



Operational Experience with IPv6: A Campus Implementation

By:
Michael Fazely
Jeffry Handal

Table of Contents

Table of Contents	2
Operational Experience with IPv6: A Campus Implementation	3
Foreword.....	3
History.....	3
Taking IPv6 to the Field - A Campus Story	4
Motivations for Adopting IPv6.....	4
Client Addressing Mechanisms	5
Routing Considerations	5
Subnetting Simplified.....	5
Link-Local Addresses Put To Use	7
Campus Network Security - A Balancing Act.....	7
Rogue RA and DHCPv6 Mitigation.....	7
ICMPv6 ACL	9
Management Subnet ACLs	9
Privacy Addressing - Friend or Foe?	9
Tunneling	10
Challenges - Learning by Doing	10
The Tagged VLANs and Identical MAC Addresses Conundrum	10
Tip.....	12
Authentication Evaluation.....	12
Bogus RAs from Unexpected/Trusted Sources	13
Tip.....	14
NETREG and DHCPv6 Options Bypass.....	14
Native IPv6 Support	15
Route Summarization.....	15
Stateful Address Configuration with DHCPv6	15
White/Black Listing	15
Closing Thoughts	16
Appendix.....	18
Appendix 1.1	18
Link-Local Addresses	18
Appendix 1.2	18
Questions About Addressing Considerations.....	18
Appendix 1.3	19
Who We Are - Louisiana State University	19
University Networking and Infrastructure	20
Acknowledgements	20
References	21

Operational Experience with IPv6: A Campus Implementation

Foreword

It is assumed that the reader already has basic working knowledge of IPv6, including understanding address types and notation, as well as general networking information. The authors will provide RFC numbers during the discussion of various topics in order to allow the reader to gain a greater understanding. The following document will chronicle Louisiana State University's efforts in setting up a fully functioning IPv6 presence on a large enterprise campus. The entire process from planning through implementation and troubleshooting will be discussed in detail. It is the goal of the authors to provide a comprehensive guide in order to bolster readers' IPv6 knowledge and point out potential pitfalls and/or caveats they may encounter.

This document has been compiled to further promote the adoption of IPv6 and to be used as a lessons learned reference. As the flagship university of the state of Louisiana, we are helping fulfill our mission of research, education, and promotion of new technologies and protocols. Finally, another driver for this document is to continue building on the legacy of technology leadership being developed by the members of the University Networking and Infrastructure (UNI) group within Louisiana State University's (LSU) Information Technology Services (ITS).

History

The history of IPv6 at LSU began after hosting the 2008 Internet2 Member meeting in New Orleans as simple curiosity. As we were preparing to deliver a successful event, there were many things that had to be put into play: multicast with CERN, IPv4 routing, wireless setup of Access Points (APs), and a random request for IPv6 with the caveat it may not be used. Regardless, LSU applied for and received a /48 address block from ARIN in August of 2008. Even though IPv6 ended up not being used during the Member meeting, our attention was focused on exploring it.

After some initial research and testing, LSU rapidly deployed IPv6 in a dual-stack network infrastructure within the Computing Services building only. The intent of this small scale deployment was to gain experience with running a dual-stack network and work out potential bugs before exposing users on campus to any issues. In September of 2010, the dual-stack network was expanded to the entire campus in an effort to give our campus community the opportunity to begin testing and using IPv6 for themselves.

In June of 2011, LSU participated in World IPv6 Day by advertising A and AAAA records on the main LSU website along with several department websites. This was largely a success. The exception will be discussed in detail in the section entitled "Tagged VLANs and Identical MAC addresses." As a tangent, it is curious to note, LSU experienced a brief 20 minute outage on the IPv4 space for the entire campus that day. Surprisingly, IPv6 routing remained operational. Thereby, all sites available on World IPv6 Day were available and clients' outage impact was highly minimized, especially if going to Facebook and some Google resources.

Shortly after World IPv6 day, we realized some caveats with a /48. ISPs would not be accepting advertisements smaller than a /48. This was a real hindrance for a campus where multihoming and having a backup provider for our campus are critical components of the network. Additionally, we concluded that a /48 was not "large enough" for our campus. LSU applied for a new /40 address space and returned the previously attained /48.

Since the initial address space was only a /48, it didn't allow for other LSU sites to be advertised out a backup Internet drain in the event that main campus lost connectivity. The /40 would allow for each site to have more than one /48 and thus would make LSU immune to prefix size limitations. The /40 prepared LSU to keep growing the number of network devices well into the future. Additionally, if other LSU satellite campus were to be tied into our network, we have plenty of IPv6 to distribute.

Most recently, on June 6 of 2012, LSU participated in World IPv6 Launch Day. During this event, even more users on campus were encouraged to embrace IPv6. We thus greatly increased the number of websites and network utilities that were IPv6-enabled. Several of our internal tools ride on IPv6. In the near future, we plan to deploy test wired and wireless networks that are IPv6-only with the intent to awaken curiosity and experimentation from our campus research community.

