NICS IPv6 Best Practices Guide

Recommendations for Deploying IPv6

Version: 1.3
Date: April 22, 2014
scott.hovis@nasa.gov

Document Change Log

Date	Version	Change Author	Affected Section	Change Description
1/15/2014	1.0	Scott Hovis	Change Log	Added document change log to track changes.
2/19/2014	1.1	Ben Hardy, Scott Hovis	Wireless	Added wireless section.
3/4/2014	1.2	Scott Hovis	Firewall Rules and Access Control Lists	Added example “Day 1” firewall rule set.
4/22	1.3	Scott Hovis	Host Networks	Added Deployment Considerations sub-section. Added notes to FHRP section based on lab tests.

Table of Contents

Document Change Log 1

Table of Contents 2

Table of Figures 4

Introduction 5

Purpose 5

Best Practices 5

- Order of Deployment 5
  - Minimizing User Impact 6
  - Efficient Use of Schedule 7
  - Proposed Order of Deployment 7
- Infrastructure 8
  - Securing the Management plane 8
  - Infrastructure Addressing 9
  - Infrastructure Links 9
  - First Hop Redundancy Protocols (FHRP) 10
  - Monitoring 11
- Host Networks 12
  - Address Assignment Use Cases 12
  - IPv6 Address Acquisition 13
  - Default Gateway 15
  - Web Redirection 17
  - Deployment Considerations 17
- Firewall Rules and Access Control Lists (ACLS) 17
  - Neighbor Discovery Protocol 18
  - Path MTU Discovery 18
  - Other Filter Considerations 19
  - BOGON Filters 21
  - Example “Day 1” Firewall Rule Set 22
- Routing 23
- Wireless 26
  - Cisco Wireless 27

Multicast 35
Multicast Addressing 35
Receiver/Router Signaling 36
RP Methods and MSDP 36
IOS Versions 37
Summary of Recommendations 39
References 41
Appendix A – Applicable JUNOS IPv6 43
Infrastructure 43
Securing the Management Plane 43
Infrastructure Addressing 45
Routing 45
First Hop Redundancy Protocol 46
Host Addressing 48
Firewall Rules and Access Control Lists 49

Table of Figures

Figure 1 – Dual Stack Web Connection 6
Figure 2 – Dual Stack Web Connection (with Happy Eyeballs) 7
Figure 3 – EUI-64 Address Generation 14
Figure 4 – IPv6 Neighbor Discovery 18
Figure 5 – IPv4/IPv6 Header Comparison 19
Figure 6 – IPv6 Extension Headers 20
Figure 7 – IPv6 Extension Header Example 20
Figure 8 - Example Outbound Firewall Rules 22
Figure 9 - Example Inbound Firewall Rules 23
Figure 10 – OSPF LSA Types 25
Figure 11 – IPv6 Next-Hop 26
Figure 12 - Example Wireless Architecture 27
Figure 13 - GUI WLC IPv6 Configuration 28
Figure 14 - Configuring Mobility Groups 29
Figure 15 - Configuring RA Guard 30
Figure 16 - Router Advertisement Throttling 31
Figure 17 - Address Type Distribution 32
Figure 18 - IP Clients over time 33
Figure 19 - Traffic by Client Type 33
Figure 20 - IPv6 Address Assignment Distribution 34
Figure 21 - Monitoring IPv6 Clients 34
Figure 22 - Client IPv6 Addresses 35
Figure 23 – IPv6 Multicast Addressing 36
Figure 24 - Embedded RP Example 37
Figure 25 - Feature/Version Dependencies 39

Introduction

The OMB FY14 mandate states that government agencies shall "upgrade internal client applications that communicate with public Internet servers and supporting enterprise networks to operationally use native IPv6 by the end of FY 2014." This is the intended end state to be reached in 9/2014. There are many steps that must be accomplished in the interim months.

In addition to the OMB mandate, NASA has marked the successful implementation of IPv6 for intranet networks an official Area of Emphasis (AoE) for the NICS contract. As such, the progress toward meeting the mandate is being reviewed by NICS management as an item on the MMR. The AoE project schedule has laid out several major milestones to ensure the OMB mandate is met on schedule. The schedule requires that much of the pre-implementation work towards IPv6 Intranet implementation be accomplished by the end of calendar year 2013 to allow enough time for the actual implementation tasks across the agency. This includes such tasks as:

1. Getting the Intranet addressing template approved (by Aug 31)
2. Submitting Center IPv6 Intranet plans (by September 30)
3. Getting CSO approval for Center plans (by October 31)
4. Completing inventory of Intranet equipment to ensure IPv6 capability and plan mitigations for shortfalls (TBD)
5. Provide guidance for configuring IPv6 on the Intranet (by November 30)

This document will address the last item by providing a set of best practices to use in the IPv6 Intranet deployment phase that will begin in January 2014.

Purpose

The purpose of this document is to provide a set of recommended best practices for implementing IPv6 consistently across the agency. In addition to standardization, this document will be a resource to reduce the amount of research and testing that must be done individually by each Center. This is not an all-encompassing IPv6 reference. Rather the intent is to provide enough background in each area to help the implementing engineer understand the reason for the recommendation. This document will focus on platforms and protocols listed in the Corporate Network Target Architecture (CNTA).

Best Practices

Order of Deployment

There are at least two things that should be considered when deciding in what order IPv6 implementation tasks should be accomplished. First and foremost care should be taken to minimize impact to users and secondly making most efficient use of schedule. The intent is to have all IPv6 services external to the LANs in place before the LAN assigns the first IPv6 prefix to a user facing LAN.

Minimizing User Impact

The deployment of IPv6 across the agency will be almost exclusively dual-stack. As it stands today the only deployed IPv6 systems are those that were deployed in response to the OMB 2012 public mandate. There has been no shortage of IPv4 addresses internally and virtually no early adoption of IPv6 only enclaves, which would have required tunneling and/or translation services that would have subverted agency security requirements. The introduction of IPv6 to the intranet will take time and will not be ubiquitous. Furthermore, there will still be IPv4 only resources internally and in the Internet that will require IPv4 host operation as there is no plan to deploy translation services.

There are inherent issues with dual-stack implementations. These are documented in RFC1671 and more recently in RFC6555. The issues are summarized in the question: When a system has addresses from two different protocols, how does it decide which to use?

Most modern operating systems enable dual-stack operations by default, including Windows 7 (deployed as the primary ACES seat). Unless IPv6 has been configured on the network infrastructure, the IPv6 portion of the stack is rendered useless as it cannot obtain a globally scoped address to use. However, once a globally scoped address is obtained, the OS has two protocols available for outgoing connections. When a DNS query results in both IPv4 (A record) and IPv6 (AAAA record) responses, the host will attempt an IPv6 connection first by default. In the case that the IPv6 path to the service is broken, or not completely implemented, this results in significant delays to the user despite the available, working IPv4 path.

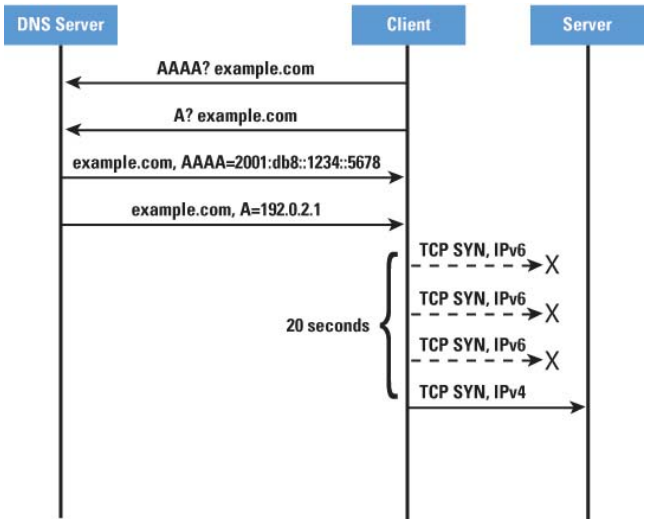

Figure 1 – Dual Stack Web Connection

Given this behavior, the order of IPv6 Intranet deployment needs to be carefully planned to avoid impacting users and creating a plethora of trouble tickets. The order of deployment should ensure that user devices do not attempt to use IPv6 for communication until the infrastructure is in place to support it end to end. The easiest way to accomplish this is to deploy IPv6 starting from the outside edge and

moving inward. Host networks should be the last to be IPv6 enabled to ensure hosts do not receive an IPv6 address too soon.

Another possible mitigation to this behavior is "Happy Eyeballs" (RFC6555), which is an attempt to solve the problem at the application level by not waiting on IPv6 timeout to initiate IPv4 communication. When both IPv4 and IPv6 addresses are returned for a service, Happy Eyeballs enabled applications initiate connections using both protocols simultaneously. Whichever connection is established first is used. Some of the latest versions of popular web browsers already incorporate this algorithm. It would be advantageous for ACES to deploy "Happy Eyeballs" capable browsers prior to IPv6 implementation.

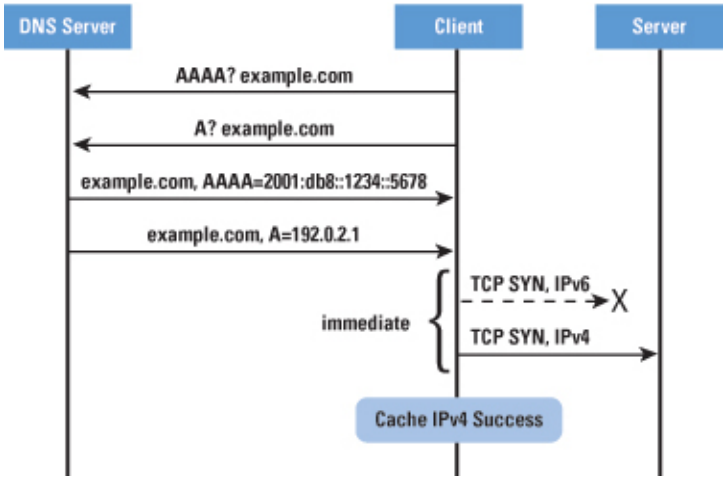

Figure 2 – Dual Stack Web Connection (with Happy Eyeballs)

Efficient Use of Schedule

Once the pre-work is completed real work of deployment will begin. As it turns out the amount of effort required by the various entities is not the same. Center LANs will have several steps to accomplish as the IPv6 deployment will be new for Intranet services. This includes upgrading code across the infrastructure for IPv6 features, in some cases upgrading hardware platforms will be necessary, configuration of IPv6 on all infrastructure links, deploying IPv6 security protections to the management plane, proxies etc.

While the Center LANs will be hit with the brunt of the IPv6 deployment effort, the network core (WAN) and other essential services such as DNS will not require as much effort because of the positioning they were required to achieve to meet the previous 2012 mandate. Agency DNS servers already function for IPv6 and the WAN infrastructure was addressed and implemented to provide connectivity to public facing services. What work remains for the Core services should be scheduled to take place while the LANs are in the process of upgrades, etc. This will facilitate the edge-inward model discussed previously.

Proposed Order of Deployment

Pre-Steps

I. Get IPv6 address plans approved.
II. Provide best-practices and guidance documentation to LANs including recommended device code versions.

Steps

1. Establish peering of Center Intranet address allocations. In most cases this is adding routes and network statements vs. actual new peering which was stood up for public services in many cases.
2. Deploy IPv6 DDI capability at the centers. This includes DNS and DHCP servers managed by the central DDI group.
3. Address Center infrastructure and ensure applicable security measures are in place on management connections.
4. Enable IPv6 on Center firewalls and implement IPv6 rule set. Test connectivity from Center to the Internet.
5. Add DHCP allocations and implement IPv6 addressing for user-facing LANs.

Infrastructure

Securing the Management plane

Enabling IPv6 in the infrastructure instantly doubles the attack surface of the network. All of the threats of the IPv4 world still exist and a new protocol stack is available to would-be attackers. Therefore the IPv6 management and control planes should be protected prior to configuring the first IPv6 address on the device. As a general rule any IPv4 controls on the device should have an IPv6 equivalent. This includes configuring IPv6 ACLs and applying them to the VTY lines and SNMP groups.

The OMB mandate does not address network management. One approach is to block all IPv6 management. This especially makes sense if the network management platforms are not yet IPv6 capable. These can be worked, if desired, after the mandated IPv6 services have been implemented. The following shows the disabling of IPv6 management in IOS.

```
! Disable HTTP connections via IPv6
no ip http server
no ip http secure-server
!
ipv6 access-list Deny_IPv6_Traffic
 remark deny all IPv6 traffic
 deny ipv6 any any
!
line vty 0 15
 ipv6 access-class Deny_IPv6_Traffic in
!
! Deny IPv6 SNMP. The 88 is the existing IPv4 ACL.
snmp-server group NETmgmt v3 noauth notify SNMPv3 access ipv6 Deny_IPv6_Traffic 88
snmp-server community netmgmt RO ipv6 Deny_IPv6_Traffic 88
```

Note: Lab testing revealed an issue with securing SNMP on Cat3k switches running 12 code. It would allow only and IPv4 or an IPv6 ACL to be applied, not both.

Infrastructure Addressing

In IPv6 subnets always have a prefix length of 64 bits with a couple of exceptions. One exception is 128 bit loopback addresses. The other exception is the 127 bit point-to-point link (RFC 6164). NASA has chosen to use the 128 bit loopback but not the 127 bit point-to-point link. This means all networks will be /64 or /128. There should be no carryover of VLSM usage from the existing IPv4 infrastructure. Address conservation is not the priority as it was in IPv4.

It is recommended to use manually configured infrastructure addresses vs. EUI-64 addressing where the host portion is derived based on the interface MAC address. If EUI-64 addressing is used on infrastructure links, equipment remove and replace activities will result in a new address. Any references to the old address in ACLs, PIM filters, etc. will have to be updated with the new router's EUI-64 based address. It is also advantageous to manually set the link-local addresses on infrastructure links rather than allowing the default EUI-64 host portions so that the next hop IPs for routes are more human readable.

