Cisco live!

What You Make Possible

TOMORROW starts here.

CISCO

IPv6 Planning, Deployment and Troubleshooting

BRKRST-2311

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

The Growing Internet Challenge…

The gap between supply and demand for IP addresses – the key Internet resource – is **widening**

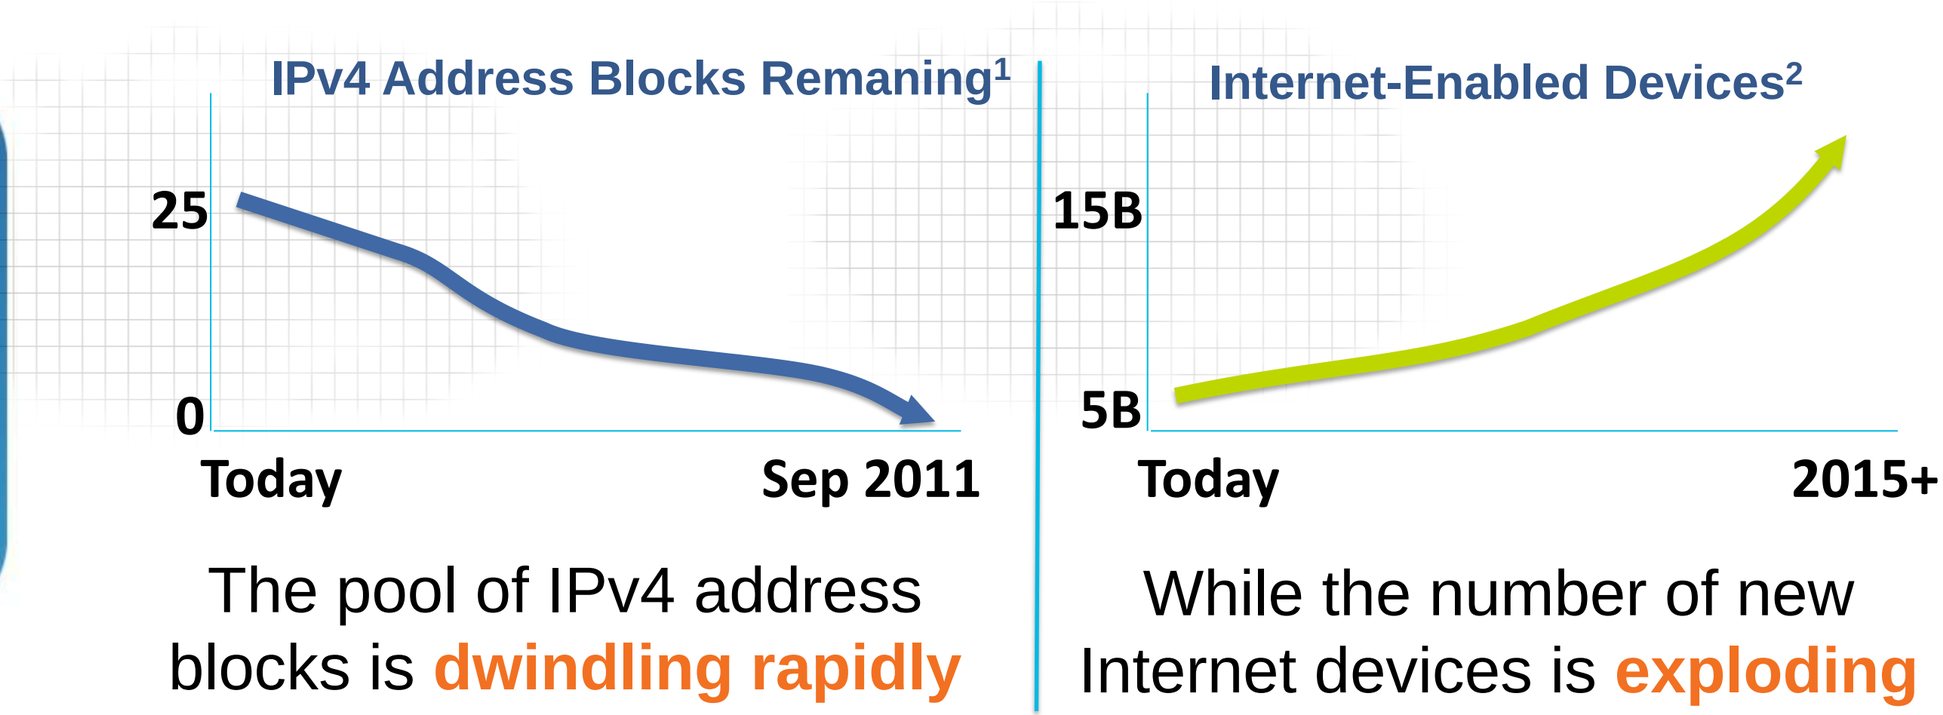

The pool of IPv4 address blocks is **dwindling rapidly**

While the number of new Internet devices is **exploding**

1 – Geoff Huston, APNIC, www.potaroo.net, tracking /8 address-blocks managed by the Internet Assigned Numbers Authority
2 – Cisco Visual Networking Index / Intel Embedded Internet Projections

IANA/RIR IPv4 Exhaustion

IANA Unallocated Address Pool Exhaustion
03-Feb-2011

Estimated Registry Exhaustion Dates

Exhaustion Date /8's	
APNIC	19-Apr-2011 9 (Actual)
RIPE	14-Sep-2012 9(Actual)
ARIN	13-Sep-2013
LACNIC	11-Jun-2015
AFRNIC	30-Aug-2019

On 14 September 2012, the RIPE NCC began to allocate IPv4 address space from the last /8 of IPv4 address space it holds. Source:RIPE

