



Enterprise IPv6 Deployment



Michael De Leo – mdeleo@cisco.com

Corporate Consulting Engineer / CTO LATAM

December 2008

Agenda

- General Concepts
- Infrastructure Deployment
- Planning and Deployment Summary

Reference Materials

“Deploying IPv6 Networks” by Ciprian Popoviciu, Eric Levy-Abegnoli, Patrick Grossetete—Cisco Press (ISBN: 1587052105)

- Deploying IPv6 in Campus Networks:
<http://www.cisco.com/univercd/cc/td/doc/solution/campipv6.pdf>
- Deploying IPv6 in Branch Networks:
<http://www.cisco.com/univercd/cc/td/doc/solution/brchipv6.pdf>
- CCO IPv6 Main Page:
<http://www.cisco.com/go/ipv6>
- Cisco Network Designs:
<http://www.cisco.com/go/srnd>
- Cisco IOS Feature Navigator – What images support IPv6?
<http://www.cisco.com/go/fn> (Registration Required)

Monitoring Market Drivers

Address space depletion

<http://www.potaroo.net/tools/ipv4/>

Impact being a slow down of the Internet growth and market penetration

National IT Strategy

U.S. Federal Mandate

IPv6 Task Force and promotion councils:
Africa, India, Japan, Korea,...

China Next Generation Internet (CNGI)
project

European Commission sponsored
projects

IPv6 “on” & “preferred” by default
Applications only running
over IPv6 (P2P framework)

IP NGN

DOCSIS 3.0, FTTH, HDTV, Quad
Play

Mobile SP – 3G, WiMax, PWLAN
Networks in Motion

Networked Sensors, ie: AIRS

NAT Overlap – M&A

MSFT Vista & Server 2008

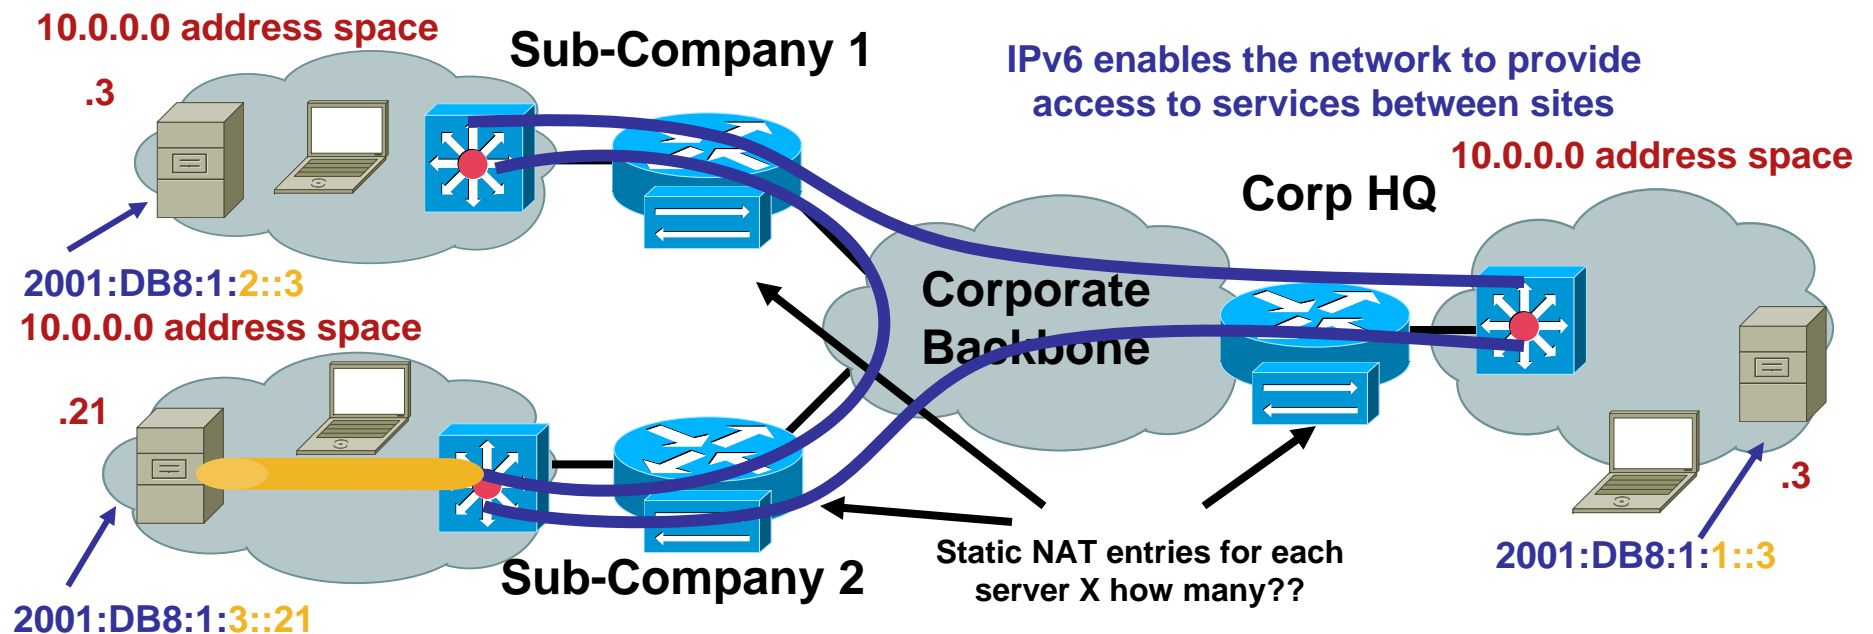
Infrastructure Evolution

Operating System Support



- Every major OS supports IPv6 today
- Top-to-bottom TCP/IP stack re-design
- IPv6 is on by default and preferred over IPv4 (considering network/DNS/application support)
- Tunnels will be used before IPv4 if required by IPv6-enabled application
 - ISATAP, Teredo, 6to4, Configured
- All applications and services that ship with Vista/Server 2008 support IPv4 and IPv6 (IPv6-only is supported)
 - Active Directory, IIS, File/Print/Fax, WINS/DNS/DHCP/LDAP, Windows Media Services, Terminal Services, Network Access Services – Remote Access (VPN/Dial-up), Network Access Protection (NAP), Windows Deployment Service, Certificate Services, SharePoint services, Network Load-Balancing, Internet Authentication Server, Server Clustering, etc...
- <http://www.microsoft.com/technet/network/ipv6/default.mspx>

NAT Overlap




- Merger and acquisition complexity force many to leave existing IPv4 address space in place vs. full integration/consolidation
- When server-to-server or client-to-server service is required then single/double static NAT translations are often required
- IPv6 can be deployed to enable service access per site and/or per application

General Concepts



Hierarchical Addressing and Aggregation



- Prefix assignment can be larger/smaller
<http://www.icann.org/announcements/announcement-12oct06.htm>
- Provider Independent proposal:
http://www.arin.net/policy/proposals/2005_1.html
- Be careful when using /127 on P2P links (See RFC 3627)

Do I Get PI or PA?

