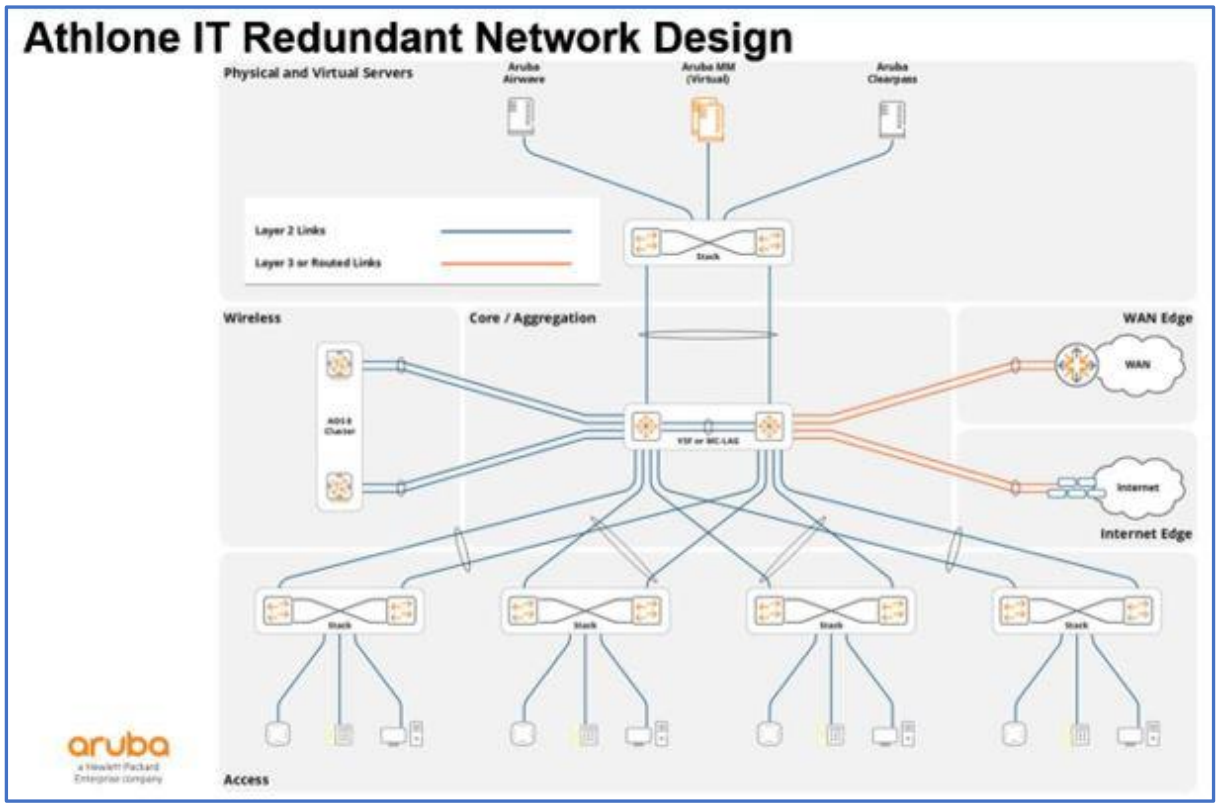

Figure 5: IPv6 WiFi-6 logical design

9.6 VLAN Design

An objective of this project was to demonstrate how IPv6 allows for simpler network design to avoid VLAN sprawl and utilise IPv6 auto-addressing features. One VLAN each for Access point deployment and network management campus wide and a single VLAN at the Core per SSID was implemented.

9.7 Wireless Controllers and Wireless Access Points

Dual physical wireless controllers were deployed in a high-availability IPv6 cluster. 802.11AX Wifi6 Access points were deployed natively over IPv6 using SLAAC and RDNSS with dual Radio on 2.4GHZ and 5GHZ. The Wireless controller provides RADIUS s proxy support over IPv6 for Eduroam.

9.8 Virtual Machine Software Provisioning

A suite of provisioning tools is hosted on the institutes Hyper V windows infrastructure including natively over IPv6.

- A centralised management platform manages all configuration and licensing of the physical controllers on the institute LAN Microsoft Hyper-V virtual infrastructure.
- Real-time visibility and Client traffic analysis of the RF status of Wireless infrastructure provides heat maps for troubleshooting.
- A Network Access Control solution identifying who and what's connected to the network utilising policy-based access control.