Infrastructure Links

The following commands should be configured on router to router interfaces. The "ipv6 nd ra suppress all" command suppresses router advertisements on the link to prevent hosts from learning the prefix and obtaining addresses. The "all" keyword was introduced in IOS 15.x is not available in IOS 12.x releases so an upgrade may be required. Without this keyword the router suppresses periodic Router Advertisements (RAs) but will still send RAs in response to host Router Solicitations (RS) allowing host address configuration. If the all keyword is not available and an upgrade to capable code is not possible, additional configuration, such as clearing the A-bit, should be done to prevent the use of SLAAC on the link. These steps are discussed in the host networks section.

```
int G0/0
 ipv6 address 2001:db8:1000::1/64
 ipv6 address fe80::1 link-local
 ipv6 nd ra suppress all
 ipv6 nd prefix 2001:db8:1000::/64 300 300 no-autoconfig
 no ipv6 redirects
 no ipv6 unreachables
```

The link local address can be manually set on the router as shown rather than using the default EUI-64 link local address. This has two benefits: 1) easier identification of the router as seen in routing neighborships, next hops, etc and 2) the address will not change as the result of hardware replacement, which could impact ACLs. The same link-local address can be used on multiple interfaces as long as they are not connected to the same link/subnet. The last two commands are IPv6 equivalents for common IPv4 commands to disable ICMP redirects and unreachable messages.

In addition, the following global configuration command should be used (if not on by default) to protect the control plane by limiting the number of IPv6 ICMP error messages the router will generate. This example limits ICMP error messages to 10 per second. This command is on by default in some versions. Use the command "show ipv6 int <interface>" to determine the current setting.

```
ipv6 icmp error-interval 100 10
```

First Hop Redundancy Protocols (FHRP)

As in IPv4, FHRPs are available in IPv6 to provide automatic failover of router interfaces that serve as IPv6 gateways. See this link for Cisco IPv6 feature/IOS version mappings.

There are differences between IPv4 and IPv6 host/router operations that drive changes to FHRP operations in IPv6. One difference is the distribution of gateway information to hosts. While default gateways are typically configured on IPv4 hosts statically or provided by DHCP, IPv6 hosts learn about router gateways through periodic or solicited router advertisements (RAs). Another difference is the addition of link-local scoped addresses in IPv6. In most cases, hosts will use the link-local address of the gateway rather than the global scoped address. Because of these differences there are two subtle changes to the configuration of FHRPs in IPv6.

- Enabling an IPv6 FHRP suppresses RAs containing the physical address of the router so that IPv6 hosts only learn the VIP as a gateway. When configured, only the active gateway sends RAs. The FHRP RAs are affected by other RA configuration settings such as interval, priority, etc. Note that RAs should not be suppressed on an interface participating in a FHRP. Doing so will prevent the FHRP from functioning correctly and hosts will not receive the VIP as a gateway address.
- The configured virtual IP address (VIP) is by default a link-local address (FE80::/10).
    - There is a provision for HSRP in newer IOS versions to allow a global scoped VIP. The application of this feature is for situations where the FHRP is used between routers as opposed to gateway redundancy for host networks. In this situation the VIP may need to be sent upstream in routing updates and therefore a link-local address would not be sufficient. See this link for more information.

Hot Standby Router Protocol (HSRP) (version 2), Gateway Load Balancing Protocol (GLBP), and Virtual Router Redundancy Protocol (VRRP) (version 3) are currently supported in IOS for IPv6 operations. Other than the VIP configuration, other options have largely stayed the same. Note that existing HSRPv1 configurations for IPv4 will have to be modified to version 2. The conversion will cause a momentary outage for IPv4 hosts. GLBP is not supported on the Cat3k switches.

Note the following restrictions for configuring an IPv6 FHRP:

- When enabling an IPv6 FHRP on an interface with existing FHRP configuration for IPv4, the IPv6 configuration must use a separate group number from the IPv4. For example if the IPv4 HSRP configuration is configured as standby group 1, the IPv6 HSRP group will have to be something other than 1.
- Although link local addresses can typically be reused on multiple interfaces, this does not apply to link local VIPs for FHRPs. If manually assigning the FHRP VIP, the VIP must be unique from interface to interface but must match between routers sharing the VIP.
- HSRP and VRRP for IPv6 require newer versions of the protocols to be enable. HSRP must be version 2, VRRP must be version 3 (see examples)

```
!
! GLBP example
!
!
interface fastethernet 0/0
 ipv6 address 2001:DB8:0001:0001:/64
 glbp 10 ipv6 autoconfig
!  The link-local VIP address can be manually set if desired but must match between routers.
 glbp 10 timers 5 18
 glbp 10 priority 254
 glbp 10 preempt delay minimum 60
 glbp 10 authentication md5 key-chain AuthenticateGLBP
 glbp 10 weighting 110 lower 95 upper 105
 glbp 10 weighting track 1 decrement 10
!
! HSRPv2 example
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:CAFE:2100::BAD1:1010/64
 standby version 2
 standby 201 ipv6 autoconfig
 standby 201 priority 120
 standby 201 preempt delay minimum 30
 standby 201 authentication ese
 standby 201 track Serial0/1/0 90
!
! VRRPv3 example
!
! Enable VRRPv3
fhrp version vrrp v3
!
interface FastEthernet0/0
 vrrp 1 address-family ipv6
   ! address can be global or link-local
   address FE80::1 primary
   timers advertise 100
   prempt
   exit-vrrp
```

Monitoring

Other aspects of the infrastructure that need to be evaluated are the various monitoring and troubleshooting tools. These vary from center to center so it is not feasible to provide a complete list of tools and required versions to ensure IPv6 support. Therefore, this section should serve as a reminder to check locally deployed tools such as those in the following list. Note that IPv6 management is not required as part of the FY14 mandate, so IPv6 capability may not be required in some management tools.

- NAM modules
- Netflow (version 9 required for IPv6 flow records)
    - Evaluate router platforms for feature support.
    - Evaluate netflow collector tools for support receiving and processing IPv6 flow information.
- Sniffers (Infinistream, Riverbed, etc)
    - The capture length may need to be increased to allow for the increased header size in IPv6.
- SNMP monitoring tools (What's up gold, Spectrum, etc)
- Ciscoworks

Host Networks

To ensure proper operation on an IPv6 enabled network, all Agency internal hosts shall be capable of the operations and features included and referenced in RFC 4294. Failure to comply with this minimal set of IPv6 standard operations could result not only in a broken or sub-optimal IPv6 deployment, but could also lead to a degraded user experience despite a stable and working IPv4 infrastructure. A few of the main topics listed in RFC 4294 are:

- Dynamic Host Configuration Protocol (DHCP) – The ability to acquire an IPv6 address using DHCP.
- Neighbor Discovery – The ability to resolve Layer 3 to Layer 2 address for IPv6 using ICMP types 135 and 136.
- Router Discovery - The ability to send Router Solicitations and receive Router Advertisements (both solicited and unsolicited). The inability to receive/process RAs can lead to broken communication with hosts on the same link.
- Path MTU Discovery/Processing fragmentation headers - The ability to process ICMPv6 "packet-too-big" messages and to process fragmentation headers is essential in IPv6 since IPv6 packets cannot be fragmented in transit by routers. All fragmentation is required to be done on the sending host for IPv6.

In addition to these requirements, it is preferable that all client web browsers implement RFC 6555, "Happy Eyeballs" to improve user experience. RFC 6555 capable versions of most popular web browsers are already available.

Some types of clients will require different methods of address acquisition on CSO-managed networks. These are described below in the following use cases.

Address Assignment Use Cases

Clients

Agency clients shall be configured to acquire an IPv6 address dynamically via DHCP from the CSO managed DHCP server. This is independent of the addressing mechanism used for IPv4. The hosts should be configured such that they do not generate temporary addresses. This is to ensure the ability to track

hosts by the DHCP assigned IPv6 addresses. For the same reason, SLAAC and stateless DHCP should not be used on NICS managed, internal client LANs.

Servers and Printers

In addition to clients, there are other types of devices that will require static IPv6 addresses. Two such cases are servers and printers. There are different ways to derive the host portion of an IPv6 address to use for static assignment listed below. Different situations may call for one method or the other. However, the manual EUI-64 method is only listed for completeness. It should not be used as it ties the IPv6 address to a MAC address and thus to specific hardware. If this method is used, the remove/replace or refresh of the device would require changes to DNS, firewall rules, etc. or would lead to confusion when the IPv6 address and MAC no longer match.

- Pick any unused address from the range.
- Randomly generate the host portion
- Embed the host's IPv4 address as the host portion
- Convert the host's IPv4 address (or some portion of the address) to hex
- Manually execute EUI-64 process and assign statically.

If any application in use on the Agency network will be IPv6 enabled and the application requires firewall exceptions, the parties responsible for the application will be required to submit corresponding firewall exceptions for the IPv6 addresses of the application. This should be accomplished prior to IPv6 deployment on the network to prevent problems with the application during/after deployment.

IPv6 Address Acquisition

Host address acquisition in IPv6 is more complicated than in IPv4. In addition to new address acquisition methods, there are differences in how hosts acquire default gateway information. The different methods of each will be discussed briefly, but focus will be mainly on DHCPv6. The majority of hosts in the agency will acquire IPv6 addresses via DHCP from the CSO managed DHCP/DDI service.

An understanding of Router Advertisements (RAs) is needed before discussing the various address acquisition methods. RAs are ICMPv6 type 134 messages that are sent periodically by IPv6 enabled routers on each IPv6 enabled interface. RAs are also sent asynchronously as the result of a Router Solicitation message from a host. The RA includes flags used by hosts in address acquisition and default routing information. The following is a breakdown of RA flags as described in RFC 4861:

- Managed Address Configuration Flag (M-bit) – "When set, it indicates that addresses are available via Dynamic Host Configuration Protocol (DHCPv6). If the M flag is set, the O flag is redundant and can be ignored because DHCPv6 will return all available configuration information."
- Other Configuration Flag (O-bit) – "When set, it indicates that other configuration information is available via DHCPv6. Examples of such information are DNS-related information."
- Autonomous Address Configuration Flag (A-bit) – When set, it indicates that the associated prefix can be used for Stateless Address Auto-Configuration (SLAAC).

SLAAC

Stateless Address Auto-Configuration (SLAAC) occurs when a host receives an RA containing an IPv6 prefix with the A-bit set. The host performs SLAAC by forming a unique host identifier and appending it to the advertised prefix resulting in a unique IPv6 address. The uniqueness of the address is verified by Duplicate Address Detection (DAD). There are several methods for forming the unique host identifier: EUI-64, Cryptographically Generated Address (CGA), random (used in Microsoft temporary addresses). EUI-64 is the most common and is based on the MAC address of the host. An EUI-64 host identifier is created by adding the hex string FFFE in the middle of the MAC address and flipping the 7th most significant bit (see figure below).

The benefit of SLAAC is the ease of deployment – no manual config or DHCP server required for addressing and most hosts and IPv6 routers are setup for SLAAC by default. However, without the DHCP server, additional configuration information such as DNS server, domain, etc. cannot be transferred to the host. The hybrid solution, stateless DHCP, is discussed in the next section. Another drawback of SLAAC in an enterprise is the inability to track addresses, especially if Microsoft temporary addresses, which are on by default, are enabled. Microsoft temporary addresses prevent host tracking by periodically changing the host portion of the IPv6 address rather than using the EUI-64 identifier based on MAC address. For these reasons SLAAC will not be used on the majority of NICS managed networks. There may be corner cases, such as lab networks, where SLAAC would be acceptable. In such cases the client addresses would not be easily tracked and would require manual configuration of DNS, etc.

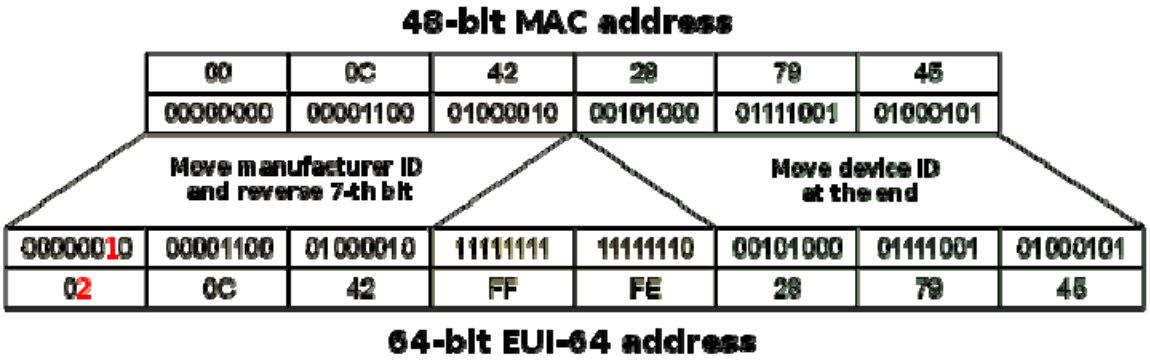

Figure 3 – EUI-64 Address Generation

Stateless DHCPv6

SLAAC by itself does not provide the host with other useful information such as DNS servers, domain name, etc. Stateless DHCP is an augmentation of SLAAC that allows hosts to receive configuration information such as DNS servers, domain name and other non-IP address information from a DHCP server while using SLAAC to derive an IPv6 address. To inform a host to perform stateless DHCP the router is configured to send RAs with the O-bit set and a prefix with the A-bit set. The O-bit tells the host to contact the DHCP server for other configuration information. The concerns in SLAAC with host address tracking and temporary addresses remain with Stateless DHCPv6. Given the requirement for a DHCP server and the added issues with address tracking, DHCPv6 (stateful) is the preferred method of client addressing on NICS managed networks.

DHCPv6 (Stateful)

Stateful DHCPv6 is analogous to normal DHCP in IPv4. The host contacts the DHCP server for both IP addressing and other configuration items. To inform a host to perform DHCP, the router is configured to send RAs with the M-bit set. This tells the host to contact the DHCP server. One benefit of using DHCP is the ability to track addresses. Note that different OS's behave differently with regard to the other bits (O and A) when the M bit is set. For example some OS's that receive the M-bit and a prefix with the A-bit will acquire a stateful address via DHCP and generate a SLAAC address. In IPv6 an interface can have multiple IPv6 addresses.