- It depends 😊
- PI space is great for ARIN controlled space (not all RIRs have approved PI space)
- PA is a great space if you plan to use the SAME SP for a very long time OR you plan to NAT everything with IPv6 (not likely)
- More important things to consider—do you get a prefix for the entire company or do you get one prefix per site (what defines a site?)

Link Level - Prefix Length Considerations

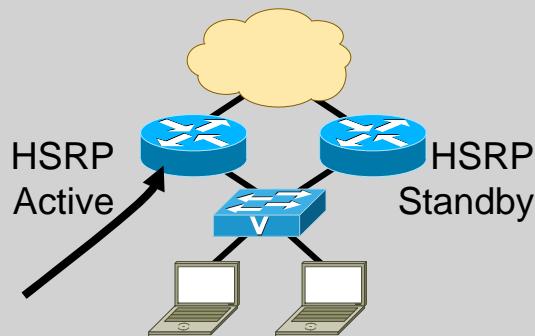
64 bits	< 64 bits	> 64 bits
<ul style="list-style-type: none">▪ Recommended by RFC3177 and IAB/IESG▪ Consistency makes management easy▪ MUST for SLAAC▪ Significant Address space loss	<ul style="list-style-type: none">▪ Enables more hosts per broadcast domain▪ Considered bad practice▪ 64 bits offers more space for hosts than the media can support efficiently	<ul style="list-style-type: none">▪ Address space conservation▪ Special cases:<ul style="list-style-type: none">/126—valid for p2p/127—not valid for p2p (RFC3627)/128—loopback▪ Complicates management▪ Must avoid overlap with specific addresses:<ul style="list-style-type: none">Router Anycast (RFC3513)Embedded RP (RFC3956)ISATAP addresses

Interface-ID Selection

Network Devices


- Reconnaissance for network devices—the search for something to attack
- Use random 64-bit interface-IDs for network devices
 - 2001:DB8:CAFE:2::1/64—Common IID
 - 2001:DB8:CAFE:2::9A43:BC5D/64—Random IID
 - 2001:DB8:CAFE:2::A001:1010/64—Semi-random IID
- Operational management challenges with this type of numbering scheme

First-Hop Router Redundancy



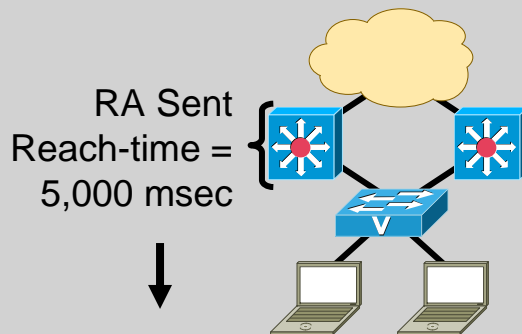
HSRP for v6

- Modification to Neighbor Advertisement, Router Advertisement, and ICMPv6 redirects
- Virtual MAC derived from HSRP group number and virtual IPv6 link-local address



GLBP for v6

- Modification to Neighbor Advertisement, Router Advertisement—GW is announced via RAs
- Virtual MAC derived from GLBP group number and virtual IPv6 link-local address



Neighbor Unreachability Detection

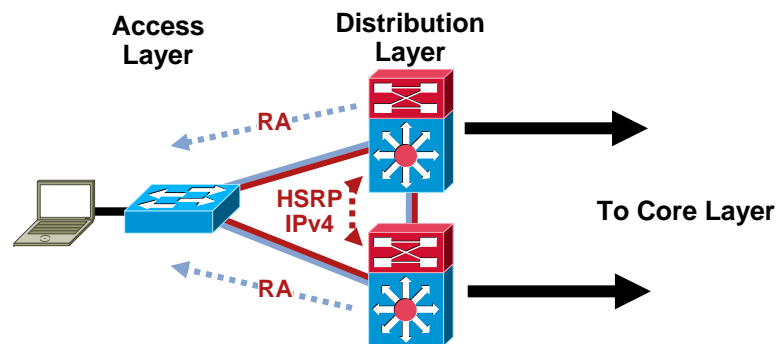
- For rudimentary HA at the first HOP
- Hosts use NUD “reachable time” to cycle to next known default gateway (30s by default)

First-Hop Redundancy

- When HSRP, GLBP and VRRP for IPv6 are not available
- NUD can be used for rudimentary HA at the first-hop (today this only applies to the Campus/DC...HSRP is available on routers)
`(config-if)#ipv6 nd reachable-time 5000`
- Hosts use NUD “reachable time” to cycle to next known default gateway (30 seconds by default)
- Can be combined with default router preference to determine primary gw:
`(config-if)#ipv6 nd router-preference {high | medium | low}`

```
Default Gateway . . . . . : 10.121.10.1  
                          fe80::211:bcff:fec0:d000%4  
                          fe80::211:bcff:fec0:c800%4
```

```
Reachable Time           : 6s  
Base Reachable Time      : 5s
```




HSRP for IPv6

- Many similarities with HSRP for IPv4
- Changes occur in Neighbor Advertisement, Router Advertisement, and ICMPv6 redirects
- No need to configure GW on hosts (RAs are sent from HSRP Active router)
- Virtual MAC derived from HSRP group number and virtual IPv6 Link-local address
- IPv6 Virtual MAC range:
0005.73A0.0000—0005.73A0.0FFF
(4096 addresses)
- HSRP IPv6 UDP Port Number 2029 (IANA Assigned)
- No HSRP IPv6 secondary address
- No HSRP IPv6 specific debug

Host with GW of Virtual IP

```
#route -A inet6 | grep ::/0 | grep eth2
::/0          fe80::5:73ff:fea0:1          UGDA 1024 0          0 eth2
```

