TOMORROW starts here.

CISCO

Cisco live!

IPv6 Planning, Deployment and Operation Considerations

BRKRST-2311

Alvaro Retana
Distinguished Engineer, Cisco Services

Agenda

- IPv6 Market Trends
- IPv6 Planning Steps
- IPv6 Addressing
- Transition Mechanisms
- IPv6 Co-existence Considerations
- Management and Operations
  – IPv6 DNS
- IPv6 Security
- Action Plan




IPv6 Market Trends

IPv6 Adoption Accelerating Worldwide

IPv4 Address Exhaustion

APAC and RIPE are out (allocating from Last /8
RIPE /22 allocation only if IPv6 address has been allocated

ARIN + LATNIC ~ Feb 2015
http://ipv6.he.net/statistics/

IPv6-Capable Devices

8 billion by 2016 (40% of all devices)*

80%
of Internet core networks are IPv6-ready

IPv6 Users

Steady growth around the globe*

2.75% globally

IPv6 Content

Google
facebook
YAHOO!

~44% of Internet content*

*US and Europe as Oct, 2012

* Source: Cisco Visual Networking Index (VNI), IPv6 adoption stats : http://6lab.cisco.com/stats

Evolving Internet ….

Connecting Things

- Devices – Phones, TV/Entertainment Systems, Game Consoles, Refrigerators, Cars, Power Meters
- Sensors - Oil Rigs, Smart Grid, Bio Sensors

Communicating

- Machine to Machine
- Vehicle to Vehicle
- Vehicle to Infrastructure

Impacting Business

- Healthcare
- Manufacturing
- Retail
- Energy
- Financial Service

Changing User Experience

- Safety
- Convenience
- Health
- Productivity

http://www.rita.dot.gov/

International Civil Aviation Organization

Internet of Things Philosophy

Drivers

Ubiquitous computing
Intelligence in things at the edge (Fog)

Ubiquitous use of IP
Convergence of proprietary protocols

Ubiquitous connectivity
Radio, Cellular, Fixed

Architectural Philosophy

From	To
Interaction with capable devices via proprietary/closed systems	Distributed intelligence & actions across standardised networks & interfaces

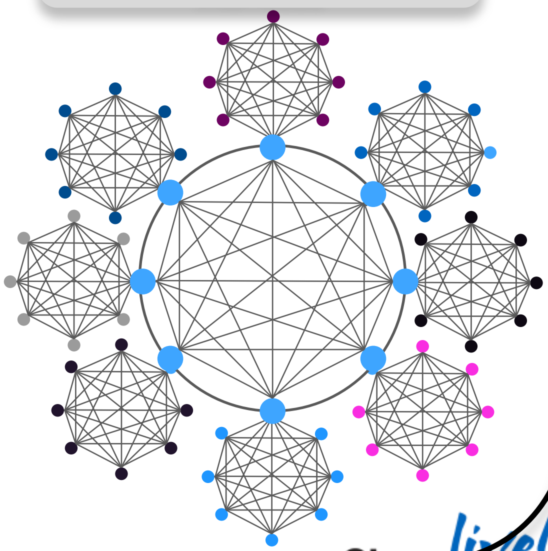

AMI IEEE 802.15.4g RF Mesh Architecture

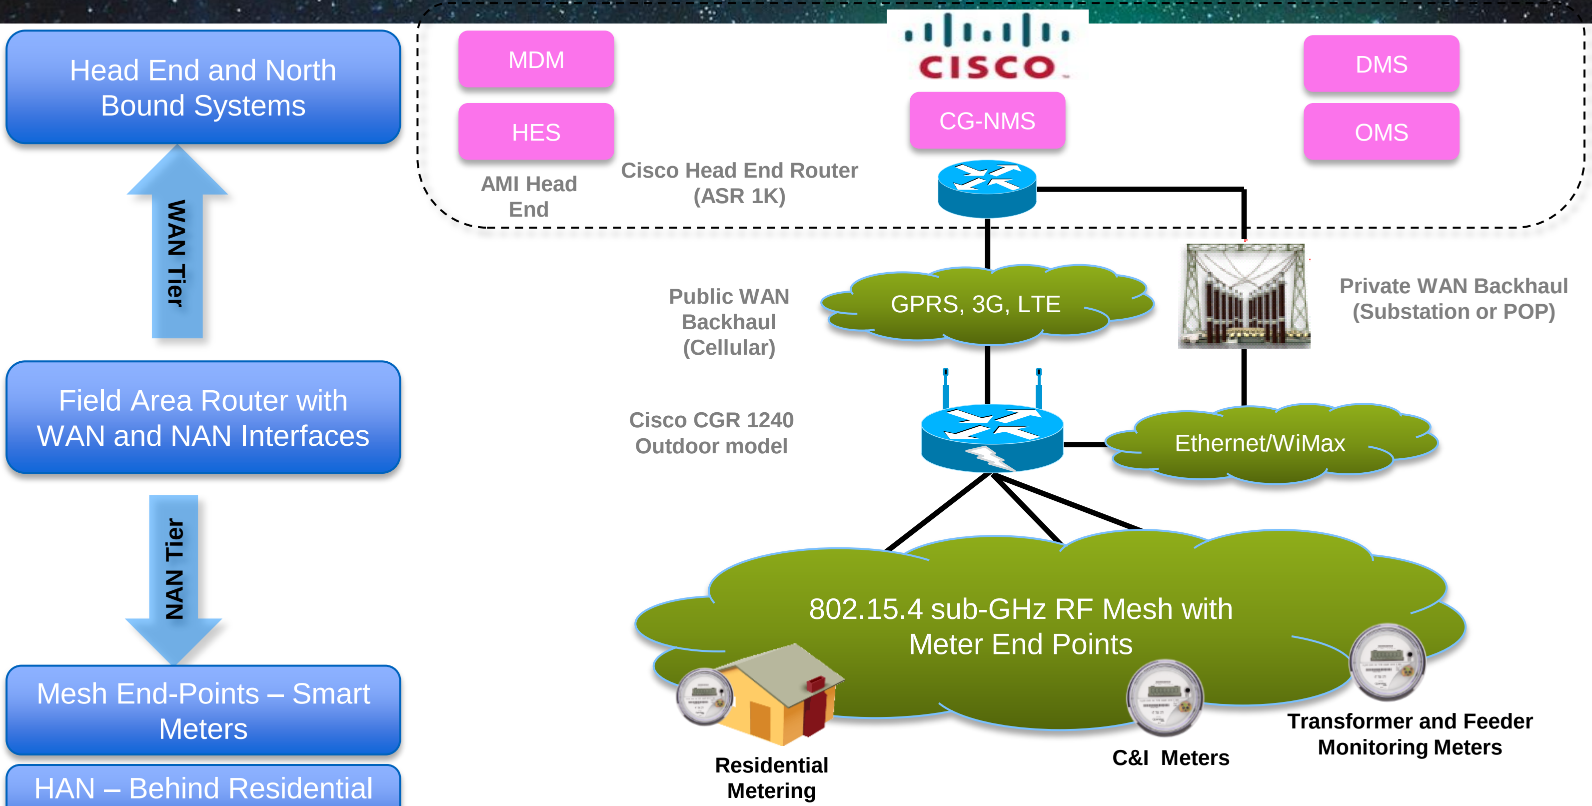

Neighbourhood Area Network based on IPv6 Wireless Mesh

IPv6 Global Deployment To Users

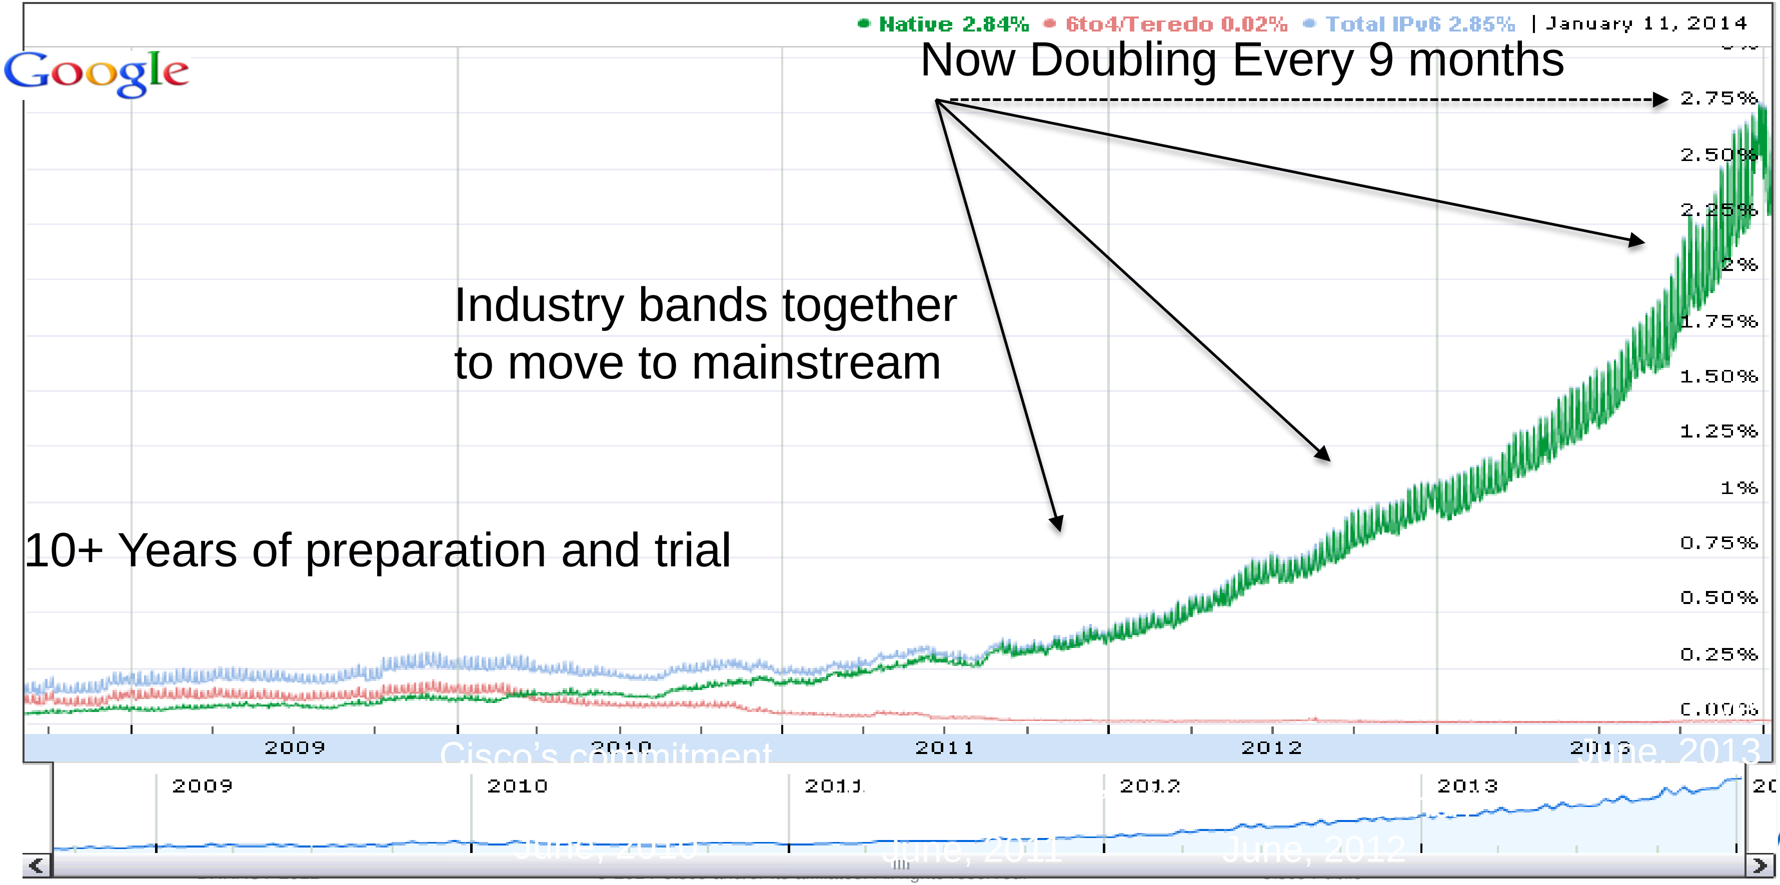

Where Are IPv6 Users Coming From?

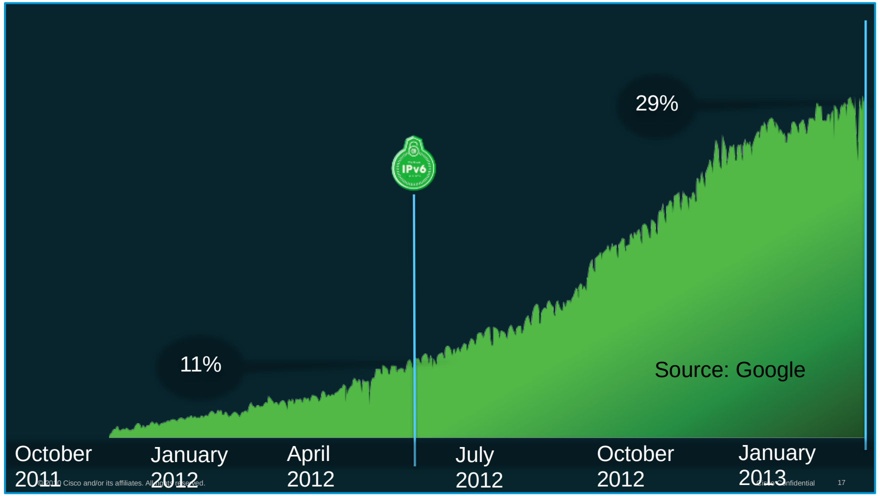

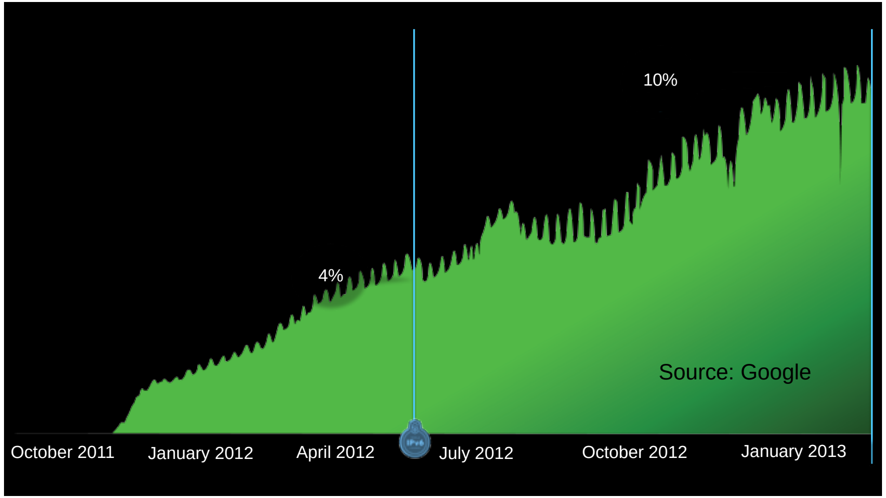

Participating Network	ASN(s)	IPv6 traffic
ATT	6389, 7018, 7132	8.39%
Free	12322	42.94%
KDDI	2516	9.31%
RCS & RDS	8708	16.06%
Verizon Wireless	6167, 22394	23.63%
Comcast	7015, 7016, 7725, 7922, 11025, 13367, 13385, 20214, 21508, 22258, 33287, 33489, 33490, 33491, 33650, 33651, 33652, 33653, 33654, 33655, 33656, 33657, 33659, 33660, 33661, 33662, 33664, 33665, 33666, 33667, 33668, 36733	1.64%

January 16th 2014

ATT	13.53%
Free	34.28%
KDDI	8.94%
RCS &RDS	23.80%
Verizon Wireless	40.03%
Comcast	20.61%
Telefonica Peru	4.60%
Deutsche Telekom AG	15.50%

Ubiquitous IPv6 Access

Adoption Metrics

- Google is seeing about 8% of traffic from Cisco using IPv6
- Performance is increasing significantly