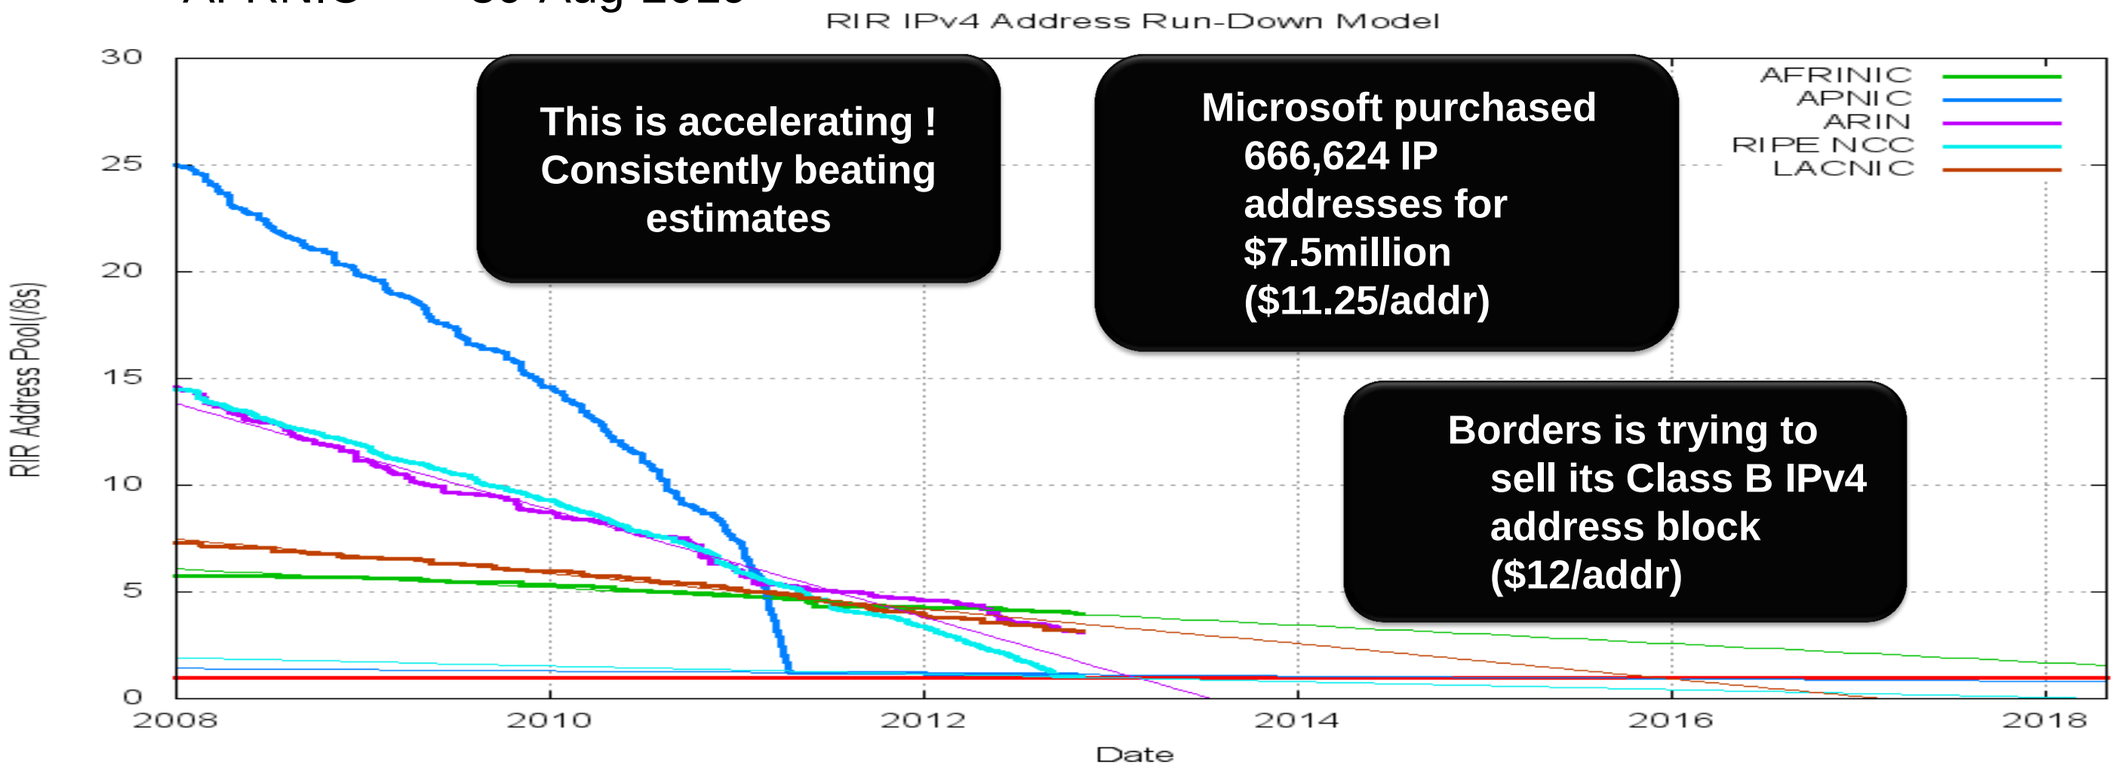

Source: Geoff Huston, APNIC

Market Factors Driving IPv6 Deployment

IPv4 Address Run-Out

2011

National IPv6 Strategies

US Federal/Civilian, US DoD, China NGI, EU

OECD

IPv6

IPv6 OS, Content & Applications

Google

facebook

Microsoft

Infrastructure Evolution

SmartGrid, SmartCities
DOCSIS 3.0, 4G/LTE, IPSO

comcast

verizon

at&t

CISCO

IPv6 Delivers Growth Across Verticals

Utilities

- Advanced Metering
- Smart Grids
- Home Energy Management
- Plug-In Electric Vehicles

Financial

- Web Banking
- E-Commerce
- Fraud Prevention
- Risk management

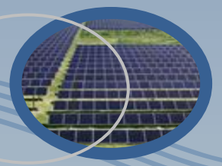

Government / Public Sector

- US Federal Mandate
- European Commission
- China and Japan Next Generation Internet (CNGI) project
- UA, EU, India, Singapore, Australia

Health Care

- Home care
- Wireless asset tracking
- Imaging
- Mobility

Consumer

- Mobility 4G/LTE
- Next generation set-top boxes
- Internet gaming
- IP-enabled appliances

Failure to Act Will Impact Business

IPv6 Estimated Adoption Timeframes

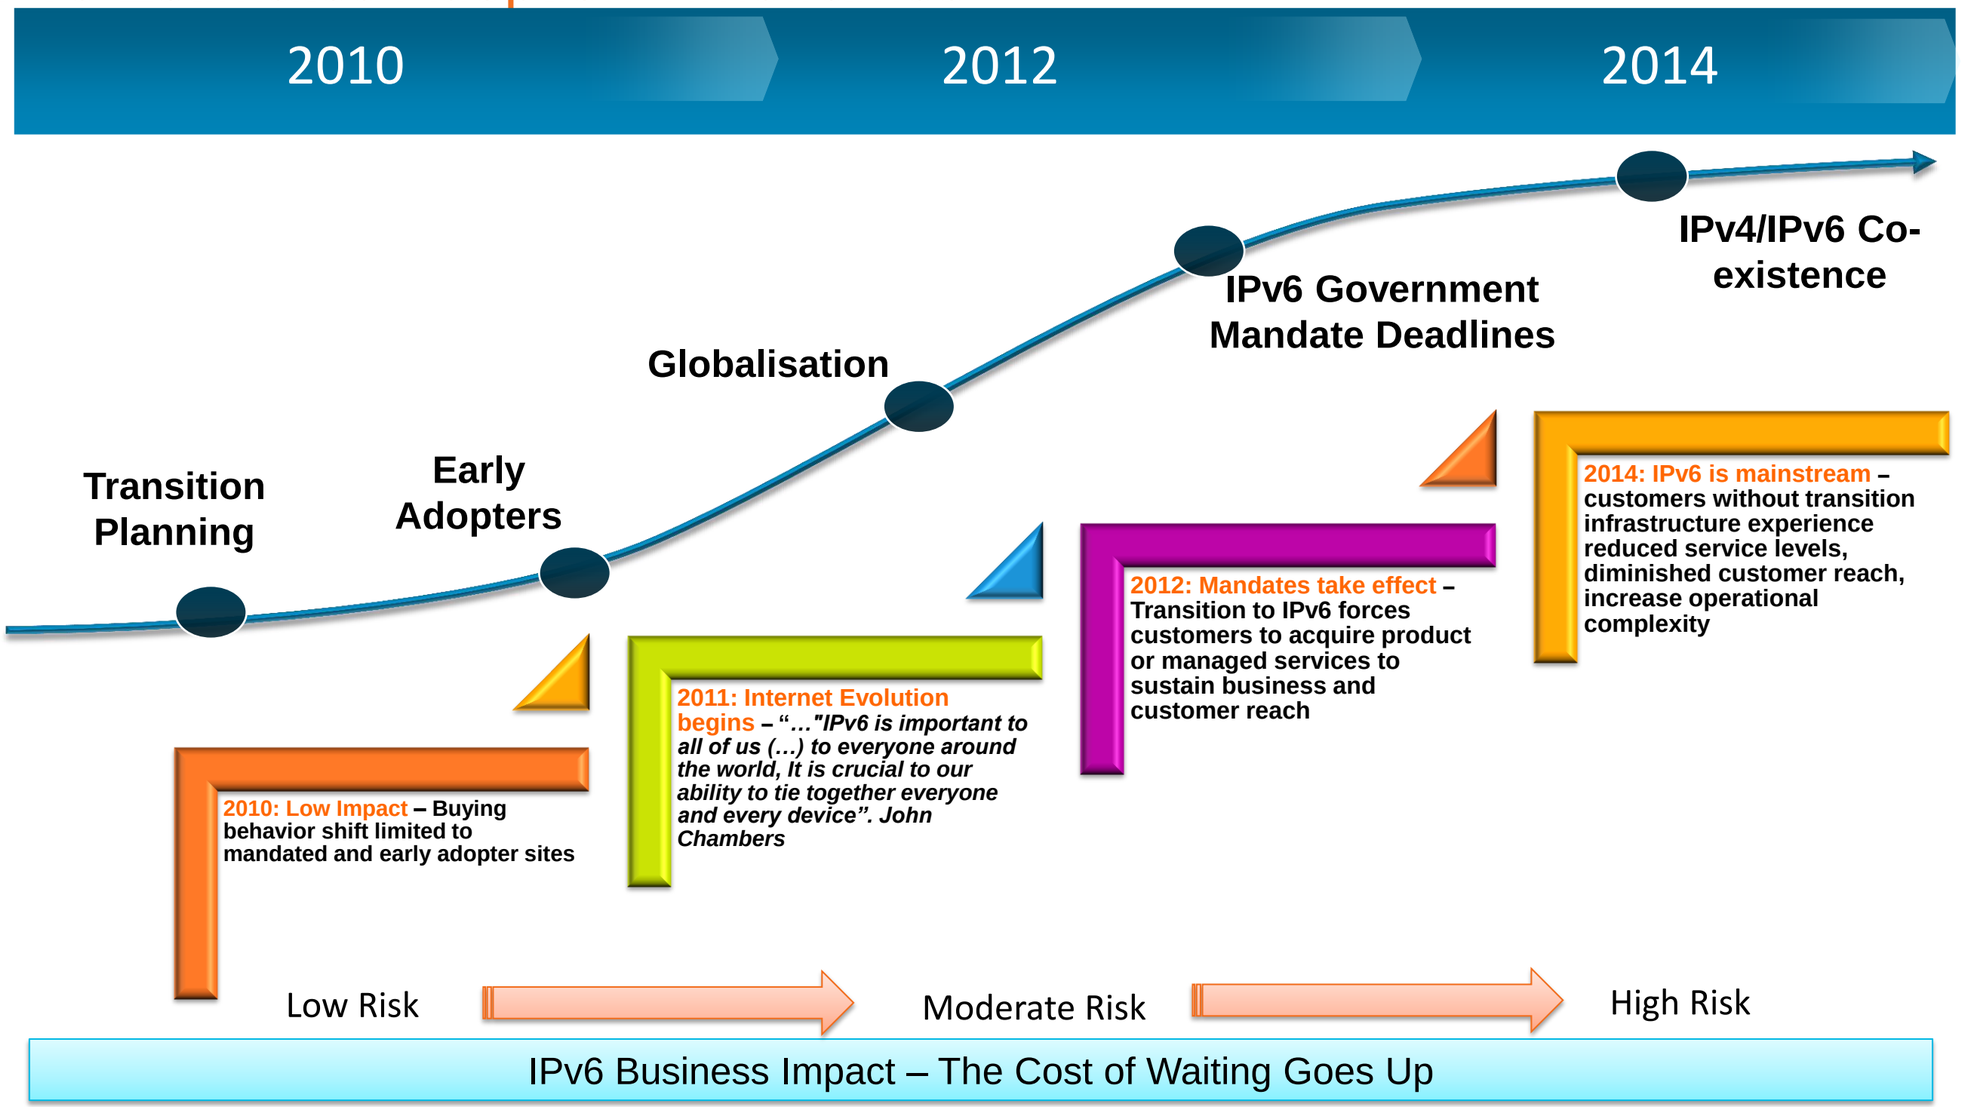

…. enterprises must prepare now for the IPv6 Internet or suffer business impacts due to connectivity failures.