Stateful DHCPv6 will be used by the majority of host LANs in the agency. This is necessary to meet security requirements for tracking addresses. In addition it alleviates the inefficient and error prone process of manually assigning 128bit IPv6 addresses. The following example shows how to ensure DHCPv6 addresses are used by enabling the M-bit while turning off the A-bit, which is on by default. The DHCP relay command is analogous to the helper address command in ipv4 to relay DHCP operations to an off-link DHCP server. In the example, the RA interval is also decreased from the default of 200 seconds to 30 seconds.

```
! Stateful DHCPv6 interface example
!
interface FastEthernet0/0
 ipv6 address 2001:DB8:CAFE:2100::1/64
 ipv6 address fe80::1 link-local
 ipv6 nd managed-config-flag
 ipv6 nd prefix 2001:DB8:CAFE:2100::/64 300 300 no-autoconfig
 ipv6 dhcp relay destination 2001:DB8:CAFE:11::9
 ipv6 nd ra interval 30
```

Note that clearing the A-bit using the "no-autoconfig" keyword is preferred to preventing the advertisement of the prefix using the "no-advertise" keyword. The reason is the impact to on-link operations. When a prefix is advertised in the RA, the receiving hosts designate the prefix as "on-link", which means that any hosts with an address in the prefix can be reached directly through typical neighbor discovery processes. If the prefix is not advertised, the host does not designate the prefix as "on-link" and will not attempt neighbor discovery for other hosts in the prefix. This means all packets that could be sent directly on the same link are rather forwarded to the router. This results on the router needlessly handling local traffic. Depending on the router and host configuration, ICMP redirects may or may not be sent/processed.

Default Gateway

In addition to dictating host address acquisition, RAs are used by hosts to determine what routers to use as default gateways. Unlike IPv4, default gateway information is not distributed by DHCP servers in IPv6. Hosts learn about gateways by hearing RAs. This introduces a couple problems: 1) on a host with multiple or redundant gateways, which should the host use and 2) what happens when RAs are introduced by an unintended gateway, either accidentally or maliciously?

For the case of redundant routers refer to the previous FHRP section. If the routers on a link are not a FHRP group, the router preference field in the RA could be manipulated to have hosts prefer one router over the other. The cisco default is medium. The following example shows the router preference being changed to high.

```
interface FastEthernet0/0
 ipv6 address 2001:DB8:CAFE:2100::1/64
 ipv6 address fe80::1 link-local
 ipv6 nd router-preference high
```

RFC 6104 introduces the problem of rogue RAs misleading hosts into using the wrong gateway. This could be the result of an attack or simple misconfiguration. Either way it could cause connectivity problems for all hosts on the LAN. RFC 6105 proposes RA Guard as a possible solution to the problem. In short, RA guard is a way to define legitimate RA sources and block others from being sent into the network. This requires the implementing platform to listen for RAs on all ports and only forward those that meet the configured criteria. More information on the cisco feature can be found at this link. IOS 15.x code is required for this feature on some platforms.

Note: Lab testing uncovered a Cisco internal bug (**CSCtd32401**) that affects the 6500 platform running a SUP720 with either the PFC3C or PFC3CXL. The bug prevents RA guard from working while MLD snooping, which is on by default, is enabled.

Implementing RA guard on non-router facing ports is recommended. The example below shows the simplest configuration. This configures the port as a "host" port through which no RAs are allowed to be injected into the network. There are additional features in the documentation that allow for the definition of RA profiles so that RAs can be allowed conditionally, depending on the contents of the RA.

```
! If this platform is affected by bug CSCtd32401 disable MLD snooping
! no ipv6 mld snooping
!
interface FastEthernet0/0
   ipv6 nd ra guard
```

There are other solutions, namely Secure Neighbor Discovery (SeND), that address this problem by authenticating RAs and other ND messages. However, the fact that it is not implemented by Microsoft OS's renders it useless for this discussion.

If there are any sites that currently implement DHCP snooping for IPv4, DHCPv6 snooping is on Cisco's roadmap. Refer to this link for feature roadmap information.

Web Redirection

Many Centers use redirection for proxies and/or content filters. It appears that IPv6 capability for WCCP is first implemented in IOS 15.1(1)SY1. If WCCP is used at your center you may need to evaluate an upgrade path to 15.1 or 15.2 code. In addition, work with whoever manages the proxies and content filters to see if they support IPv6. This needs to be evaluated early as it may take some time to upgrade the various components needed for content filtering.

Deployment Considerations

The Cisco default RA settings (A-bit = 1, M-bit = 0) conflict with the desired NICS settings (A-bit = 0, M-bit = 1). The default RA setting would result in all IPv6 enabled hosts to acquire SLAAC and/or temp addresses which are not trackable via DHCP logs. The previous examples in this section have shown how to configure the NICS recommended settings, however, the examples show the end-state router configurations as would be seen in show running-config output.

If the commands are actually entered in this order it will cause problems. As soon as the IPv6 address is assigned to an interface, the router will begin sending RAs with default settings causing hosts to derive SLAAC and/or temp addresses with a valid and preferred lifetime of 30 days. For this reason some care should be taken to prevent this scenario from playing out. There are two methods of preventions:

- Configure RA suppression before beginning IPv6 interface config. Remove RA suppression only after config is complete. Example:

```
interface GigabitEthernet0/0
 ipv6 nd ra suppress all  ! Enable RA suppression
 ipv6 address 2001:db8:café:2001::1/64
 ipv6 nd prefix 2001:db8:cafe:2001::/64 300 300 no-autoconfig
 ipv6 nd managed-config-flag
 <other IPv6 configuration>
 no ipv6 nd ra suppress all  ! Disable RA suppression
```

- Configure M-bit and A-bit settings first, before configuring an IPv6 address or using the “ipv6 enable” command on an interface.

```
Interface GigabitEthernet0/0
 ipv6 nd ra suppress all  ! Enable RA suppression
 ipv6 address 2001:db8:café:2001::1/64
 ipv6 nd prefix 2001:db8:café:2001::/64 300 300 no-autoconfig
 ipv6 nd managed-config-flag
 <other IPv6 configuration>
 no ipv6 nd ra suppress all  ! Disable RA suppression
```

Firewall Rules and Access Control Lists (ACLS)

There are a couple differences in IPv6 that drive changes to traffic filtering. One difference is the neighbor discovery (ND) (RFC 4861) process which is analogous to IPv4 Address Resolution Protocol (ARP). The other is the requirement for Path MTU (PMTU) Discovery (RFC 1981).

Neighbor Discovery Protocol

Neighbor Discovery (ND) is the IPv6 equivalent to ARP. Since there is no concept of broadcast in IPv6, ND relies on link local scoped multicast for transport. Rather than implementing a separate protocol such as ARP, ND operates via IPv6 ICMP messages (types 135 and 136).

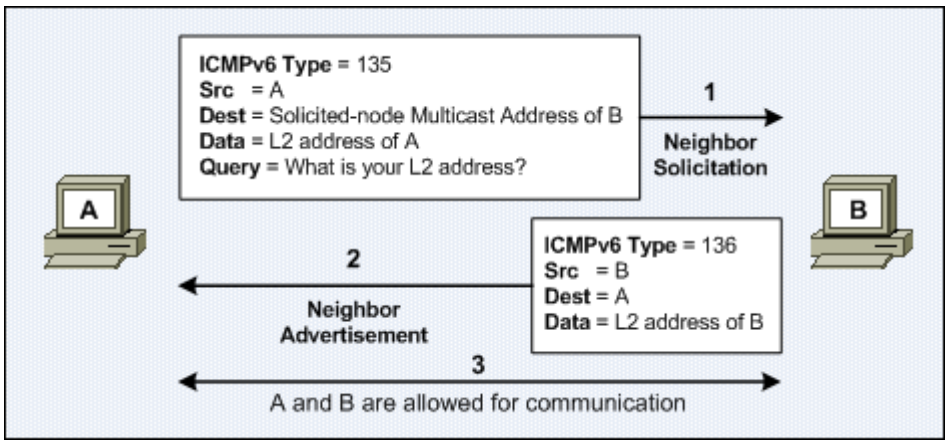

Figure 4 – IPv6 Neighbor Discovery

The host that initiates the ND sends an ICMPv6 type 135 message to the solicited node multicast address of the other host. The solicited node multicast is used in place of an all-nodes broadcast. A solicited node multicast address is created by taking the last 24 bits of a unicast or anycast address and appending them to the prefix ff02:0:0:0:0:1:ff00::/104. Therefore, a host always knows the multicast address a destination host will be listening to.

The use of ICMPv6 messages as an alternative to a separate protocol such as ARP means that any ACLs or firewall filters must allow ICMPv6 types 135 and 136 for protocol operation. Note that Cisco IPv6 ACLs by default contain implicit rules to allow IPv6 ND messages (Cisco). They are applied before the implicit deny at the end of the access-list. However, if an explicit deny is configured by the user, these rules are no longer applicable. Therefore, any IPv6 ACL with an explicit "deny all" at the end or with a deny for the link-local address range (FE80::/10) should manually account for permitting ND messages higher in the ACL to ensure protocol operation. Example of permitting ND before explicit deny:

```
ipv6 access-list IPv6_IN
 permit icmp any any nd-na
 permit icmp any any nd-ns
 deny ipv6 any any
```

Path MTU Discovery

Another difference is the requirement to allow certain ICMP message types end to end. In the IPv4 world ICMP was often completely blocked at the border. Unlike IPv4, IPv6 packets cannot be fragmented by routers in transit. Only the sending host is able to perform fragmentation (RFC 2460 Sec 5) in IPv6. In order for this to work IPv6 relies on PMTU Discovery which was optional in IPv4.

To allow PMTU Discovery to operate, the following ICMP message types must be allowed by firewalls and ACLs. In the case of firewalls, these messages should still be statefully inspected to ensure they are the result of an actual data flow. The last type in the list is not required for PMTU discovery but could still be useful in some circumstances.

- ICMP type 2 – packet too big
- ICMP type 3 – time-exceeded
- ICMP type 4 – parameter-problem (optional)

Other Filter Considerations

Source Routing and Extension Headers

A comparison of the IPv4 and IPv6 headers reveals that that the IPv6 header contains several fewer fields and options. The increase in address length makes the overall IPv6 header larger. All non-essential fields were removed from the main header and moved into extension headers. Extension headers are additional headers that are added to a packet only when necessary.

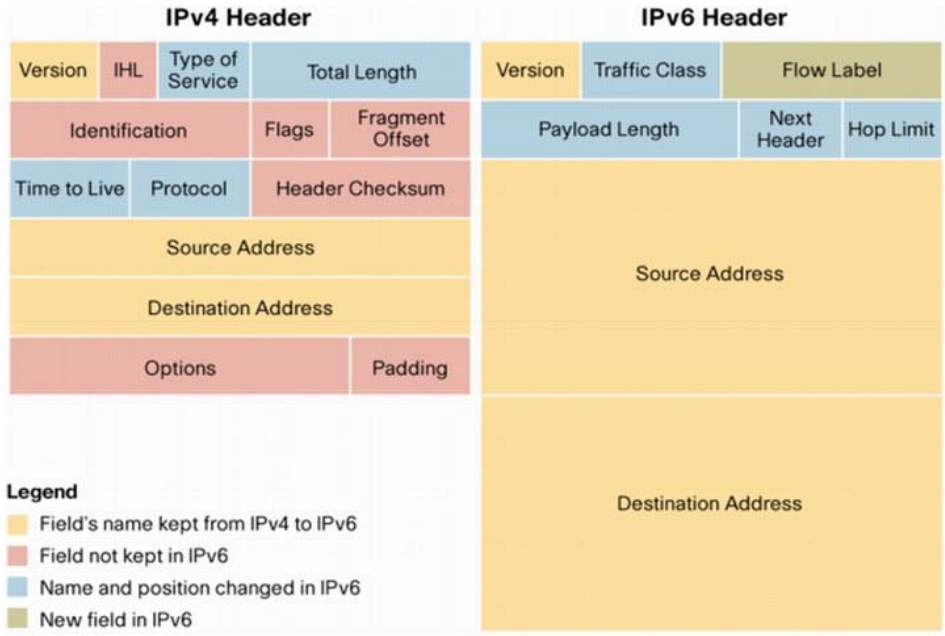

Figure 5 – IPv4/IPv6 Header Comparison

Extension headers are an intrinsic part of IPv6. The different types of extension headers are listed in the table below. Notice the Fragmentation Header. Anytime a sending host has to fragment a packet, the fragmentation header must be used as there are no fragmentation flags/offset fields in the standard IPv6 header.

Order	Header Type	Next Header Code
1	**Basic IPv6 Header**	-
2	**Hop-by-Hop Options**	0
3	**Destination Options (with Routing Options)**	60
4	**Routing Header**	43
5	**Fragment Header**	44
6	**Authentication Header**	51
7	**Encapsulation Security Payload Header**	50
8	**Destination Options**	60
9	**Mobility Header**	135
	No next header	59
Upper Layer	TCP	6
Upper Layer	UDP	17
Upper Layer	ICMPv6	58

Figure 6 – IPv6 Extension Headers

As shown in the figure below, when extension headers are employed, the standard header points to and describes the type of the next header. This process is repeated by each subsequent extension header until the end of the chain is reached.

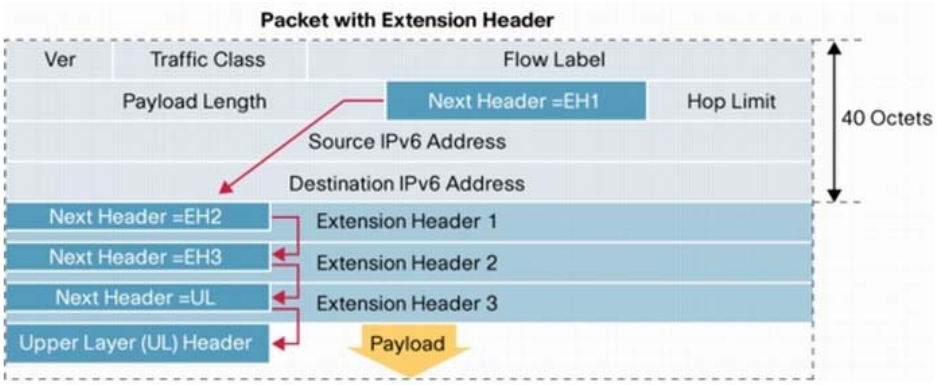

Figure 7 – IPv6 Extension Header Example

While Extension headers are necessary for protocol operation, they can also be used as an attack vector. A malicious host could attempt to hide information by chaining a large number of headers together to go beyond the inspection bounds of ACLs and firewalls or cause increased utilization by causing filtering devices to examine all of the additional header information. Therefore, it is recommended to block any truly unnecessary headers.