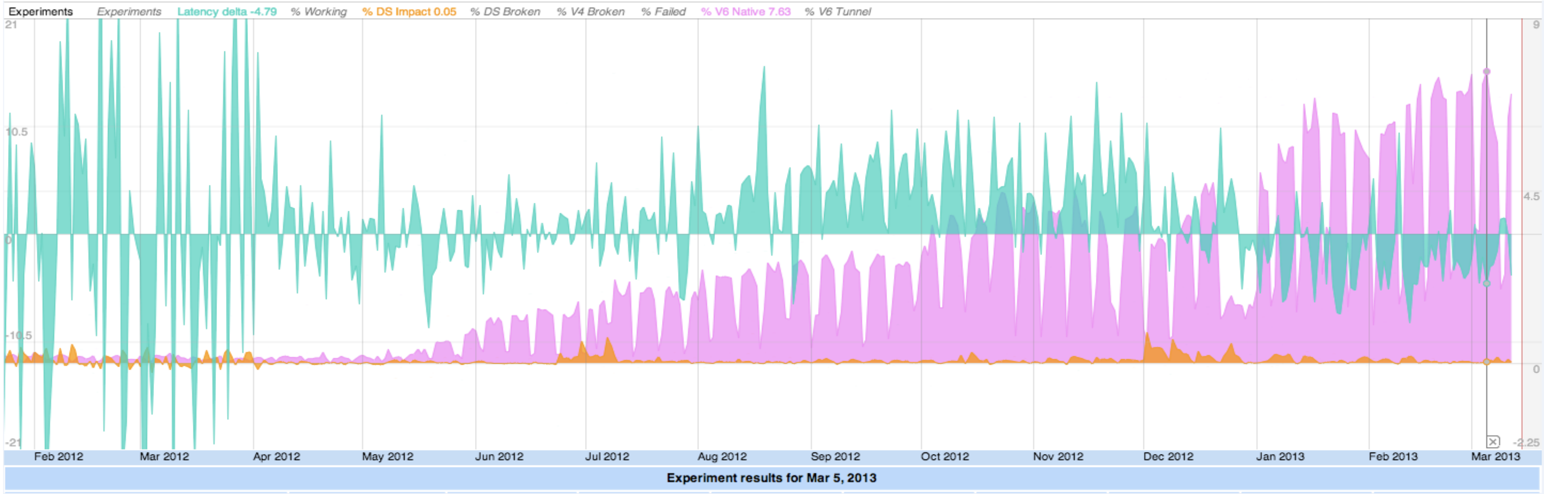

→ 13%
1/16/2014

Source: Google

Deployment Concerns

ISP Concerns

- Difficult to add/support new IPv4 customers
- Have to deal w/ smaller IPv4 address blocks
- Difficult to plan for IP NGN Services
- Business Continuity could be impacted
- Re-use of private address blocks

Enterprise Concerns

- What does IPv4 address depletion mean for us?
- How complex is IPv6 migration? What are the potential challenges?
- How should we go about migrating/transiting to IPv6?
- What are the key benefits of migrating to IPv6?

IPv6 is inevitable. Not migrating to IPv6 is not an option

General Observations

- Service Providers do not seem to consider IPv6 unless…
  - A lack of IPv4 space hinders their progress or there is consumer demand
  - However, IPv6 underpins SP transformation - collaboration, content delivery, mobility, video, cloud, m2m
- Enterprises will not ask for IPv6 unless…
  - They have an application requirement to drive it
  - Their presence on the Internet is compromised by lack of IPv6 access
  - The price of an IPv4 address exceeds the hardware cost to route it
- Consumers are generally ambivalent
  - Do not/should Not care whether IPv4 or IPv6 broadband delivery




IPv6 Planning

The Scope of IPv6 Deployment

Planning and coordination is required from many across the organisation, including …

- ✓Network engineers & operators
- ✓Security engineers
- ✓Application developers
- ✓Desktop / Server engineers
- ✓Web hosting / content developers
- ✓Business development managers
- ✓…

Moreover, **training will be required** for all involved in supporting the various IPv6 based network services

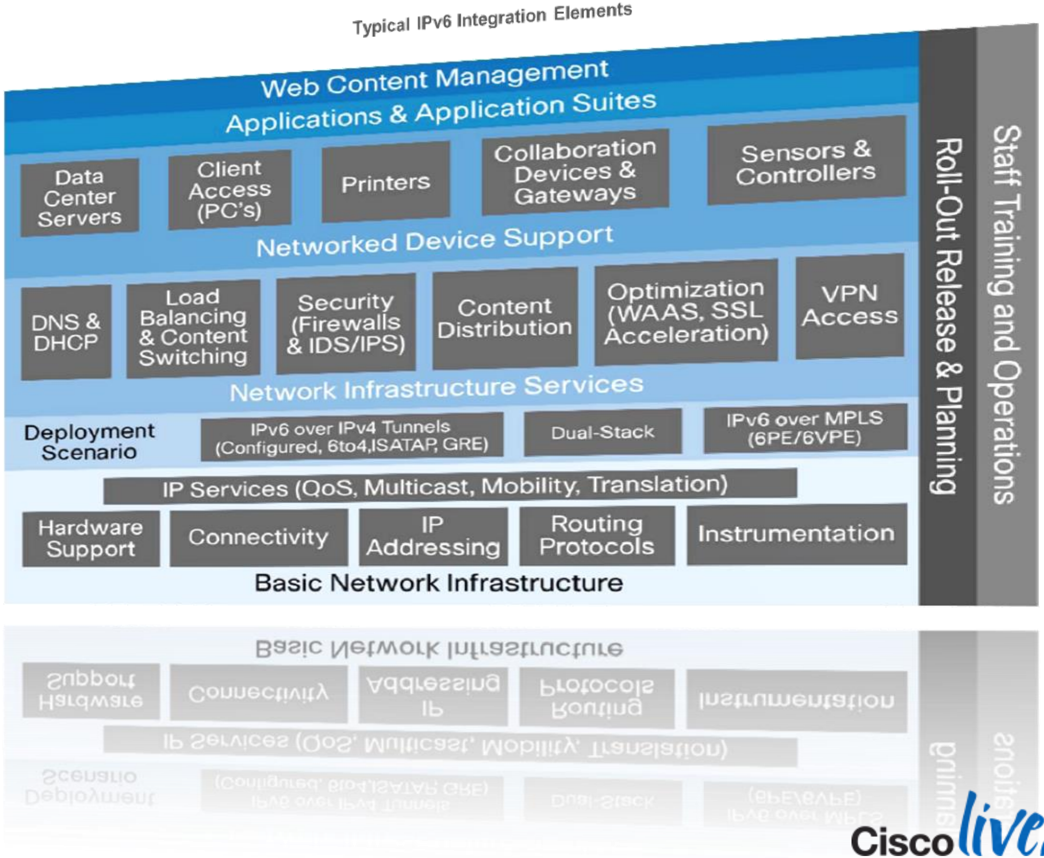

IPv6 Integration Planning

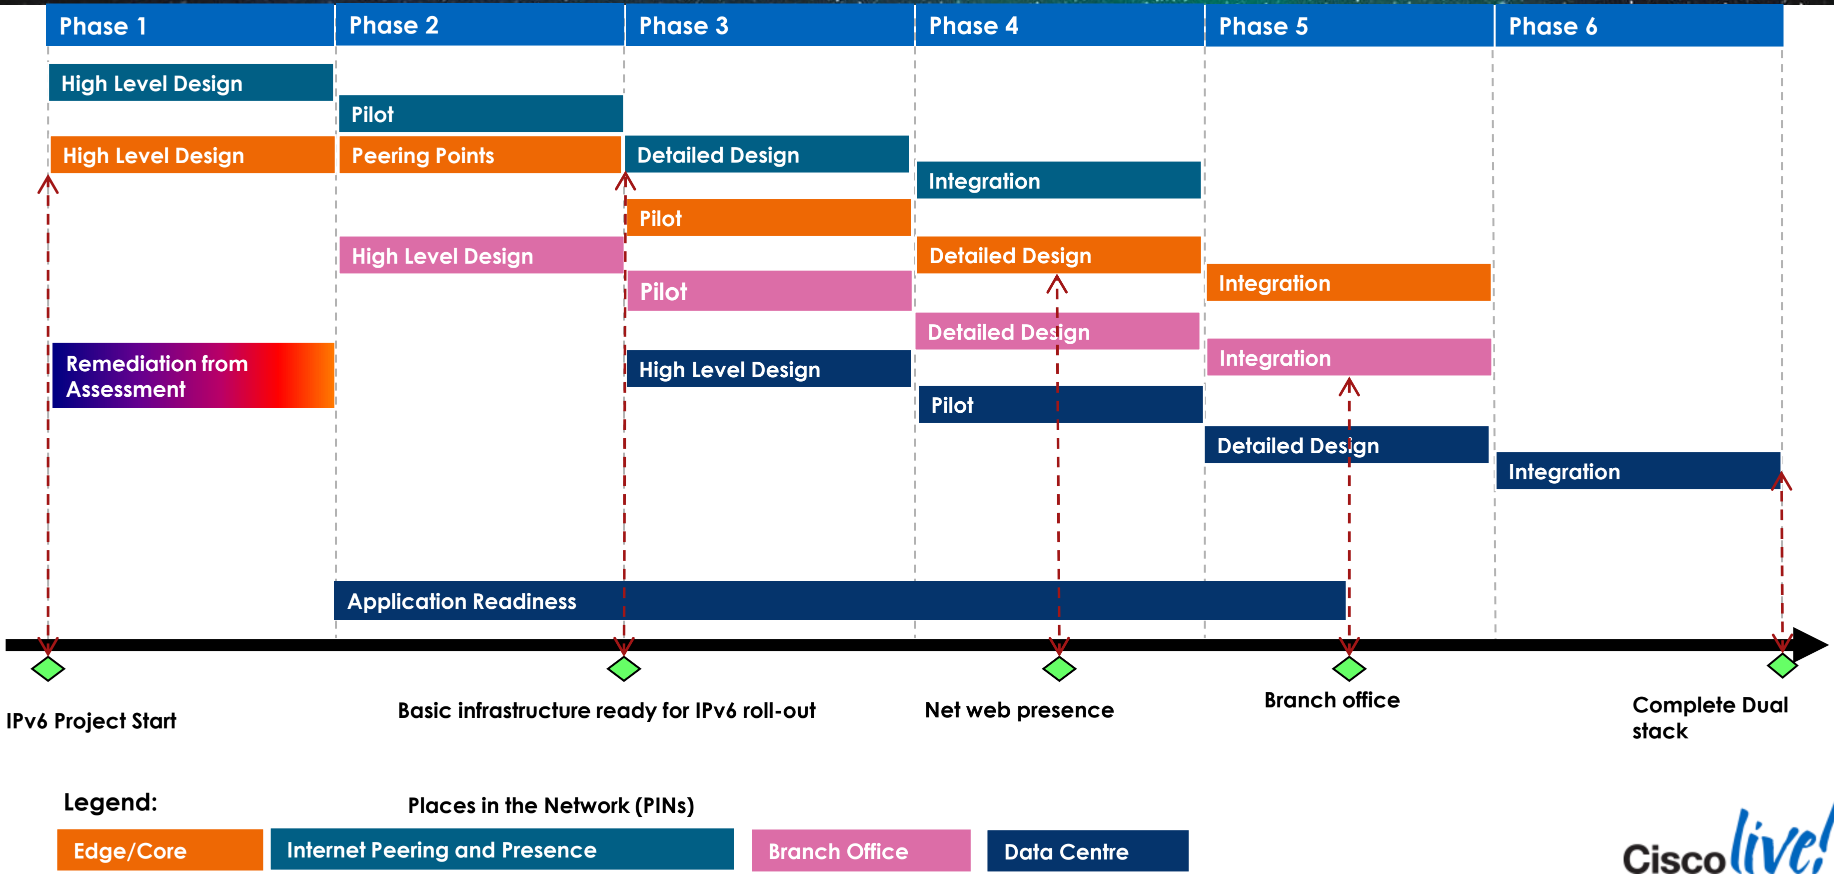

High Level Lessons Learned

- Cross functional effort across the IT Stack
  - Starts with networking team taking the lead
  - Early engagement of security team, infrastructure and application teams follow
- Business case for IPv6 Internet Presence is simpler to articulate
- Business case for IPv6 on internal corporate network takes more work
- Absorb the IPv6 effort into existing network lifecycle management process
- Security concerns and mitigation
- Operational readiness
  - Training and knowledge of operations staff
  - Network management and tooling, Configuration (automate where you can)
- Planning is key, so is early hands-on experience with IPv6

Internal Network: Where do I Start ?

- Life-Cycle management, depends on Timing and Use case
- Native/Dual-Stack where you can, Tunnels where you must
- Security – Visibility – Management
- IPv6 Host Configuration.

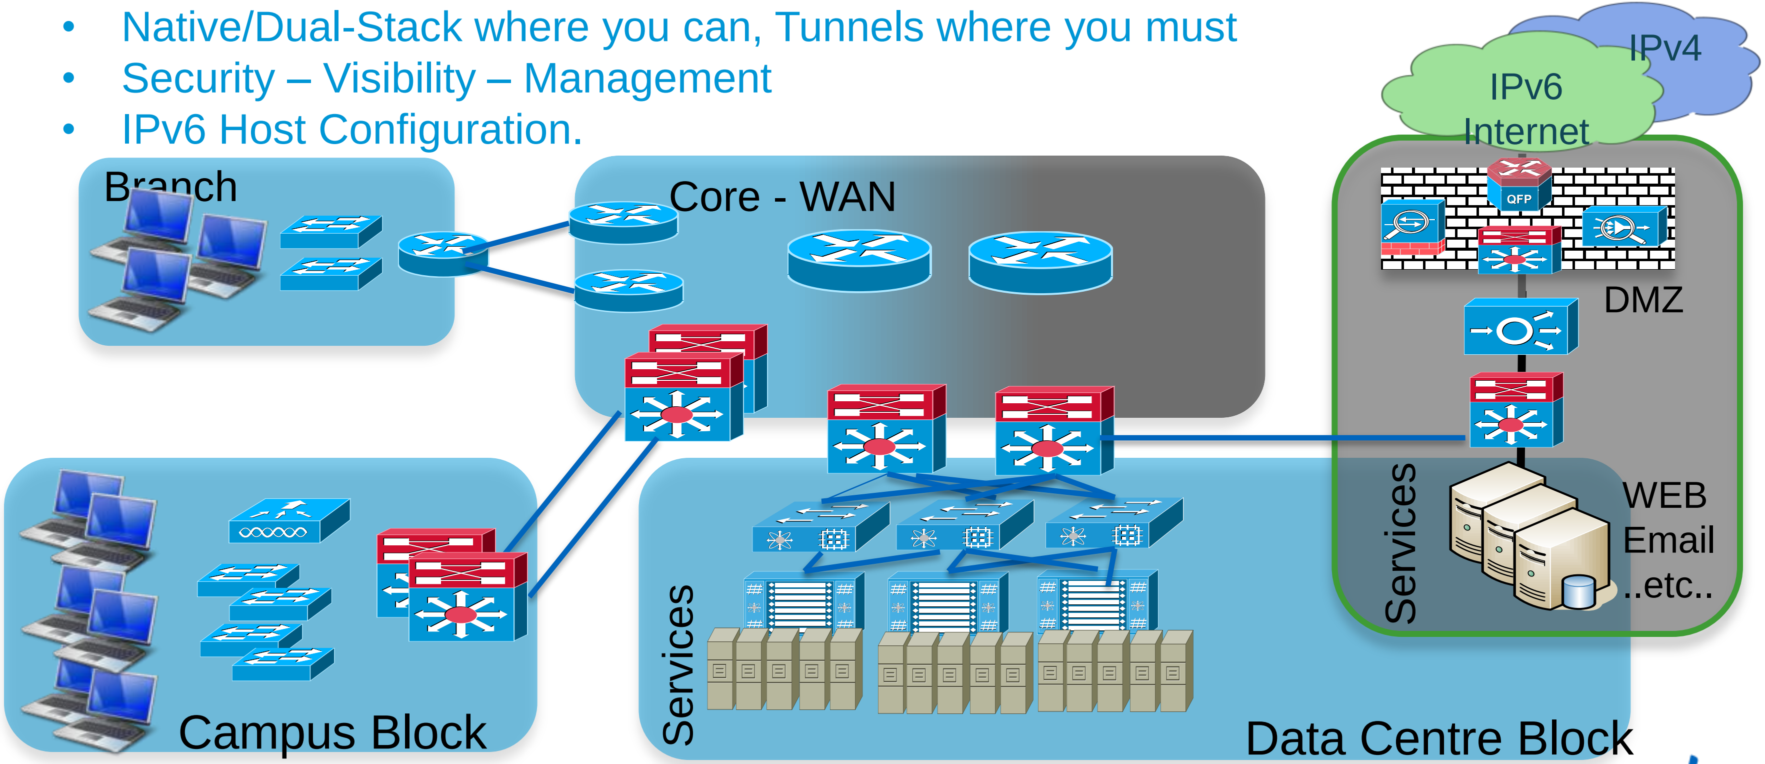

Core to Edge !