Gartner – 5th April 2012

Impact of IPv4 Run-Out on Service Providers

- Difficult to add/support new IPv4 customers
- May find smaller IPv4 address blocks
  - Results in increasing routing table bloat
- Difficult to plan for IP NGN Services
  - For example, mobility and sensor networks
- Business Continuity could be impacted
  - No IPv6 Internet presence for customer with IPv6 only access
- Re use of private address blocks
  - Result is network management and operations difficulties

Enterprise Concerns

- What does IPv4 address depletion mean for us?
- What will happen if we don’t migrate?
- How complex is IPv6 migration? What are the potential challenges?
- How should we go about migrating/transiting to IPv6?
- Who has migrated/transited to IPv6? Why?
- 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

What Are the Current Issues?

- Urgency
  - Most companies see IPv6 as “in the indefinite future”
  - Have missed the fact of IPv4 address exhaustion
- Business Case
  - While company’s business partners do not require IPv6, they generally do not see the need
  - Unless – they offer services that depend on address availability.
- Education
  - IPv6 ≠ IPv4, and absent business case many have not taken the time to figure out the differences
- Vendor Support
  - Issues with Load managers, residential CPEs, legacy equipment
- Application Support
  - Application and Content Providers have the same business case issues

Innovator’s Dilemma

Business Continuity

IPv6 Planning

IPv6 Adoption – Customer Momentum

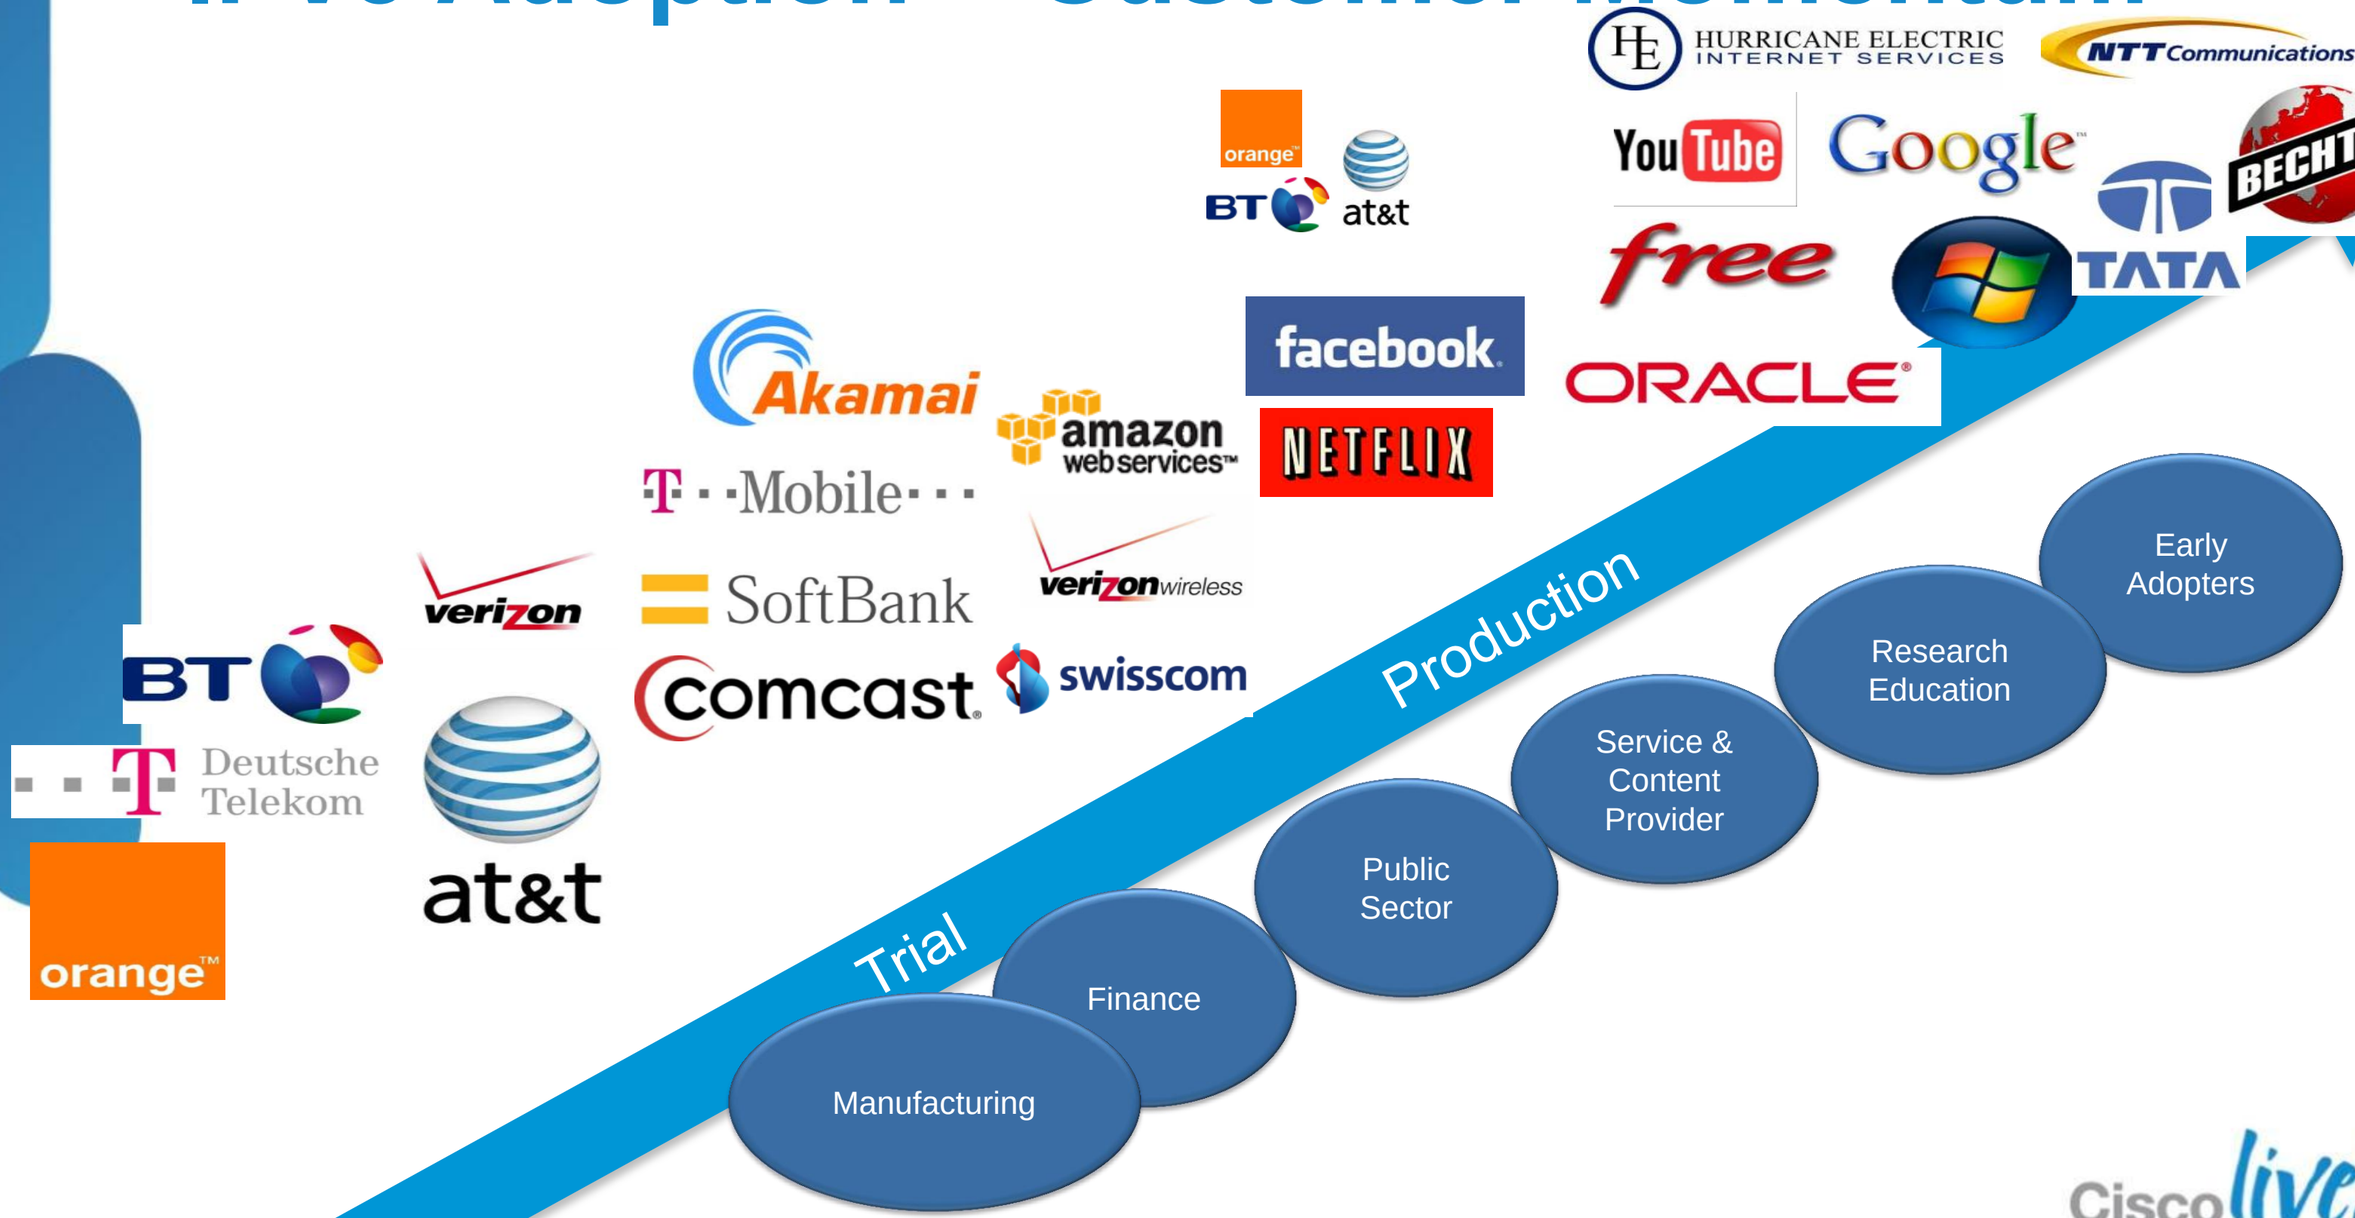

Cisco's IPv6 Web Presence Business Case

- Top level business drivers for IPv6
  - IPv6 leadership and mindshare
  - IPv6 product and solution readiness
  - Growth/continuity in the face of IPv4 address exhaustion
- Web presence business drivers
  - IPv4 address exhaustion NOW
  - Growing number of IPv6 endpoints on Internet
  - Cisco does business with customers and partners on the Internet
  - Act now or suffer business impacts due to connectivity failures
  - Act now by making www.cisco.com IPv6 accessible

Cisco's IPv6 Web Presence Roadmap

2002-2009 | 2010 | 2011 | 2012 | 2013 | 2014

IPv6 Internet Presence (Outside-In)

- www.ipv6.cisco.com
- www.cisco.com accessible over IPv6
- Entire cisco.com platform accessible over IPv6

Ubiquitous IPv6 User Access (Inside-Out)

- On-demand tunnel services
- Dual stack global core
- Dual stack user access pilot
- Dual stack user access (prod)
- Dual stack DC and apps

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

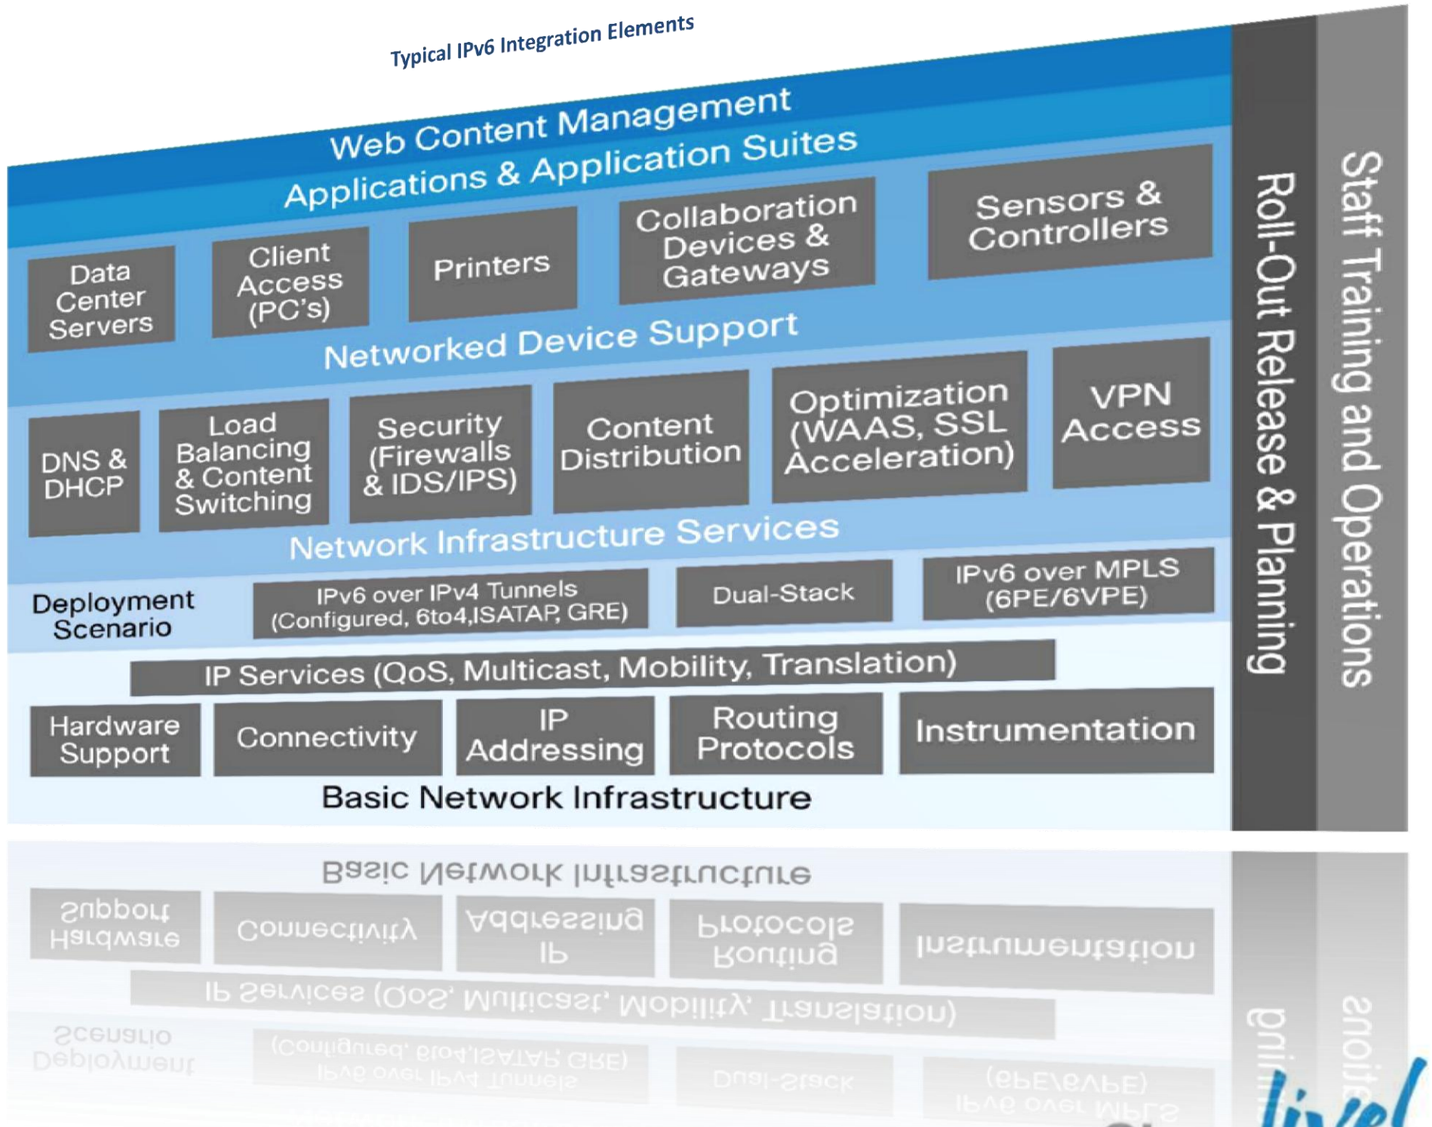

How Do We Get There from Here?

- IPv4 & IPv6 will coexist for the foreseeable future
  - No D-Day / Flag Day.
- Education & Careful Planning are crucial.
  - How long does it take in your environment?
  - What impacts are there to the existing network?
- IPv4 & IPv6 implementations must be scalable, reliable, secure and feature rich.