Taking IPv6 to the Field - A Campus Story

Motivations for Adopting IPv6

On February 3, 2011, IANA (Internet Assigned Numbers Authority) assigned its last remaining IPv4 address space. [1] This event shouldn't have come as a surprise as it had been known for quite some time that available IPv4 space was quickly running out. With available IPv4 space exhausted, it became paramount that LSU finish its IPv6 deployment and continue to explore its intricacies.

With the limited amount of public IPv4 address available to LSU (two /17's and one /16) it became difficult to satisfy the addressing needs of the increased number of wireless devices on campus. Students, faculty, and staff now have one or more wireless devices on their person at any given time. To combat the issue of a growing number of devices, one could argue that NAT would be a viable solution, but it also can add complexity to a network. If a user on campus performs a nefarious act, like the downloading/sharing of copyrighted material, it can prove difficult to track that user down if they are using a private address that is NATed.

LSU's ITS strives to provide cutting edge services and technology to campus. With this goal and the amount of research departments on campus, it is essential that ITS continues to be a step ahead of the users. Exploration and implementation of IPv6 also fit in line with the mission statement of University Networking & Infrastructure (UNI) (a department within ITS): The mission of UNI is to design, implement and manage solutions which provide for the University's voice, data, and video communications. Services provided by UNI include installation and maintenance of metallic and optical cabling, network architectural design, software specification and configuration; and the specification, operation and maintenance of transmission hardware and software.

As the flagship university of Louisiana, it is a goal of LSU to not only be a leader within the state for IPv6 and other technology services, but to also be at nation's forefront when it comes to these services. Through LSU's success in these endeavors, we hope to educate and assist other Universities or entities with issues that they may have.

And finally, we are moving forward with IPv6 in an attempt to keep the Internet open and free through network transparency. Network transparency can be achieved by providing end-to-end connectivity, which IPv6 purposefully does. Also, implementing IPv6 on our network allows us to avoid the pitfalls of carrier-grade NAT (CGN).

Client Addressing Mechanisms

In regards to the topic of address configuration, LSU explored stateful address configuration with DHCPv6 and stateless address autoconfiguration (SLAAC) with EUI-64. SLAAC was initially chosen in favor of DHCPv6 for its ease of use and rapid deployment. Using SLAAC with EUI-64 allows clients to configure an IPv6 address themselves using the neighbor discovery protocol (NDP) instead of a DHCP server. The neighbor discovery protocol operates by using link-local and multicast addresses for communication (covered in a later section). It should be noted LSU's long term address configuration goal is to move to a stateful address configuration with DHCPv6 as it will allow for better control of address deployment and user tracking.

Routing Considerations

Initial discussions of an IPv6 presence on campus dealt with deployment, transition mechanisms, hardware/software compatibilities, and potential negative effects. As discussed in the brief history section, it was decided that deployment would initially be performed on a small scale within the Computing Services building, but that still left the transition mechanism up for debate.

Options for transition mechanisms include dual-stack and various types of tunneling mechanisms. After evaluating 6to4, ISATAP, Teredo, and other tunneling options, a dual-stack implementation seemed to be the best fit for LSU's current needs and long term goals. A dual-stack implementation builds IPv6 onto an already existing IPv4 infrastructure, hence associating IPv6 subnets with IPv4 subnets on interfaces. In LSU's network implementation, OSPF is the routing protocol of choice. Naturally, this option leads to running two separate instances of OSPF: one for IPv4 and one for IPv6 (OSPFv2 and OSPFv3 respectively).

It was also essential that all currently deployed network gear be evaluated for IPv6 compatibility. Fortunately, much of the networking gear on LSU's campus only needed software upgrades in order to enable IPv6. To this day, however, there remain certain IPv6 features not supported on a small number of network devices. Those wanting to implement IPv6 should review manufacturers' specifications for each device running on the network in order to determine the software release needed for desired IPv6 functionality. If there are devices which are not supported, engage your vendor for a road map of features to see when, or if, features will be available. Other groups within the organization will also need to be cognizant of this when ordering hardware or software products that are IPv6 capable.

Subnetting Simplified

During the initial testing phase, the campus was allocated sequential /64's throughout the campus. With this iteration, using the original /48 space left the following lessons learned:

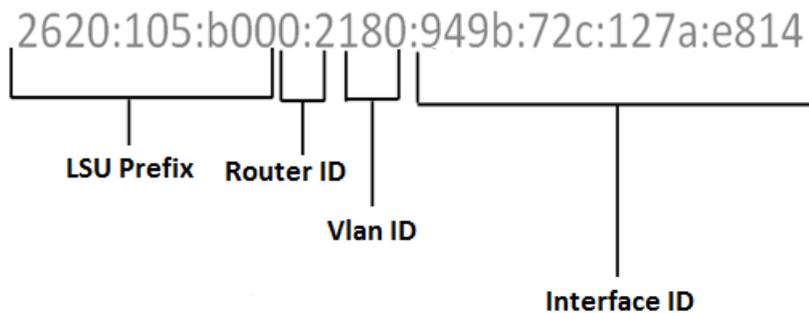
- Realization of how big the IPv6 space is. Sequential allocation did not make a dent on the /48 we owned. Most of it remained unallocated.
- Sequential allocation without clear tie-ins to the existing IPv4 space made it impossible to troubleshoot.