Orderly Transition – Slow to dual-Stack all the way to user

- Dual-Stack Core – Network based Tunnel to connect island
- ISATAP for IPv6 services to users… Design gotchas
- Dual-Stack selected part of DC (server front-end)

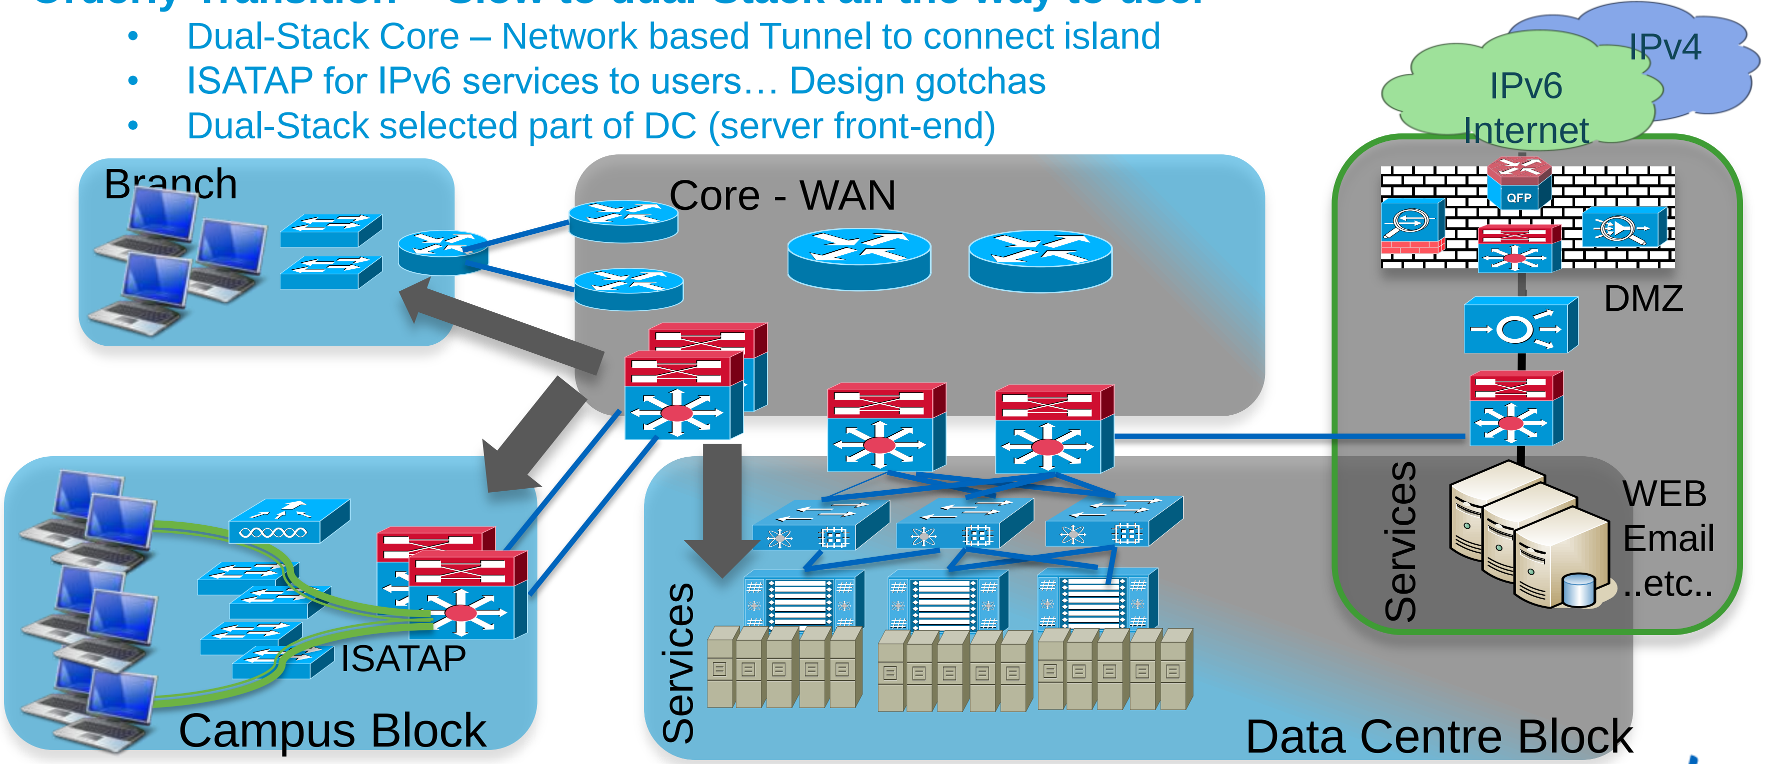

Edge to Core!

End User and Service first - Challenging but Doable

- First Hop Security
- Network based Tunnel to connect Islands
- Dual-Stack selected part of DC (server front-end)

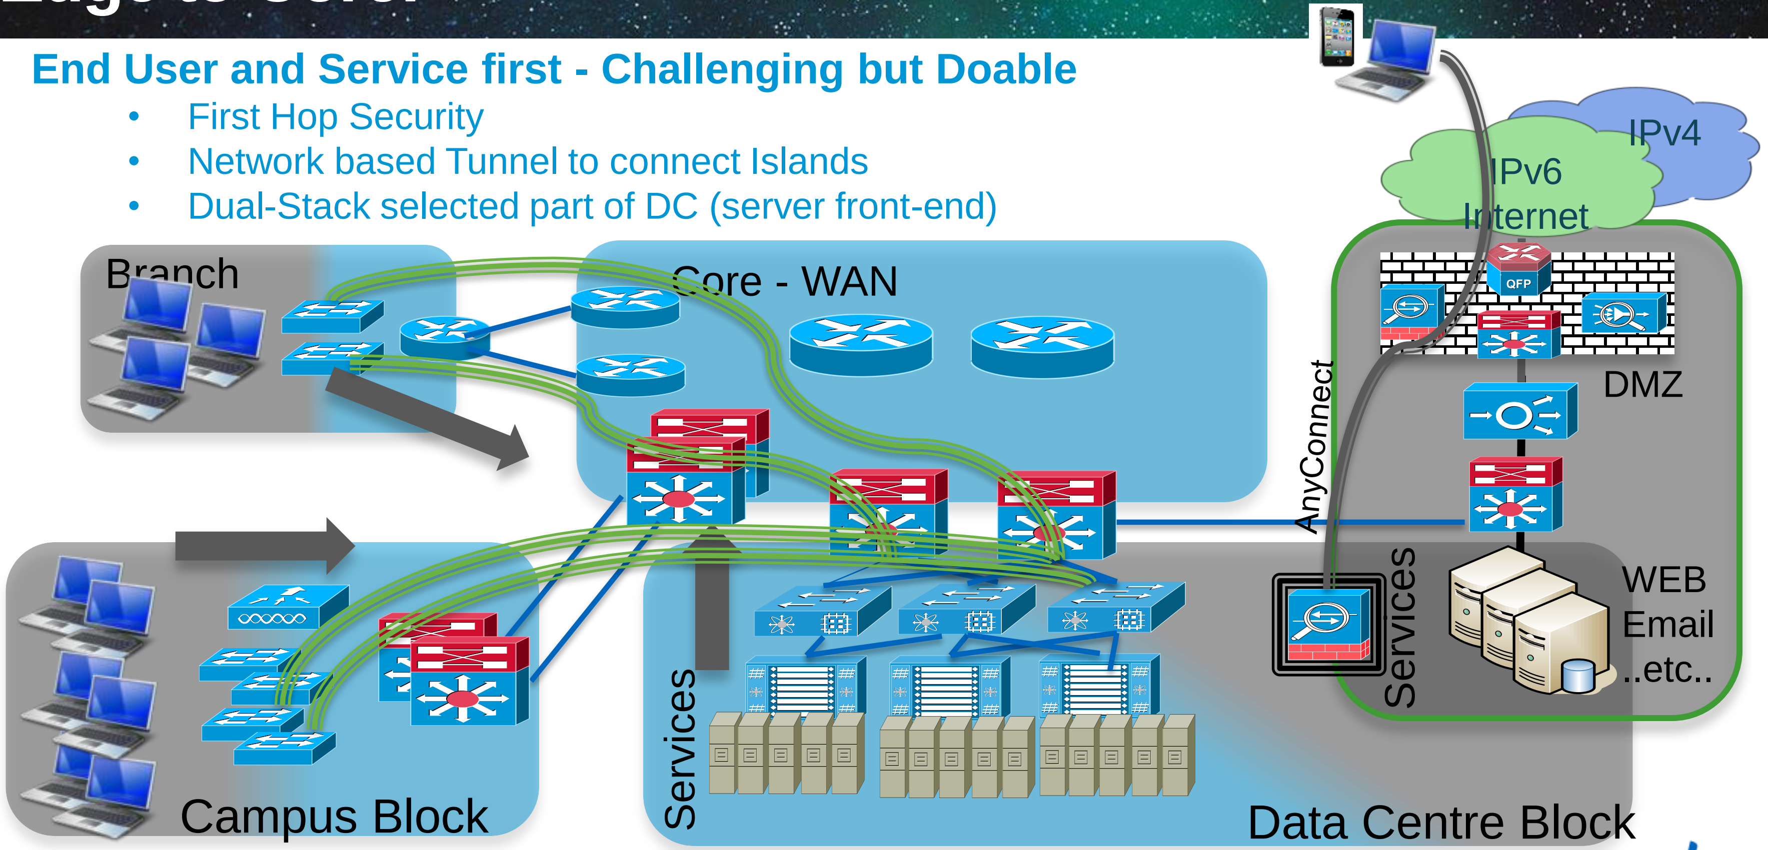

Questions to Ask Your Service Provider

http://docwiki.cisco.com/wiki/What_To_Ask_From_Your_Service_Provider_About_IPv6

- SP Deployment Type
  - Dual Stack, Native or Overlay (if so what kind of overlay) ?
  - What kind of SLA are provided for the services ? Do you post metrics online ?
- What kind of services are offered
  - Internet Services
  - Layer 2 or Layer 3 VPN's
  - IPv6 Multicast support or plans ?
  - DNS Services over v4 or V6 ?
- Visibility and footprint to the IPv6 Internet
  - Peering arrangements
- Service availability on nodes
- Acceptance Policy
  - Prefix length acceptance?
  - Provider Independent or Provider Assigned acceptance
  - Do your Peering partners have similar policy to yours?
  - What prefix length do your upstream providers accept ?
- Provisioning
  - Is there a self service portal ?
  - Routing add and deletes
  - When do you plan on providing v6 services as a default offering ?
- Charging model
  - Do you charge for IPv6 ?

IPv6 Assessments

IPv6 Readiness Considerations

- Network Hardware & Software Readiness
  - Check all network hardware for correct memory/software
  - Can device support IPv6 and needed features?
  - Is device in Critical Path ? Is IPv6 forwarded in the HW path?
- Establish upgrade Plan
  - Is new hardware needed ??
  - Does software need to be upgraded to get a certain features ?
  - How many devices will need to be upgraded ?
  - Resource budgeting, Maintenance windows etc .
- Ensure new procurements of hardware/software is IPv6 capable
- Identify components that will remain on IPv4
  – Could be for many reasons technical, business or cost

Readiness Assessment

- A key and mandatory step to evaluate the impact of IPv6 integration
- May be split in several phases
  - Infrastructure – networking devices and back end systems
  - Hosts, Servers and applications
- Must be as complete as possible to allow upgrade costs evaluation and planning
  - Hardware type, memory size, interfaces, CPU load,…
  - Software version, features enabled, license type,…, forwarding path, known limitations, best practices, etc.
- Difficult to complete if a set of features is not defined per device's category for a specific environment
  - IPv6-capable definition, knowledge of the environment and applications, design goals
- Break Network into Places in the network for a more accurate assessment
  - Should Map directly into your IPv6 Network Architecture strategy, Cost analysis and time lines

Assessment Example

- Break the project down into phases
- Determine place in the network (PIN), platforms, features that are needed in each phase
- Work with your vendor to address the gaps

	ISR G1/G2	ASR 1000	6500 (Sup 720)	3750
Phase I (Initial Deployment - Infrastructure Only)				
IPv6 Neighbor Discovery	12.2(2)T	12.2(33)XNA	12.2(17a)SX1	12.2(25)SEA
IPv6 Address Types— Unicast	12.2(2)T	12.2(33)XNA	12.2(17a)SX1	12.2(25)SEA
ICMPv6	12.2(2)T	12.2(33)XNA	12.2(17a)SX1	12.2(25)SEA
EIGRPv6	12.4(6)T	12.2(33)XNA	12.2(33)SXI	12.2(40)SE
SSH	12.2(8)T	12.2(33)XNA	12.2(17a)SX1	12.2(25)SEE
Phase II (Internet Edge Enablement)				
Multiprotocol BGP Extensions for IPv6	12.2(2)T	12.2(33)XNA	12.2(17a)SX1	-
NetFlow for IPv6 Unicast Traffic	12.3(7)T	12.2(33)XNC	12.2(33)SXH	-
RFC 4293 IP-MIB and RFC 4292 IP-FORWARD-MIB (IPv6 Only)*	15.1(3)T	12.2(33)XNA	12.2(50)SY	12.2(58)SE
IPv6 over IPv4 GRE Tunnels	12.2(4)T	12.2(33)XNA	12.2(17a)SX1	-
NAT64 - Stateful	-	15.1(3)S	-	-
Phase III (Access Edge Enablement)				
IPv6 RA Guard	-	-	12.2(33)SXI4	-**
HSRP for IPv6 (HSRPv2)	12.4(4)T	15.1(3)S	12.2(33)SXI	12.2(46)SE
HSRP Global IPv6 Address	-	-	12.2(33)SXI4	-
DHCPv6 Relay Agent	12.3(11)T	-	12.2(33)SXI	12.2(46)SE
* Must include HW switched packets				
** 12.2(46)SE does support PACL				

Commonly Deployed IPv6-enabled OS/Apps

Operating Systems

- Windows 7
- Windows Server 2008/R2
- SUSE
- Red Hat
- Ubuntu
- The list goes on

Virtualisation & Applications

- VMware vSphere 4.1
- Microsoft Hyper-V
- Microsoft Exchange 2007 SP1/2010
- Apache/IIS Web Services
- Windows Media Services
- Multiple Line of Business apps

Most commercial applications won't be your problem – it will be the custom/home-grown apps that are difficult

Coexistence Strategy

Don’t Forget the Applications

While infrastructure is everyone’s initial focus, nothing happens until the applications use the new API. IPv4-only apps will remain IPv4-only, and these legacy apps will fail when presented with an IPv6-only infrastructure.

Line Number : 39 Type :STRUCTURE
Name: sockaddr
Migration Tip: 1. If you are using struct sockaddr to allocate storage, you need to change sockaddr to sockaddr_in6

Dual Stack Affecting IPv4 Applications

- Slowness because of IPv6 Path brokenness
  - Need registry fix to override the default behaviour to chose IPv6 stack
  - Happy Eyeballs
- Embedded IPv4 addresses
- Path MTU
  - Fragmentation and Reassembly… adds latency.
- Address representation and logging
  - Scripts that match on address
  - IP Address Logging - Database Structure: Is the database is structured to accommodate the IPv6 addresses?