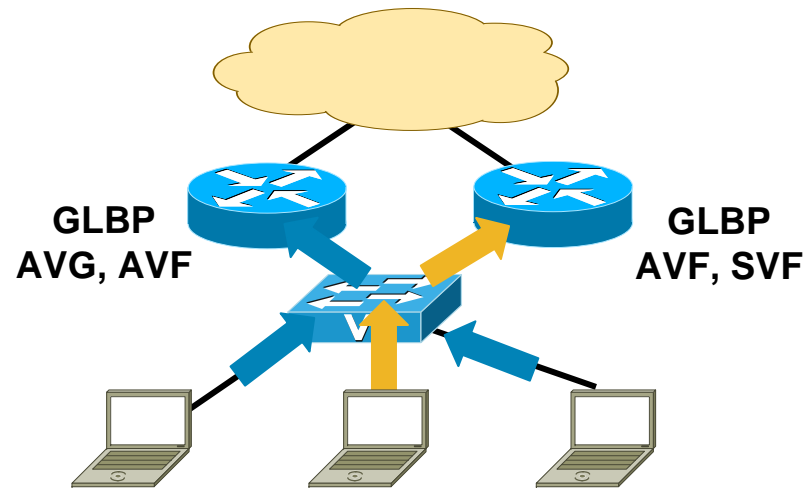


```
interface FastEthernet0/1
  ipv6 address 2001:DB8:66:67::2/64
  ipv6 cef
  standby version 2
  standby 1 ipv6 autoconfig
  standby 1 timers msec 250 msec 800
  standby 1 preempt
  standby 1 preempt delay minimum 180
  standby 1 authentication md5 key-string cisco
  standby 1 track FastEthernet0/0
```

GLBP for IPv6

- Many similarities with GLBP for IPv4 (CLI, Load-balancing)
- Modification to Neighbor Advertisement, Router Advertisement
- GW is announced via RAs
- Virtual MAC derived from GLBP group number and virtual IPv6 Link-local address

AVG=Active Virtual Gateway
AVF=Active Virtual Forwarder
SVF=Standby Virtual Forwarder

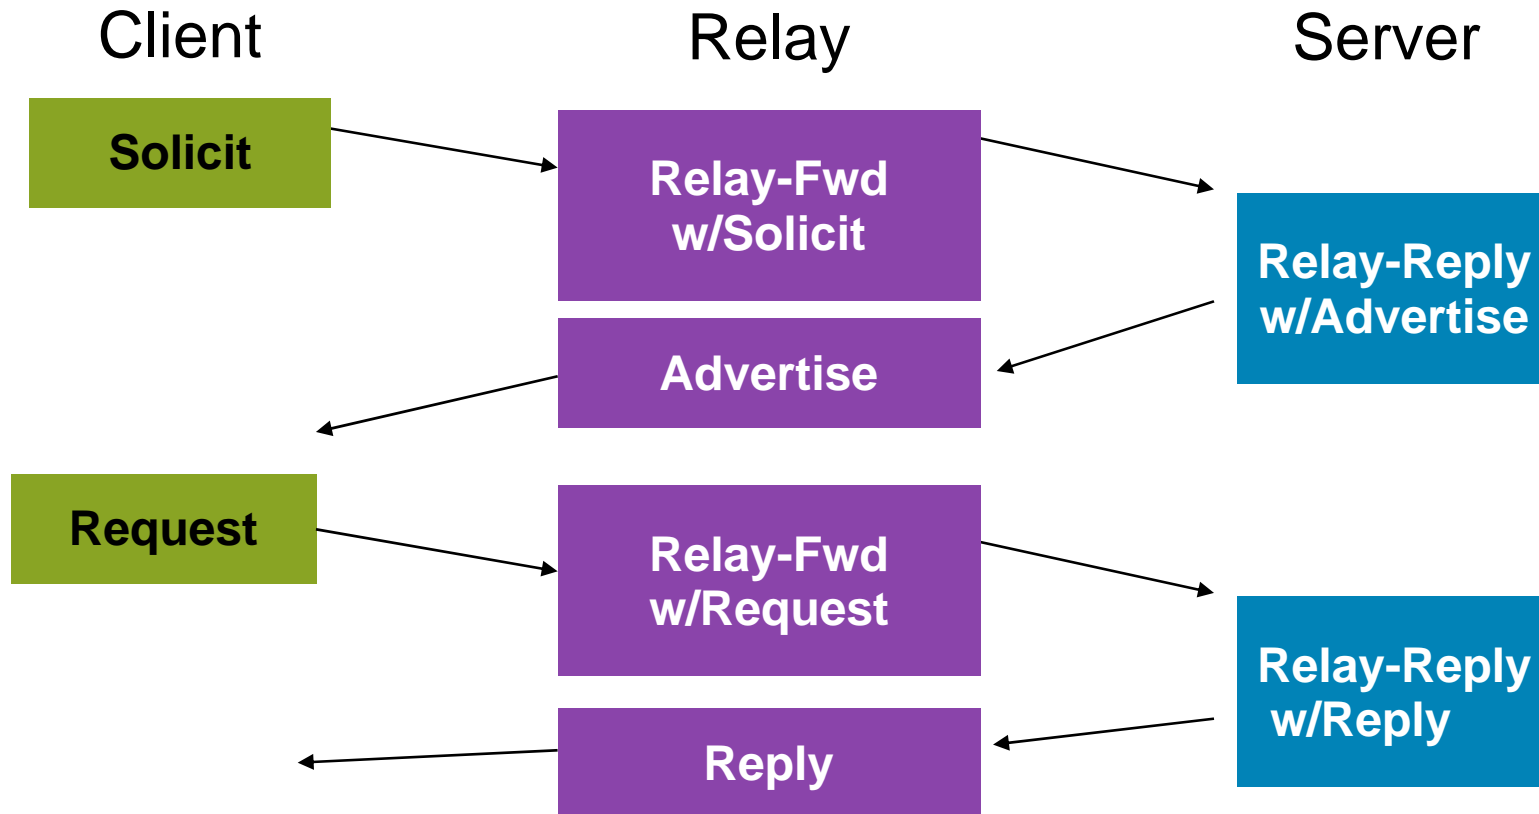


```
interface FastEthernet0/0
  ipv6 address 2001:DB8:1::1/64
  ipv6 cef
  glbp 1 ipv6 autoconfig
  glbp 1 timers msec 250 msec 750
  glbp 1 preempt delay minimum 180
  glbp 1 authentication md5 key-string cisco
```