One extension header that should be blocked is the routing-header type 0, which was intended for source routing. RFC 5095 officially deprecates the type 0 routing header stating that it creates such a

severe threat that it should be blocked until it is removed from all host implementations. Most versions of code have a global switch for disabling source routing. For other versions of code this should be blocked by ACL. Routing header type 2 is not the same threat as type 0, however, it is used for IPv6 mobility which is not currently planned for implementation. Both type 0 and type 2 routing headers are currently filtered at the WAN/Internet border.

```
! Global command (if supported)
no ipv6 source-route
!
! otherwise apply via ACL
!
Ipv6 access-list BLOCKRH
 deny ipv6 any any routing-type 0  ! blocks type 0 routing header
 ! deny ipv6 any any routing        ! alternatively block type 0 and type 2 (mobility)
 permit ipv6 any any
!
interface FastEthernet0/0
 ipv6 traffic-filter BLOCKRH in
!
interface FastEthernet0/1
 ipv6 traffic-filter BLOCKRH out
```

Tunneling Protocols

Many host platforms support IPv6 tunneling. IPv6 host tunneling provides an attack vector that would allow inside and outside hosts to communicate without their conversation being subjected to infrastructure filters. Since the agency is following the OMB mandate to build a native IPv6 infrastructure, there is no need for host-to-host or site-to-site tunneling. As such the various tunneling methods should be blocked. Some are identified by prefix while others by protocol/port.

- Deny 6to4 prefix (2002::/16)
- Deny Teredo prefix (2001::/32)
- IPv4 filters for IPv6 tunneled traffic
    - Deny IPv4 protocol 41
    - Deny IPv4 UDP port 3544 (Teredo)

BOGON Filters

This section is included for reference as BOGON filtering is most affectively applied at the enterprise edge and has been configured on the WAN/Internet peering routers. A BOGON is a prefix that should never appear in the internet routing table. Any traffic from such a prefix is invalid. Cymru is an organization that maintains a list of IPv4 and IPv6 BOGONs. The internet usage of IPv6 is still sparse enough that the IPv6 BOGON list is based on explicit permits and a blanket deny where IPv4 is the opposite. Example of BOGONs are the documentation address (2001:0DB8::/32), deprecated Site-local address space and all other unassigned IPv6 address space. See this link for the current list: http://www.team-cymru.org/ReadingRoom/Templates/IPv6Routers/ios.html.

Example "Day 1" Firewall Rule Set

Based on the filtering information presented above, the following table provides an example that could be used as a starting point for developing a "Day 1" firewall rule set. This table is not represented as complete or as addressing every scenario or security requirement of the various firewall rule approval authorities in the Agency.

Outbound Rules				
Source	**Destination**	**Protocol**	**Action**	**Notes**
Any	Any	IPv6 with source-route header	Deny	Deny all packets with source-route extension header.
Any	2002::/16	IPv6	Deny	Deny all packets to 6to4 tunneling prefixes.
Any	2001::/32	IPv6	Deny	Deny all packets to Teredo tunneling prefix.
Any	Any	IPv4/UDP port 3544	Deny	Deny IPv4 teredo packets carrying IPv6 tunneled data.
Any	Any	IPv4/41	Deny	Deny IPv4 packets carrying IPv6 in a manually configured v6in4 tunnel. * Any intended tunnels would need specific exception.
Fe80::/10	Any	IPv6/ICMP Types 135,136	Permit	Permit link local traffic required for neighbor discovery.
2001:04D0:xxxx:2000::/52 2001:04D0:xxxx:4000::/52	Any	IPv6/HTTP	Permit	Allow center intranet wired and wireless clients to execute IPv6 web traffic. * Could exclude traffic to other Centers depending on SOC readiness.
2001:04D0:xxxx:2000::/52 2001:04D0:xxxx:4000::/52	Any	IPv6/DNS	Permit	Allow center intranet wired and wireless clients to execute IPv6 DNS queries.
2001:04D0:xxxx:0/52	Any	IPv6/ICMP types 2,3	Permit	Allow routed infrastructure to send ICMP responses required for Path MTU discovery.
Specific Host	Specific Host	IPv6/BGP	Permit	Permit routing communication from internal routers to external routers
Any	Any	IPv6	Deny	Deny all other IPv6 traffic outbound

Figure 8 – Example Outbound Firewall Rules

Inbound Rules				
Source	**Destination**	**Protocol**	**Action**	**Notes**
Any	Any	IPv6 with source-route header	Deny	Deny all packets with source-route extension header.
Any	2002::/16	IPv6	Deny	Deny all packets to 6to4 tunneling prefixes.
Any	2001::/32	IPv6	Deny	Deny all packets to Teredo tunneling prefix.
Any	Any	IPv4/UDP port 3544	Deny	Deny IPv4 teredo packets carrying IPv6 tunneled data.
Any	Any	IPv4/41	Deny	Deny IPv4 packets carrying IPv6 in a manually configured v6in4 tunnel. * Any intended tunnels would need specific exception.
Fe80::/10	Any	IPv6/ICMP Types 135,136	Permit	Permit link local traffic required for neighbor discovery.
Any	2001:04D0:xxxx:2000::/52 2001:04D0:xxxx:4000::/52	IPv6/HTTP	Stateful Permit	Allow return web traffic to center intranet wired and wireless clients. * Could exclude traffic to other Centers depending on SOC readiness.
Any	2001:04D0:xxxx:2000::/52 2001:04D0:xxxx:4000::/52	IPv6/DNS	Stateful Permit	Allow return DNS responses to center intranet wired and wireless clients.
Any	2001:04D0:xxxx:2000::/52 2001:04D0:xxxx:4000::/52	IPv6/ICMP types 2,3	Permit	Allow ICMP responses required for Path MTU discovery. Some firewalls may support stateful allowance of these base on connection attempts from inside.
Specific Host	Specific Host	IPv6/BGP	Permit	Permit routing communication from external routers to internal routers
Any	Any	IPv6	Deny	Deny all other IPv6 traffic inbound

Figure 9 – Example Inbound Firewall Rules

Routing

Interior Gateway Protocol (IGP) routing for IPv6 is largely the same as in IPv4 with a few key differences. The CSO Corporate Network Target Architecture (CNTA) lists OSPF as the standard IGP. It is assumed that any new routing instances stood up for IPv6 will adhere to the standard so this section will focus on OSPFv3. Note that OSPFv3 will run concurrently with the existing IGP.

The first difference in configuring OSPFv3 for IPv6 vs. IPv4 is the absence of the network statement under the router process configuration. In IPv6 an interface level command is used to enable the protocol on an interface rather than matching the interface IP address with a network statement. The same is true for EIGRP for IPv6. Also, the IPv6 OSPF process does not have to be explicitly configured. It is automatically added to the configuration when OSPF is configured on the first IPv6 interface. Note that the router must have at least one up/up IPv4 enabled interface to start the OSPF process for IPv6 as the OSPFv3 RID is still in dotted decimal format. If not, the OSPF RID must be manually configured before the process will start.

```
! Enable IPv6 routing.
!
! Cat 3k switches require modifying the SDM template and a reboot prior to enabling ipv6 routing.
! sdm prefer dual-ipv4-and-ipv6 routing
!
ipv6 unicast-routing
ipv6 cef  [distributed] ! distributed keyword required on some platforms
!
interface FastEthernet0/0
 ip address 155.2.58.5 255.255.255.0
 ipv6 address FC00:2:0:58::5/64
 ipv6 ospf 100 area 58
 ipv6 ospf authentication ipsec spi 500 md5 [key]
!
ipv6 router ospf 100
 router-id 150.2.5.5
 log-adjacency-changes
 area 58 range FC00:2::/56
 redistribute eigrp 100 include-connected
```

Note the use of the “include-connected” keyword in the example above. Unlike IPv4, IPv6 redistribution does not automatically include connected routes on which the source protocol is running when sending routes into the destination protocol. By default, only the routes actually learned through the source protocol are sent into the destination protocol. The “include-connected” keyword command reverts back to the IPv4 behavior of automatically including connected routes running the source protocol. Also note the absence of the keyword “subnets which is no longer required for protocols being redistributed into OSPF.

Other OSPF concepts such as summarization are configured under the process just like in IPv4. An example is the area range command in the example above. The various types of OSPF areas (normal, stub, totally stubby, NSSA, totally NSSA) are conceptually the same in IPv6 as IPv4. Some of the underlying Link State Advertisements (LSA) types did change in IPv6 but are largely transparent in standard implementations. Specifically, the link/prefix information that was represented in OSPFv2/IPv4 in the Router LSA (type 1) has been split into two separate LSA types in IPv6. The link information is still carried in a Router LSA (type 1). The prefixes are sent in a separate, new LSA type 9 which is called an Intra-Area Prefix LSA. This enhances the scalability of OSPFv3 by allowing prefixes to be added/removed

without triggering SPF calculations. Another new Link LSA (type 8) was added to relay information between routers attached to the same link such as link local addresses.

OSPFv3 LSAs		OSPFv2 LSAs	
LS Type	**Name**	**Type**	**Name**
0x2001	Router LSA	1	Router LSA
0x2002	Network LSA	2	Network LSA
0x2003	Inter-Area Prefix LSA	3	Network Summary LSA
0x2004	Inter-Area Router LSA	4	ASBR Summary LSA
0x4005	AS-External LSA	5	AS-External LSA
0x2006	Group Membership LSA	6	Group Membership LSA
0x2007	Type-7 LSA	7	NSSA External LSA
0x0008	Link LSA		*No Corresponding LSA*
0x2009	Intra-Area Prefix LSA		*No Corresponding LSA*

Figure 10 – OSPF LSA Types

Another difference to note is the use of link-local addresses for next-hops. Link-local addresses are addresses from the range FE80::/10 that are valid only for on-link communication. Each interface that is IPv6 enabled will have a link-local address whether or not it has been given a globally scoped address. The link-local address can be configured or defaults to EUI-64 for the host portion of the address. It is used for on-link communications such as neighbor discovery. IGP messages make use of link-local addresses and use them as the next-hop for learned routes. As shown below, a link-local address is non-deterministic without also including the interface for which the link-local address is applicable. As previously mentioned it would be beneficial to manually set the host portion of infrastructure link-local addresses so that it easier to determine what router the next hop address is pointing to.

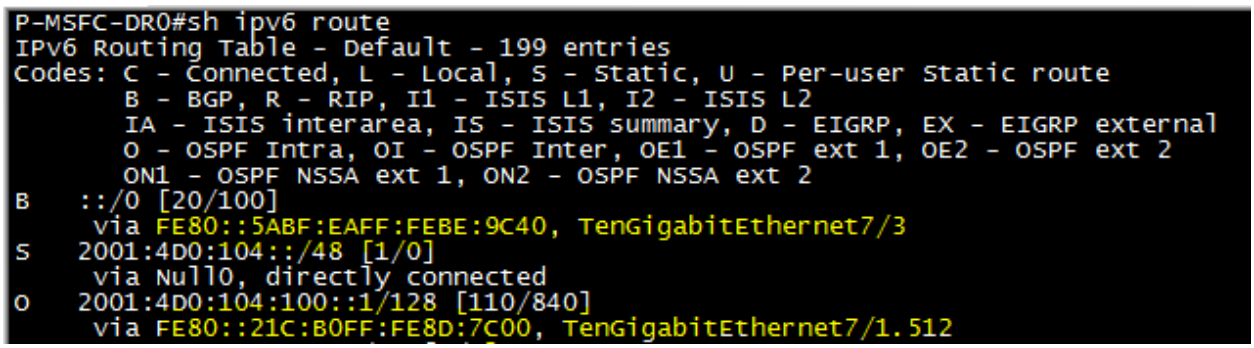

Figure 11 – IPv6 Next-Hop

OSPFv3 authentication is not included until IOS 15.1 for some platforms such as 6500s. It is, however, recommended that authentication be used. This should be taken into account when deciding whether code upgrades are required prior to IPv6 deployment.

NBMA interfaces are rare these days but it is worth mentioning that frame relay inverse arp does not work for IPv6. All layer 3 to layer 2 mappings must be done statically. And since next hop addresses are usually link-local, the link-local addresses must be statically mapped in addition to the globally scoped addresses.

```
Interface Serial0/0/0
 ip address 155.2.0.5 255.255.255.0
 ipv6 address FE80::5 link-local
 ipv6 address 2001:2:0:1234::5/64
 ipv6 ospf 100
 frame-relay map ipv6 2001:2:0:1234::4 504 broadcast
 frame-relay map ipv6 FE80::4 504 broadcast
```

Wireless

Cisco and Aruba wireless solutions support IPv6 clients with the same features as IPv4 clients, such as mobility, security, guest access, QoS, and endpoint visibility. This includes IPv6-only and Dual-stack hosts. The information in this section is based on the following products/models.

- Cisco
  - Wireless LAN controller release v7.2 or higher
  - Wireless Services Module 2 (WiSM2)
  - 5500 series Wireless Controller
  - 1130, 1140, 3500, and 3600 series APs
- Aruba
  - Aruba Wireless LAN Controller release 6.1.4.1-FIPS or higher
  - 6000 series Wireless Controllers w/ M3 line cards
  - 3600 series Wireless Controllers
  - 61, 65, 105, 135, 124, and 175 series APs

To enable wireless IPv6 client connectivity, IPv6 routing and a host address assignment mechanism (DHCPv6) must be configured in the network. Wireless networks should ensure that only CSO-managed DHCP servers are used for address assignment.

For Aruba, IPv6 is enabled by default in the current OS. The following sections provide detailed instructions for configuring IPv6 on a Cisco deployment.

Cisco Wireless

In a Cisco centralized wireless deployment the wireless LAN controller (WLC) must L2 adjacency to an IPv6 router and the VLAN needs to be tagged when packets enter the controller. The APs do not require IPv6 connectivity because all traffic is encapsulated in an IPv4 CAPWAP tunnel between the AP and the controller. Alternatively, autonomous access points (AP) in Work Group Bridge mode can also support IPv6 WGB clients.