IPv4 Address Audit

- Assess how the existing IPv4 address space is used
- Useful information for
  - IPv6 integration
  - IPv4 address consolidation
  - Reclaiming unused address space
- Use existing tools
  - IPAM
  - ARP tables
  - Routing tables
  - DHCP logs

→ Better visibility into how the existing Address space is used → Can better answer when IPv6 is critical

IPv6 Addressing

IPv6 Address Space

- Possible Options
  - Get one large global block from local RIR and subnet out per region
  - Get a separate block from each of the RIR you have presence in
- Which route to go ?
  - Depends on specific business case
  - Enterprise that have a heavy consumer interaction using a block from each RIR will help avoid DNS and routing hacks to lead clients to regional Data Centres
- Do I Get PI or PA?
  - PI space is great for organisations who want to multihome to different SPs changing ARIN policy on block sizing
  - PA is a great space if you plan to use the same SP for a very long time or you plan to NAT/Proxy everything with IPv6 (not likely)
- Building the IPv6 Address Plan
  - Hierarchy is key
  - Cisco IPv6 Addressing White Paper
    - http://www.cisco.com/en/US/docs/solutions/SBA/February2013/Cisco_SBA_BN_IPv6AddressingGuide-Feb2013.pdf

PI Space Concerns

- Concerns around prefix announcement from other regions
  - Will providers accept prefixes from other regions?
- Concerns around prefix lengths
  - What length prefix will providers accept?
  - How do I do traffic engineering?
  - What about providers upstream peers?
- Bottom line is to have a detailed conversation w/ your provider or peering partner about what their policies are
  - http://www.us.ntt.net/support/policy/routing.cfm#v6PeerFilter

IPv6 Address Considerations

- Many ways of building an IPv6 Address Plan
  - Regional Breakdown, Purpose built or Generic buckets, Separate per business function, M&A or divestment focused
  - No matter which method you use look for ways to have some structure and Hierarchy
  - Don’t worry too much about potential inefficiencies
    - IPv6 is much larger space and trade IPv4 conservation mentality for Operational benefits

- Prefix length selection
  - P2P links, Host LAN, Small LAN interconnecting network elements

- Addressing hosts
  - SLAAC, DHCP (stateful), DHCP (stateless), Manually assigned

Infrastructure - Type of Address

- Global Unicast vs. Unique Local Address for Infrastructure

Global Unicast Address	Unique Local Address (ULA)
Use of Global address space – requiring a registered address block	Free – you could use FC00::/8 or FD00::/8
No need for Address translation or Proxy for host trying to reach to Internet	Requires translation from Private to public address – there is no scalable translation solution giving V4 type NAT/PAT
Operationally Simplistic because managing only one type of space	Management becomes complicated – have to manage private and public spaces
Could gain the same security as using ULA, if filtering is done correctly at the edge	No Security benefit of using Private space – the infrastructure could still get under attack if optimal security not in place at the Edge
Global Reachability means even connectivity to islands spread out connected via Internet	No Global reach meaning islands connected over Internet have to be administered in isolation

Recommendation: Use Global Unicast Addresses

P2P Links IPv6 Address Selection

/64	/126	/127
Ping Pong could occur if a packet sent to an un-specified address	Theoretically Optimal but still could result in a ping pong loop	Old RFC 3627 and 5375 recommends using against /127 due to Subnet-Router anycast but newer RFC 6164 Recommends using /127
Common use with overall consistency to other LAN blocks – IOS devices have a fix for Ping pong loops	Common use keeping IPv4 type of conservation mentality – IOS devices have a fix for Ping pong loops	Cisco devices disables Subnet-Router Anycast upon configuration of a /127 address
Also, mandated by RFC 4443 to send a Code 3 Destination Unreachable message to the neighbour router	Also, mandated by RFC 4443 to send a Code 3 Destination Unreachable message to the neighbour router	Most vendor equipment does not use subnet-router anycast
Use this style, if operational focus to keep the same length across the board	Use this style, if operational focus to keep the v4 /30 type addressing semantics	Use this style, if operational focus to keep the v4 /31 type addressing semantics

Recommendation: Use what makes sense within the context of your organisation

/48 Prefix Breakdown Example

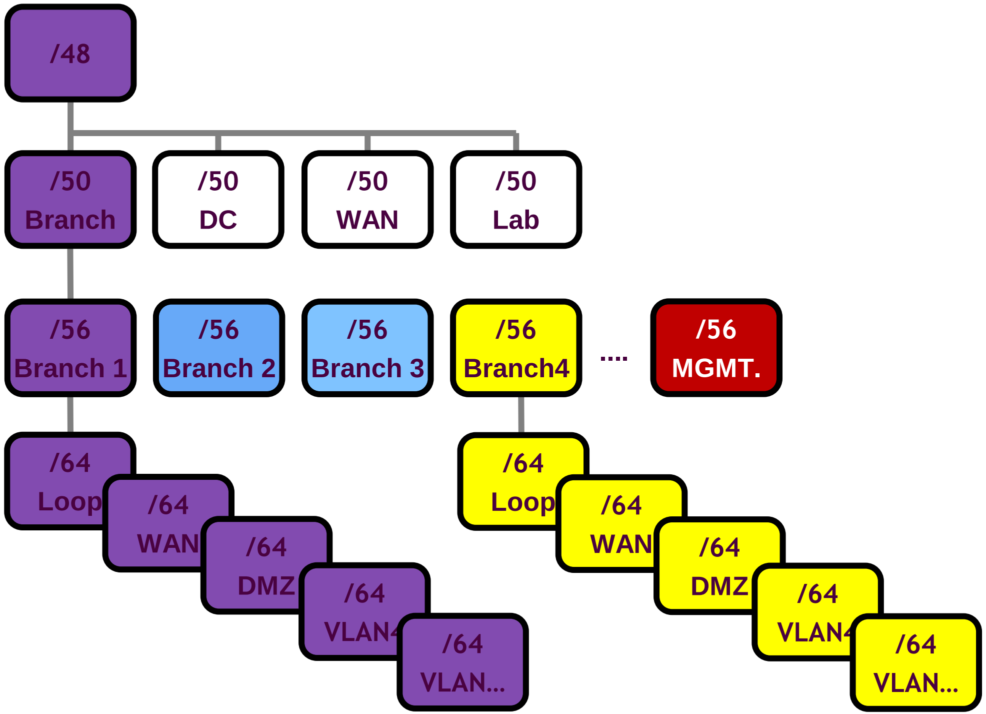

- High Level addressing plan. Indicative only. Can be modified to suit needs
- /48 = 65536 x /64 prefixes
- Break up into functional blocks (4 x /50 in this case)
- Each functional block simplifies security policy
- Assumes up to 64 Branch networks
- Each Branch has access to 256 /64 prefixes for WAN, DMZ, & VLAN use

Address Plan Example

- Given 2001:db8:1000/40 by ARIN
  - Choose to stick with ARIN only assigned block
- Use a /44 block per region
  - Potential for 16 regional blocks
  - Follow regional registry breakdown
    - 2001:db8:1010::/44 for North America (reserve next block for expansion)
  - Use one regional block for Data Centres
- Break up North America into sites
  - Define site scopes
    - /52 for large sites (2048 subnets), /56 for mid-size sites (256 subnets), /60 (16 subnets)for small sites
    - Can assign contiguous blocks
  - Use /48's from the Data Centre block for NA data centres

Address Plan Example

- Template addressing
  - Build information into the address
  - Stay w/ /64 subnets for any segment that will have end systems attached
  - Example 2001:0db8:1010:1000::/52
    - Already know that this is a North American site
    - Site #'s map to physical site locations (Site 1 = Atlanta GA)

- Use the site bits to identify specific locations and/or functions w/in the site
    - 2001:0db8:1010:1xyz::/52
    - **X** = building(or floor); **Y** = organisation; **Z** = subnet function (e.g. servers, users, DMZ, wireless, voice, etc.)

- Short numbers: less chance of transcription errors for loopbacks
  - Compare: 2001:db8:1111:1:1:1:1/128 with 2001:db8:1234:1111::1/128

- Split address block into two example of a /32
  - /33 for internet Enabled devices /33 for Internal Restricted devices.
  - Helps with Route Identification and makes filtering on edge easier.

Host Address Assignment

	Manual	Stateless	DHCPv6
Pros	Address is stable Controlled assignment Well understood process	Scales well Time to deploy Widely implemented	Well understood process Controlled assignment Time to deploy
Cons	Does not scale Time to deploy	No control on assignment process Not well understood Lack of management Privacy concerns	Implementation in OS Must design for HA

- The choice of assignment depends on the existing processes and the adaptability of that process
- Remember that the methods are not mutually exclusive - all three can be used
- Regardless of choice must still control the stateless address assignment of addresses

What about NAT?

- A couple of versions of address translation related to IPv6
  - NAT-PT
    Original specification
    Deprecated
  - NPTv6
    Stateless translation method
    Only manipulate the prefix
  - NAT66
    Stateful translation
    Not specified in RFC
  - NAT64
    Translation between IPv6 and IPv4 address families
    Stateless and stateful methods available

- Where should NAT be applied?
  - NAT66
    Address hiding ???
    That’s the way we do IPv4???
    It provides security???
    Multi-homing
  - NAT64
    Boundaries between IPv4 only and IPv6
    Highly successful in getting quick IPv6 access
    Cannot be the final state
    Must move towards full IPv6 integration

Importance of IP Address Management Tools

- Spreadsheets do not scale and are not auditable
- Tools should allow customers to manage IP address space consistent with their management methods.   Having a single source helps.

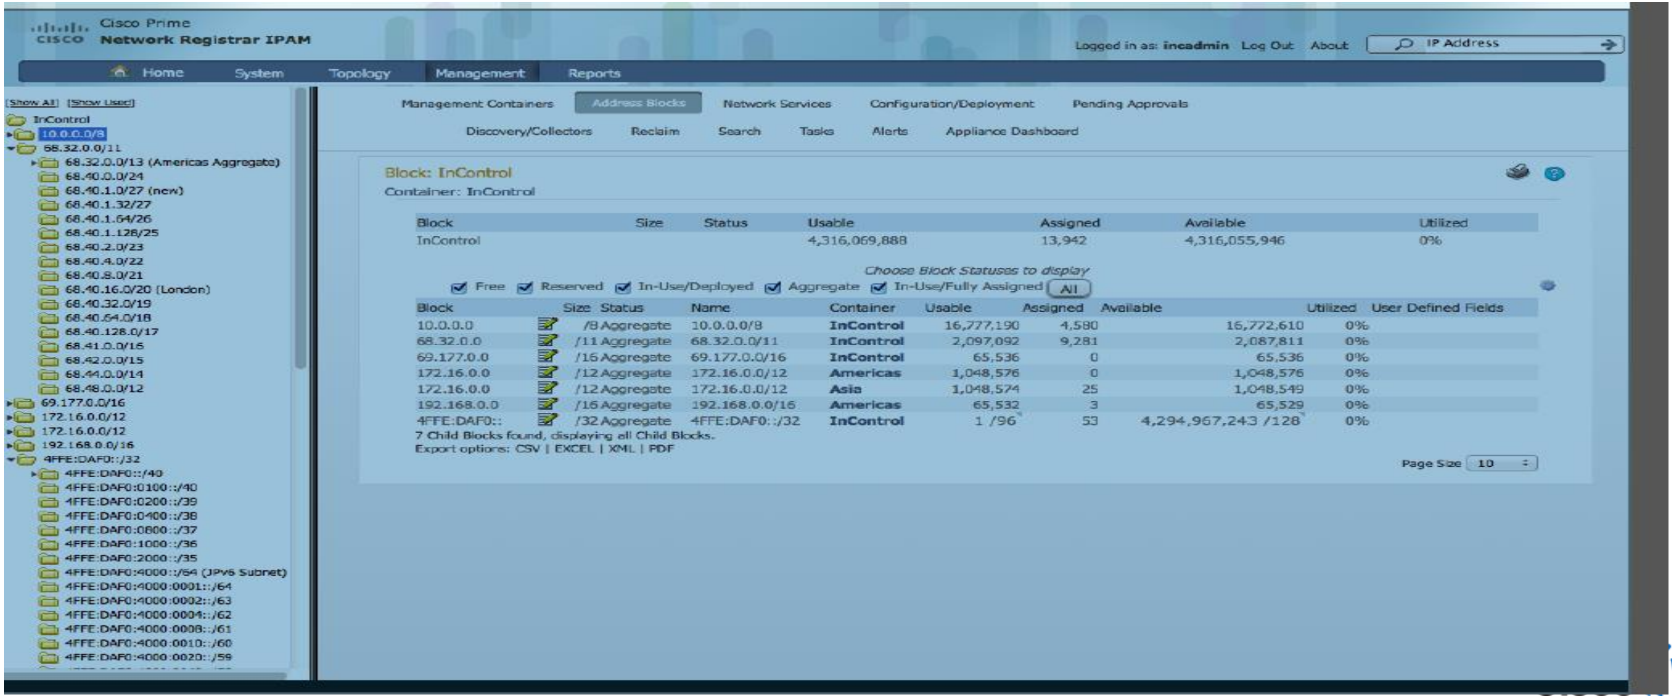




Transition Mechanisms

IPv6 Co-Existence Solutions

Dual Stack

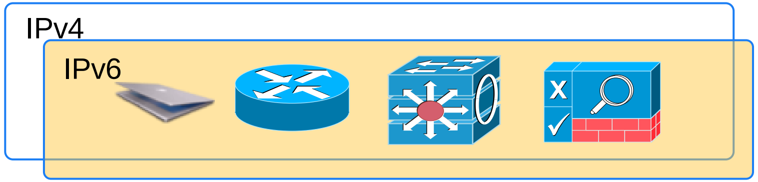

Recommended Enterprise Co-Existence Strategy

Tunnelling Services

Connect Islands of IPv6 or IPv4

Translation Services

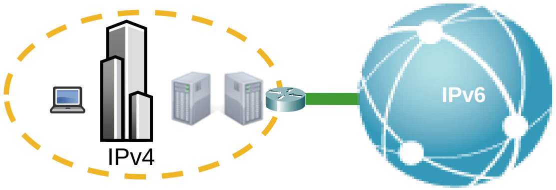

Connect to the IPv6 Community

Considerations

- IPv6 allows you to architect a new network frugally
  - In parallel with and over existing IPv4 infrastructure
  - Minimal capital outlay
  - Implement where it is needed
- Consider Routing co-existence
  - ISIS for IPv6
  - OSPF for IPv4
- Consider addressing
  - How will you allocate your IPv6 prefixes to customers
- Consider interoperability between vendors
- Consider billing systems
- Watch the standards and policies

Connecting IPv6 Sites Together

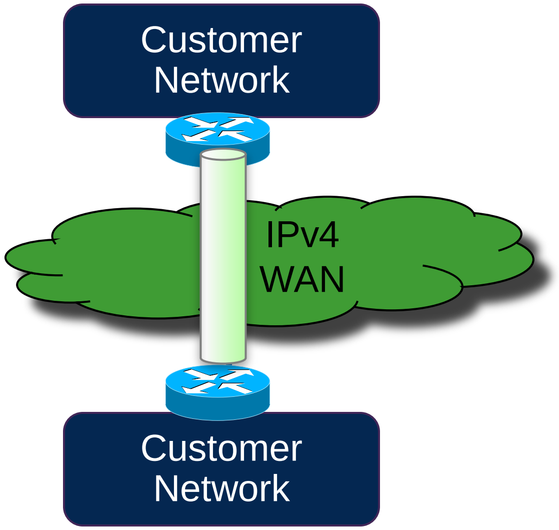

Using Tunnels

- Manually configured tunnels
- IPv6 over GRE
- LISP
- IPSec Tunnels
- Dynamic Multipoint VPN (DMVPN)