DHCPv6

- Updated version of DHCP for IPv4
- Client detects the presence of routers on the link
- If found, then examines router advertisements to determine if DHCP can or should be used
- If no router found or if DHCP can be used, then
 - DHCP Solicit message is sent to the All-DHCP-Agents multicast address
 - Using the link-local address as the source address

DHCPv6 Operation



- All_DHCP_Relay_Agents_and_Servers (FF02::1:2)
- All_DHCP_Servers (FF05::1:3)
- DHCP Messages: Clients listen UDP port 546; servers and relay agents listen on UDP port 547

Stateful/Stateless DHCPv6

- Stateful and Stateless DHCPv6 Server

Cisco Network Registrar:

<http://www.cisco.com/en/US/products/sw/netmgmtsw/ps1982/>

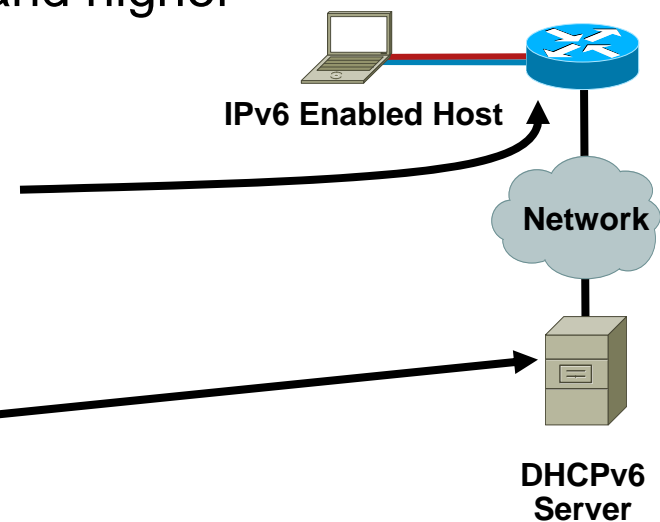
Microsoft Windows Server 2008:

<http://technet2.microsoft.com/windowsserver2008/en/library/bab0f1a1-54aa-4cef-9164-139e8bcc44751033.msp?mfr=true>