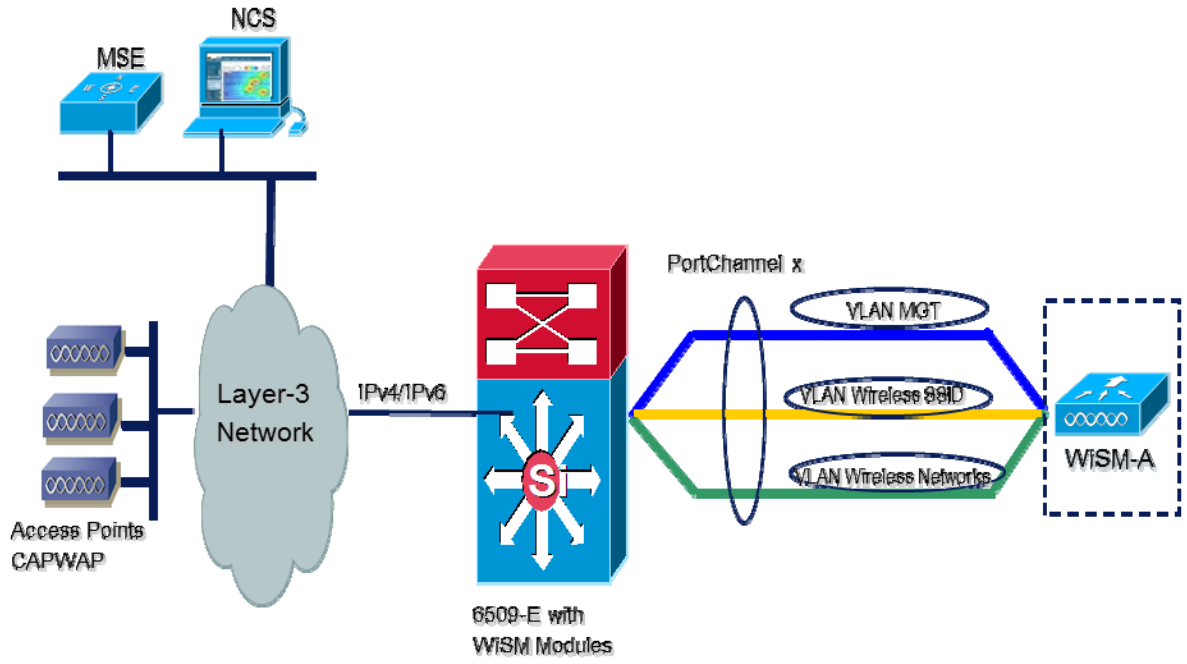

Figure 12 – Example Wireless Architecture

IPv6 must be enabled globally on the Cisco wireless controller. You can enable/disable IPv6 from the CLI or the GUI.

GUI WLC IPv6 Configuration

1. Choose Controller > General.
2. From the Global IPv6 Config drop-down list, choose Enabled or Disabled.
3. Click Apply.
4. Click Save Configuration

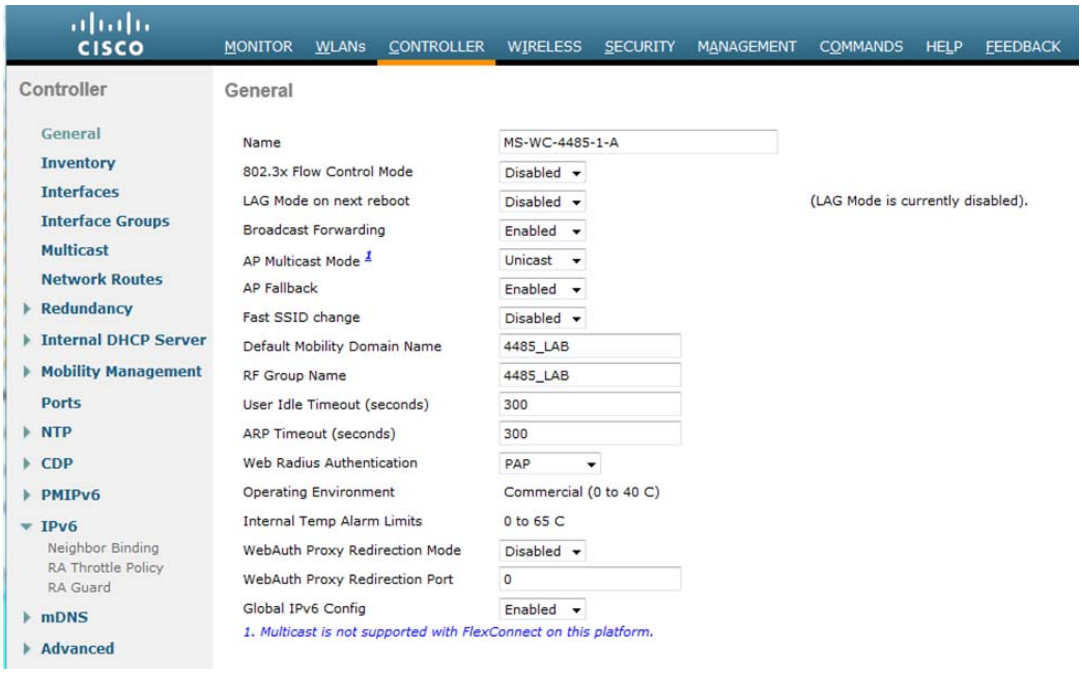

Figure 13 – GUI WLC IPv6 Configuration

CLI WLC IPv6 Configuration

Enable or disable IPv6 globally by entering this command:

```
config ipv6 {enable | disable}
```

IPv6 Wireless Client Roaming/Mobility

In order to support IPv6 client roaming/mobility, some ICMPv6 messages such as Neighbor Solicitation (NS), Neighbor Advertisement (NA), Router Advertisement (RA), and Router Solicitation (RS) must be handled differently. The controllers keep track of IPv6 clients by intercepting the ICMPv6 messages. These messages are then converted from multicast to unicast and delivered individually per client. This unique solution ensures that Neighbor Discovery and Router Advertisement packets are not leaked across VLANs. Clients receive specific Neighbor Discovery and Router Advertisement packets ensuring correct IPv6 addressing and avoiding unnecessary multicast traffic. The configuration for IPv6 mobility is the same as for IPv4 mobility and requires no separate software on the client side to achieve seamless roaming.

1. If both controllers have access to the same VLAN the client was originally on, the roam is simply a Layer 2 roaming event where the client record is copied to the new controller and no traffic is tunneled back to the anchor controller.

2. If the second controller does not have access to the original VLAN the client was on, a Layer 3 roaming event will occur, meaning all traffic from the client must be tunneled via the mobility tunnel (Ethernet over IP) to the anchor controller.

a. In order to ensure the client retains its original IPv6 address, the RAs from the original VLAN are sent by the anchor controller to the foreign controller where they are delivered to the client using L2 unicast from the AP.

b. When the roamed client goes to renew its address via DHCPv6 , the RS, NA, and NS packets continue to be tunneled to the original VLAN so the client will receive an IPv6 address that is applicable to that VLAN.

The only required configuration is that the controllers must be part of the same mobility group/domain.

Go to the **Controller** tab > **Mobility Groups**, and add each controller by MAC Address and IP address into the group. This must be done on all controllers/WiSM in the mobility group.

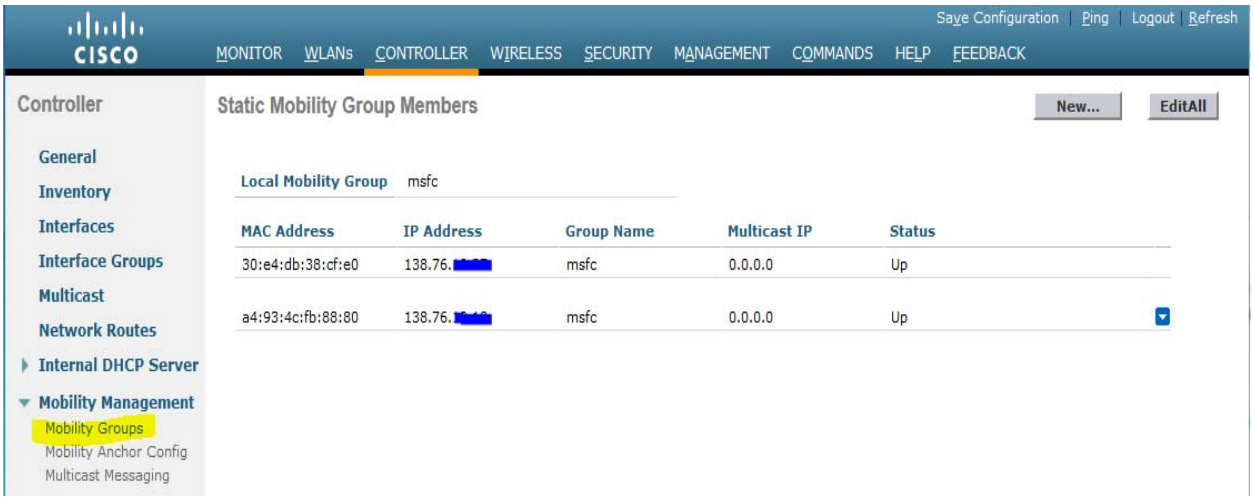

Figure 14 – Configuring Mobility Groups

Security for IPv6 Wireless

Router Advertisement Guard

The RA Guard feature increases the security of the IPv6 network by dropping RAs coming from wireless clients. Without this feature, misconfigured or malicious IPv6 clients could announce themselves as a router for the network, often with a high priority which could take precedence over legitimate IPv6 routers.

By default, RA Guard is enabled at the AP (but can be disabled at the AP) and is always enabled on the controller. Dropping RAs at the AP is preferred as it is a more scalable solution and provides enhanced per−client RA drop counters. In all cases, the IPv6 RA will be dropped at some point, protecting other wireless clients and upstream wired network from malicious or misconfigured IPv6 clients.

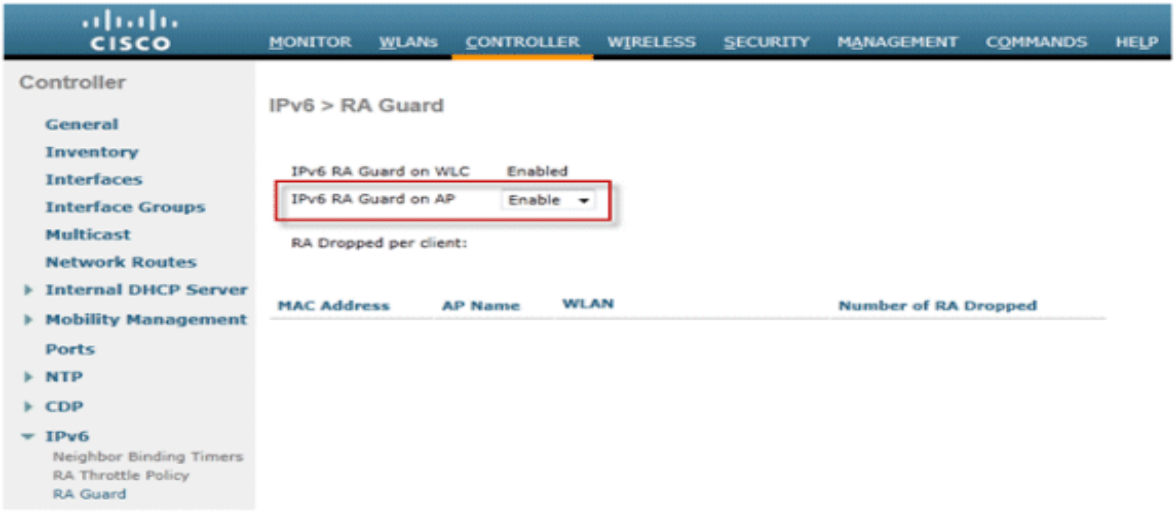

Figure 15 – Configuring RA Guard

DHCPv6 Server Guard

The DHCPv6 Server Guard feature prevents wireless clients from handing out IPv6 addresses to other wireless clients or wired clients upstream. In order to prevent DHCPv6 addresses from being handed out, any DHCPv6 advertise packets from wireless clients are dropped. This feature operates on the controller, requires no configuration and is enabled automatically.

IPv6 Source Guard

The IPv6 Source Guard feature prevents a wireless client spoofing an IPv6 address of another client. This feature is analogous to IPv4 Source Guard. IPv6 Source Guard is enabled by default but can be disabled via the CLI.

IPv6 Address Accounting

The NOC wireless management system records all IPv6 addresses in use by each client and historically logs them each time the client roams or establishes a new session. These records can be configured in the wireless management system to be held for up to a year.

IPv6 Access Control Lists

IPv6 Access Control Lists (ACLs) can be used to identify traffic and permit or deny it. IPv6 ACLs support the same options as IPv4 ACLs including source, destination, source port, and destination port (port ranges are also supported). Pre−authentication ACLs are also supported to support IPv6 guest authentication using an external web server. The wireless controller supports up to 64 unique IPv6 ACLs with 64 unique rules in each. The wireless controller continues to support an additional 64 unique IPv4 ACLs with 64 unique rules in each for a total of 128 ACLs for a dual−stack client.

IPv6 ACLs are applied on a per−WLAN/SSID basis and can be used on multiple WLANs concurrently.

Optimization for IPv6 Clients

Neighbor Discovery Caching

The IPv6 neighbor discovery protocol (NDP) utilizes NA and NS packets in place of Address Resolution Protocol (ARP) in order to allow IPv6 clients to resolve the MAC address of other clients on the network. The NDP process can be very chatty as it initially uses multicast addresses to perform address resolution; this can consume valuable wireless airtime as the multicast packets are sent to all clients on the network segment.

In order to increase the efficiency of the NDP process, neighbor discovery caching allows the controller to act as a proxy and respond back to NS queries that it can resolve. Neighbor discovery caching is made possible by the underlying neighbor binding table present in the controller. The neighbor binding table keeps track of each IPv6 address and its associated MAC address. When an IPv6 client attempts to resolve another client's link layer address, the NS packet is intercepted by the controller which responds back with a NA packet.

Router Advertisement Throttling

Router Advertisement Throttling allows the controller to enforce rate limiting of RAs headed towards the wireless network. By enabling RA throttling, routers which are configured to send RAs very often (for example, every three seconds) can be trimmed back to a minimum frequency that will still maintain IPv6 client connectivity. This allows airtime to be optimized by reducing the number of multicast packets that must be sent. In all cases, if a client sends an RS, then an RA will be allowed through the controller and unicast to the requesting client. This is to ensure that new clients or roaming clients are not negatively impacted by RA throttling.