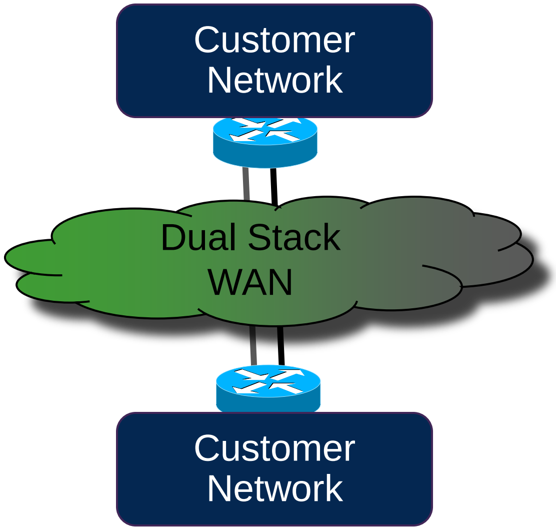

Dual Stack IPv4/IPv6

- Dual Stack CPEs
- Dual Stack Headquarters
- Dual Stack WAN

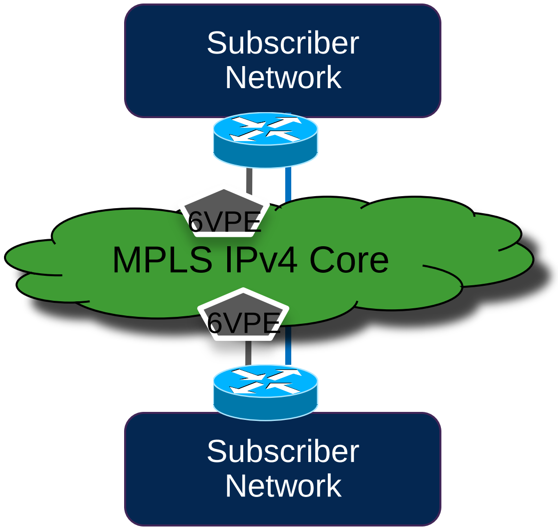

6VPE Service

- Dual Stack IPv4 / IPv6
- 6VPE VPN Service

SP IP Network Transition Options

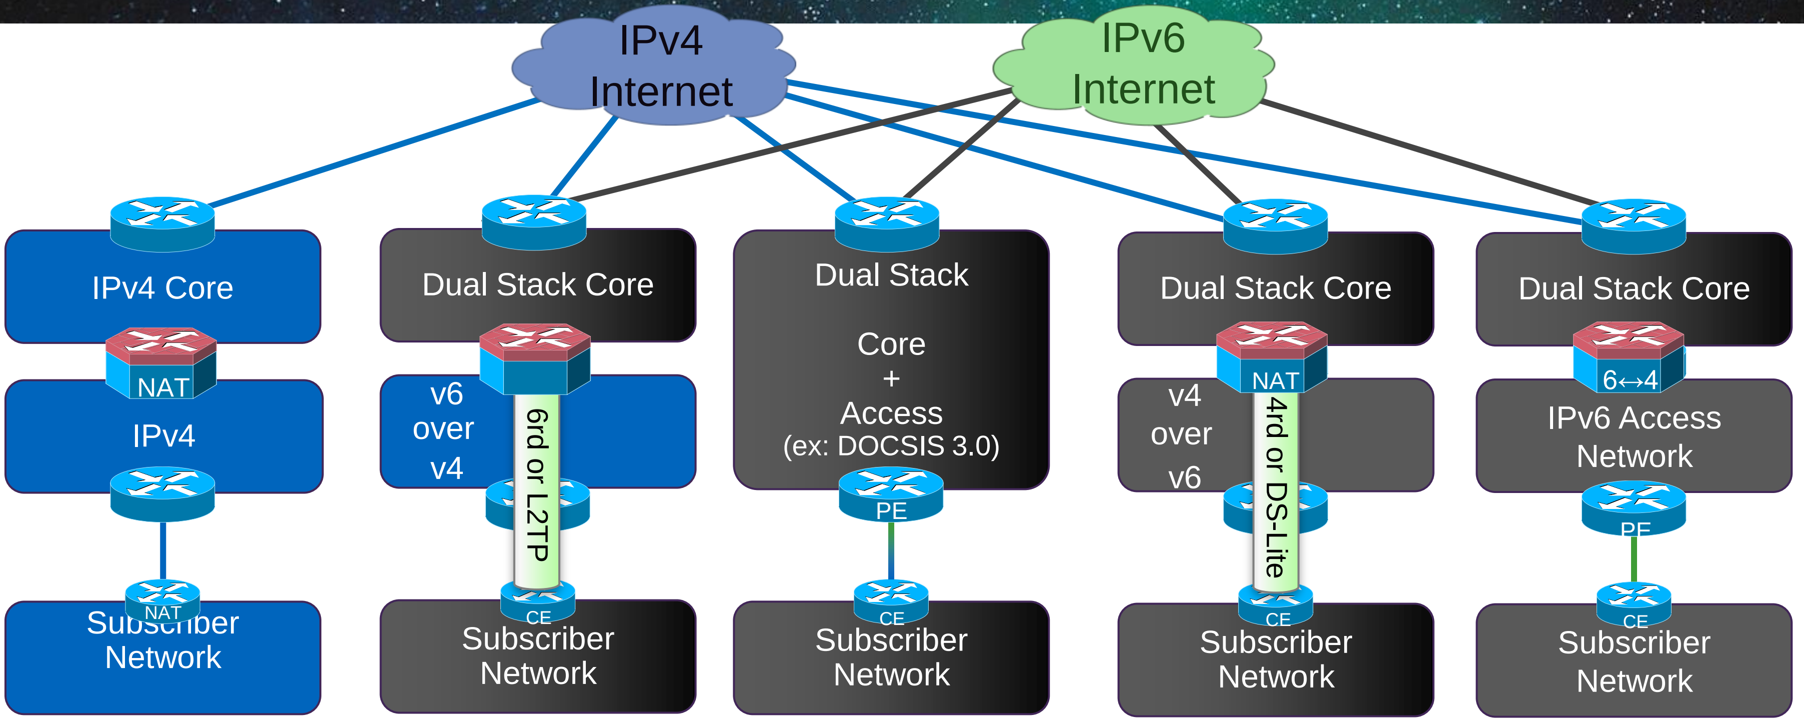

NAT444

6 Rapid Deployment (6rd
L2TP
Softwires

Broad Band Connectivity
Dual Stack Core
DOCSIS Access

IPv4 via IPv6
Using DS-Lite (w/NAT44)

AFT64

IPv6 Data Centre Network Architecture

Distribution/Core
- Dual Stack
- Routing protocols (OPSFv3, ISISv6, BGPv6..)
- IPv6 Mcast
- IPv6 security: classification, ACL & policing,CoPP
- BFD
- Flexible Netflow
- 6VPE
- ECMP
- Interface stats
- uRPF

L2/L3 Boundary

Towards Access
- Dual Stack
- HSRPv6/VRRPv3
- BFD
- SVI
- Snooping (MLDv2)
- IGMPv3
- First Hop Security (RA guard)
- PACL/VACL
- IPv6 Management