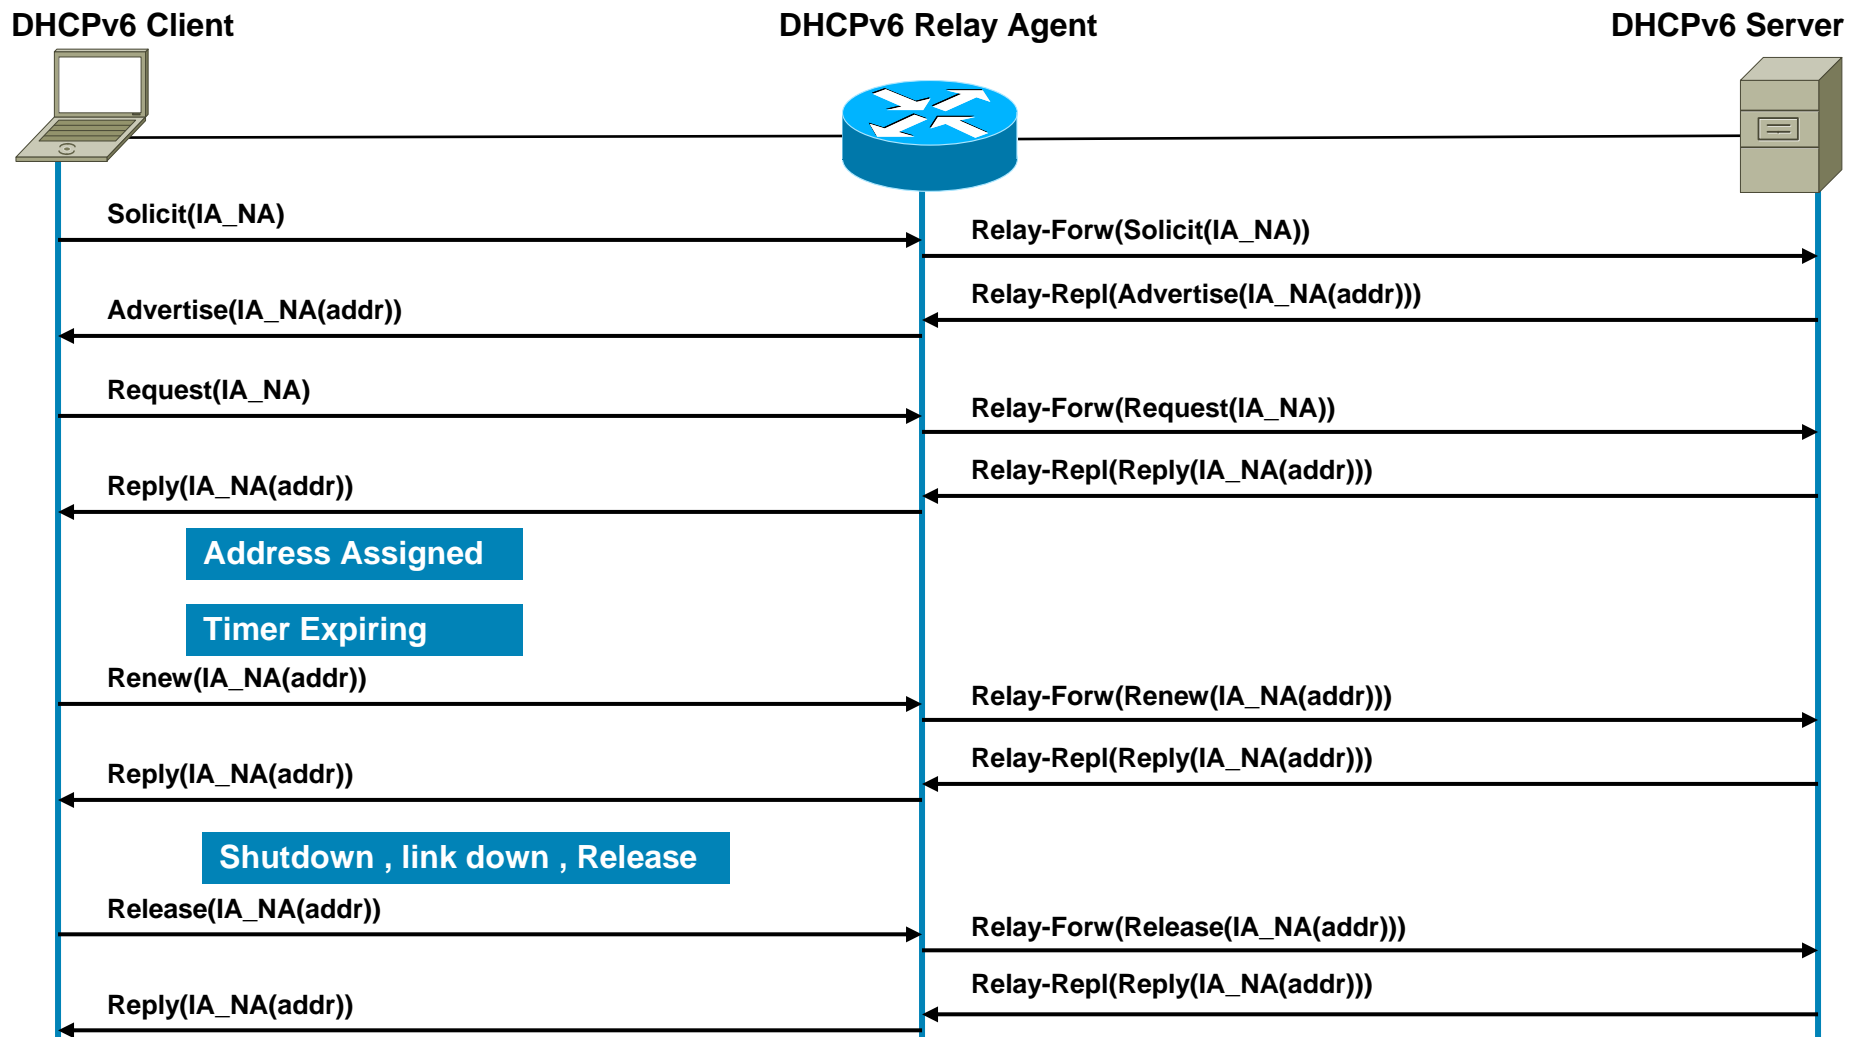
Dibbler: <http://klub.com.pl/dhcpv6/>

- DHCPv6 Relay—12.3(11)T/12.2(28)SB and higher

```
interface FastEthernet0/1
description CLIENT LINK
ipv6 address 2001:DB8:CAFE:11::1/64
ipv6 nd prefix 2001:DB8:CAFE:11::/64 no-advertise
ipv6 nd managed-config-flag
ipv6 nd other-config-flag
ipv6 dhcp relay destination 2001:DB8:CAFE:10::2
```



Basic DHCPv6 Message Exchange



CNR/W2K8—DHCPv6

CISCO SYSTEMS **Network Registrar - Local** [About](#) | [Help](#) | [Logout](#)
Name: admin
Host: shmcfarl-srv-1:1234

[Home](#) | [Administration](#) | [Servers](#) | [Clusters](#) | [Routers](#) | **[DHCP](#)** | [DNS](#) | [Hosts](#) | [Address Space](#)

[Scopes](#) | [Scope Templates](#) | [Reservations](#) | **[Prefixes](#)** | [Links](#) | [Options](#) | [Policies](#) | [Clients](#) | [Client-Classes](#) | [VPNs](#) | [Networks](#) | [Failover](#) | [DNS](#) | [LDAP](#) | [Extensions](#) | [Traps](#) | [DHCP Server](#)

List DHCPv6 Prefixes

Name	Address	Link	DHCP Type	Policy	Leases	Reservations
bld1-acc1-vlan11	2001:db8:cafe:11::/64	VLAN11	DHCP	bld1-policy	🔗	🔗

DHCP

- lhr0-01
 - IPv4
 - IPv6
 - Scope [2001:db8:cafe:10::] Local
 - Scope [2001:db8:cafe:11::] VLAN11
 - Address Leases
 - Exclusions
 - Reservations
 - Scope Options**
 - Server Options

Option Name	Vendor	Value
00023 DNS Recursive Name Server IPV6 Address List	Standard	2001:db8:cafe:10::4
00024 Domain Search List	Standard	cisco.com

IPv6 General Prefix

- Provides an easy/fast way to deploy prefix changes
- Example: 2001:db8:cafe::/48 = General Prefix
- Fill in interface specific fields after prefix

"ESE ::11:0:0:0:1" = 2001:db8:cafe:11::1/64


```