- Troubleshooting should always be on your radar. The sequential allocation made it hard to tell support personnel what to use on their firewalls.

After acquiring the /40 space, the addressing plan was revisited. We threw out all prior experiences and notions we had from IPv4 addressing. We cleaned the slate and decided to work from the ground up. A few decisions that were made include:

- Set aside portions of the space for main campus, residential halls, remote sites, and future campus sites. For such purpose, we assigned /44's.
- Each /44 was then further subdivided into /52's.
- A /52 was assigned per distribution router. This would allow for growth and uniformity on a per router instance. NOTE: Typically a router corresponds to a building in our setup.
- Within the /52, the last two digits were designated as a ROUTER ID to help identify a building or area.
- VLAN IDs were tied into the addressing scheme as well. This was done by placing the VLAN ID immediately following the ROUTER ID. A 3 digit VLAN ID number is used.

The end result of all the above internal rules leads to a very informative address that aids in troubleshooting. Now, all we need to do is educate our NOC, Help Desk, and users to identify the intelligent components built into the address.



NOTE: Since LSU has a /40, only the first 40 bits of the address will remain constant:

2620 = 0010011000100000 = 16 bits
 0105 = 0000000100000101 = 16 bits
 B0 = 10110000 = 8 bits

Link-Local Addresses Put To Use

For a breakdown of the link-local architecture see Appendix 1.1 Point-to-point link-locals were selected for routing. They were cleverly assigned in such a way that matched our gear's interfaces. In a way, the address embeds documentation on how physical copper/fiber connections are hooked up to each other. For instance, we assign link-locals in the following fashion:

FE80::(Device ID):(Slot ID):(Port):(1 if core or 2 if remote)

The only exception to this previously mentioned convention may be found with core-to-core links as the ports do not match between links.

Even though link-locals are automatically assigned to interfaces once IPv6 is enabled, we override that automatic assignment and use our own link-local addresses for VLAN interfaces. We will dig deeper into this extra step's rationale in a later section. These addresses look like this:

FE80::(OSPF process ID):(Router ID):(VLAN ID):1

This convention holds true for every VLAN interface on campus.

Campus Network Security - A Balancing Act

Proper network security should continue to be a goal for any organization. With the deployment of IPv6, the network becomes only as secure as its least secure IP stack. Thus, it was important to look into IPv6 security and address the issue accordingly.

When first looking at IPv6, it had been said that the new protocol was inherently more secure than IPv4, which was found to be a myth after further research. [2] IPv6 is just as secure as IPv4. It eliminates some issues of IPv4 but introduces new vulnerabilities. For example, IPv6 subnets are considerably larger than IPv4. This makes any type of network scanning much more time consuming and difficult to consider. On the other hand, we introduce a new threat known as rogue RAs, which will be discussed below.

The inclusion of security mechanisms may result in some lost functionality. The following section will cover the testing that was performed in an attempt to balance securing the IPv6 network without sacrificing functionality.

Rogue RA and DHCPv6 Mitigation

A rogue router advertisement (RA) threat is one of the greatest security threats in the IPv6 world and warrants our attention. RA's are an integral part of IPv6 ND [3] and thus can't be completely blocked (unless you wish to deploy addresses using DHCPv6). They provide hosts with important information including: MTU size, auto configuration flag, and the link-layer address of the router. Thus, it is important that the measures taken to block bogus RA's are examined carefully. This can be accomplished in Catalyst IOS, using one of the following command strings:

```
int gigabitethernet ##  
  ipv6 nd rguard
```

The above macro command is not available on all platforms. (Please consult vendor documentation.) The above may also be accomplished with an ACL that gets applied to an interface as shown below:

```
ipv6 access-list RAguard  
  deny icmp any any router-advertisement  
  permit ipv6 any any
```

```
int gigabitethernet ##  
  ipv6 traffic-filter RAguard in
```

Using a test network, the ACL protecting against rogue RA's was enabled in order to determine not only if it worked properly, but if it had adverse effects. The test was first performed with no RA guard and a machine was allowed to send RA's via The Hacker's Choice-IPv6 Attack Toolkit (available for Unix) [4], to three laptops (one Mac OSx Lion, two Windows 7) on the same network. Each machine received a bogus link-local as well as dozens of IP addresses. Each machine also experienced high spikes in their CPU during this RA attack. Once the ACL was applied to interfaces on the switch, the laptops no longer experienced any issues associated with a bogus RA's.

Another way of mitigating rogue-RA's in IPv6 is to set the router-preference on the SVI's to "high". The default value for devices is "medium," therefore, setting the preference to "high" will force that RA from your router to be preferred. This can be achieved in IOS through the following command:

```
interface vlan100  
  ipv6 nd router-preference High
```

Rogue DHCP servers are always a problem in any large enterprise. A local layer 2 DHCP server will usually overpower your users unsuspectingly and use its addresses. They will have preference for their addresses over yours, since the time to reach your central DHCP server is longer. The same holds for the IPv6 space. To mitigate the ease for which rogue DHCP servers appear in your environment, ports can be blocked to prevent DHCPv6 requests coming from client ports. An example of this ACL in IOS is:

```
ipv6 access-list blockrogue  
  deny udp any eq 547 any eq 546  
  permit ipv6 any any
```

Note: It is possible to combine the RA guard and rogue DHCPv6 ACLs into one.