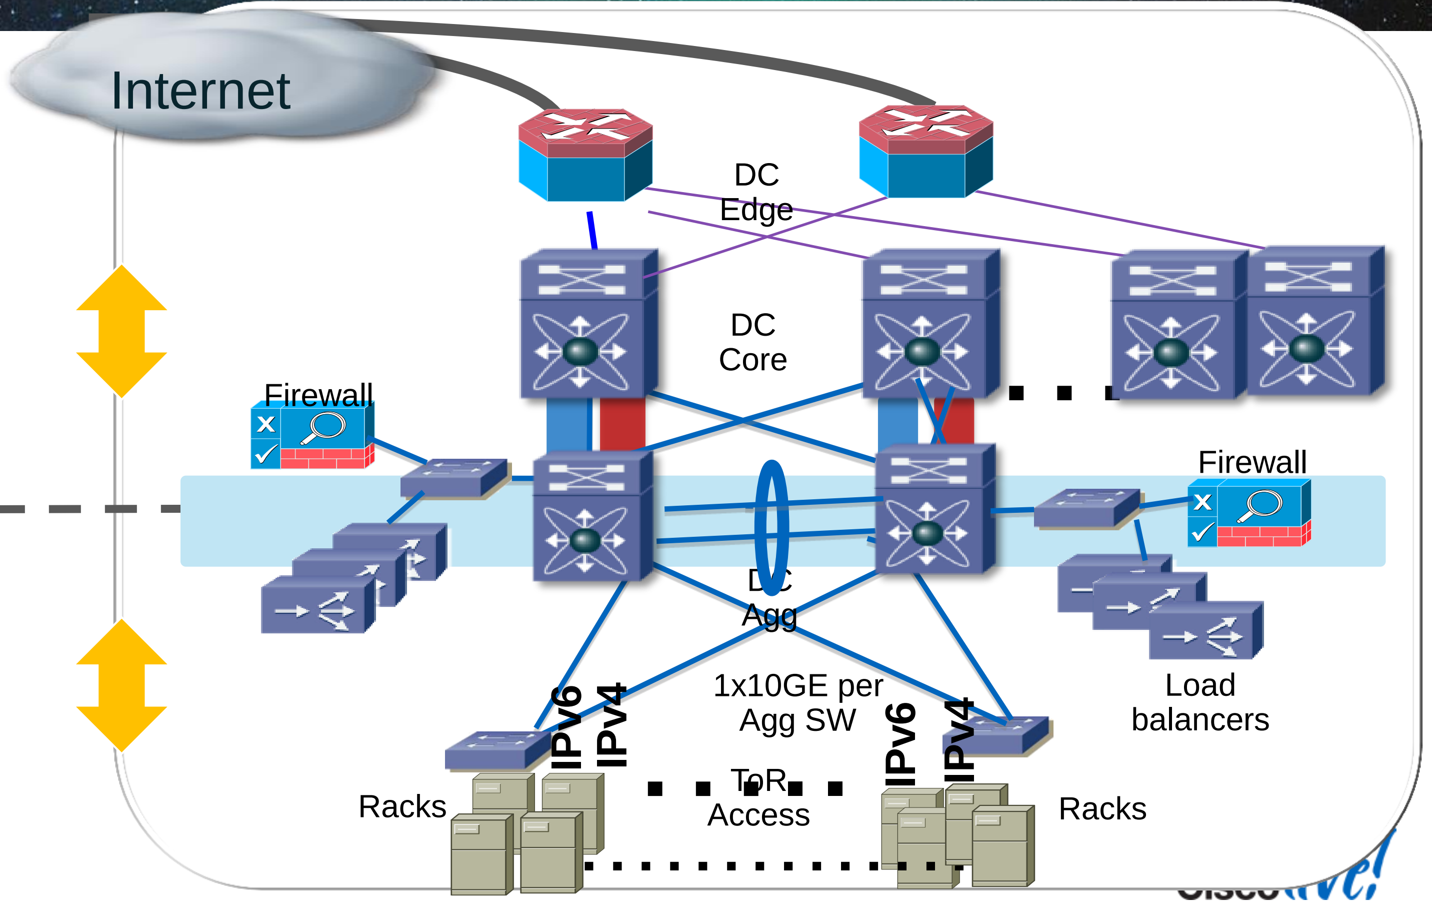

Common Deployment Models for Internet Edge

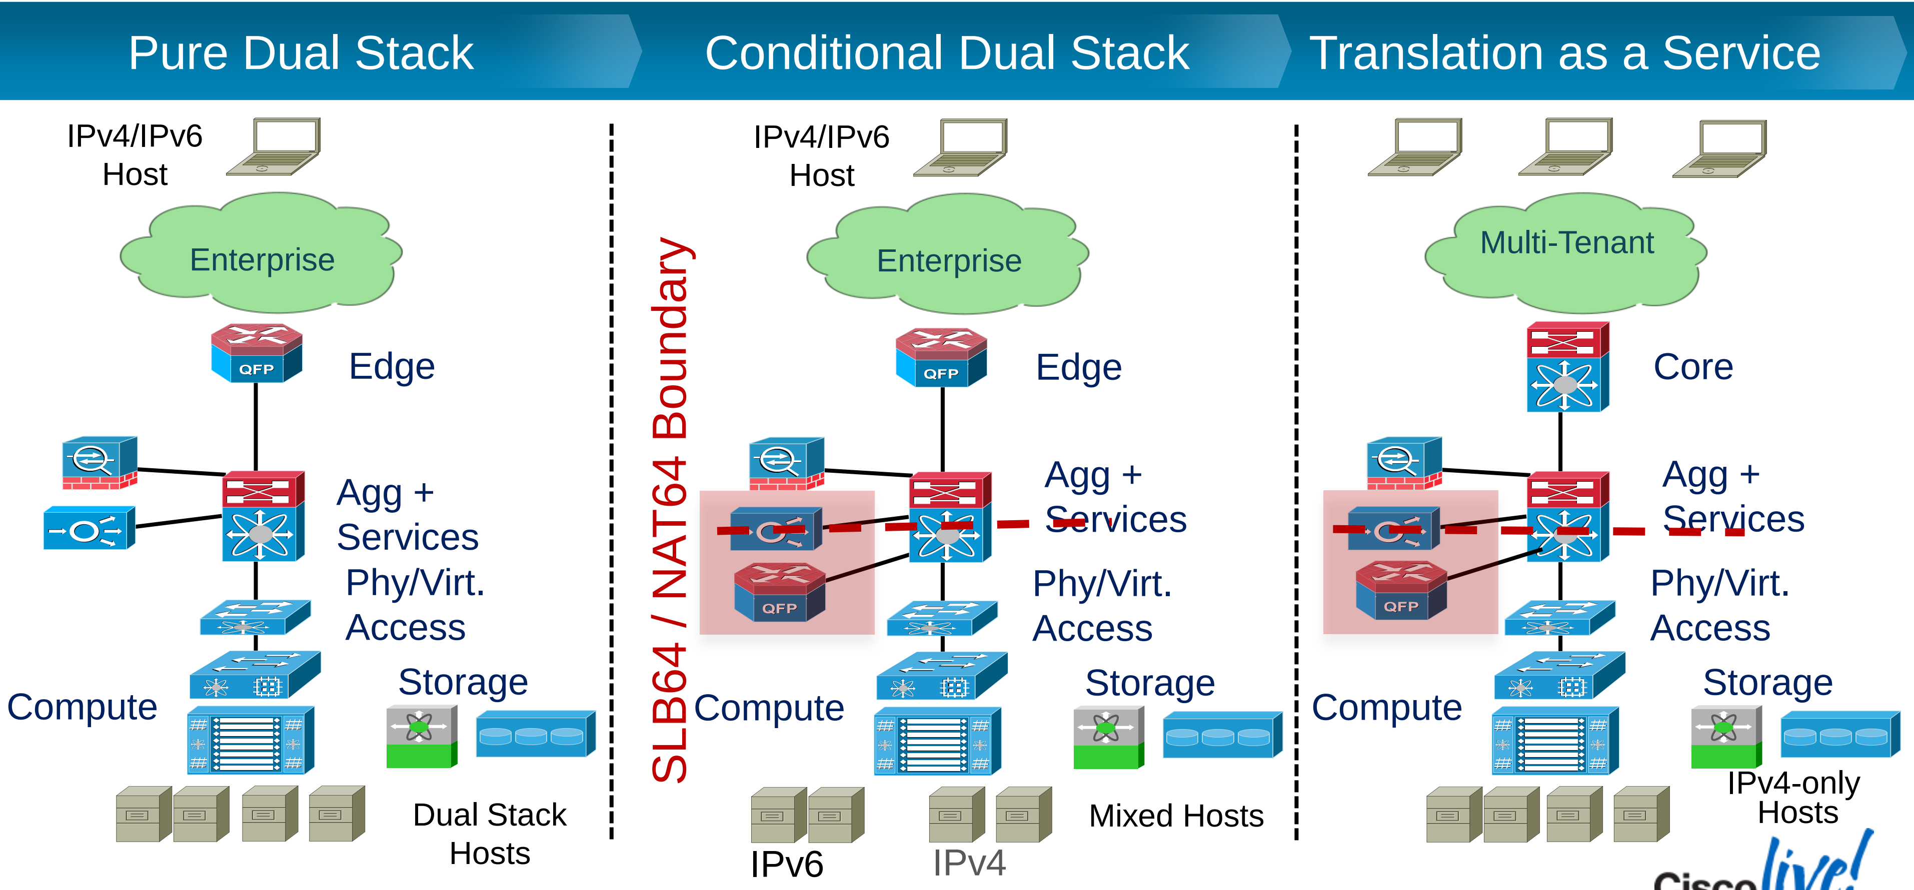

IPv6 Integration – NAT Overlap

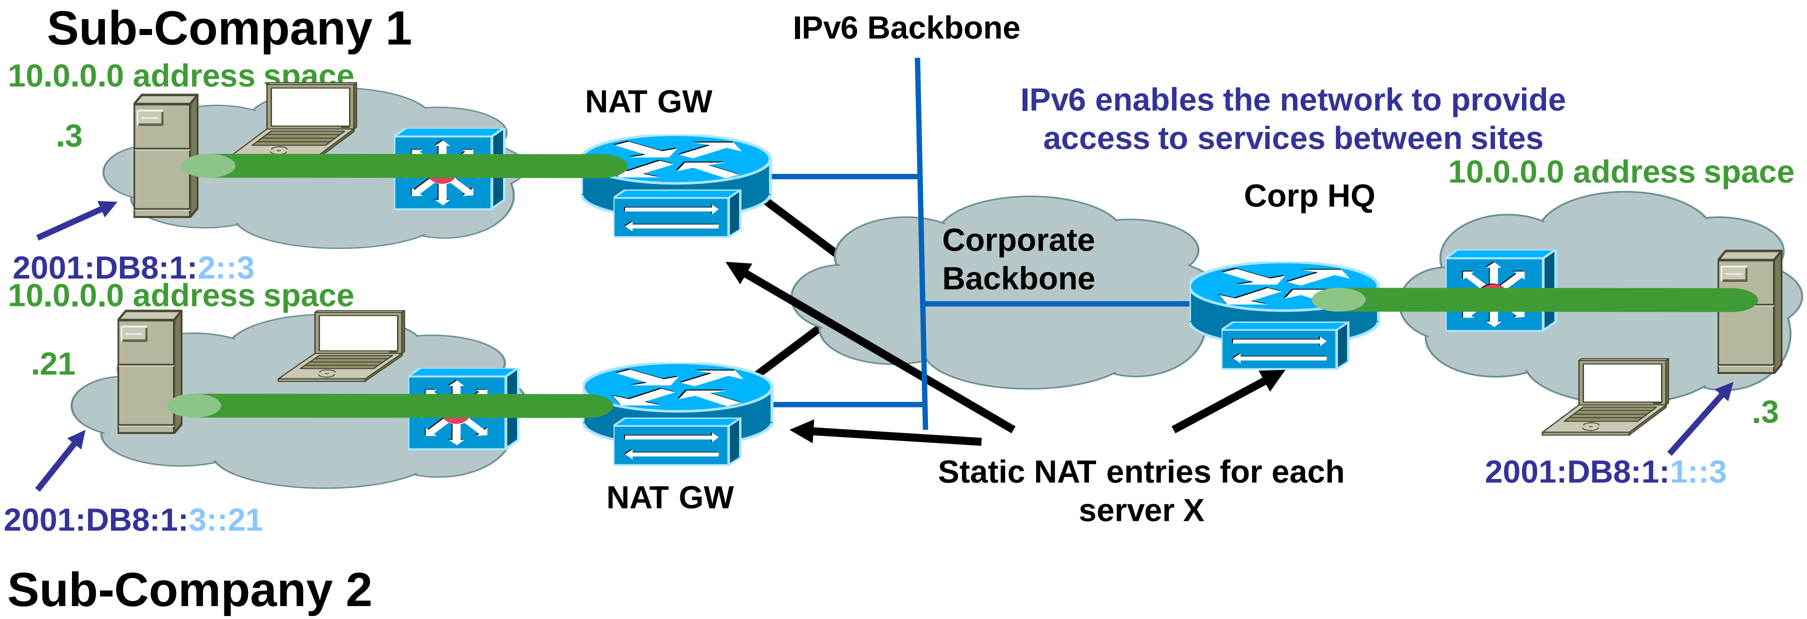

Customer requirement

- Speed up deployment of applications across the Enterprise

Business Challenges

- Merger and acquisition complexity
- Overlapping private address space

Solution

- IPv6 can be deployed to enable service access per site and/or per application

Interface Connectivity Types and Considerations

Physical or sub-interface (dot1.q)
Dual Stacked IPv6 and IPv4 on the same interface
IPv6 only on interface
IPv4 only interface
Peering Considerations

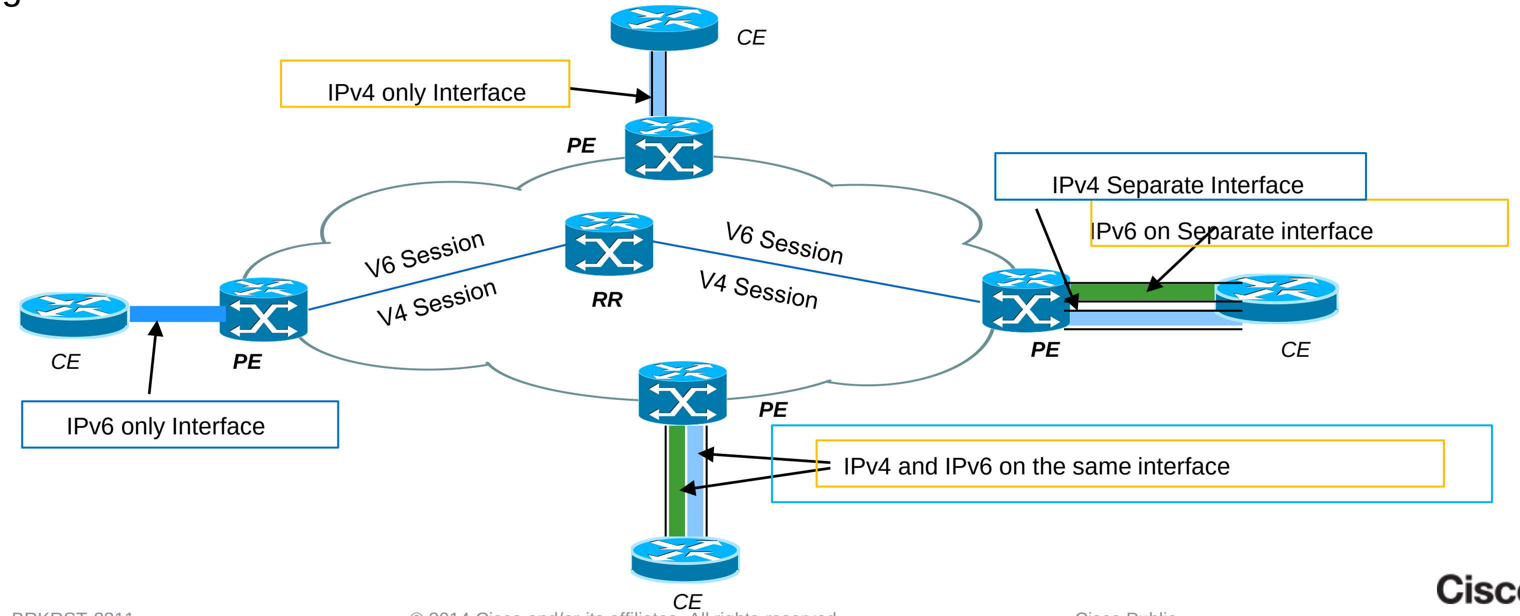

IPv6 PE to CE E-BGP Peering Options

- Separate BGP peering whenever possible.
  - Keep V4 and V6 Prefix exchange independent
- E-BGP over IPv6 native session to Link-local addressing
- If required both IPv4/IPV6 Address-families can be established over IPv4 peer.
- Requires in and outbound route-maps to manually set next hops depending on software/vendor/product implementations.

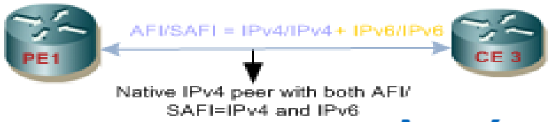

IPv6 and IPv4 Routing Protocols

RIP	RIPv2 for IPv4 RIPng for IPv6 Distinct but similar protocols with RIPng taking advantage of IPv6 specificities
OSPF	OSPFv2 for IPv4 OSPFv3 for IPv6 Distinct but similar protocols with OSPFv3 being a cleaner implementation that takes advantage of IPv6 specificities
IS-IS	Extended to support IPv6 Natural fit to some of the IPv6 foundational concepts Supports Single and Multi Topology operation
EIGRP	Extended to support IPv6 Some changes reflecting IPv6 characteristics
BGP	Extended to support IPv6 through multi-protocol extensions

IPv4 and IPv6 Co-existence

	Single Process / Single Topology	Single Process / Multi Topology	Multi Process / Multi Topology
Protocols	IS-IS ST	IS-IS MT	OSPFv2 + OSPFv3 EIGRP + EIGRPv6
IP topologies	Single (IPv4+IPv6) Congruent	Multiple Non-congruent	Multiple Non-congruent
Flooding + router/network resources	Common	Common	Multiple protocol instances on given link
SPF	Single	Multiple	Multiple (OSPF)
LS databases / topology tables	Single Large	Single Large	Multiple
Control plane **For Your Reference**	- Common - Less resource intensive - More deterministic IPv4/IPv6 co-existence	- More separation - Protocol-specific optimisation possible - More resource intensive	- Clear separation - Protocol-specific optimisation possible - More resource intensive