ipv6 unicast-routing
ipv6 cef
ipv6 general-prefix ESE 2001:DB8:CAFE::/48
!
interface GigabitEthernet3/2
ipv6 address ESE ::2/126
ipv6 cef
!
interface GigabitEthernet1/2
ipv6 address ESE ::E/126
ipv6 cef
```

```
interface Vlan11
  ipv6 address ESE ::11:0:0:0:1/64
  ipv6 cef
!
interface Vlan12
  ipv6 address ESE ::12:0:0:0:1/64
  ipv6 cef
```

Global unicast address(es):
2001:DB8:CAFE:11::1, subnet is 2001:DB8:CAFE:11::/64

IPv6 Multicast Availability


- Multicast Listener Discovery (MLD)
 - Equivalent to IGMP
- PIM Group Modes: Sparse Mode, Bidirectional and Source Specific Multicast
- RP Deployment: Static, Embedded
 - NO Anycast-RP Yet



Multicast Listener Discovery: MLD

Multicast Host Membership Control


- MLD is equivalent to IGMP in IPv4
- MLD messages are transported over ICMPv6
- MLD uses link local source addresses
- MLD packets use “Router Alert” in extension header (RFC2711)
- Version number confusion:
 - MLDv1 (RFC2710) like IGMPv2 (RFC2236)
 - MLDv2 (RFC3810) like IGMPv3 (RFC3376)
- MLD snooping



Multicast Deployment Options

With and Without Rendezvous Points (RP)


SSM, No RPs



ASM Single RP—Static definitions



ASM Across Single Shared PIM Domain, One RP—Embedded-RP



IPv6 QoS Syntax Changes

- IPv4 syntax has used “ip” following match/set statements

Example: `match ip dscp, set ip dscp`

- Modification in QoS syntax to support IPv6 and IPv4

New `match` criteria

`match dscp` – Match DSCP in v4/v6

`match precedence` – Match Precedence in v4/v6

New `set` criteria

`set dscp` – Set DSCP in v4/v6

`set precedence` – Set Precedence in v4/v6

- Additional support for IPv6 does not always require new Command Line Interface (CLI)