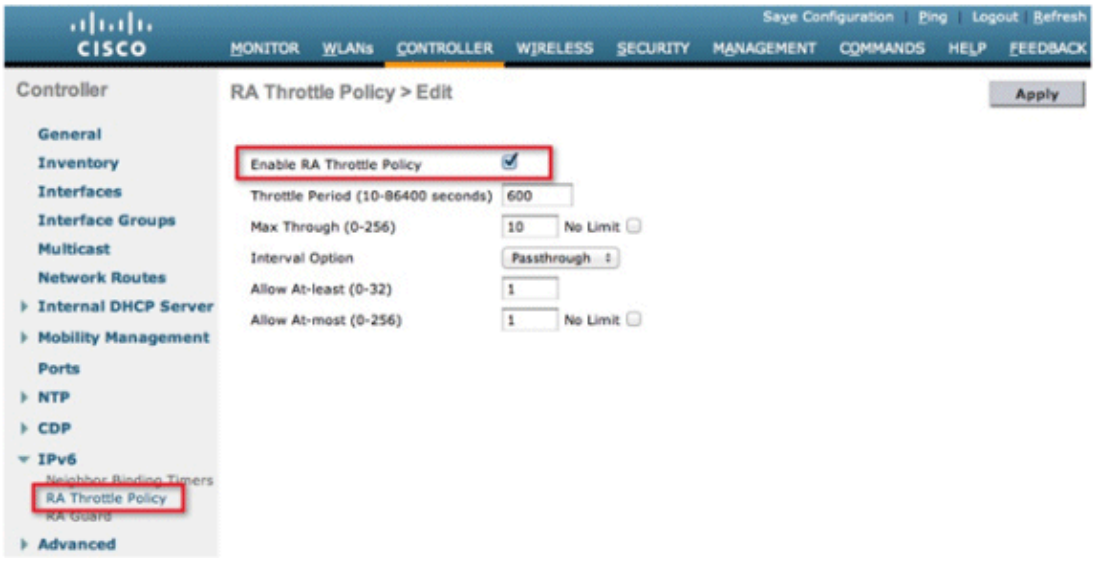

Figure 16 – Router Advertisement Throttling

Note: When RA Throttling occurs, only the first IPv6 capable router is allowed through. For networks with multiple IPv6 prefixes being served by different routers, RA throttling should be disabled.

IPv6 Guest Access

The controller internal wireless guest captive portal for IPv4 clients work in the same manner for dual−stack and IPv6−only clients. Once the guest user associates, they are placed in a WEB_AUTH_REQ run state until the client is authenticated via the IPv4 or IPv6 captive portal. The controller will intercept both IPv4 and IPv6 HTTP/HTTPS traffic in this state and redirect it to the virtual IP address of the controller. Once the user is authenticated via the captive portal, their MAC address is moved to the run state and both IPv4 and IPv6 traffic is allowed to pass.

In order to support the redirection of IPv6−only clients, the controller automatically creates an IPv6 virtual address based off of the IPv4 virtual address configured on the controller. The virtual IPv6 address follows the convention of [::ffff:<virtual IPv4 address>]. For example, a virtual IP address of 1.1.1.1 would translate to [::ffff:1.1.1.1].

Both the IPv4 and IPv6 virtual address of the controller must be defined in DNS to match the SSL certificates hostname. This ensures that clients do not receive a security warning stating that the certificate does not match the hostname of the device. The controller's auto−generated SSL certificate does not contain the IPv6 virtual address.

IPv6 Wireless Client Visibility

Cisco Prime Infrastructure management system release v1.3 and later provides IPv6 capabilities to monitor and manage clients on the wireless network. Information on types of clients that are present on the wireless network and provide insight into IPv6 specific statistics and offer the capability to drill down into IPv6 clients.

Displays the types of IP clients on the network:

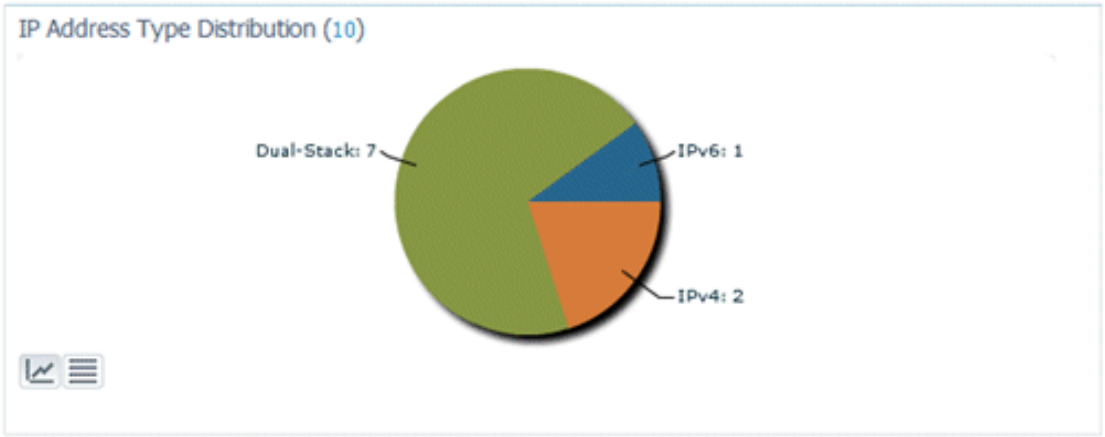

Figure 17 – Address Type Distribution

Displays the IP client type over time:

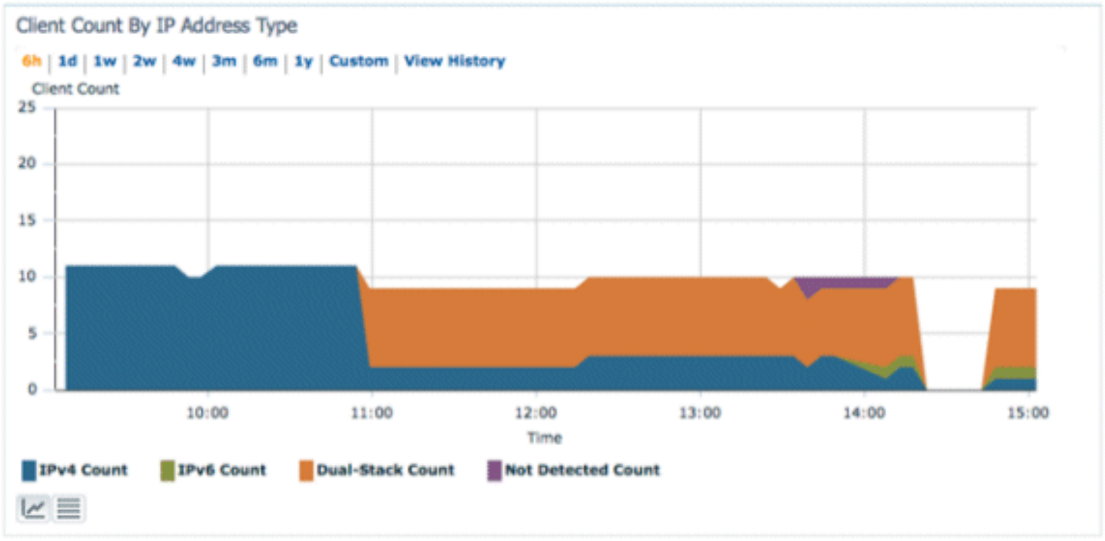

Figure 18 – IP Clients over time

Displays the traffic from each type of client. Clients in the dual−stack category include both IPv4 and IPv6 traffic:

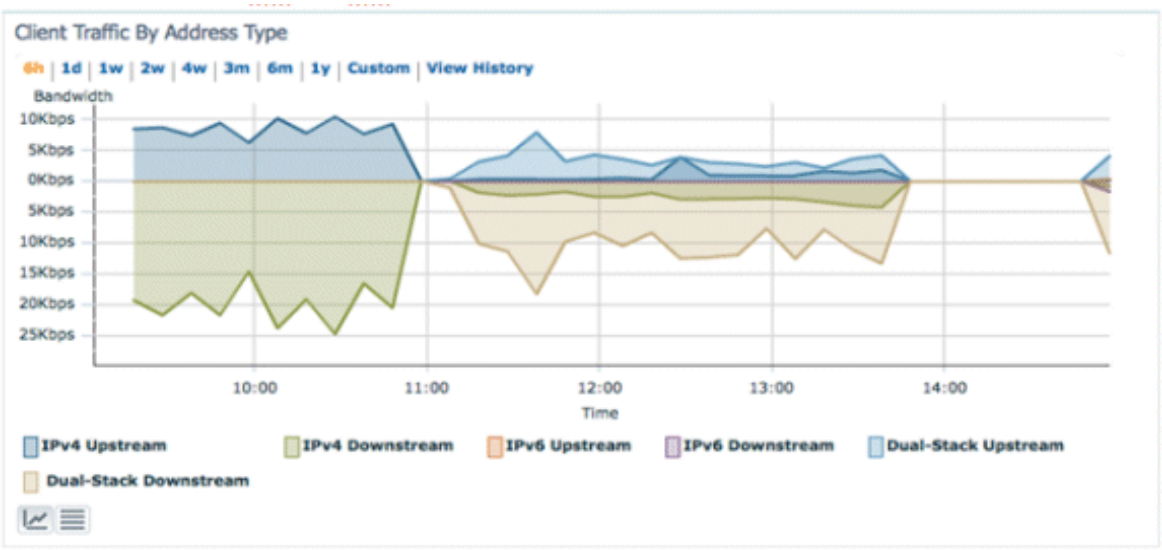

Figure 19 – Traffic by Client Type

Displays the method of address assignment for each client as one of these four categories:

- DHCPv6 – For clients with addresses assigned by a central server. (The client may also have a SLAAC address as well.)
- SLAAC or Static – For clients using stateless address auto assignment or using statically configured addresses.
- Unknown – For cases where the IPv6 address assignment cannot be discovered.
- Self−Assigned – For clients with only a Link−local address which is entirely self−assigned.

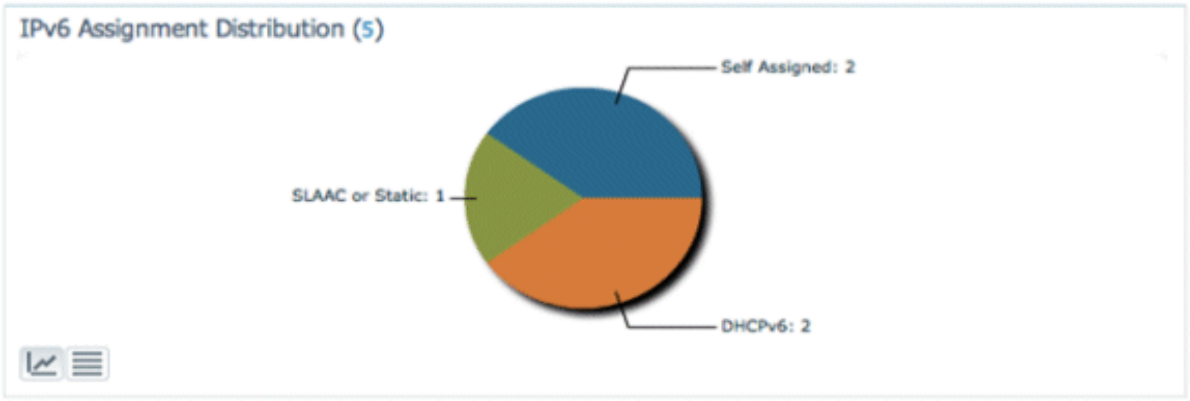

Figure 20 – IPv6 Address Assignment Distribution

Monitor IPv6 Clients

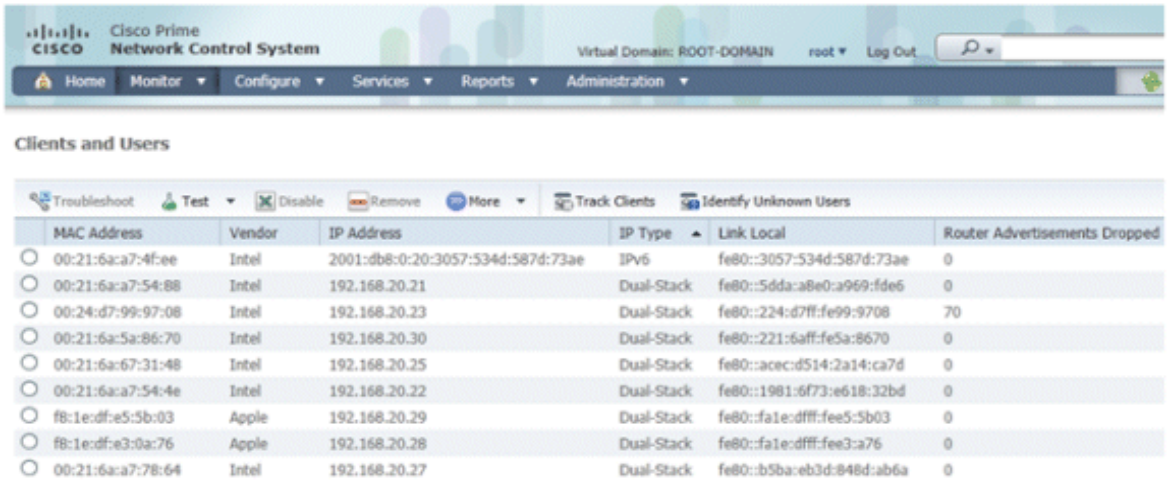

Figure 21 – Monitoring IPv6 Clients

- **IP Type** – The type of client based on what IP addresses have been seen from the client. The possible options are IPv4, IPv6, or Dual−Stack which signifies a client with both IPv4 and IPv6 addresses.
- **IPv6 Assignment Type** – The method of address assignment is detected by NCS as either SLAAC or Static, DHCPv6, Self−Assigned, or Unknown.
- **Global Unique** – The most recent IPv6 global address used by the client. A mouse−over on column contents reveals any additional IPv6 global unique addresses used by the client.
- **Local Unique** – The most recent IPv6 local unique address used by the client. A mouse over on column contents reveals any additional IPv6 global unique addresses used by the client.
- **Link Local** – The IPv6 address of the client which is self−assigned and used for communication before any other IPv6 address is assigned.
- **Router Advertisements Dropped** – The number of router advertisements sent by the client and dropped at the AP. This column can be used to track down clients that

may be misconfigured or maliciously configured to act like an IPv6 router. This column is sortable, which allows offending clients to be identified easily.

Figure 22 – Client IPv6 Addresses

In addition to displaying IPv6 specific columns, the IP Address column will show the current IP address of the client with a priority to display the IPv4 address first (in the case of a Dual−Stack client) or the IPv6 Global Unique address in the case of an IPv6−only client.

Multicast

The primary mechanics of multicast routing did not change from IPv4 to IPv6. However, there are a few key differences to consider.

- No dense mode in IPv6
- MLD vs. IGMP for receiver/router signaling
- No MSDP in IPv6
- Introduction of Embedded RP functionality in IPv6

In addition, some IOS behaviors have changed with regard to multicast. When multicast-routing is enabled for IPv6, PIM and MLD are automatically enabled on every IPv6 enabled interface. For interfaces where multicast routing is not desired use the interface command “no ipv6 pim”.

Multicast Addressing

IPv6 multicast groups are of the range FF00::/8. Bits 9-12 (from the left) are designated as flags while bits 13–16 determine the scope of the multicast group. Currently the flags are always set to 0x7 for embedded RP groups or 0x0 for all other groups. Several multicast group scopes have been defined and are shown below.