9.9 IPv6 Site Address Allocation Plan

In this wireless design, the controller authenticates the user either with 802.1x, Captive portal process. The layer 3 router interface is defined on the Wired Core per VLAN subnet. The client is bridged to the wired core where it obtains an IPv6 address either by SLAAC or DHCPV6 based on SSID selected. Eduroam and Guest users on the captive portal utilised SLAAC and IoT networks utilised DHCP.

9.10 WAN Edge and Routing

A separate External 10GBps interface was added to the institute firewall to route Wireless IPv6-only client traffic to the internet. Network Management traffic and another provisioning service on LAN and Virtual environment were routed via firewall interfaces.

Firstly, native IPv6 traffic routes directly to the IPv6 internet. NAT64 & DNS64 tunnelling services provided by the institute firewall allow IPv6-only wireless clients the ability to access IPv4 internet. The NAT64 feature worked seamlessly and as efficiently as NAT44 for wired Dual stack LAN segments.

9.11 Security with Policy-Based Access Control

Wireless security is delivered by a combination of the use of a Policy Enforcement firewall when a user is authenticated by Wireless Controller with Policy Access control assigning a role to the user at the point of access. This allows for centralised Zero trust access with a dynamic policy approach.

9.12 User Experience

Overall user experience has been very positive due to the high access speeds provided by WI-FI 6 and also the enhanced roaming capabilities of IPv6. STEAM gaming site only application where issues occurred. Also, some android clients experience intermittent issues with Captive portal guest access. The key figure when students are on campus during term 70% of all of Athlone s internet traffic is IPv6. Interestingly this drops to 60% when students are not on campus explained by the drop in use of Content Delivery networks.

9.13 IPv6 only Testbed Faculty of Engineering

The concept is to exploit IPv6 only facility and to develop an IPv6-only playground for both Undergraduates and researchers to embrace the use of IPv6 in the University sector. Interconnect all devices – Raspberry Pi's, wireless sensors, robots VR&AR headsets in a secure authenticated policy-based access with IPv6 addressing. Industrial IoT IIoT and 6LowPan sensor networks are a research focus.

Conclusion

As demonstrated and highlighted within this document, IPv6 is the only practical and viable solution to address the IPv4 address depletion issue. The IETF, which develops open standards through open processes to make the internet work better, are only pursuing these such developments now with IPv6. Technologies such as 5G, cloud, IoT and CDN all require the large Ipv6 address space for continued growth. Developments in Industrial IoT (IIoT) utilising wireless mesh and sensor networks are using IPv6 centric protocols such as 6LoWPAN [i.41] and 6TiSCH [i.42] at the forefront of these innovations. The removal of NAT from enterprise networks will bring us back to the original internet architecture i.e., endpoint-to-endpoint connectivity. This also will aid address a range of the security and identity concerns. Cloud Computing is an essential component of digital transformation and having a local IPv6 infrastructure routing to an IPv6 cloud without the need for NAT delivers a simple scalable secure architecture. IPv6 will soon hit the critical 50% adoption rate on the global internet which will only accelerate the transition to IPv6.

The experience of the IPv6 installation in TUS Athlone outlined in this document hopes to encourage other University Network teams to begin this exciting technical challenge. It will ultimately deliver a simpler, secure, and more scalable infrastructure. One of the acknowledged reasons for hesitancy to deploy IPv6 on enterprise networks is due to legacy design and applications on these Networks. The TUS Athlone switched wired network was no different, with little change and innovation in over 30 years. However, the WI-FI 6 IPv6 only network delivered a software defined networking solution combining Controller based access and policy-based access for all clients. IPv6 as outlined slotted seamlessly into this new architecture with the large /64 prefix delegations allowing for simple VLAN design. Athlone's wired access infrastructure is currently being upgraded to utilise the same controller and policy-based access for all devices. The result is a simple scalable integrated wired and wireless infrastructure optimised for secure cloud services. Athlone's IPv6 usage of 70% versus 30% IPv4 on its internet connectivity shows that when deployed the protocol is utilised and delivers the return of endpoint to endpoint connectivity and all associated benefits.