Example—WRED

Scalability and Performance

- IPv6 Neighbor Cache = ARP for IPv4

In dual-stack networks the first hop routers/switches will now have more memory consumption due to IPv6 neighbor entries (can be multiple per host) + ARP entries

ARP entry for host in the campus distribution layer:

```
Internet 10.120.2.200                2    000d.6084.2c7a  ARPA  Vlan2
```

IPv6 Neighbor Cache entry:

```
2001:DB8:CAFE:2:2891:1C0C:F52A:9DF1    4    000d.6084.2c7a  STALE V12
```

```
2001:DB8:CAFE:2:7DE5:E2B0:D4DF:97EC    16   000d.6084.2c7a  STALE V12
```

```
FE80::7DE5:E2B0:D4DF:97EC              16   000d.6084.2c7a  STALE V12
```

- Full Internet route tables—ensure to account for TCAM/Memory requirements for both IPv4/IPv6—Not all vendors can properly support both
- Multiple routing protocols—IPv4 and IPv6 will have separate routing protocols. Ensure enough CPU/Memory is present
- Control Plane impact when using tunnels—Terminate ISATAP/configured tunnels in HW platforms when attempting large scale deployments (hundreds/thousands of tunnels)

Selecting an IPv6 IGP



“IPv6 is an Evolutionary not a Revolutionary step and this is very clear in the case of routing which saw minor changes even though most of the Routing Protocols were completely rebuilt.”

IPv6 Challenges to Router Performance

Addressing Driven

- Forwarding Challenges—lookup not impacted as much as originally thought, different size prefixes typically see little difference in forwarding performance.
- Control Plane Challenges—routing table sizes:
 - IPv6 supports multiple addresses per interface (not the most significant concern at this time but it could be in the future)
 - IPv6 can have a lot more prefixes due to a significantly larger address space

IPv6 Routing Protocols

- Static Routing
- RIPng ↔ RIPv2.
- IPv6 EIGRP ↔ EIGRP
- IPv6 IS-IS
- OSPFv3 ↔ OSFPv2
- IPv6 extensions for BGP

Remember: **ipv6 unicast-routing** (for IOS not for IOS-XR)

The Questions Are the Same as for IPv4 ... Almost

- Is one routing protocol better than any other routing protocol?
- Define “Better!”



- Converges faster?
- Uses less resources?
- Easier to troubleshoot?
- Easier to configure?
- Scales to a larger number of routers, routes, or neighbors?
- More flexible?
- Degrades more gracefully?
- ...


IPv6 IGP Selection...in Theory

In **Theory**:

- The similarity between the IPv6 and IPv4 routing protocols leads to similar behaviour and expectations
- To select the IPv6 IGP, start by using the IPv4 IGP rules of thumb

IPv4 IGP Selection Rules—Topology

Rules of Thumb



IPv6 IGP Selection...in Practice

In **Practice**:

- The IPv6 IGP implementations might not be fully optimized yet so there is a bit more uncertainty
- Not all Fast Convergence optimizations might be available
- Operational experience with large scale IPv6 networks is still being developed

IGP Conclusions

- Same topology considerations as for IPv4
- Convergence time

When comparing apples to apples the convergence times are very similar

Other tools are also leveraged: Bidirectional Forwarding Detection (BFD)

There are HW and SW dependencies

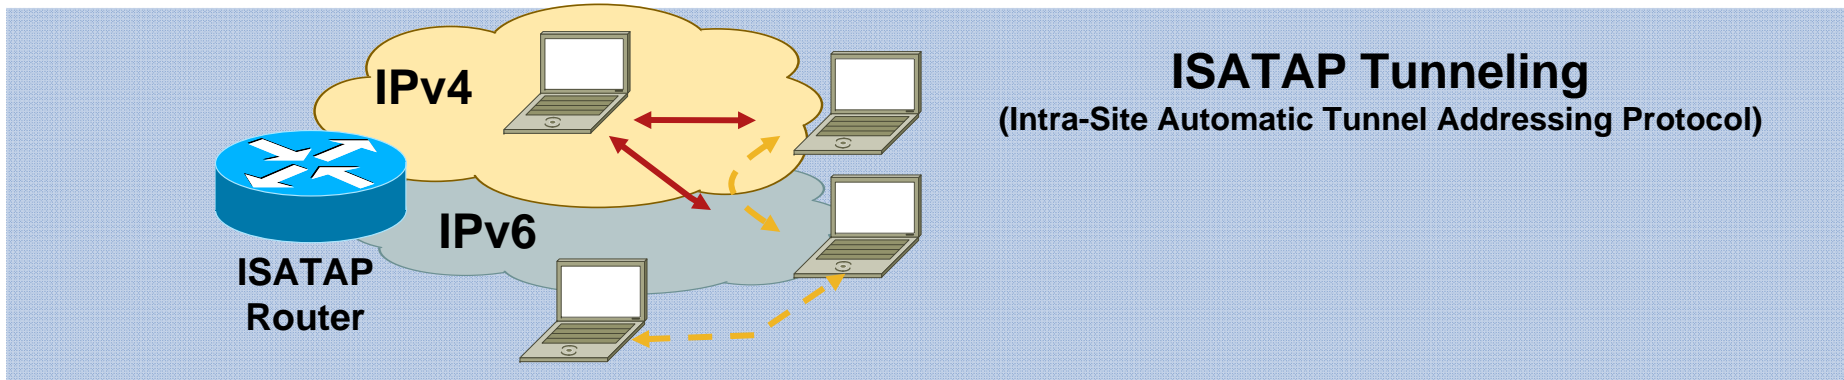
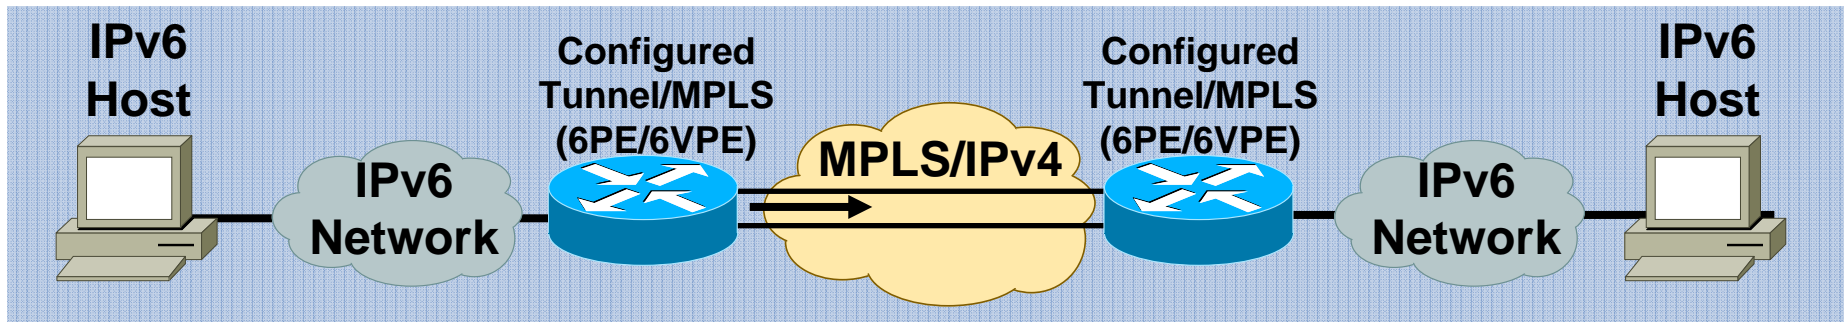
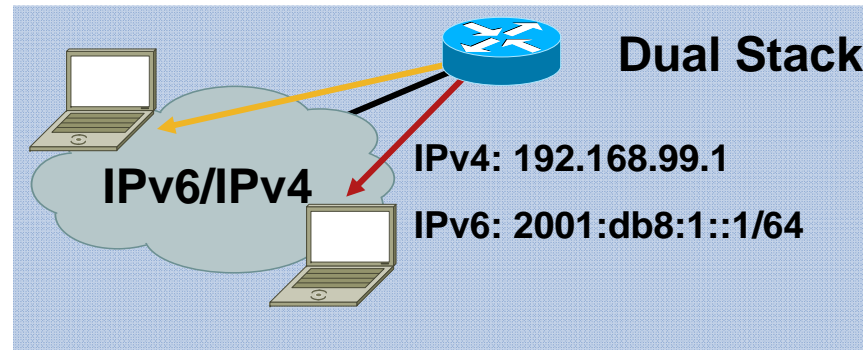
Deployment Considerations: Coexistence