Co-Existence Considerations

Scalability and Performance

- IPv6 Neighbour Cache = ARP for IPv4
  - In dual-stack networks the first hop routers/switches will now have more memory consumption due to IPv6 neighbour 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 Neighbour 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
```

- There are some implications to managing the IPv6 neighbour cache when concentrating large numbers of end systems

Neighbour Unreachability Detection (NUD) Implications

- The neighbour cache maintains mapping information
  – Neighbour's reachability state is also maintained
- Neighbours can be in one of 5 possible states
  – INCOMPLETE – Address resolution is in progress and link-layer address is not yet known.
  – REACHABLE – Neighbour is known to be reachable within last reachable time interval.
  – STALE – Neighbour requires re-resolution, traffic may flow to neighbour.
  – DELAY – Neighbour pending re-resolution, traffic might flow to neighbour.
  – PROBE – Neighbour re-resolution in progress, traffic might flow to neighbour.
- Every entry that is marked STALE in the neighbour cache will need to have it's state verified
  – Traffic will be forwarded using the STALE entry
  – NUD will use NS/NA to detect reachability
- How often NUD is run depends on the value of AdvReachableTime that is set in RA messages
  – Cisco default is 30 seconds
- **Consider CPU load for maintaining state for thousands to tens of thousands of entries!**

Neighbour Unreachability Detection (NUD) Implications

- What to do?
- Don't Panic!
  – Unless you forgot your towel
- New features to manage the neighbour cache
  – Extend the reachable time advertised in RA's(max value is 1 hour)
  – Unsolicited NA glean (more to avoid traffic disruption)
  – ND cache timers (control how long an entry is maintained in STALE state; default is 4 hours)
  – ND cache refresh (run NUD before purging STALE neighbours)
  – NUD exponential retransmit (spread out the NS packets)

Scalability and Performance

- Full internet route table
  – Ensure to account for TCAM/memory requirements for both IPv4/IPv6
  – Not all platforms 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 tunnels on platforms that use HW switching when attempting large scale deployments (hundreds/thousands of tunnels)

58

Understanding Co-Existence Implications

- Resources considerations
  - Memory (storing the same amount of IPv6 routes requires less memory than might be expected)
  - CPU (insignificant increase in the case of HW platforms, additive in the case of SW platforms)
- Control plane considerations
  - Balance between IPv4/IPv6 control plane separation and scalability of the number of sessions
- Performance considerations
  - Forwarding in the presence of advanced features
  - Convergence of IPv4 routing protocols when IPv6 is enabled

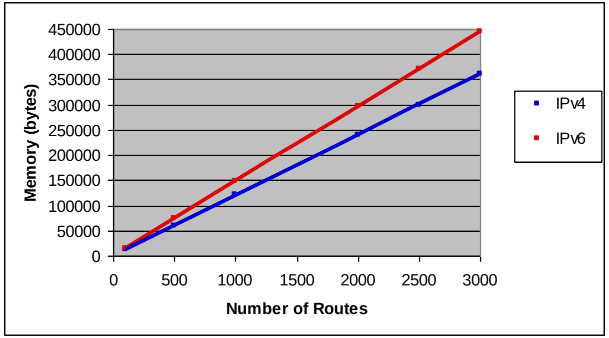

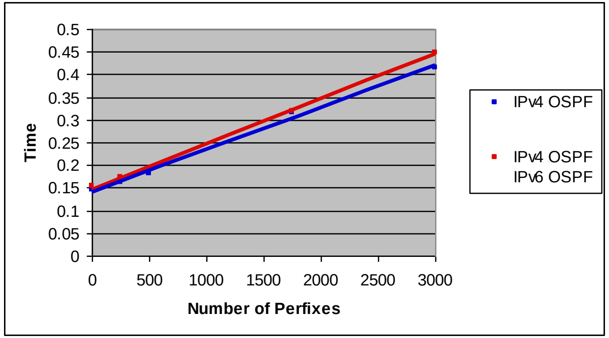

The Coexistence Twist

- IPv6 IGP impact on the IPv4 IGP convergence
- Aggressive timers on both IGPs will highlight competition for resources
- Is parity necessary from day 1?

Tuned IPv4 OSPF, Untuned IPv6 OSPF

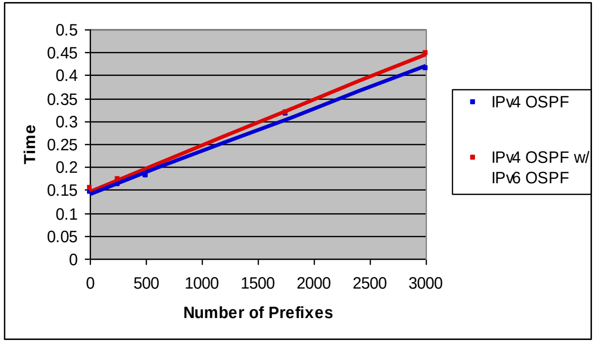

Tuned IPv4 OSPF, Tuned IPv6 OSPF

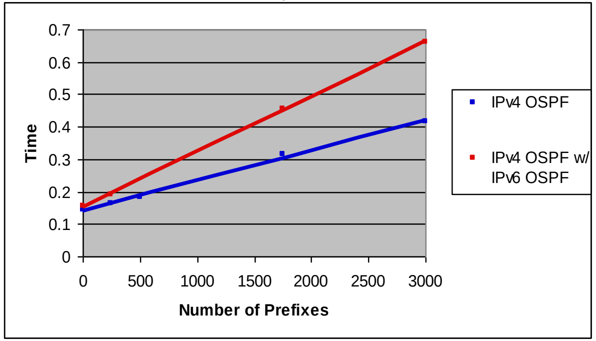

IPv6 QOS

- QOS Policy for IPv4 and IPv6 will be consistent (RFC 2460/3697)
- Ipv6 Traffic class field maps to the same dscp values as IPv6 and will be mapped to corresponding EXP values set on the network today.
- IPv6 classification will follow the same IP Precedence, Service Class, DSCP and EXP QOS Taxonomy values already defined for IPv4.
- Some devices will need additional configuration to match values on the IPv6 traffic class field.
- IPv6 will utilise the same Network Control, Voice, SIG, Gold, Bronze, Silver, Best Effort classes

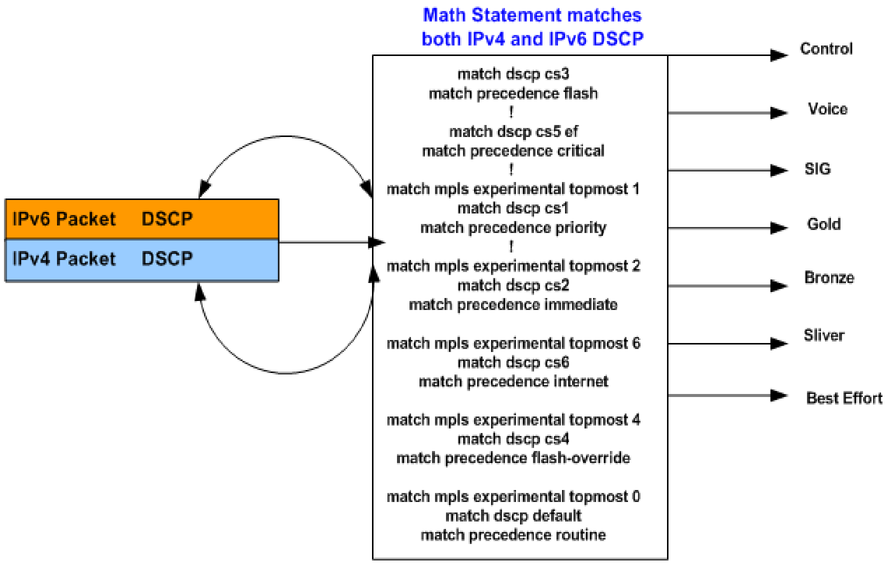

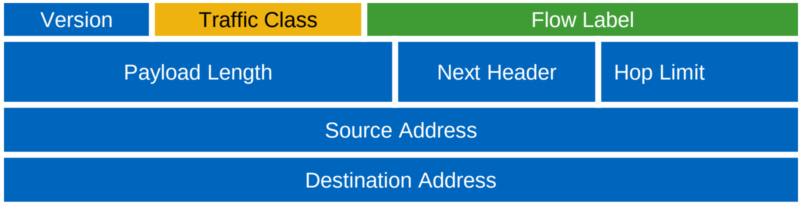




Management and Operations

Don’t Forget About Network Management

Introduction of IPv6 creates new network management challenges

- Management and design strategies for IPv6 addressing model, policies and operation
- Introduction of extended IP services: DHCPv6, DNSv6, IPAM
- Managing security infrastructures: Firewall, IDS, AAA
- Tool visibility, insight and analysis of IPv6 traffic Netflowv9, IPv6 SLA
- Troubleshooting
  - IPv4-IPv6 interaction
- Requires support in
  - Instrumentation (MIB , Netflow records, etc.)
  - NMS tools and systems
- Dual Stack Interfaces will result in tools i.e. MRTG reporting combined v4 and V6 traffic statistics.

NetFlow for IPv6

- Application Performance monitoring is a great differentiator for IPv6
- IPv6 support added as part of Flexible NetFlow (metering) and NetFlow v9 (exporting) Monitors the IPv6 traffic.
- Export is over an IPv4 Transport
- Exporting: NetFlow version 9
  - –Advantages: extensibility
    - Integrate new technologies/data types quicker (MPLS, IPv6, BGP next hop, etc.)
    - Integrate new aggregations quicker
  - –Note: for now, the template definitions are fixed
- Metering: Flexible NetFlow
  - –Advantages: cache and export content flexibility
    - User selection of flow keys
    - User definition of the records

IPv6 Traffic Visibility

IPv6 MIBs and host support

IPv6 Traffic Metering with Flexible Netflow (export over IPv4)

Response measurement with IP SLA
UDP-Jitter, UDP-Echo, ICMP Echo, TCP Connect

Tunnel detection with NBAR2

Tunnel Filtering with ASA

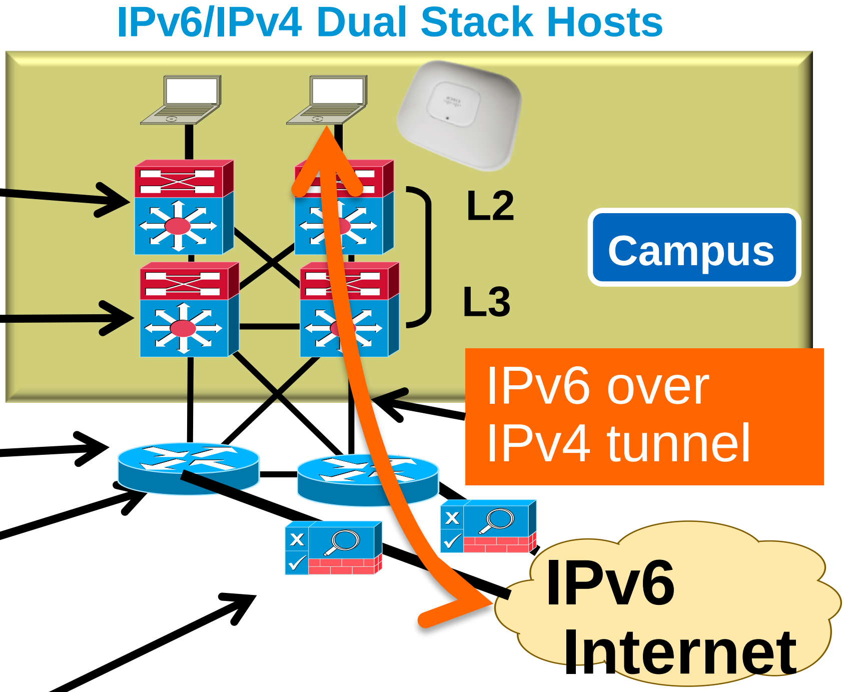

IPV6 Testing Considerations

- Create base line template that should be run as part of all IPv6 solution test.
- Template should consist of basic IPv6 RFC 2460 functionality.
- PMTU Testing is very important
- How do hosts re-act to auto-configuration?
- Are devices taking both a static and auto-configuration (Understand so that security Policy is not affected)?
- Should IPv6 RA’s be disabled how do devices re-act to that?
- Does application being used implement SAS (Source address selection) algorithm correctly?
- How do devices re-act with A and AAAA DNS records?

IPv6 Tools

- Different ways to check on what is happening
- Where's my prefix?
  - Route servers and looking glasses - http://www.bgp4.as/looking-glasses
- What's happening with traffic and adoption rates?
  - Cisco - http://6lab.cisco.com/stats/
  - Internet Society - http://www.worldipv6launch.org/measurements/
  - Google - http://www.google.com/ipv6/statistics.html
- Who's out there?
  - DNS
  - Registry whois database




IPv6 DNS

Introduction to DNS and IPv6

- Introduction of IPv6, will require use both IPv4 &IPv6 addresses in your network
- Need to add mappings from names to IPv6 addresses in parallel with the existing mapping from names to IPv4 addresses
- One example of such a mapping, using the AAAA resource record type, is shown here:
  – www.ipv6.cisco.com. 86400 IN AAAA 2001:420:80:1::5
- Mapping from a name to an IPv6 address is performed using an AAAA resource record, with the IPv6 address given as a hexadecimal address (RFC 3596)

IPv6 and DNS

	IPv4	IPv6
Hostname to IP Address	**A Record:** www.abc.test. A 192.168.30.1	**AAAA Record:** www.abc.test AAAA 2001:db8:C18:1::2
IP Address to Hostname	**PTR Record:** 1.30.168.192.in-addr.arpa. PTR www.abc.test.	**PTR Record:** 2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.1.c.0.8.b.d.0.1.0.0.2.ip6.arpa PTR www.abc.test.

Enabling DNS

- Add AAAA records in your DNS server for the hostnames of the devices that can be reached through the IPv6 protocol.

- Add pointer (PTR) records in your DNS server for the IP addresses of the devices that can be reached through the IPv6 protocol.

- Enable IPv6 access to the authoritative DNS servers.
  – Be sure that DNS servers can be accessed through IPv6.

- Enable IPv6 connectivity to the external full-service resolvers that send DNS queries to authoritative servers in the world.

AAAA Records on the Wire

1	0.000000	144.254.8.239	144.254.10.12	DNS	Standard query A ipv6.google.com
2	0.030695	144.254.10.123	144.254.8.239	DNS	Standard query response CNAME ipv6.l.google.com
3	0.058595	144.254.8.239	144.254.10.12	DNS	Standard query AAAA ipv6.l.google.com
4	0.070745	144.254.10.123	144.254.8.239	DNS	Standard query response AAAA 2a00:1450:8003::68 AAAA
5	0.071204	144.254.8.239	144.254.10.12	DNS	Standard query MX ipv6.l.google.com
6	0.087707	144.254.10.123	144.254.8.239	DNS	Standard query response

```
    Authority RRs: 4
    Additional RRs: 4
  ▷ Queries
▽ Answers
  ▽ ipv6.l.google.com: type AAAA, class IN, addr 2a00:1450:8003::68
      Name: ipv6.l.google.com
      Type: AAAA (IPv6 address)
      Class: IN (0x0001)
      Time to live: 5 minutes
      Data length: 16
      Addr: 2a00:1450:8003::68
  ▷ ipv6.l.google.com: type AAAA, class IN, addr 2a00:1450:8003::67
```

DNS as an Integration Tool

- DNS controls how people will access the application or service
  - Who wants to remember 2001:420:1101:1::a?
- Control when the service is available
  - AAAA record in DNS means service is available
- Control who receives the AAAA record
  - Whitelist who gets the AAAA response
- Control how the service is accessed
  - Separate domain

    ipv6.cisco.com vs cisco.com

DNS as an Integration Tool

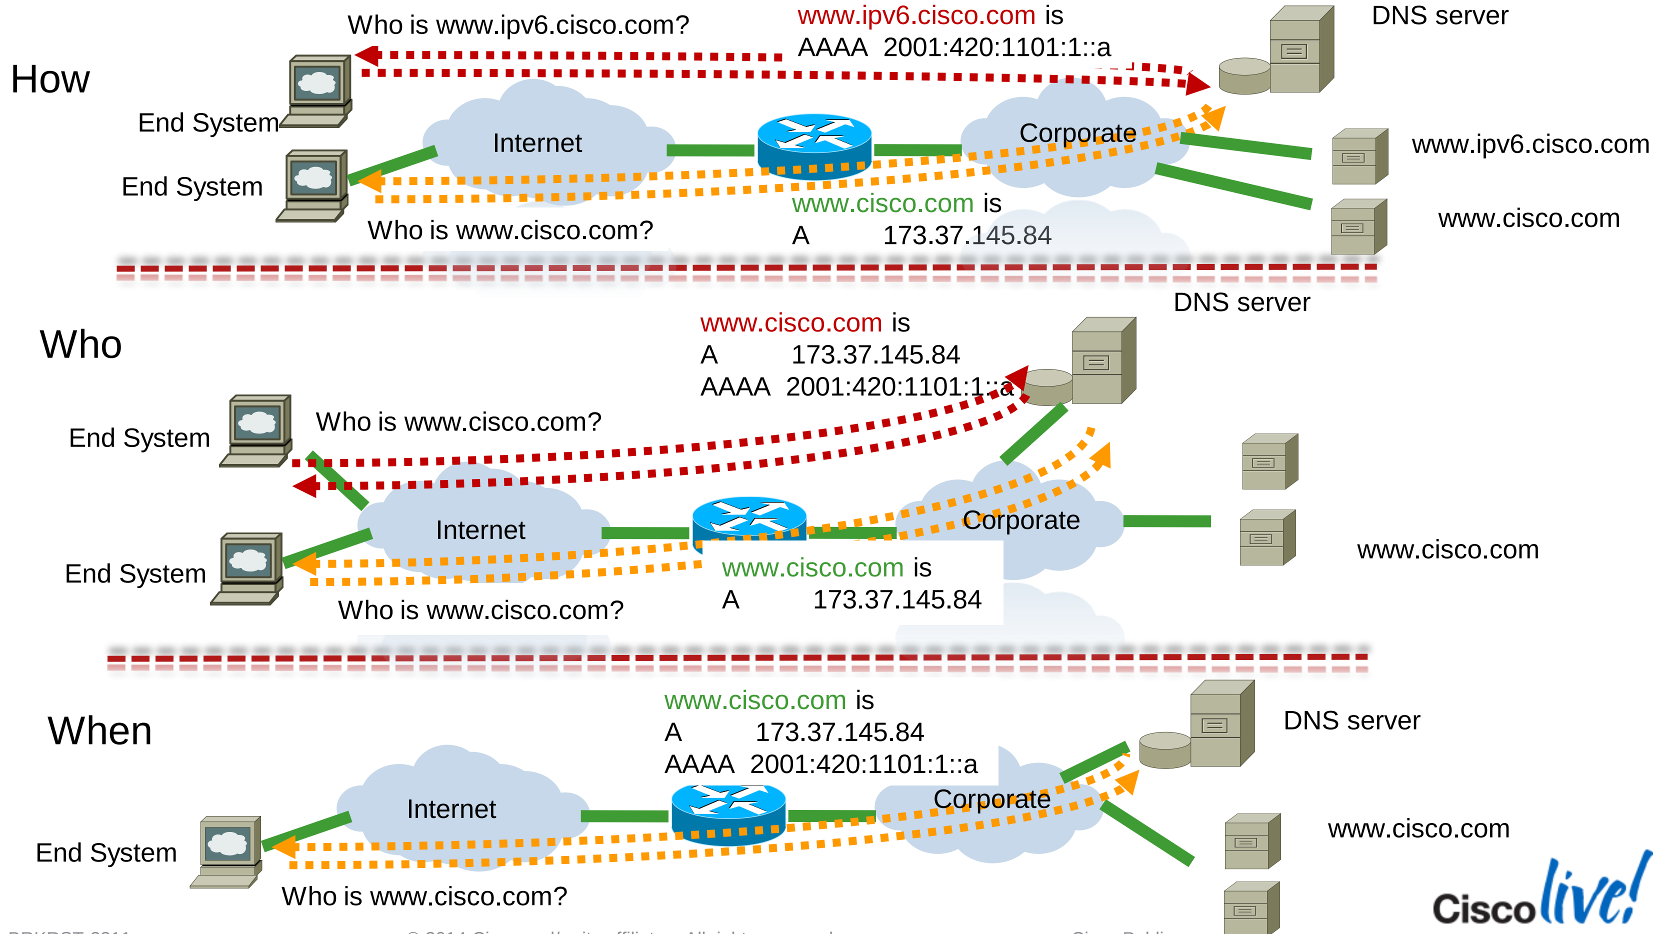




IPv6 Security

Security Considerations

Fate Sharing

- Dual Stack host with IPv6 Enabled IPv6 by Default be impacted by external services over tunneling.

Internet Edge Security

- Peering
- Ingress and Egress filters
- Bogon filters

Infrastructure Security

- Device Hardening
- Control Plane Policing
- DoS Detection
- ACLs
- EH processing

First Hop Security

- Secure Neighbor Discovery (SeND)
- Port ACLs
- Router Advertisement Guard
- NDP Inspection
- ND Cache limits

Dual Stack Host Considerations

- Host security on a dual-stack device
  - Applications can be subject to attack on both IPv6 and IPv4
  - **Fate sharing**: as secure as the least secure stack...
- Host security controls should block and inspect traffic from both stacks
  - Host intrusion prevention, personal firewalls, VPN clients, etc.

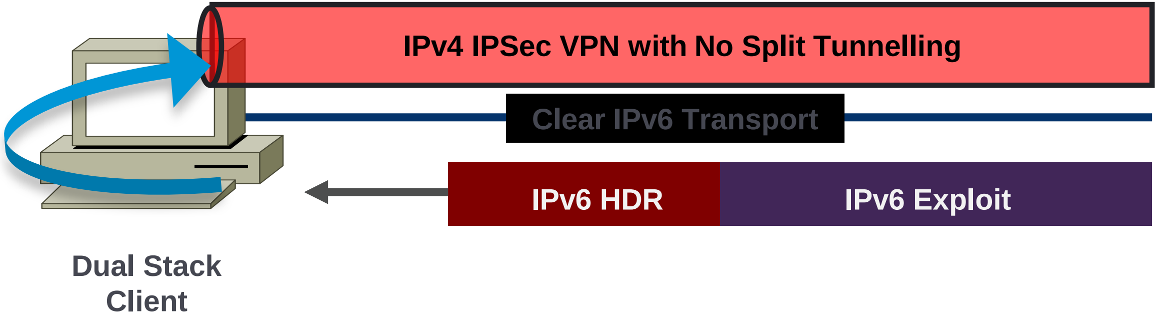

Does the IPSec Client Stop an Inbound IPv6 Exploit?

Infrastructure Security
Management Plane

- SSH, syslog, SNMP, NetFlow all work over IPv6
- Dual-stack management plane
    - More resilient: works even if one stack is down
    - More exposed: can be attacked over IPv4 and IPv6
- RADIUS over IPv6 is recent but IPv6 RADIUS attributes can be transported over IPv4
- As usual, infrastructure ACL is your friend as well as out-of-band management