Similarly, tests for evaluating DHCPv6 rogue servers were conducted using the Hacker's Choice-IPv6 Attack Toolkit. Testing and sniffing how these events occur provides great insight into how IPv6 works. It is easier than we think.

ICMPv6 ACL

ICMPv6 is an integral part of IPv6 and not only allows for the host to discover other hosts on the network, but also provides several other mechanisms that allow for built-in error checking and ensuring that IPv6 runs accordingly. ICMP plays a critical role in IPv6 and cannot be blocked as we normally do for ICMP in v4. You will not get away with it here. With that being said, there are some potential security risks associated with specific ICMPv6 messages.

RFC 4890 [6], entitled "Recommendations for Filtering ICMPv6 Messages in Firewalls" gives a detailed breakdown of each ICMPv6 message along with security considerations and steps recommended to mitigate potential threats. An ACL may be derived from this RFC to address potential security risks, but as with any security measure, it should be reviewed closely to assess how it impacts your particular environment.

Management Subnet ACLs

As in IPv4, it is important to limit access to your network devices by allowing only specific subnets to be able to reach them. With IPv6, you can no longer "hide" your management IPs by using RFC 1918 [5] IPv4 addresses. All of a sudden, with an IPv6 address on your devices, they could potentially be reachable from anywhere in the world using IPv6.

ACLs become your friend again. Securing your devices can be done through the use of an IPv6 access-lists. A sample ACL for access management in IOS can be found below:

```
ipv6 access-list ipv6-acl-95
  permit ipv6 2620:105:B000:<subnet>::/64 any
```

```
line vty 0 15
  ipv6 access-class ipv6-acl-95
```

Note: This is not the case for all platforms, as some support dual-stack natively. Check the vendor documentation to see if this step is necessary.

Privacy Addressing - Friend or Foe?

Several software vendors have implemented privacy state addressing [7] to accompany stateless autoconfiguration using EUI-64. The intent is to prevent the host's MAC address from being known from the interface identifier. This address is generated by the host concatenating the prefix supplied by the router through an RA with a randomly generated interface ID. This randomly generated privacy address will change periodically making it even more difficult to track a user.

Having privacy addressing is a double-edged sword. On one hand it protects users from having their MAC addresses known by outside entities. On the other hand it makes user tracking difficult for network or security administrators. For some environments, this is a problem. User tracking is a must when having to respond to DMCA [8] violation requests.

For machines that we have control over, i.e. machines on the Windows Domain environment, it was decided to disable privacy state addressing through netsh commands made using a group policy.

```
netsh interface ipv6 set privacy state=disabled store=active
netsh interface ipv6 set privacy state=disabled store=persistent
netsh interface ipv6 set global randomizeidentifiers=disabled store=active
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```

Tunneling

As discussed earlier in the implementation section, there are several transition mechanisms that can be used in networks. The most common tunneling mechanisms used are Teredo, 6to4, and ISATAP. Since LSU chose a dual-stack infrastructure, there is no need for clients on campus to use any tunneling services, therefore we have taken steps to block tunneling. These steps have been implemented as follows:

1. Blocking tunneling protocol 41 at the campus firewall.
2. Blocking the relay address 192.88.99.1 at the campus firewall.
3. On the Windows environment, by issuing netsh commands using a group policy and blocking IP protocol 41 and UDP port 3544 outbound at our border.

```
netsh interface teredo set state disable
netsh interface 6to4 set state disabled
netsh interface isatap set state disable
```

Challenges - Learning by Doing

The Tagged VLANs and Identical MAC Addresses Conundrum

As discussed in the history section, a problem was encountered after World IPv6 Day when A and AAAA records were advertised for certain LSU websites. Users in a particular department on campus began to experience latency while attempting to load their main website. This web server was hosted locally in their building router. The issue was blatantly apparent as many users within the building had their browsers homepage set to the department website. For off site pages, not on their router, no issues were encountered. It was determined that clients within the department who tried to visit the website, would attempt to use IPv6 first and would timeout, before reverting to IPv4. This preference for IPv6 is a default behavior of Windows clients, which were the only clients experiencing the latency at the time. It was later determined that connectivity within the building became erratic.

To explain the cause of the issue above, a brief overview of port configuration is necessary. Every switchport on campus is set up with an untagged VLAN for data and a voice VLAN command. This setup allows for the rapid deployment of VoIP phones in the event of a disaster scenario. This network implementation will also allow disaster personnel to attach VoIP phones without the aid of a network team to configure ports on-the-fly. With this configuration, Windows clients will receive an address from the untagged data VLAN and then later an address from the tagged VoIP VLAN. While configured with the untagged address, network connectivity works properly, but upon receiving the tagged address, IPv6 network connectivity

ceases within the building. This scenario only presents itself on routers that have the same MAC address and thus same link-local address on each switch virtual interface (SVI) Figure 1.