Start Here: Cisco IOS Software Release Specifics for IPv6 Features

http://www.cisco.com/univercd/cc/td/doc/product/software/ios123/123cgcr/ipv6_c/ftipv6s.htm

IPv6 Coexistence



Planning and Deployment Summary



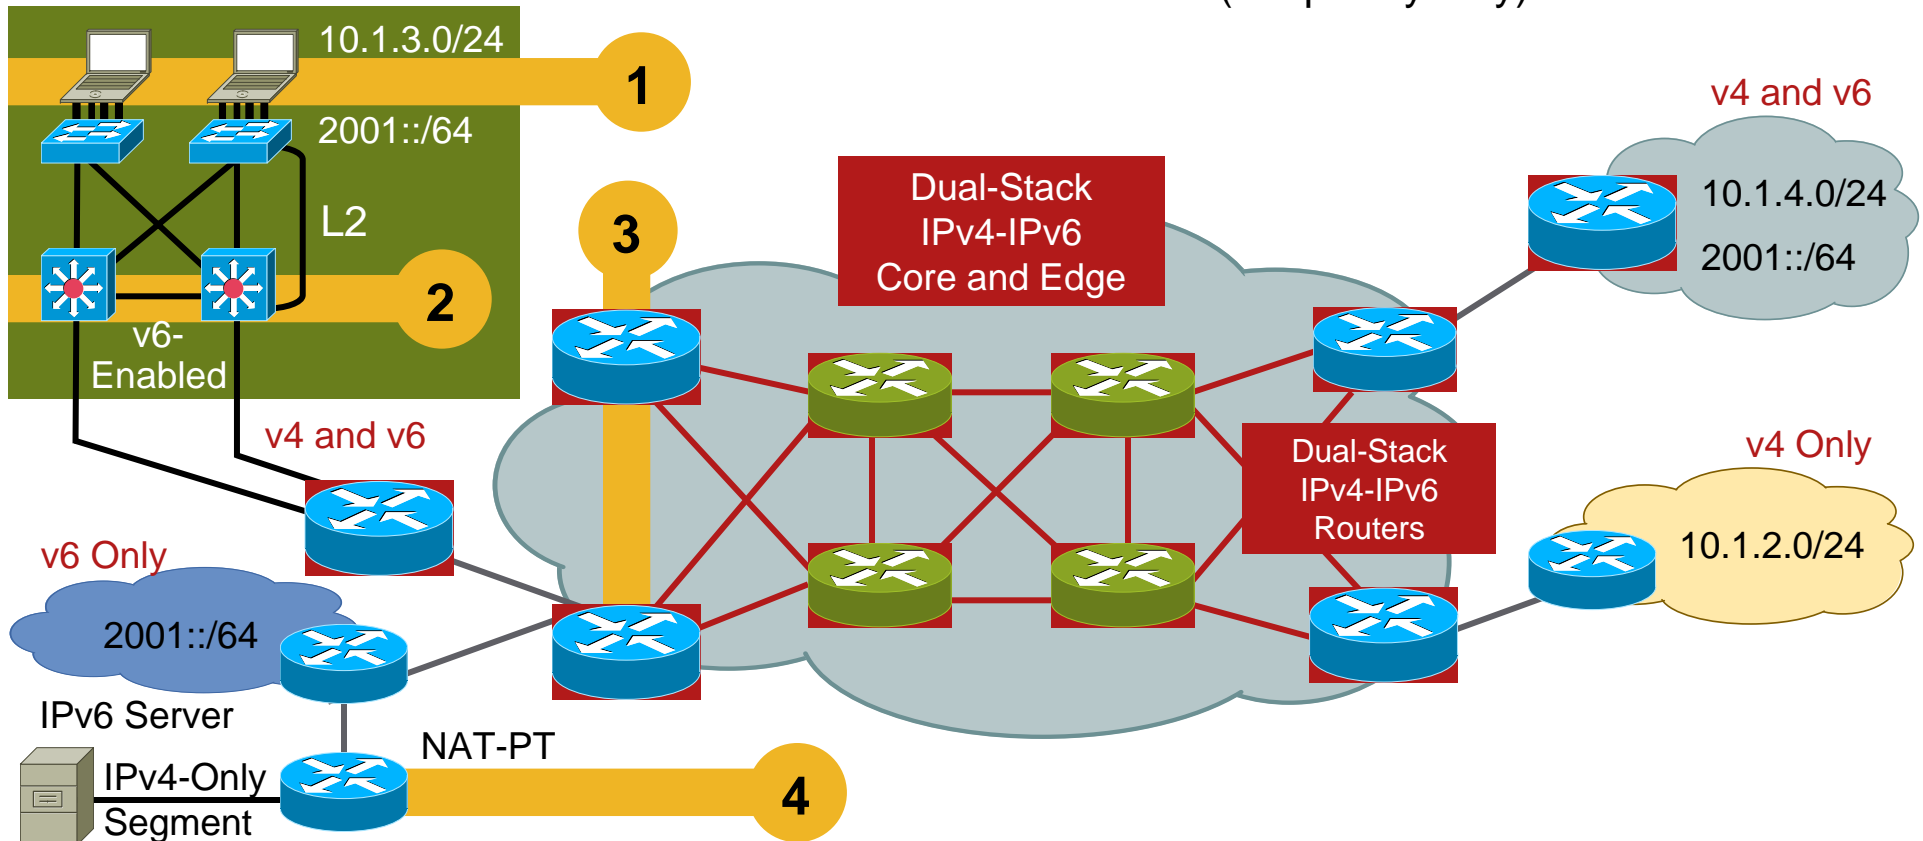
IPv6 Integration Outline

Pre-Deployment Phases	Deployment Phases
<ul style="list-style-type: none">▪ Establish the network starting point▪ Importance of a network assessment and available tools▪ Defining early IPv6 security guidelines and requirements▪ Additional IPv6 “pre-deployment” tasks needing consideration	<ul style="list-style-type: none">▪ Transport considerations for integration▪ Campus IPv6 integration options▪ WAN IPv6 integration options▪ Advanced IPv6 services options

Integration/Coexistence Starting Points

Example: Integration Demarc/Start Points in Campus/WAN

- 1 Start dual-stack on hosts/OS
- 2 Start dual-stack in campus distribution layer (details follow)
- 3 Start dual-stack on the WAN/campus core/edge routers
- 4 NAT-PT for servers/apps only capable of IPv4 (temporary only)



Pre-Deployment Checklist

Other Critical Network Planning Requirements

- ✓ **Establish starting point, network assessment, security guidelines**
- ✓ **Acquire IPv6 address block and create IPv6 addressing scheme**
- ✓ **Create and budget for an IPv6 lab** that closely emulates all network elements (routers, switches, hosts, OS)
- ✓ **Upgrade DNS server to support IPv6**
- ✓ Establish **network management** considerations (hardware, MIBs required for v6, etc.)
- ✓ **Routing and multicast protocol and selection/evaluation** process (align with IPv4 choice is possible)
- ✓ Consider options for centralized ISATAP router (see campus example)
- ✓ Evaluate **IPv6-capable transport services** available from current Service Provider (SP)

Link support to timeline needed, not before

Does L3 VPN service support QoS? Dual-homing? Security at NAP?

Transport Deployment Options for Integration

Applied to Campus, WAN, Branch, and Other

- Campus (also applies to Data Center)
 - Dual-stack (IPv4/v6 enabled on all L3 devices—core/distr/access)
 - Hybrid (combination dual-stack, tunnels, ISATAP)
 - Services block (dedicated for IPv6 ISATAP tunnel termination)
- WAN (used for core or branch interconnect)
 - Dual-stack core/edge
 - WAN L2 transport (IPv4/v6 over ATM/FR, PPP/HDLC, T1/T3, OC-x)
 - Metro Service (Ethernet, point-to-point, point-to-multipoint)
- VPN/transport considerations
 - Self-deployed MPLS VPNs: PE to PE (VPN or non-VPN service)
 - SP Offering L3 VPN service: CE to CE (encryption? QoS? multicast?)
 - Overlay 6 over 4 IPsec: site-to-site, VPN client-based using ISATAP
 - IPv6 over WiFi (802.1x is not required to be supported over IPv6)
- Other service options
 - Broadband, internet (as transport), remote access supporting IPv6

General IPv6 Requirements

Considered in Each Place in the Network

- General Coexistence

IPv4 and IPv6 coexist with no impact on performance

Flexible integration tools

- Routing

High-performance IPv6-aware routing protocols

- QoS

Identify and prioritize traffic based upon a wide-variety of criteria

Contiguous over campus, WAN, branch

SP offered

- IP Multicast

Optimize traffic utilization with a broad range of deployment types

- Security

User-based policy enforcement

Stress Host-based features

Privacy extensions

Monitoring and reporting

- Mobility

Access to applications and services while in motion

Design into core infrastructure for IPv4 and IPv6

Each Category Applied to Campus, WAN, Branch, Other

The Scope of IPv6 Deployment

Web Content Management

Applications & Application Suites

Data Center Servers

Client Access (PC's)

Printers

Collaboration Devices & Gateways

Sensors & Controllers

Networked Device Support

DNS & DHCP

Load Balancing & Content Switching

Security (Firewalls & IDS/IPS)

Content Distribution

Optimization (WAAS, SSL acceleration)

VPN Access

Networked Infrastructure Services

Deployment Scenario

IPv6 over IPv4 Tunnels
(Configured, 6to4, ISATAP, GRE)

Dual-Stack

IPv6 over MPLS
(6PE/6VPE)

IP Services (QoS, Multicast, Mobility, Translation)

Hardware Support

Connectivity

IP Addressing

Routing Protocols

Instrumentation

Basic Network Infrastructure

Staff Training and Operations

Roll-out Releases & Planning

Conclusion

- Start learning now—Books, presentations, your own pilot lab
- Create a virtual team of IT representatives from every area of IT to ensure coverage for OS, Apps, Network and Operations/Management
- Microsoft Windows Vista and Server 2008 will have IPv6 enabled by default—Understand what impact **any** OS has on the network
- Things to consider:
 - Full parity between IPv4 and IPv6 is the goal, but not a reality today
 - Watch the standards and policies

Q and A