Strategy That Reflects This…

Starting with Edge Upgrades Enable IPv6 Service Offerings Now

World IPv6 Launch, June 6th 2012

http://www.worldipv6launch.org/

THIS TIME IT IS FOR REAL

Major Internet service providers (ISPs), home networking equipment manufacturers, and web companies around the world are coming together to permanently enable IPv6 for their products and services by 6 June 2012.

AKAMAI
COMCAST
FREE TELECOM
KDDI
TIME WARNER CABLE

AT&T
D-LINK
GOOGLE
LIMELIGHT
XS4ALL

CISCO
FACEBOOK
INTERNODE
MICROSOFT BING
YAHOO!

DO YOUR PART
JOIN THE LAUNCH!

We welcome web companies, ISPs, and home router vendors to join the cause and spread the word and follow along.

Typical Project Steps

How do we proceed?

Key Question: Where Do I Start?

- Based on Timeframe/Use case
- Core-to-Edge – Fewer things to touch
- Edge-to-Core – Challenging but doable
- Internet Edge – Business continuity

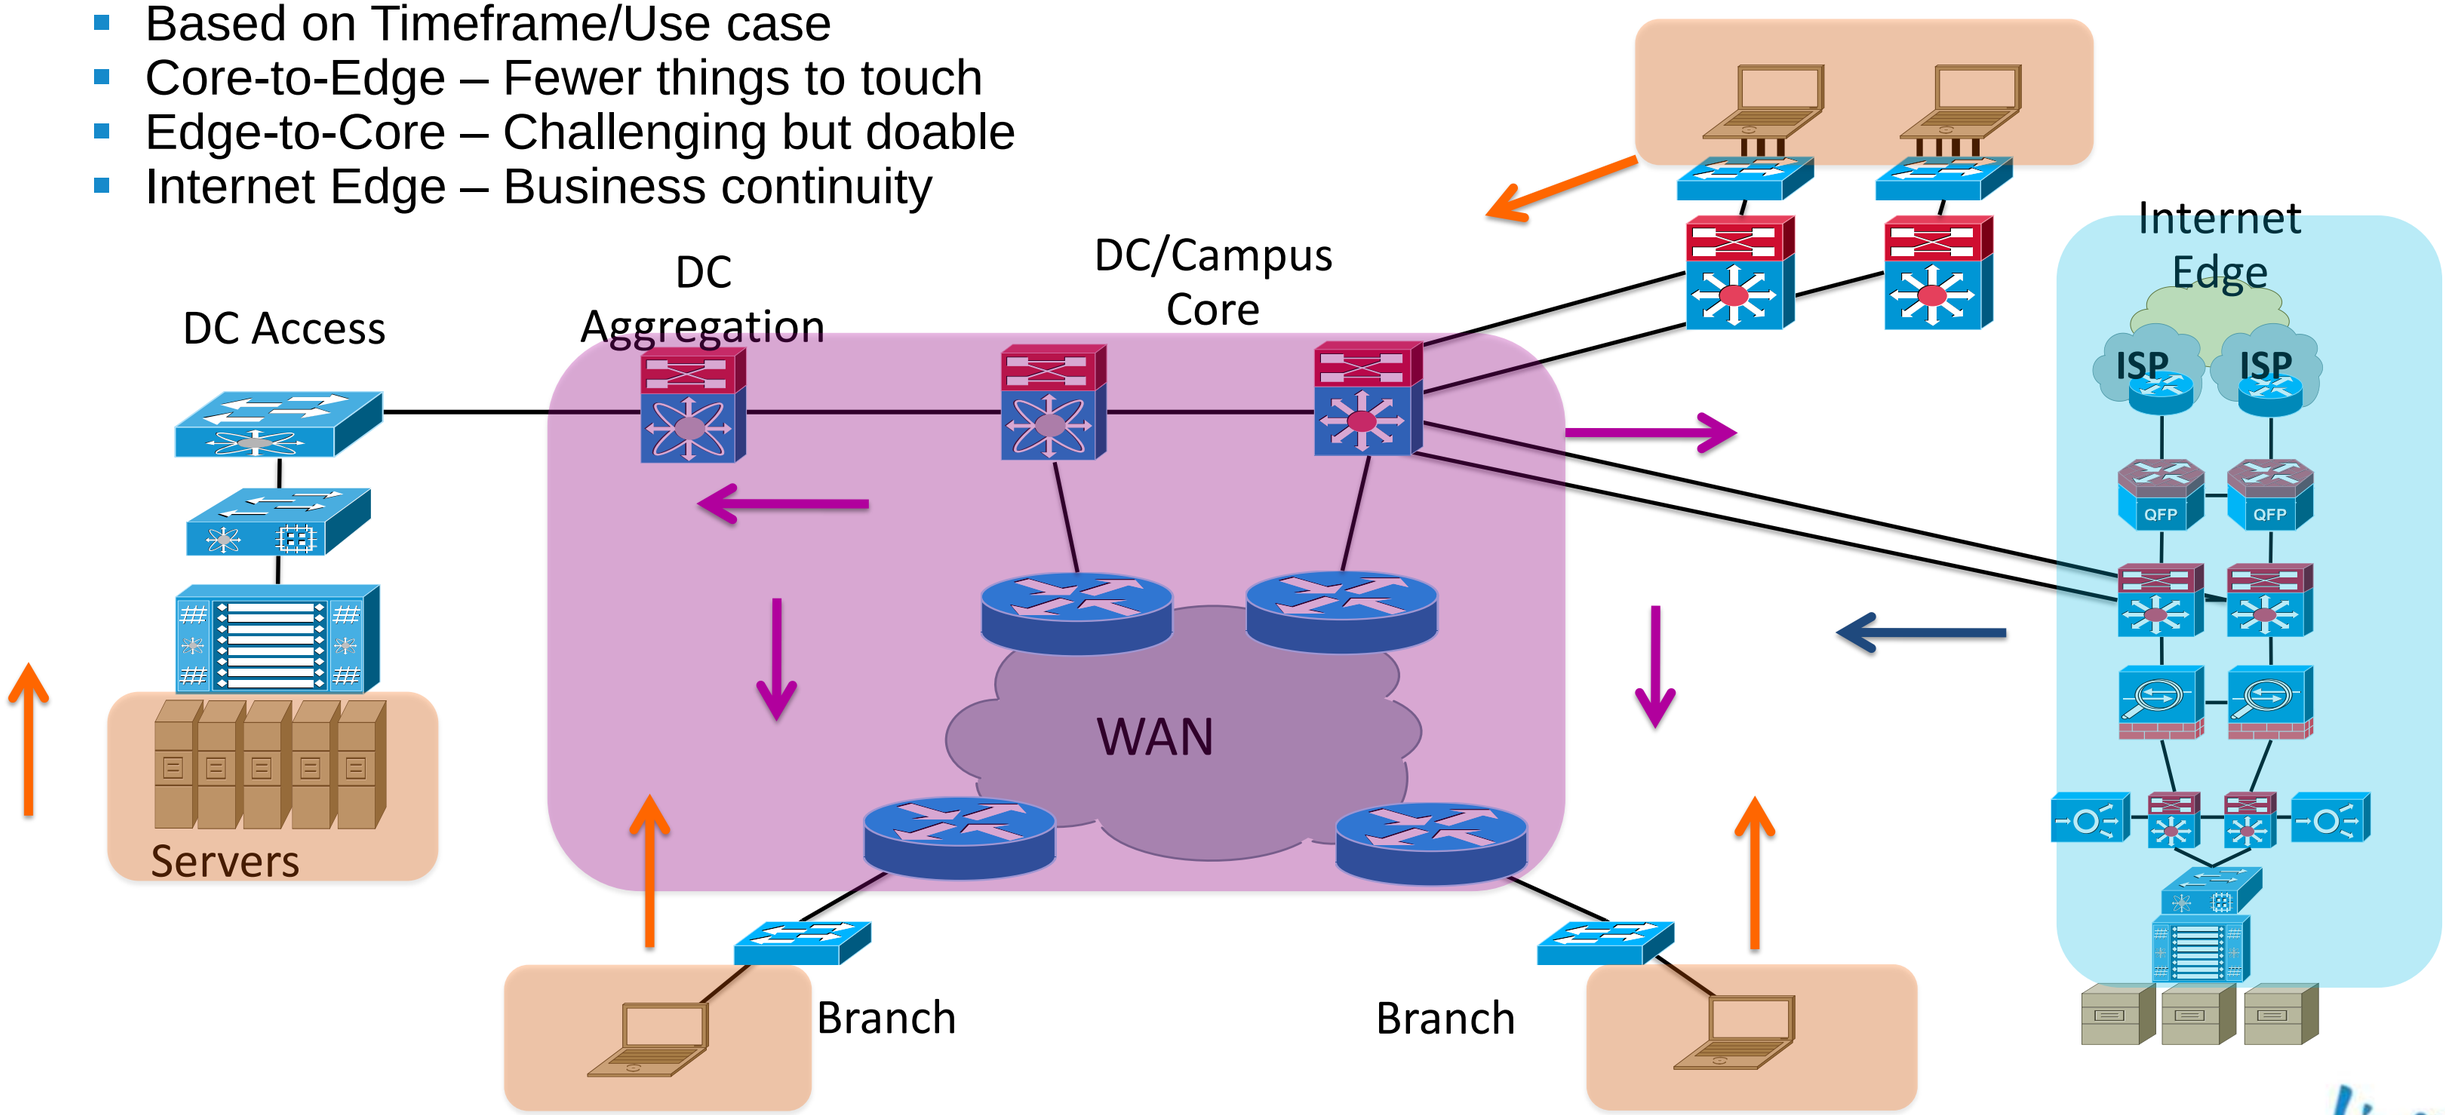

Questions to Ask Your Service Provider

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
  - Available over 802.1Q or VLANs ?
  - Separate or Same VRF's ?
- 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