```
ipv6 access-list VTY
  permit ipv6 2001:db8:0:1::/64 any

line vty 0 4
  ipv6 access-class VTY in
```

In IOS-XR: The command is
'`access-class VTY ingress`',
And
The IPv4 and IPv6 ACL must have the same name

IPv6 First Hop Security

IPv6 Device Tracking
Revoke network access for inactive devices

IPv6 PACL
Filter traffic on Layer 2 ports

IPv6 RA Guard
Stops false router advertisement threats

IPv6 NDP inspection
Prevents neighbour discovery spoofing attacks

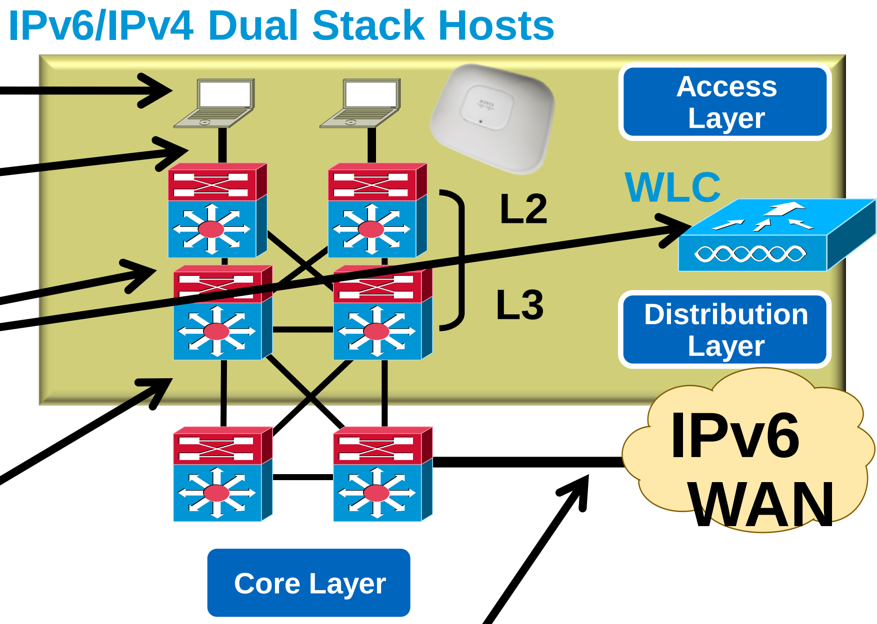

IPv6 uRPF
Blocks spoofed traffic in hardware

Control Plane Policing

- Control Plane Policing can be applied to IPv6
- Adapt what's in place today to accommodate IPv6
  - Routing protocols
  - Management protocols
- Remember the extended functionality of ICMP
- Monitor carefully to see what shows up in the logs
- Remember the default rules at the end of all IPv6 ACLs

  permit ipv6 any any nd-na

  permit ipv6 any any nd-ns

  deny ipv6 any any

  - They apply to any CoPP policy that uses ACLs to match

```
policy-map COPPr
 class ICMP6_CLASS
   police 8000
 class OSPF_CLASS
   police 200000
 class class-default
   police 8000
!
control-plane cef-exception
 service-policy input COPPr
```

Routing Protocol Authentication

Control Plane

- BGP, ISIS, EIGRP no change:
  - MD5 authentication of the routing update
- OSPFv3 has changed and pulled MD5 authentication from the protocol and instead rely on transport mode IPsec (for authentication and confidentiality)
- Or New Alternative is Authentication trailer for OSPFv3 (Refer to RFC 6506)
- IPv6 routing attack best practices
  - Use traditional authentication mechanisms on BGP and IS-IS
  - Use IPsec to secure protocols such as OSPFv3

```
interface Ethernet0/0
 ipv6 ospf 1 area 0
 ipv6 ospf authentication ipsec spi 500 md5
    1234567890ABCDEF1234567890ABCDEF
```

```
interface Ethernet0/0
 ipv6 authentication mode eigrp 100 md5
 ipv6 authentication key-chain eigrp 100 MYCHAIN

key chain MYCHAIN
 key 1
key-string 1234567890ABCDEF1234567890ABCDEF
accept-lifetime local 12:00:00 Dec 31 2006 12:00:00 Jan
    1 2008
send-lifetime local 00:00:00 Jan 1 2007 23:59:59 Dec 31
    2007
```

No crypto maps, no ISAKMP: transport mode with static session keys

Infrastructure Security
Data Plane

- Same as in IPv4, on real P2P without NDP, if not for me, then send it on the other side... Could produce looping traffic
- Classic IOS and IOS-XE platforms implement RFC 4443 so this is not a threat
  - on 76xx see CSCtg00387 (tunnels)
  - IOS-XR see CSCsu62728

Solution:

1. Use /127 on P2P link (see also RFC 6164) Or
2. Use infrastructure ACL
3. Use only link-local addresses

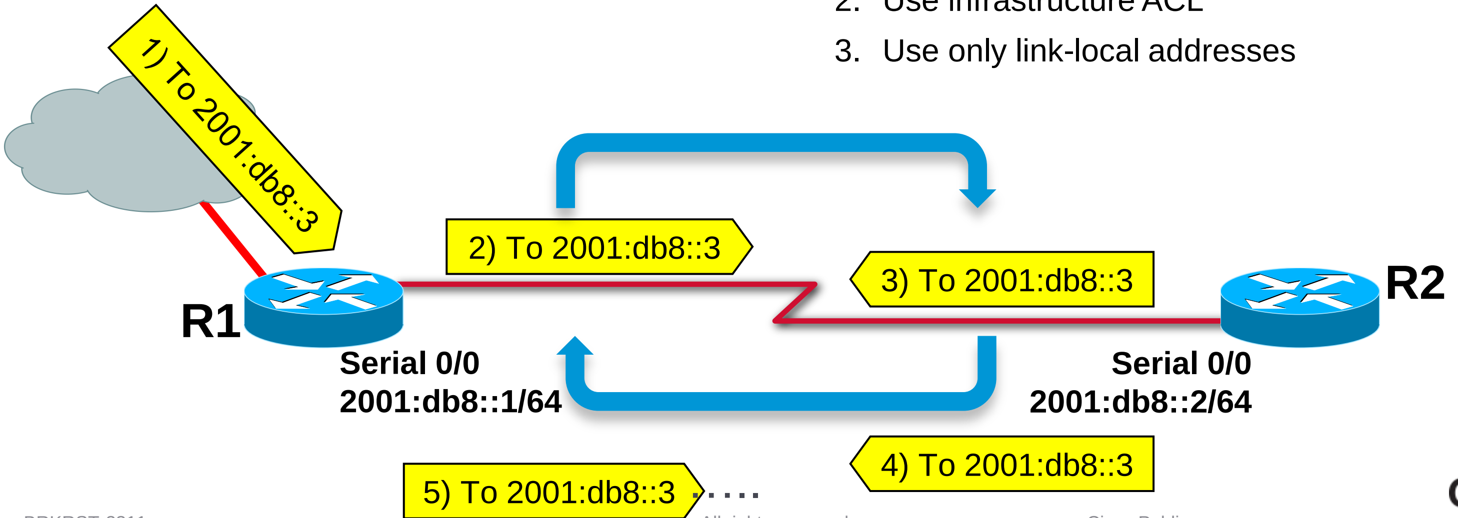

Perimeter Security: Anti-Spoofing and Bogon Filters

- Similar to IPv4, IPv6 has Bogons
- Anti-spoofing in IPv6 same as IPv4
  - => Same technique for single-homed edge= uRPF

```
ipv6 access-list NO_BOGONS
    remark Always permit ICMP unreachable (PMTUD)
    permit icmp any any unreachable
    remark Permit only large prefix blocks from IANA
    permit ip 2001::/16 any
    permit ip 2002::/16 any
    permit ip 2003::/18 any
    permit ip 2400::/12 any
    permit ip 2600::/10 any
    permit ip 2800::/12 any
    permit ip 2a00::/12 any
    permit ip 2c00::/12 any
    Remark implicit deny at the end
```

For Full list of Bogons:

http://www.team-cymru.org/Services/Bogons/fullbogons-ipv6.txt

Inter-Networking Device with uRPF Enabled

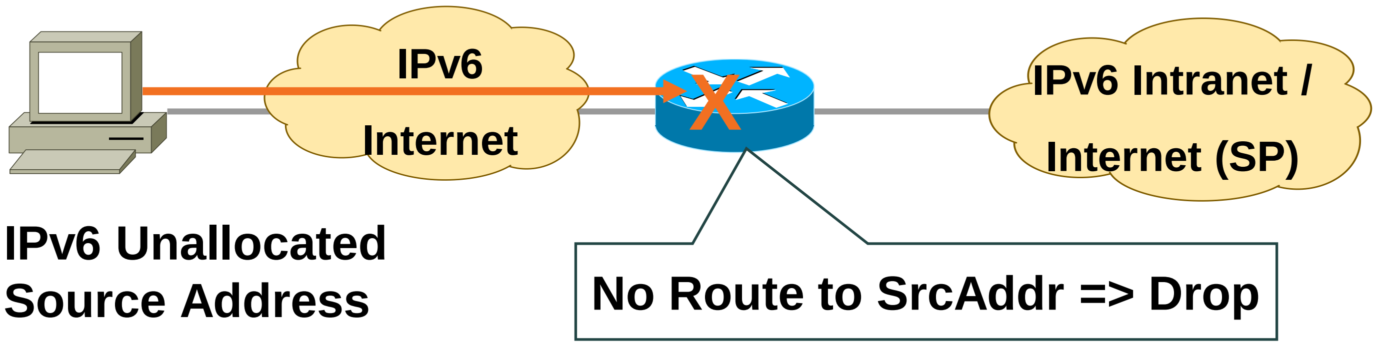

Remote Triggered Black Hole (RTBH)

- RFC 5635 RTBH is easy in IPv6 as in IPv4
- uRPF is also your friend for blackholing a source
- 100::/64
  - RFC 6666 has a specific discard ONLY prefix announced by IANA (100::/64)
  - added the prefix to the "IANA IPv6 Special Purpose Address Registry”
- Consult the following RTBH CCO Resource:
  - http://www.cisco.com/web/about/security/intelligence/ipv6_rtbh.html

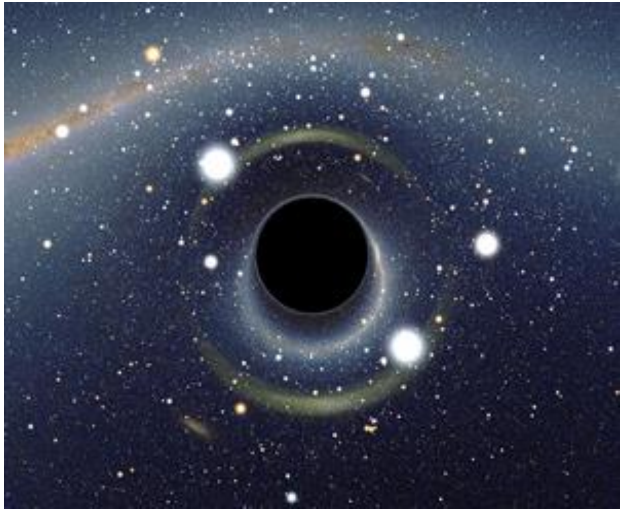

Source: Wikipedia Commons

Conclusion

- Start now and position for growth
- Next Steps:
  - Assess, Plan, Design Trial, Train, Roll out
- Map out opportunities to be IPv6 ready in planned technology refresh cycles
  - Reference IPv6 Ready Logo, USGv6 and RIPE-501
- Adapt IPv4 best practices for IPv6
- IPv6 is not identical to IPv4 so a review of the current architectures is necessary to understand the possible impact of integrating IPv6
- Education is key!

http://www.cisco.com/go/ipv6

Recommended Reading

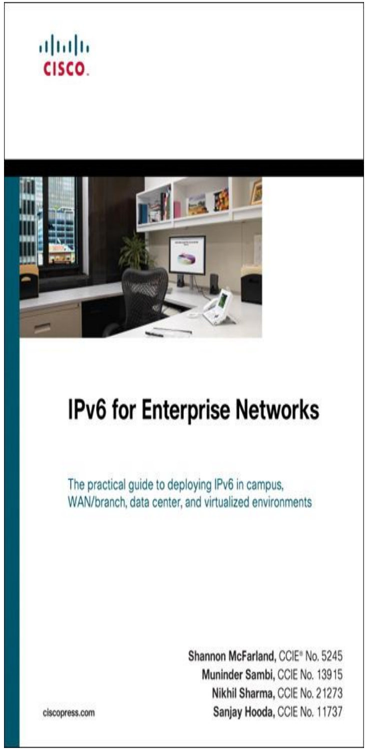

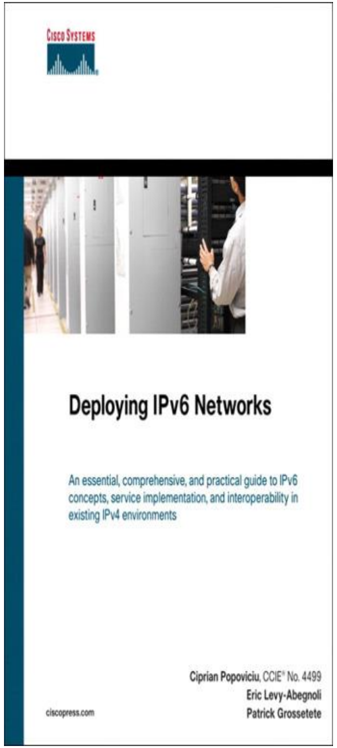

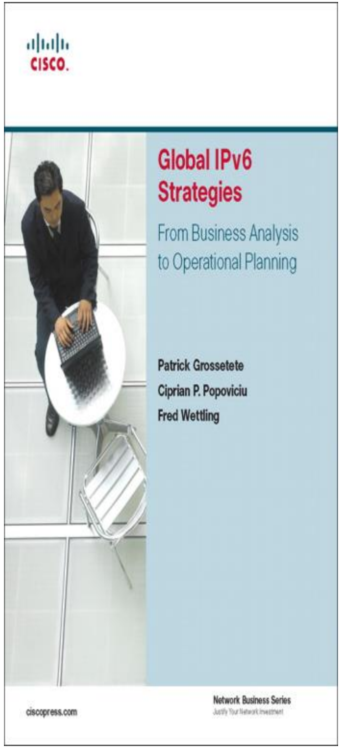

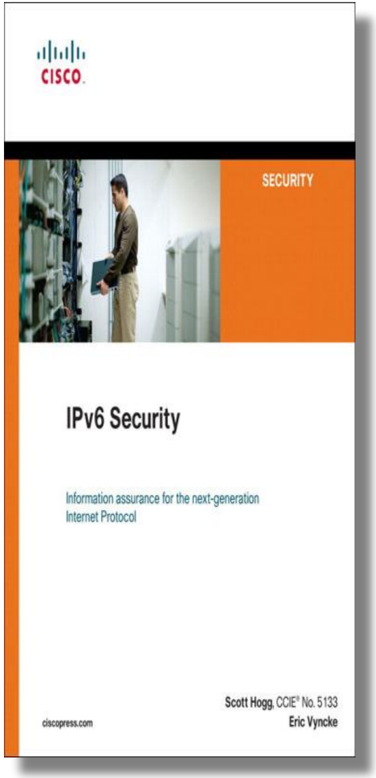




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations. www.CiscoLiveAPAC.com

CISCO TM