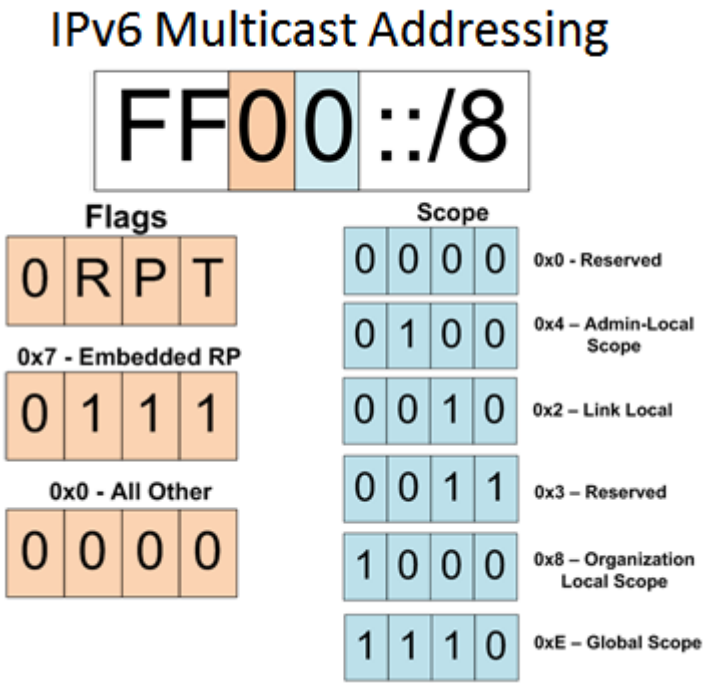

Figure 23 – IPv6 Multicast Addressing

Receiver/Router Signaling

There are no major changes to the Protocol Independent Multicast (PIM) protocol which is used for router to router signaling. There are minor changes to the receiving host to router signaling protocol. Multicast Listener Discovery (MLD) protocol replaces Internet Group Management Protocol (IGMP) for receiver/router signaling in IPv6. Unlike IGMP which is a unique protocol, MLD is implemented as a set of ICMPv6 message types 130-132, 143. The two protocols implement roughly equivalent functionality.

- MLDv1 = IGMPv2 – Basic functionality to join a multicast group (IGMPv1) with the enhanced "fast leave" functionality to quickly stop sending multicast on a LAN after the last receiver "leaves" (IGMPv2).
- MLDv2 = IGMPv3 – Adds support for Source-Specific Multicast by allowing joins for specific sources of multicast.

RP Methods and MSDP

Multicast Source Discovery Protocol (MSDP) has two main functions in IPv4 multicast. The primary function is to interconnect PIM Sparse Mode (PIM-SM) domains. Rendezvous Points (RPs), which serve to connect multicast senders and receivers, use MSDP to peer with RPs from other domains and exchange information about available multicast groups and sources. MSDP is used in the agency today to tie Center RPs to the backbone RPs which are in turn tied to Internet RPs and other Center RPs. In the current configuration MDSP is essential for any inter-domain multicast (between centers).

The other function of MSDP is redundancy through the Anycast-RP functionality. Anycast-RP provides RP redundancy by configured more than one router with the same RP unicast address. The multiple RPs are then tied together with MSDP to make sure all RPs are aware of all multicast group/source pairs.

However, despite its central role in IPv4 multicast, MSDP has not been implemented for IPv6. This necessitates a change in the enterprise multicast architecture. Instead of each Center relying solely on their own RP, which in IPv4 would be aware of external multicast through MSDP, RPs need to be scoped so that internal multicast points to the internal RP while inter-domain and external multicast points to the backbone RP. This in turn means that all Centers should agree on what multicast addresses to use for internal vs. inter-domain/external. For this reason, work is underway to generate a multicast address plan for the agency. As for replacing the RP redundancy aspect of MSDP for IPv6, a version of Anycast-RP for IPv6 which does not rely on MSDP was introduced in IOS 15.1 code.

Another new aspect of IPv6 multicast is the concept of embedded RP multicast addresses (RFC 3956). The concept is to encode the unicast address of the RP address into the multicast group address so that the RP does not have to be known or configured ahead of time. This is particularly useful for multicast with external entities with which there is no prior RP relationship. Embedded RP groups are denoted with the flag bits set to 0111 or 0x7. The following example shows how to derive the RP address from an embedded RP multicast address.

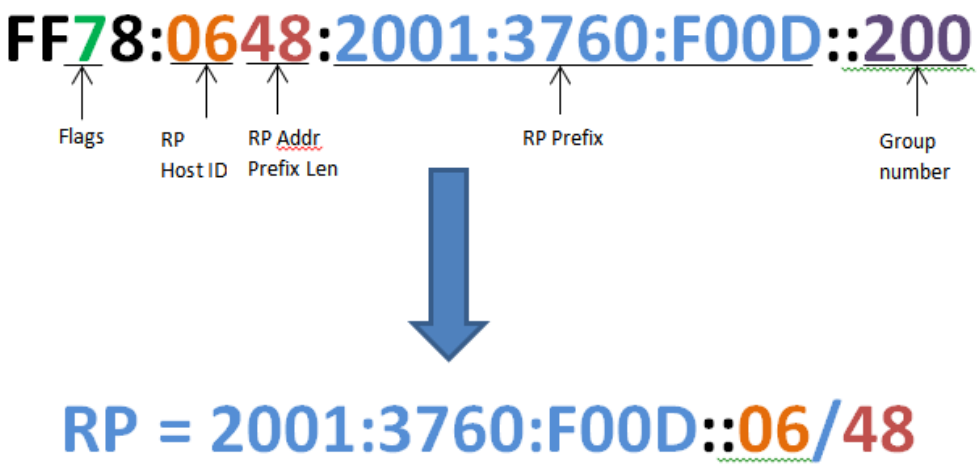

Figure 24 – Embedded RP Example

There are two methods for RP dispersion in IPv6: 1) manual configuration and 2) Bootstrap Router (BSR). The AutoRP protocol has not been implemented in IPv6. Manual configuration is the recommended method. RP information rarely changes and since scoping will be used this alleviates the need to send BSR messages across firewalls/ACLs.

IOS Versions

As mentioned in various sections of this document there are IOS version dependencies for some of the recommended IPv6 features to be deployed. Examples are OSPFv3 authentication, RA Guard, Anycast-RP for PIMv6, and WCCP. Given these dependencies and the number of platforms deployed across the

agency, Cisco will be engaged via the ESA to provide minimum code level recommendations per platform based on the recommended features that are planned to be deployed. Any features that are not listed in this document that are planned for implementations need to be made known for inclusion in this effort.

The table below is a summary of features that require IOS 15.x. The reference core/distribution platform is a 6500/Sup720 and the access platforms are a 3560G and 3560x. Lab testing was limited to these platforms. Refer to vendor documentation for specifics on other platforms.

Feature	Required Code	Purpose	Work Around	Where (core, dist, access, border)	Disposition
"all" keyword for command "ipv6 nd suppress"	15.x	Recommended use on infrastructure links to prevent unintended hosts from soliciting a router and getting a response. This would allow the host to configure an address.	Clear the A-bit on the infrastructure prefix to prevent SLAAC.	All	Preferred
VRRP	15.x	First Hop Redundancy Protocol. Would be used for feature parity for any LANs that use VRRP as the IPv4 FHRP.	Use HSRP or GLBP which are available in older code for IPv6.	Core/Dist	Optional
Netflow v9/IPv6 suport	depends on platform	Provide export of netflow records for IPv6 traffic.	N/A	Core/Dist	Optional
Basic RA guard – access layer	15.x	Block unintended RAs from being injected into the network. This is important since IPv6 hosts are not configured with a "default gateway" but instead learn the gateway through RAs heard on the LAN. This feature is present in 12.x code for Core/Distribution platforms but requires 15.x code for the Access layer.	N/A	access	Required
RA guard advanced options	15.x	Core/dist platforms support basic RA guard in 12.x code. 15.x code on these platforms provides RA guard options to conditionally filter RAs based on contents vs. basic RA guard which statically permits or denies all RAs on a port.	Use basic RA guard w/o options.	Core/Dist	Optional
WCCP for IPv6	15.x	Provides redirection of web flows to the proxy/content filter. Currently used at 6 centers.	None	Border	Required
OSPFv3 Authentication (for 6500 platform)	15.x	Allows authentication of OSPFv3 messages.	Don't implement OSPFv3 authentication.	Core/Dist	Preferred
RP redundancy without MSDP	15.x	Allows IPv6 RP redundancy.	Implement with no redundancy or not at all.	Core/Dist	Optional

SNMP ACL for ipv4 and ipv6	15.x	Protect management plane	Do not address incapable platforms with an IPv6 address to ensure they cannot be managed via IPv6.	Access	Required

Figure 25 – Feature/Version Dependencies

Summary of Recommendations

	Recommendation	Reference Section
1	Deploy IPv6 from the outside edge inward. Do not enable IPv6 on any host networks until the path from the host network to the internet is fully functional including DNS, web redirection, etc.	Order of Deployment
2	The first step in deploying IPv6 on the infrastructure is to secure the management plane. Address the following: - Apply IPv6 ACLs to VTY lines - Apply IPv6 ACLs to SNMP communities/groups	Infrastructure: Securing the Management Plane
3	Configure IPv6 ICMP error message rate limiting on routers.	Infrastructure: Infrastructure Links
4	Manually assign the link-local addresses on infrastructure links.	Infrastructure: Infrastructure Links
5	Suppress Router Advertisements on infrastructure links. Also disable redirects and unreachables.	Infrastructure: Infrastructure Links
6	All IPv6 interface addresses should be either /64 or /128 (loopbacks).	Infrastructure: Infrastructure Addressing
7	Use RA suppression or specific command ordering when implementing IPv6 on host facing interfaces to prevent the router from sending default RAs during the configuration process and causing hosts to obtain SLAAC/temp addresses.	Host Networks: Deployment Considerations
8	Use HSRP, VRRP, or GLBP for host networks.	Infrastructure: First Hop Redundancy Protocols.
9	Verify existing monitoring/troubleshooting tools are capable and configured for IPv6 operations.	Infrastructure: Monitoring

10	Configure host facing router interfaces to ensure DHCP operations (M-bit) and to not allow SLAAC (disable A-bit per prefix).	Host Networks: DHCPv6 (Stateful)
11	Configure host facing router interfaces to advertise a router-preference of High in RAs.	Host Networks: Host Default Gateway
12	Implement RA guard or port ACLs to ensure the legitimate IPv6 gateway is not usurped by mis-configuration or maliciously. This may require code upgrades.	Host Networks: Host Default Gateway
13	If WCCP is used, evaluate code rev needed to implement WCCP for IPv6.	Host Networks: Web Redirection
14	Verify Center proxies and content filters are IPv6 capable. Code upgrades may be required.	Host Networks: Web Redirection
15	When writing ACLs and/or firewall rules, make sure ICMPv6 types that are required for correct protocol operation are allowed. - Neighbor Discovery (types 135 and 136) - Path MTU Discovery (types 2 and 3)	Firewall Rules and ACLS: Neigbor Discovery Protocol, PATH MTU Discovery
16	Disable source routing by filtering packets with the routing (type 0) extension header. Type 2 can be filtered as well if IPv6 mobility is not is use.	Firewall Rules and ACLS: Source Routing and Extension Headers
17	Filter all IPv6 tunneling methods by ACL or firewall. This includes some IPv4 filters to deny IPv6 over IPv4 tunnel methods.	Firewall Rules and ACLS: Tunneling Protocols
18	Implement OSPFv3 as the IPv6 IGP.	Routing
19	If IPv6 multicast is planned consider the following: - The lack of MSDP in IPv6 means RP scoping will be required for inter-domain/external multicast. - RP redundancy will require a new feature that impelments Anycast-RP without MSDP. - IGMP is replaced by MLD which runs over ICMPv6	Multicast
20	IOS versions required will depend on what features are required. Centers will be asked to help determine this by providing platform/feature information.	IOS Versions

References

- OMB IPv6 Mandate - http://www.whitehouse.gov/sites/default/files/omb/assets/egov_docs/transition-to-ipv6.pdf
- RFC 1671 – IPng White Paper on Transition and Other Considerations - http://tools.ietf.org/html/rfc1671
- RFC 6555 - http://tools.ietf.org/html/rfc6555#section-3.2
- NIST – Guidelines for the Secure Deployment of IPv6 - http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf
- Cisco - Improving User Experience with IPv6 and SCTP - http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_13-3/133_he.html
- RFC 6164 - http://tools.ietf.org/html/rfc6164
- Cisco IOS IPv6 Feature Mapping - http://docwiki.cisco.com/wiki/Cisco_IOS_IPv6_Feature_Mapping
- Cisco – Configuring First Hop Redundancy Protocols in IPv6 - http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/15-2mt/ip6-fhrp.html#GUID-9E309D6A-17FA-40F9-8675-73C0C95E68CD
- VRRPv3 - http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp_fhrp/configuration/15-sy/fhrp-vrrpv3.html
- RFC 4861 – Neighbor Discovery for IP version 6 (IPv6) - http://tools.ietf.org/html/rfc4861
- IPv6 EUI-64 - http://wiki.mikrotik.com/wiki/File:Ipv6eui64.png
- RFC 6104 - Rogue IPv6 Router Advertisement Problem Statement - http://www.rfc-editor.org/rfc/rfc6104.txt
- RFC 6105 – IPv6 Router Advertisement Guard - http://www.rfc-editor.org/rfc/rfc6105.txt
- Cisco – Implementing First Hop Security in IPv6 - http://www.cisco.com/en/US/docs/ios-xml/ios/ipv6/configuration/12-2sy/ip6-first-hop-security.html#GUID-B82A9A22-FC81-4AC2-B094-3EF89EE875AD
- Cisco – WCCPv2 IPv6 support (IOS 15M&T Configuration Guide) - http://www.cisco.com/en/US/docs/ios-xml/ios/ipapp/configuration/15-sy/iap-15-sy-book_chapter_01000.html
- RFC 1981 – Path MTU Discovery for IP version 6 - http://tools.ietf.org/html/rfc1981
- IT Cert Notes – IPv6 Neighbor Discovery - http://www.itcertnotes.com/2011/11/ipv6-neighbor-discovery-protocol.html
- Cisco – Implementing Traffic Filters and Firewalls for IPv6 Security - http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6-sec_trfltr_fw.html#wp1072522
- RFC – Internet Protocol, Version 6 (IPv6) Specification - http://tools.ietf.org/html/rfc2460#section-5
- Cisco – IPv6 Extension Headers Review and Considerations - http://www.cisco.com/en/US/technologies/tk648/tk872/technologies_white_paper0900aecd8054d37d.html
- RFC 5095 – Deprecation of Type 0 Routing Headers in IPv6 - http://tools.ietf.org/html/rfc5095