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				

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

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

IPv6 Addressing

Building the IPv6 Address Plan

- Build on the lessons learned from how the IPv4 plan was developed and implemented
  - Does it make sense to follow the current IPv4 assignment model?
- Must be proportional to current usage and expected growth
- Changing ARIN policy on block sizing
- Hierarchy is key
  - Do you get a prefix for the entire company or do you get one prefix per site (what defines a site?)
- Cisco IPv6 Addressing White Paper
  - http://www.cisco.com/web/strategy/docs/gov/IPv6_WP.pdf

IPv6 Address Considerations

- What type of addressing should I deploy internal to my network?
  - ULA (Unique Local Address)
  - ULA + Global
  - Global
  - Global + LL (Link Local) infra
- Do you need a Provider Independent (PI) or a Provider Assigned (PA) space?
  - What about presence across regional registries?
- How big of space do I need or could I get?
  - Consider number of sites, current geographic footprint, Address concentration, Growth projection, Analyse Current address aggregation challenges

ARIN Policy: https://www.arin.net/policy/nrpm.html

ULA, ULA + Global or Global

- What type of addressing should I deploy internal to my network?
  - ULA-only
    - IPv6 to IPv6 translation boundary needed
    - NPTv6 (Network Prefix Translation) is relatively new and not widely deployed
  - ULA + Global allows for the best of both worlds but at a price
    - Much more address management with DHCP, DNS, routing and security
    - Source Address Selection (SAS) does not always work as it should
  - Global-only
    - Recommended approach but the old-school security folks that believe topology hiding is essential in security will bark at this option
  - Global + LL infra LL only
    - Use LL on network infrastructure and source everything from global address
    - Described in IETF draft http://wiki.tools.ietf.org/html/draft-behringer-lla-only-00

32

Global-Only

Recommended

- Global is used everywhere
- No issues with SAS
- No requirements to have NAT for ULA-to-Global translation—but, NAT may be used for other purposes
- Easier management of DHCP, DNS, security, etc.
- Only downside is breaking the habit of believing that topology hiding is a good security method ☺