```

Ulan100 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7100::1
2620:105:B000:7100:213:5FFF:FE89:8000
Ulan102 [down/down]
FE80::213:5FFF:FE89:8000
2620:105:B000:7102::1
2620:105:B000:7102:213:5FFF:FE89:8000
Ulan103 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7103::1
2620:105:B000:7103:213:5FFF:FE89:8000
Ulan104 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7104::1
2620:105:B000:7104:213:5FFF:FE89:8000
Ulan105 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7105::1
2620:105:B000:7105:213:5FFF:FE89:8000
Ulan200 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7200::1
2620:105:B000:7200:213:5FFF:FE89:8000
Ulan204 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7204::1
2620:105:B000:7204:213:5FFF:FE89:8000
Ulan300 [up/up]
FE80::213:5FFF:FE89:8000
2620:105:B000:7300::1
2620:105:B000:7300:213:5FFF:FE89:8000
  
```

Figure 1 - All SVIs have the same link local address because they use the same virtual MAC address for their interfaces.

If the link-local addresses are different on each SVI, this will not cause an issue, as the client will receive multiple link-local addresses for its default gateway and will be able to route properly.

```

IPv6 Route Table
-----
Active Routes:
  If Metric Network Destination Gateway
  10 276 ::/0 fe80::6600:f1ff:fe9d:6c9f
  1 306 ::1/128 On-link
  10 28 2620:105:b000:b315::/64 On-link
  10 276 2620:105:b000:b315:223:aef:fe98:6ccc/128 On-link
  10 28 2620:105:b000:b500::/64 On-link
  10 276 2620:105:b000:b500:223:aef:fe98:6ccc/128 On-link
  10 276 fe80::/64 On-link
  10 276 fe80::223:aef:fe98:6ccc/128 On-link
  1 306 ff00::/8 On-link
  10 276 ff00::/8 On-link
-----
Persistent Routes:
None
  
```

Figure 2 - Default route ::/0 only points to one link local.

The issue described above was only encountered on Windows clients. This behavior is a result of how VLAN tagging is handled according to Microsoft’s NDIS [9] document. If VLAN tags are not stripped, then SVI’s having the identical MAC addresses will not have an issue as shown with Mac and Linux operating systems. A potential workaround to this issue will be discussed in the section entitled “Potential Solutions to Issues”.

Tip

Assuming that turning off IPv6 on the VoIP VLAN interface is not a viable option for your network, there are several potential workarounds that were either presented by vendors or during the troubleshooting phase of this issue. The first would be to suppress routing advertisements on the voice VLAN. This will allow for phones to still receive an IPv4 address, but won't allow for the distribution of IPv6 addresses to clients.

The second would be to enter the driver configurations for every client to disable priority VLAN tagging. This option will vary depending on the network card and driver version installed on the client. This solution is not recommend for a large campus environment. For example, a registry edit may be performed with Broadcom adapters in order to prevent the issue. This can be done by entering the registry edit, searching for TxCoalescingTicks in HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Classes, and creating string value PreserveVlanInfoInRxPacket with the value set to 1. After a reboot of the client, the client will no longer receive the tagged voice VLAN.

Lastly, an alternative solution was discovered by the networking design group that created a link-local address for each SVI. This would allow for multiple link-locals to be learned by the clients, since the link-locals would now vary between SVI's instead of being identical. This was the most viable solution to implement that allows all desired functionality to persist.

The table below can be used as a quick reference to possible fixes and their impact:

	Group Policy change	IPv6 on phones	IPv6 on clients	Router configuration change
RA suppression	No	No	Yes	Yes
Driver changes	Yes	Yes	Yes	No
Registry edit	Yes	Yes	Yes	No
Link-locals on SVI	No	Yes	Yes	Yes

Authentication Evaluation

In order to find a method of authentication for IPv6, LSU looked into 802.1X as well as MAC-based authentication. During this testing and evaluation phase, LSU encountered several issues that will be detailed below.

In testing 802.1X [10], an IEEE standard for port-based network access control, LSU hoped to be able to tie a specific user to a port at specific times. This seemed more attractive than MAC-based authentication as any person could gain access to a machine registered under another user and perform nefarious acts.

802.1X was initially deployed on a couple dozen clients within the Computing Services Building for testing. Client operating systems included: Windows XP, Windows 7, along with various versions of Mac OSx and Linux. Windows clients on the domain were configured using group policy updates, but Mac clients had to be manually configured. In order to configure 802.1X for a Mac client, one of three profiles (login window, user, system) had to be chosen. The system profile was the only profile that would allow for a client to be connected to the network even without a user being logged on. This particular configuration was the only viable option for LSU as users need to be able to access workstations remotely. The problem with a system profile, however, is that the username and password are stored in the network preferences and thus have to be changed if administrators of active directory (AD) require periodic password changes. Users that forgot to update their password in the network preferences would lose connectivity as soon as the password change takes effect in active directory.

In several cases, client workstations are attached as a pass-through via VoIP phones. Thus 802.1X had to be enabled on the phones so the clients were then able to authenticate. Supplicants on VoIP phones had to be manually configured to allow for authentication. Also, the MAC addresses and passwords for the phones had to be entered into the authentication server (e.g. freeRADIUS).

While dealing with configuration on the client, it was determined that Windows XP clients had to be upgraded to Service Pack 3 and Windows Server 2003 clients had to be upgraded to SP2 in order to support 802.1X. Linux distributions tested, included Fedora, Ubuntu, Mepis, and Gentoo. Both Mepis and Gentoo needed a supplicant to be manually configured for authentication to be possible. A list of the software releases that were tested for each distribution can be found below:

- Fedora-13/i686
- Ubuntu 9.04 and 10.04
- Mepis 8.5

Bogus RAs from Unexpected/Trusted Sources

While deploying management IPv6 addresses, an irregularity was noticed in a building on campus. There were three layer 2 switches that had IPv6 addresses placed on the management VLAN interface, but only two of the three switches were reachable via their IPv6 address from outside the building.

First, the configuration was double checked to make sure that the address wasn't entered incorrectly. Once it was determined that the address had been properly entered, an attempt was made to ping the device from another switch within the same VLAN. This attempt was successful as well as a separate attempt from the distribution switch for that building. However, no devices outside of the building were able to ping the address. The issue was complicated by the fact that each distribution switch and all core routers had the address in their routing tables.

It was determined that the switch in question needed a second look and debugging (e.g. *debug ipv6 icmp*) was turned on in order to troubleshoot the issue. An excerpt of the log can be found below:

```
Feb 22 13:17:41 CST: ICMPv6: Received echo request, Src=2620:105:B000:7915::1,  
Dst=2620:105:B000:7915::72  
Feb 22 13:17:41 CST: ICMPv6: Sent echo reply, Src=2620:105:B000:7915::72,  
Dst=2620:105:B000:7915::1
```

....

Feb 22 13:17:57 CST: ICMPv6: Received Unreachable code 0, Src=2620:105:B000:7915::192, Dst=2620:105:B000:7915::72

The first two lines shown are ICMPv6 requests and replies between the building's distribution switch and the switch itself, but the last line shown seemed out of place. It belongs to a separate switch in the building and the error given in the log is generated by the router when there is no route to a particular destination. (RFC 2463) Upon further investigation, it was found that *ipv6 unicast-routing* had been accidentally enabled on the other layer 2 switch (2620:105:B000:7915::192) and that switch was sending out RA's to other devices within the building.

Tip

In order to prevent rogue devices from sending out bogus RAs, one of two actions have to be taken. If your hardware vendor has support for it, enable RA guard on each hosts' switchport interface. If your hardware does not yet support RA guard, the same goal can be accomplished by writing an ACL that denies ICMP type 134.

NETREG and DHCPv6 Options Bypass

When a machine comes into our network for the first time they must register their devices. This consists in capturing the users machine MAC address and their username through a portal known as NETREG[11]. Digging further into how this works, when an unregistered machine connects to the network, they get a 10.39.x.x address that has no connection to the world – only internal resources at LSU. In addition, these devices are restricted to a DNS view that only points to the NETREG portal to force users to register. Once registration is complete, our DHCP servers assign them a valid, public IPv4 address.

When IPv6 came around, an attempt was made to distribute our anycast IPv6 DNS address out to users by using the DHCPv6 options via the “other config flag.” This configuration was achieved by issuing the following IOS commands:

```
ipv6 dhcp pool IPv6_DHCP_Pool  
  dns-server <anycast address>  
  domain-name <university>.edu
```