- Cymru BOGON Reference - http://www.team-cymru.org/Services/Bogons/
- Cymru IPv6 BOGON Prefix-List - http://www.team-cymru.org/ReadingRoom/Templates/IPv6Routers/ios.html
- NICS Corporate Network Target Architecture - https://sharepoint.msfc.nasa.gov/sites/cso/crew/Architecture%20Reference%20Documentation/Forms/AllItems.aspx?RootFolder=%2Fsites%2Fcso%2Fcrew%2FArchitecture%20Reference%20Documentation%2FReference%20Architecture&FolderCTID=0x01200071CB59CDFA6C6743B80B757C80C4D11B&View=%7BF6E1DA32%2DD6F8%2D4AA7%2DA6CC%2D35451CD44069%7D&InitialTabId=Ribbon%2ELibrary&VisibilityContext=WSSTabPersistence
- RFC 5340 – OSPF for IPv6 - http://tools.ietf.org/html/rfc5340
- OSPF LSA Types - http://adrian-brayton.blogspot.com/2010/10/ipv6-ospf-v3-lsa-types.html
- RFC 2710 – Multicast Listener Discovery Version 2 (MLDv2) for IPv6 - http://tools.ietf.org/html/rfc3810
- RFC 3618 – Multicast Source Discovery Protocol (MSDP) - http://tools.ietf.org/html/rfc3618
- RFC 4610 – Anycast –RP Using Protocol Independent Multicast (PIM) - http://tools.ietf.org/html/rfc4610
- RFC 4601 – Protocol Independent Multicast – Sparse Mode (PIM-SM) – http://tools.ietf.org/html/rfc4601
- RFC 3810 – Multicast Listener Discovery Version 2 (MLDv2) for IPv6 - http://tools.ietf.org/html/rfc3810
- RFC 4294 – IPv6 Node Requirements - http://tools.ietf.org/search/rfc4294
- Anycast-RP Capability for IPv6 - http://www.cisco.com/en/US/docs/ios-xml/ios/ipmulti_pim/configuration/15-sy/imc_basic_ipv6.html#GUID-00CCAD1F-B427-4FDF-8885-1F2229A44DD1
- RFC 3956 – Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address - http://www.ietf.org/rfc/rfc3956.txt
- NICS IPv6 Address Acquisition Presentation - https://sharepoint.msfc.nasa.gov/sites/cso/nisn_organization/engineering/Shared%20Documents/IPv6/IPv6_Address_acquisition_abridged_20130829.pptx
- Cisco Design Zone for IPv6 - http://www.cisco.com/en/US/netsol/ns817/networking_solutions_program_home.html

Appendix A – Applicable JUNOS IPv6

Infrastructure

Securing the Management Plane

The simplest method for securing the JUNOS management plane is to delete all unneeded services in the JUNOS "system services" configuration syntax. Only protocols that are identified in the site's security plan should be secured using firewall filters (i.e. SNMP, SSH); the rest should be disabled.

In JUNOS, SNMP (all versions) is not enabled by default. Enabling SNMP to a JUNOS device requires that the community string is assigned a read-only attribute and is associated with valid clients that can access the device's MIB table. Note: there is no need to enable IPv6 SNMP at this time assuming the device is actively monitored through the IPv4 infrastructure.

```
snmp {
        community <string> {
                authorization read-only;
                clients {
                        2001:DB8:1::2;
                        2001:DB8:2::2;
                }
        }
}
```

In addition to limiting SNMP, several unneeded services should be removed from the "system services" configuration hierarchy.

```
[edit system services]                          ## Some applications may not be configured
root@junos# delete rsh
root@junos# delete rlogin
root@junos# delete ftp
root@junos# delete finger
root@junos# delete telnet
root@junos# delete web-management http
```

After all non-essential services are disabled, the Juniper's routing engine should be protected using a firewall filter permitting authorized IPv6 hosts to access the device and prohibiting all others. The recommended method for RE protection is to use a default-deny method which requires the network engineer to explicitly permit what IPv6 address can access the device. The firewall filter must take into consideration routing protocols (BGP, OSPF, etc...), access protocols (SSH, J-Web, etc...), management (SNMP, NTP, DNS, etc...), and operational/troubleshooting (ICMPv6, traceroute, etc...). The firewall filter must be applied to the node's loopback0 interface (lo0.0). All traffic from transit interfaces must go through the loopback interface to reach the node's CPU. The firewall filter must include ICMPv6 parameters to allow neighbor discovery messages otherwise the filter will break the Juniper node.

```
firewall {
        family inet6 {
                filter IPv6_PROTECT_RE {
                        term IPv6_ICMP {
                                from {
                                        next-header icmp6;
                                        icmp-type [ echo-reply echo-request neighbor-advertisement neighbor-
solicit packet-too-big time-exceeded ];
                                }
                                then {
                                        policer POLICER_500KBPS;
                                        accept;
                                }
                        }
                        term IPv6_IBGP {
                                from {
                                        source-prefix-list {
                                                IPv6_IBGP_NEIGHBORS;
                                        }
                                        next-header tcp;
                                        destination-port bgp;
                                }
                                then accept;
                        }
                        term IPv6_OSPF {
                                from {
                                        source-prefix-list {
                                                IPv6_OSPF_NEIGHBORS;
                                        }
                                        next-header ospf;
                                }
                                then accept;
                        }
                        term IPv6_SSH_HOSTS {
                                from {
                                        source-prefix-list {
                                                IPv6_SSH_HOSTS;
                                        }
                                        next-header tcp;
                                        port ssh;
                                }
                                then accept;
                        }
                        term DENY {
                                then {
                                        log;
                                        discard;
                                }
                        }
                }
        }
        policer POLICER_500KBPS {
                if-exceeding {
                        bandwidth-limit 500k;
                        burst-size-limit 1500;
                }
                then discard;
        }
}
```

Infrastructure Addressing

Juniper devices should follow the same infrastructure addressing scheme listed earlier in this document.

Routing

OSPFv3 is the version of OSPF that support IPv6 interior routing. OSPFv3 differs from OSPFv2 (IPv4) by removing protocol route authentication and relying on the IPv6 stack authentication. OSPFv3 uses IPSEC for security. All Juniper IPv6 should be using OSPFv3 as the standard IGP.

First step is to configure a basic OSPFv3 configuration. This includes verifying that the router-id is defined in the routing-options. OSPFv3 will use the same router-id that is used for OSPF.

```
routing-options {
        router-id 192.168.1.254;
}
protocols {
        ospf3 {
                area 0.0.0.0 {
                        interface lo0.0 {
                                passive;
                        }
                        interface ge-0/0/1;
                }
        }
}
```

After the basic OSPFv3 configuration is configured, the next steps are to configure the IPSEC security association and bind the association with the OSPFv3 interface.

```
[edit]
root@junos# edit security ipsec security-association OSPFv3_CORE

[edit security ipsec security-association OSPFv3_CORE]
root@junos# set description “OSPFv3 Neighbor Authentication”

[[edit security ipsec security-association OSPFv3_CORE]
root@junos# set mode transport

[edit security ipsec security-association OSPFv3_CORE]
root@junos# set manual direction bidirectional protocol ah

[edit security ipsec security-association OSPFv3_CORE]
root@junos# set manual direction bidirectional spi 256

[edit security ipsec security-association OSPFv3_CORE]
root@junos# set manual direction bidirectional authentication algorithm hmac-md5-96

[edit security ipsec security-association OSPFv3_CORE]
root@junos# set manual direction bidirectional authentication key ascii-text <complex password>
```

The below configuration associations the IPSEC security-association with the applicable OSPFv3 interface.

```
[edit security ipsec security-association OSPFv3_CORE]
root@junos# top edit protocols ospf3 area 0.0.0.0

[edit protocols ospf3 area 0.0.0.0]
root@junos# set interface ge-0/0/1 ipsec-sa OSPFv3_CORE
```

Please refer to the primary paper for information on the use of link-local addresses for OSPFv3 next-hop.

First Hop Redundancy Protocol

JUNOS support Virtual Router Redundancy Protocol (VRRP) for IPv6 first hop redundancy. VRRP allows hosts on a LAN to use redundant routing platforms without additional configurations on the hosts. VRRP works by each router sharing an IPv6 address that is used for a host's default route. At any time, one router is the active (Master) router and provides gateway services to the hosts; other routers are backups and use timers to take control from the master if needed.

VRRP requires the routers to transmit router advertisements (ICMPv6) which are disabled by default. To enable VRRP, the first steps include enabling router advertisements and to identify the prefix that should be advertised as part of the RA. The RA is needed since it provides a default route for the IPv6 host. The RA may need to be tweaked depending on the addressing mechanism used (SLAAC versus stateful/stateless DHCPv6). In the case of DHCP, additional flags will be configured to limit SLAAC addressing while still advertising a default gateway.

```
[edit]
root@junos# edit protocols router-advertisement interface ge-1/0/1.100

[edit protocols router-advertisement interface ge-1/0/1.100]
root@junos# set prefix 2001:DB8:0:9::/64
```

Several commands are used to configure the global IPv6 address and to specify the link-local address (makes configurations more readable). The following commands configure link-local, global, logical link local, and logical global (logical will be the shared address between multiple VRRP routers). Additionally, priorities are assigned for VRRP mastership.

```
[edit]
root@junos# edit interfaces ge-1/0/1 unit 100 family inet6

[edit interfaces ge-1/0/1 unit 100 family inet6]                    ## Configure link local
root@junos# set address FE80:DB8:0:9::2/64

[edit interfaces ge-1/0/1 unit 100 family inet6]                    ## Configure global; associate VRRP
root@junos# edit address 2001:DB8:0:9::2/64 vrrp-inet6-group 10     ## VRRP Group 10

[edit interfaces ge-1/0/1 unit 100 family inet6 address 2001:DB8:0:9::2/64 vrrp-inet6-group 42]
root@junos# set virtual-inet6-address 2001:DB8:0:9::1
root@junos# set virtual-link-local-address FE80:DB8:0:9::1

[edit interfaces ge-1/0/1 unit 100 family inet6 address 2001:DB8:0:9::2/64 vrrp-inet6-group 42]
root@junos# set priority 250
root@junos# set preempt
root@junos# set accept-data
```

The output below can be copied and modified for the redundant router. Configuration changes should include specifying a lower priority and changing the redundant router's link local and global address to a unique address. Virtual IPv6 address should remain the same.

```
protocols {
        router-advertisement {
                interface ge-1/0/1.0 {
                        prefix 2001:DB8:0:9::/64;
                }
        }
}
interfaces {
        ge-1/0/1 {
                unit 0 {
                        family inet6 {
                                address FE80:DB8:0:9::2/64;
                                address 2001:DB8:0:9::2/64 {
                                        vrrp-inet6-group 10 {
                                                virtual-inet6-address 2001:DB8:0:9::1;
                                                virtual-link-local-address FE80:DB8:0:9::1;
                                                priority 250;
                                                preempt;
                                                accept-data;
                                        }
                                }
                        }
                }
        }
}
```

Host Addressing

IPv6 host addressing is a radical change from IPv4. Besides the obvious difference of a 128-bit versus 32-bit address, IPv6 relies on several addresses to include link-local addresses, global addressing, and various multicast addresses. Within the area of global IPv6 addresses, a host can have multiple IPv6 addresses acquired through various methods. Basic to all of these methods is the router advertisement (RA). For more information on RA, refer to the "Host Networks" section of this paper or Appendix B which discusses IPv6 address acquisition.

Router point-to-point links require that the inet6 address family is activated and configured. These "point-to-point" links are preconfigured and have no need for any router advertisements (which are disabled by default. Below is a sample configuration for two devices connected through ge-0/0/0. In JUNOS, there is the ability to use EUI-64 for addressing the router interface; however, this does make the interface address dependent on the physical hardware.

```
## Router A
interfaces {
        ge-0/0/0 {
                unit 0 {
                        family inet6 {
                                address 2001:DB8:0:9::1/64;
                        }
                }
        }
}
## Router B
interfaces {
        ge-0/0/0 {
                unit 0 {
                        family inet6 {
                                address 2001:DB8:0:9::2/64;
                        }
                }
        }
}
```

Redirects should also be turned off on each interface. This can be accomplished using a global system command or using specific interface syntax if redirects may be needed on some interfaces.

```
[edit]
root@junos# set system no-redirects-ipv6                    ## Global for all interfaces

[edit interfaces]
root@junos@ set ge-0/0/0 unit 0 family ipv6 no-redirects    ## Interface specific
```

Pushing IPv6 addresses to end systems in the LAN require addressing through stateless auto address configuration (SLAAC) or stateless/stateful Dynamic Host Configuration Protocol v6 (DHCPv6) – unless you really want to manually assign every address.

End-system addressing can be complex depending on addressing requirements. The easiest addressing mechanism for IPv6 is SLAAC which requires minimal configurations on the host and the router. In JUNOS, enabling SLAAC requires that a router advertisement (RA) is generated for each interface supporting end-systems. While no decision has been made on NASA addressing requirements, most likely hosts will be addressed using stateful DHCPv6. The below configuration enables RAs on an interface and configures options to enable a host to negotiate (stateful) an address and pull other configuration options through a DHCPv6 server.

```
protocols {
        router-advertisement {
                interface ge-1/0/1.200 {
                        other-stateful-configuration;
                        managed-configuration;
                }
        }
}
```

Firewall Rules and Access Control Lists

Care must be taken when applying a firewall ruleset to an interface on a JUNOS device since that firewall rule may break ICMPv6 messages that are critical to IPv6 operations. As shown in previous sections, ICMPv6 rules must exist for the neighbor discovery (neighbor solicitation / announcement) messages to function properly enabling an interface to communicate with the link-local IPv6 address. The below firewall term will allow neighbor discovery messages to pass. If a firewall filter has any explicit deny configurations, the below term must be configured to allow for IPv6 ICMP neighbor discovery. Ping and other ICMPv6 messages must also be explicitly permitted (see earlier example).

```
family inet6 {
        filter IPv6_ICMP_ND {
                term ALLOW_ICMP_ND {
                        from {
                                next-header icmp6;
                                icmp-type [ neighbor-advertisement neighbor-solicit ];
                         }
                        then accept;
                }
        }
}
```