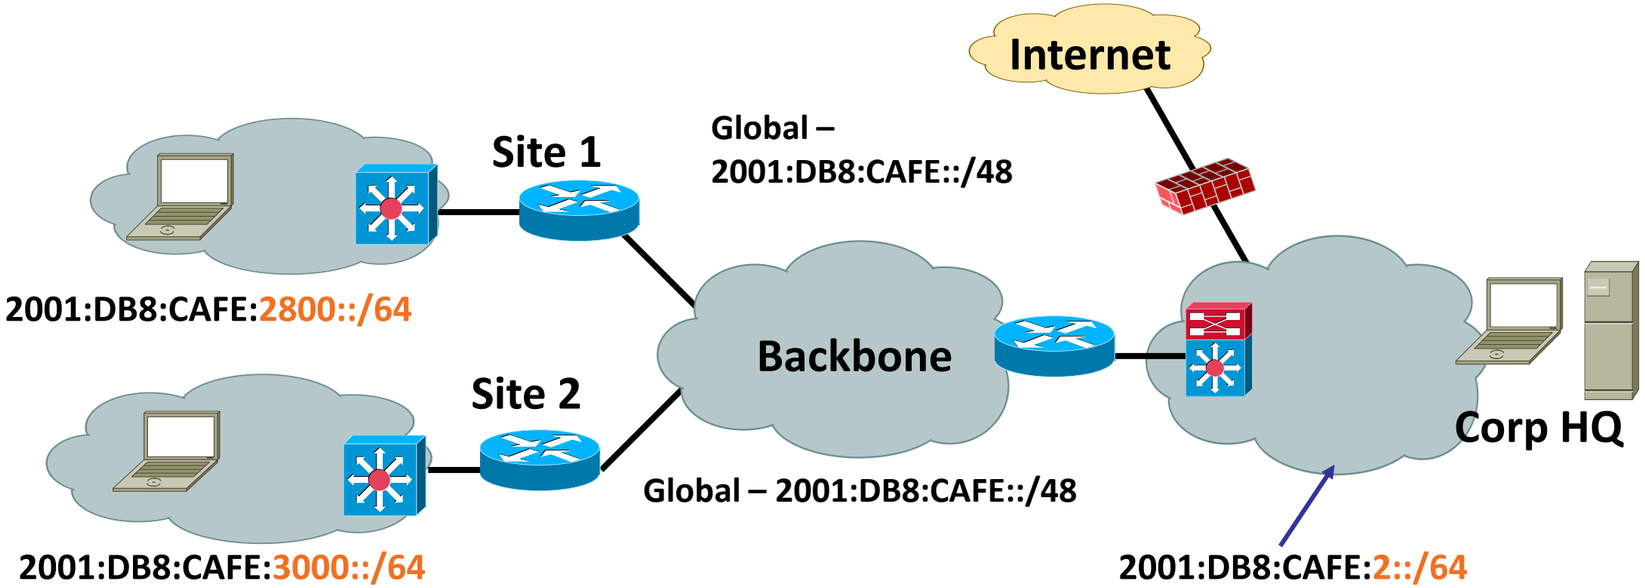

PI Space Concerns

- 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

34

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

/48 Prefix Breakdown Example

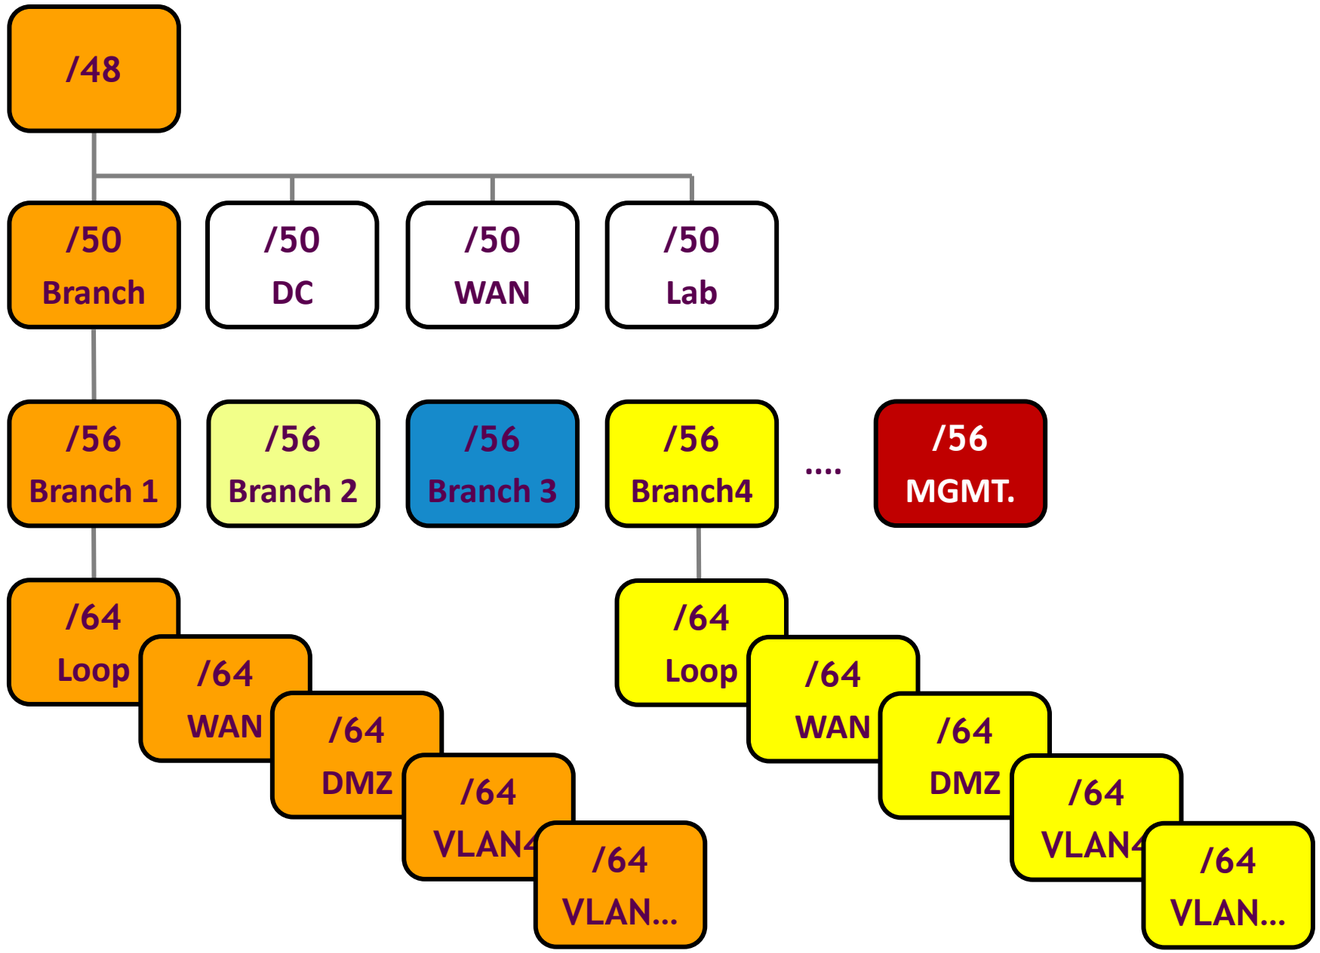

- High Level addressing plan. Indicative only. Can be modified to suit needs
- /48 = 65536 x /64 prefixes
- Break up into functional blocks (4 x /50 in this case)
- Each functional block simplifies security policy
- Assumes up to 64 Branch networks
- Each Branch has access to 256 /64 prefixes for WAN, 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)
  - Can 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.
- IPv6 Address Management – How are you going to manage these blocks?

Transition Mechanisms

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

IPv6 Co-Existence Solutions

Dual Stack

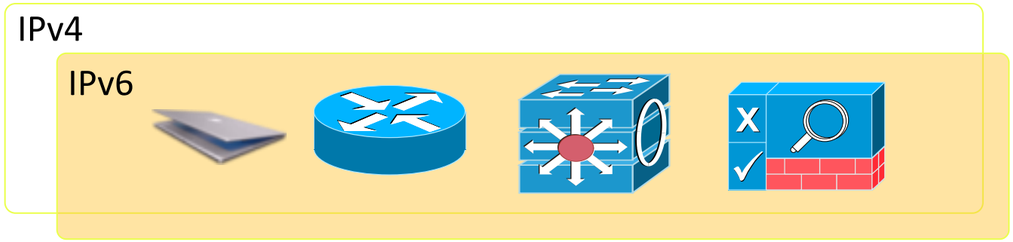

Recommended Enterprise Co-Existence Strategy

Tunnelling Services

Connect Islands of IPv6 or IPv4

Translation Services

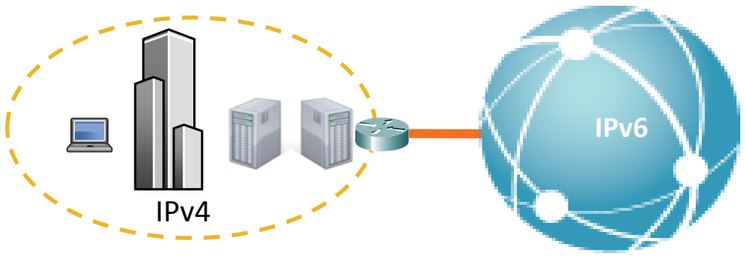

Connect to the IPv6 Community

Dual Stack

Backbone

- All P + PE routers are capable of IPv4+IPv6 support
- Two IGPs supporting IPv4 and IPv6
- Memory considerations for larger routing tables
- Native IPv6 multicast support
- All IPv6 traffic routed in global space
- Good for content distribution and global services (Internet)

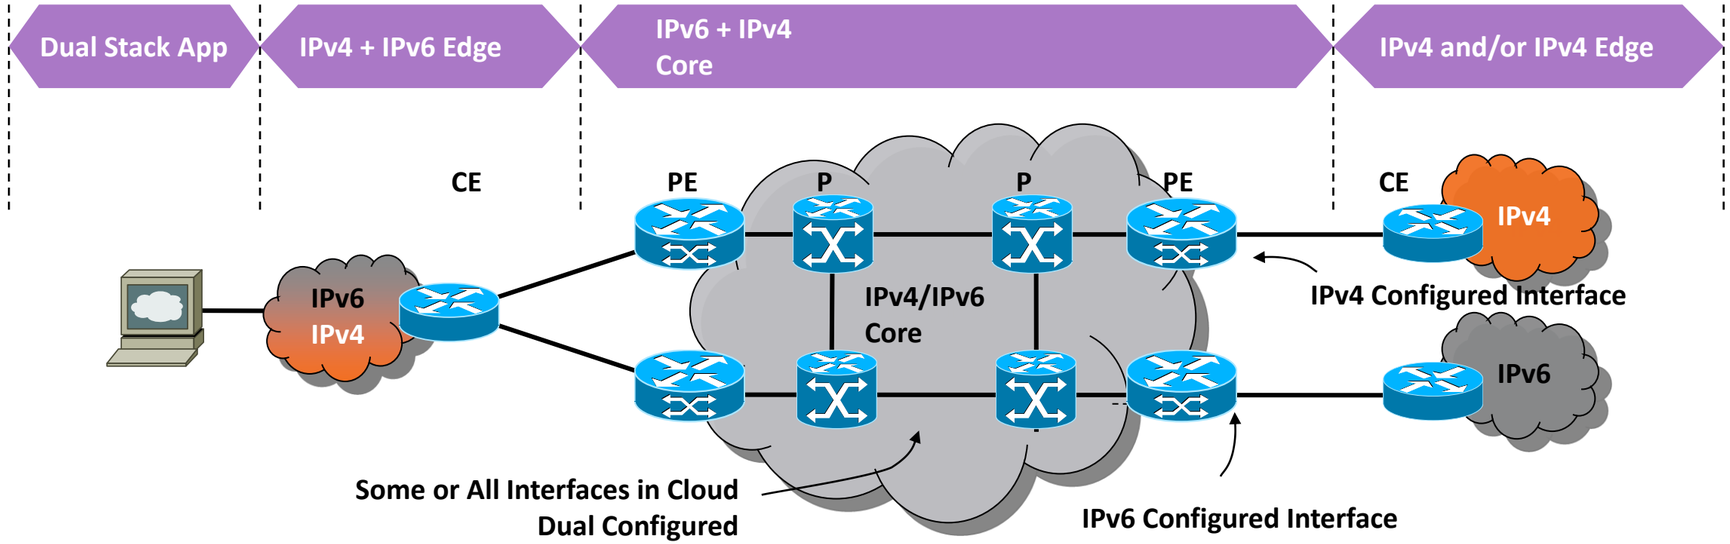

Dual Stack Application Approach

- Dual stack in a device means
  - Both IPv4 and IPv6 stacks enabled
  - Applications can talk to both
  - Choice of the IP version is based on DNS and application preference
- Dual stack at edge does not necessarily mean dual stack backbone