```
interface vlan100  
  ipv6 dhcp server IPv6_DHCP_Pool  
  ipv6 nd other-config-flag
```

This configuration was later removed from our distribution routers after it was determined that users would be able to bypass NETREG. The main reason for it was because the DNS view cannot be reproduced. There is no such thing as a private address in IPv6 equivalent to RFC1918. (The closest equivalent in functionality to RFC1918 are unique local addresses (ULAs) described in RFC4193.) Therefore, instead of their traffic redirecting to the registration page, because of the availability of IPv6 DNS, their requests would resolve proper addresses. Websites would be unreachable because their IPv4 stack would not be able to reach the site unless it was a dual-stacked site. This would deter user experience and the expected redirect would never take place. Unfortunately, NETREG has not been ported to the IPv6 space yet.

Therefore, we rely on IPv4 DNS servers to provide IPv4 and IPv6 name resolution until further notice.

Native IPv6 Support

It should be noted that while deploying IPv6, not all devices will support IPv6 commands natively. For example, some Cisco devices require you to enter the following command before being able to enable IPv6 unicast-routing (Layer 3) or enter an IPv6 management address and turn on RA guard (Layer 2):

sdm prefer dual-ipv4-and-ipv6 routing (for Layer 3)

sdm prefer dual-ipv4-and-ipv6 default (for Layer 2)

Note: This is not the case for all platforms, as some support dual-stack natively. Check the vendor documentation to see if this step is necessary.

Route Summarization

Implementing IPv6 is a great stride forward in your organization, but you have to remember, we are limited by the laws of physics. Once you start turning up IPv6, everything takes up more space in memory. This will put more demand and strain on your gear. One way to extend the useful life of your existing gear without replacing it is summarization on your Layer 3 devices. What this does is reduce the size of your routing table which in turn, takes up less memory on your routing device. For example, half way doing our summarizing, this is what our stats look like:

<i>Route Source</i>	<i>Networks</i>	<i>Overhead</i>	<i>Memory (bytes)</i>
<i>ospf 2055</i>	<i>760</i>	<i>88800</i>	<i>94240</i>

Number of prefixes:

/0: 1, /8: 1, /40: 1, /42: 1, /52: 63, /64: 587, /128: 114

Again, proper design of your IPv6 space allows for ease of summarizing. In our case, /52's summarize many /64's. Instead of having over 1500 routes, we can condense our network down to less than 200.

Stateful Address Configuration with DHCPv6

If stateful address configuration is chosen using DHCPv6, it should be noted that Windows XP and Mac OSx versions prior to Lion do not have a native DHCPv6 client. More will be discussed on this topic in part 2 of this paper.

White/Black Listing

Early on, Internet giants like Google, Facebook, and others began using IPv6 internally in their networks. As more of the world started to catch on, they began advertising their AAAA

records to the world but in an obscure manner. This started the practice known as whitelisting[12].

With the advent of World IPv6 Launch, the behavior of these Internet giants changed. Without warning, they began practicing blacklisting [13]. Blacklisting, we all know, is not a new practice. Blacklisting to prevent DNS servers from receiving AAAAs is. LSU and other big universities running IPv6 on their networks quickly made this list. Without warning, we stopped receiving AAAAs. After many days of investigating, we ran into a Google website [14] that keeps track of which DNS servers are not dishing out AAAAs to. Due to Google citing privacy, little to no information has come out of them. Our investigations have led us to the following conclusions:

- You get on the list if Google detects any latency that will affect the user experience.
- A DNS server makes the list if enough of the clients this DNS server manages are slow in responding to traffic flows over IPv6 as it compares to IPv4. We suspect some kind of Java script gets loaded in the background randomly when going to different Google sites.
- As you get lower in the list, that means you are on your way out of the list.
- The list is updated daily.

With no guidance available, some of the problem areas we have encountered when going over to IPv6 are:

- Two buildings with misconfigured user VLANs for IPv6. This would not have been a problem if we did not block tunneling protocols, but we do.
- One router running out of memory to hold all IPv6 routes for LSU. Summarization solved this issue.
- RA guard macro command bug on certain platforms prevent solicitations and other IPv6 traffic from coming into the port. Command has been removed and an ACL has been put in its place to accomplish the same action.

Therefore, the message from Google is if you want to use IPv6, make sure it is working correctly. The bottom line is good user experience.

Another important fact about the Google blacklist is that other Internet giants are looking at it and using it to see who they block or not. From first hand experience and email communications, we found Facebook was indeed using the list to update their blacklists for AAAAs distribution.

Closing Thoughts

The spread and eventual adoption of IPv6 is inevitable. As technology continues to evolve, we are becoming more connected than ever. There is a future where everything, even our clothes will have sensors and will want to connect to something. How do we accomplish such feat? IPv6 could be a start. Therefore, it is vital that there be a network infrastructure ready to support the IP needs of new technology. IPv6 has growing pains just like any new technology, but they are necessary in order to move forward.

It is important to note that immersing one's self by being hands on is the best way to gain a strong understanding of how IPv6 works. Starting off with a small test network to become familiar with addresses and adding a few different clients will go a long way in the learning process. It is one thing to read about bugs with vendors and security issues; to be able to recreate occur and he or she will be ready to combat them in the real world.

Although, there is a tendency to continue to follow current IPv4 implementations when deploying IPv6, the writers of this paper encourage the reader to take IPv6 deployment as an opportunity to reevaluate their network. If there are certain design aspects that need to be changed or looked into, it is better that they be changed before IPv6 is deployed. With a new Internet protocol, the last thing you need is a bad legacy design decision to be prevalent.

There is still much to be learned and there are quite a few features of IPv6 that aren't being used just yet (e.g. jumbograms). Nevertheless, it is important for the networking community to continue to embrace IPv6 to keep the Internet open and free for the future.

Appendix

Appendix 1.1

Link-Local Addresses

According to RFC 4291 Section 2.5.6, link-local addresses are broken up as follows:

The first 10 bits – 111111010

The next 54 bits – All zeros

The final 64 bits – Interface ID

The hexadecimal equivalent of 111111010 is FE80, thus all link-local addresses begin with FE80.

The 54 bits of all zeros is abbreviated with ::

Appendix 1.2

Questions About Addressing Considerations

Upon receiving your new address space, you'll have several decisions to make before deployment is possible.

Where do I start?

Reference RFC6177 - "IPv6 Address Assignment to End Sites." Of course, nobody should believe anything anybody says about IPv6 addressing policy, but the available IETF/IANA/RIR recommendations are about as much clarity as can be obtained.

How long should you ponder on an addressing scheme?

You may or may not agree that it is better to proceed with an addressing plan than it is to agonize over getting it right on the first try. This is something several folks are fond of pointing out. For example, Alan Whinery from the University of Hawaii and Ron Broersma with DREN.

How would you remember an IPv6 address?

Simple answer: DNS. DNS reverses become more essential, as well as more complicated. Having coherent DNS per-hop naming in your traceroutes is (even more) useful with IPv6.

If you're planning on a dual-stack implementation, are you assigning subnets in with any tie in to the IPv4 space?

This question sparked some debate within our group. Initially, the group was divided between the two ideas proposed: assign addresses sequentially or assign addresses with clues tied into the IPv4 space it resides on. On one hand, it can be argued that starting off at the beginning of an address space is better as you don't waste any addresses. It can also be argued that it is better to make the IPv6 subnets recognizable by incorporating the

VLAN IDs of the associated SVI within the address. Even though you will inevitably skip addresses, it will make the subnets easier to work with and make troubleshooting issues easier. Therefore, devise a scheme that makes sense for your campus network. Build clues into it that will aid in troubleshooting.

Are you going to assign your entire IPv6 allocation from the start or will you leave a portion aside for future expansion?

It is understandable that taking a crystal ball approach to address assignment may be difficult, but it may prove to be important at a later date with IPv6. Much like with the inception of IPv4, there are many saying that we will never use the amount of addresses allocated, but it is impossible to know exactly what leaps or bounds technology may take. For this reason, it may prove beneficial in the future if some of the address space is left unassigned. Assume that your organization will have growth in terms of IPv6 address consumption even though you do not see it today.

Is your campus contained to one geographical location or are you responsible for multiple sites?

Campus size and how they are connected can heavily skew your decisions. As discussed previously, ISPs will not peer with sites smaller than a /48, thus it is important to pre-plan your address deployment or at least leave a portion of the space to the side to address remote sites at a later date. Request space higher than a /48 for your campus environment. As a university, you are considered like an ISP.

Appendix 1.3

Who We Are - Louisiana State University

Since 1860, LSU has served the people of Louisiana, the region, the nation, and the world through extensive, multipurpose programs encompassing instruction, research, and public service. The University brings in more than \$150 million annually in outside research grants and contracts, a significant factor for the Louisiana economy. [<http://www.lsu.edu/visitors/>]

Quick Facts [<http://www.lsu.edu/visitors/quickfacts.shtml>]

- LSU was founded by the Louisiana General Assembly in 1853 under the name Louisiana State Seminary of Learning and Military Academy and was located near Pineville, La., with the first session beginning Jan. 2, 1860.
- LSU is one of only 30 universities nationwide designated as a land-grant, sea-grant and space-grant institution.
- LSU is accredited by the Commission on Colleges of the Southern Association of Colleges and Schools to award bachelor's, master's, doctoral, and professional degrees.
- LSU includes 10 senior colleges and schools, in addition to specialized centers, divisions, institutes, and offices.
- As of the spring of 2009, LSU's enrollment is more than 26,000 students, including more than 1,400 international students and over 4,000 graduate students.
- LSU has more than 1,500 faculty members and a staff of more than 5,000.
- LSU Libraries contain more than 3.2 million volumes.

- The School of the Coast and Environment (formerly CCEER) was designated as the first Coastal Marine Institute by the Minerals Management Service of the U.S. Department of the Interior.

University Networking and Infrastructure

The mission of University Networking & Infrastructure (UNI) is to design, implement and manage solutions which provide for the University's voice, data and video communications. Services provided by UNI include installation and maintenance of metallic and optical cabling, network architectural design, software specification and configuration; and the specification, operation and maintenance of transmission hardware and software. [15]

Acknowledgements

The authors would like to extend their gratitude to the following persons for their contributions to this paper. Their time and dedication to this document was essential in its completion.

Patrick Cavell
Brian Middleton
Gary Mumphrey
Craig Callender
Dustin Mouton
Michael Sparks
Alan Whinery

References

- [1] <https://www.arin.net/announcements/2011/20110203.html>
- [2] <http://www.ipv6now.com.au/primers/IPv6Myths.php>
- [3] <http://tools.ietf.org/html/rfc4861>
- [4] <http://www.thc.org/thc-ipv6/>
- [5] <http://tools.ietf.org/html/rfc1918>
- [6] <http://www.ietf.org/rfc/rfc4890.txt>
- [7] <http://tools.ietf.org/html/rfc4941>
- [8] http://en.wikipedia.org/wiki/Digital_Millennium_Copyright_Act
- [9] <http://msdn.microsoft.com/en-us/library/windows/hardware/hh463937%28v=vs.85%29.aspx>
- [10] <http://www.networkworld.com/news/2010/0506whatisit.html>
- [11] <http://netreg.sourceforge.net/>
- [12] <http://en.wikipedia.org/wiki/Whitelist>
- [13] <http://en.wikipedia.org/wiki/Blacklisting>
- [14] http://www.google.com/intl/en_ALL/ipv
- [15] <http://itsweb.lsu.edu/UNI/>