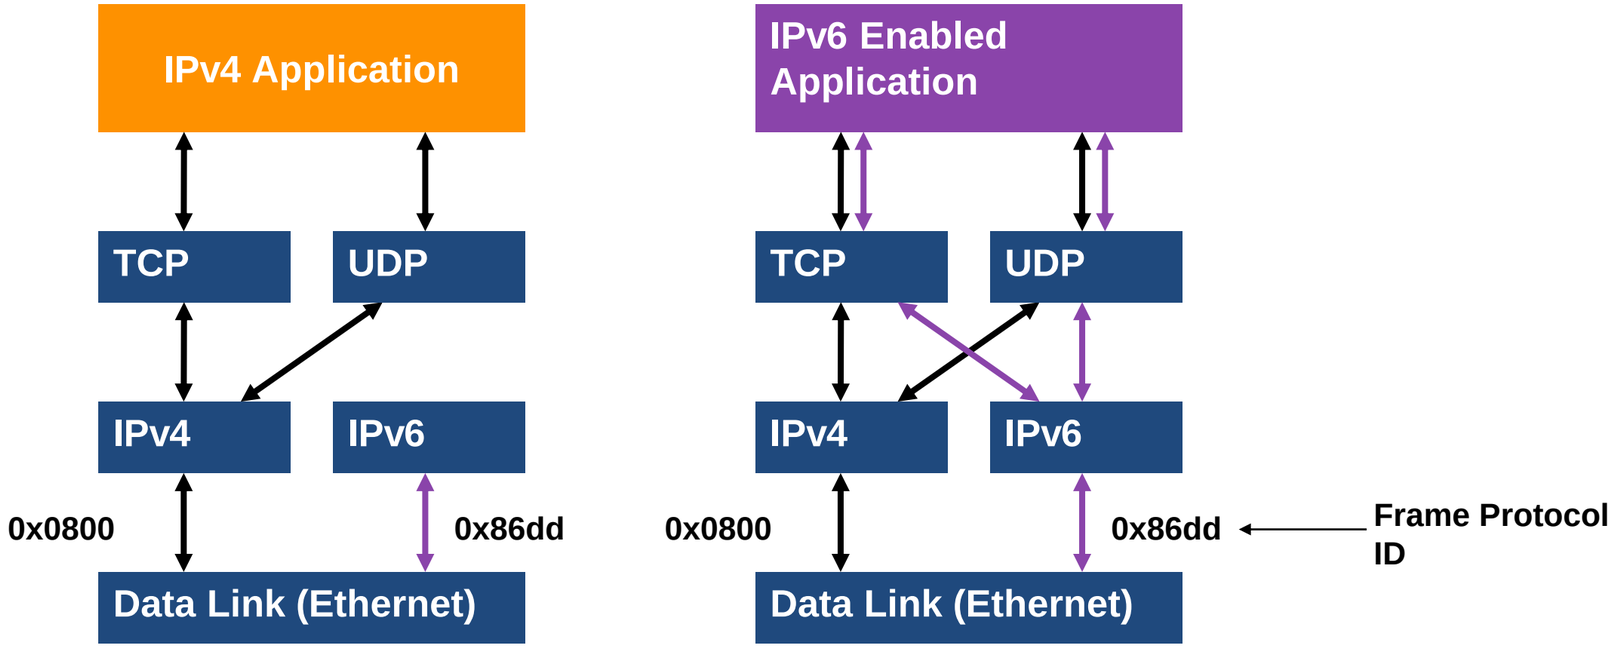

Tunnels

Manual Using RFC 4213 and GRE

- 6in4 was one of the first transition mechanisms developed for IPv6
  - Static P2P tunnel, IP protocol type = 41, no additional header, NAT breaks
  - IPv4 end point address must be routable
  - IPv6 prefix configured on tunnel interface
- An IPv6 in GRE tunnel solution also exists (not shown)
  - uses 0x86DD to identify IPv6 payload
- Usual manual point to point tunnel caveats apply

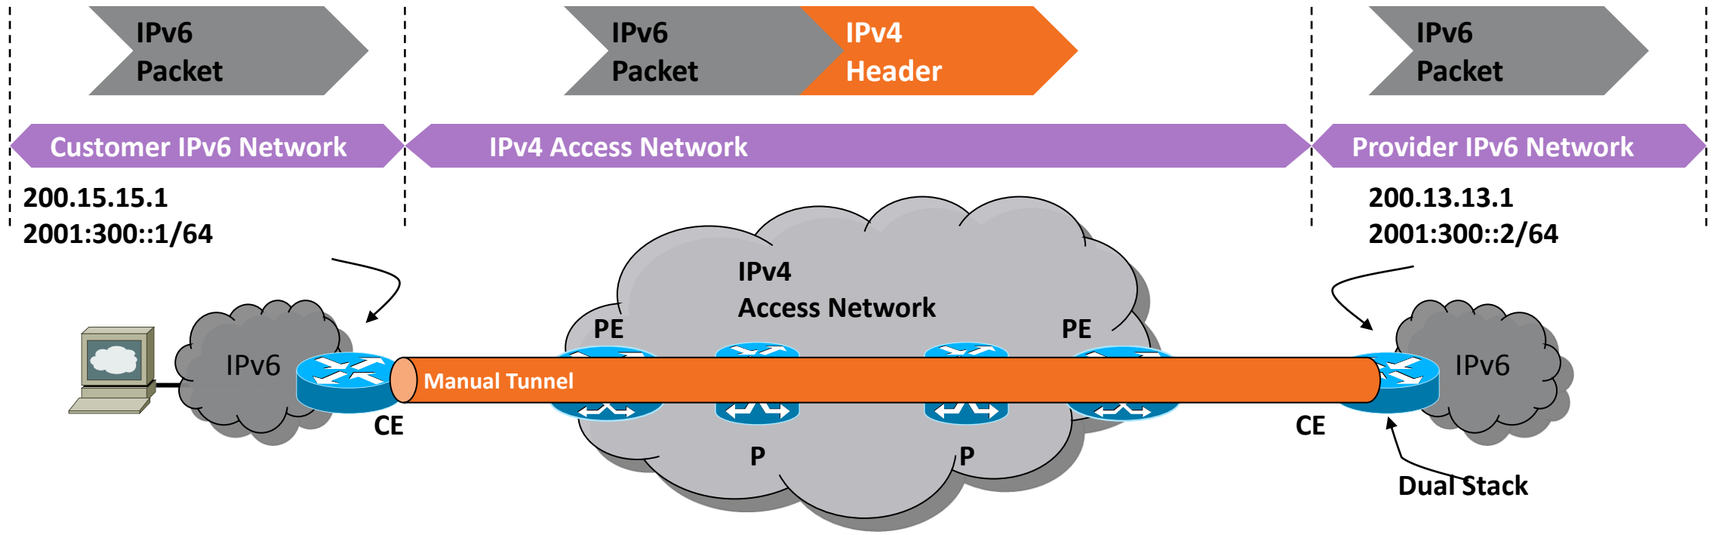

Tunnels
6 to 4 (RFC 3056)

- Automatic tunnel method using 2002:IPv4::/48 IPv6 range
  - IPv4 embedded in IPv6 format e.g.. 2002:c80f:0f01:: = 200.15.15.1
- No impact on existing IPv4 or MPLS Core (IPv6 unaware)
- Tunnel endpoints have to be IPv6 and IPv4 aware (Dual stack)
- Transition technology – not for long term use
- Intrinsic linkage between destination IPv6 Subnet and IPv4 gateway interface
  - IPv4 Gateway = Tunnel End point

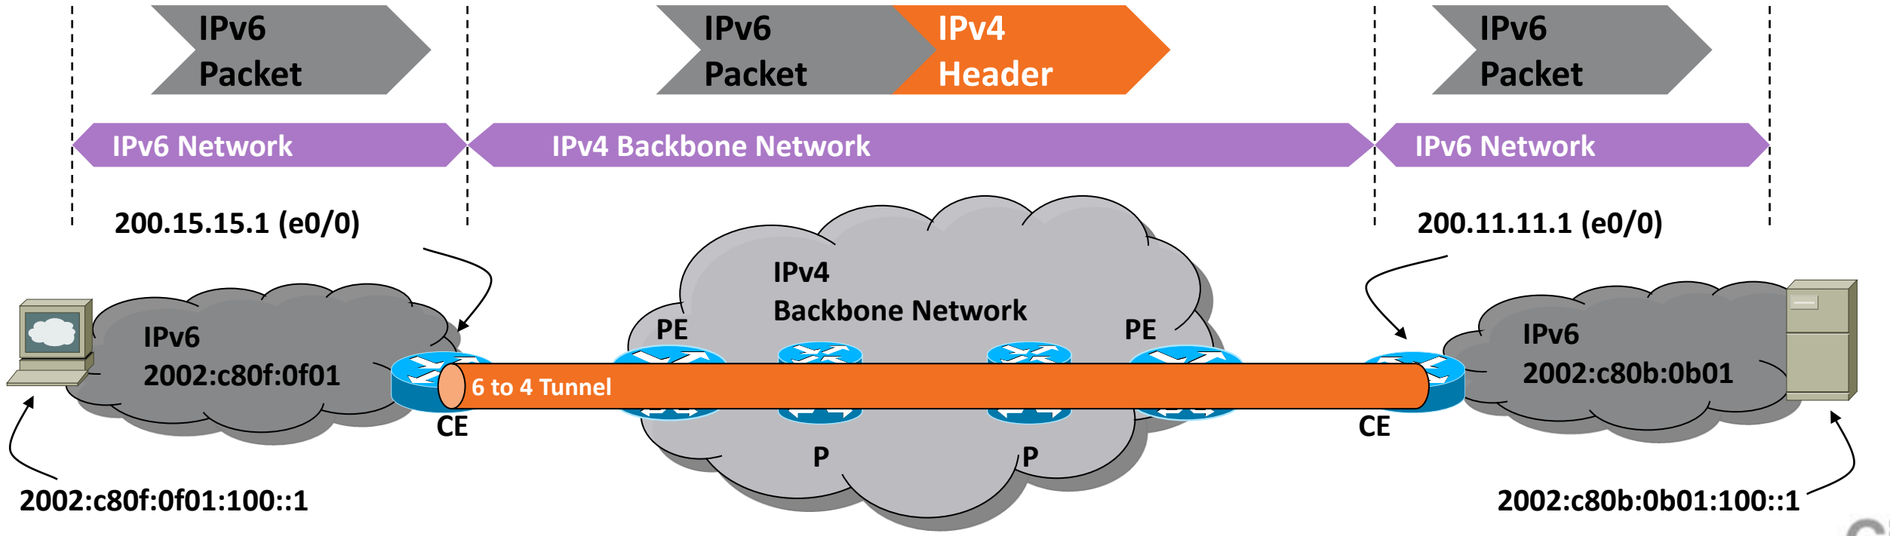

Tunnels

6rd

- No DHCPv6, Neighbour Discovery, etc. to deploy in access network
- IPv6 addressing automatically created from IPv4 addressing, synced with IPv4 lease
- 6rd Gateway provides access to IPv6 core network
- Simple, stateless, automatic IPv6-in-IPv4 encap and decap function on RG
- Based on 6to4 (RFC 3056) and RFC 5569 + new draft from Cisco coming soon

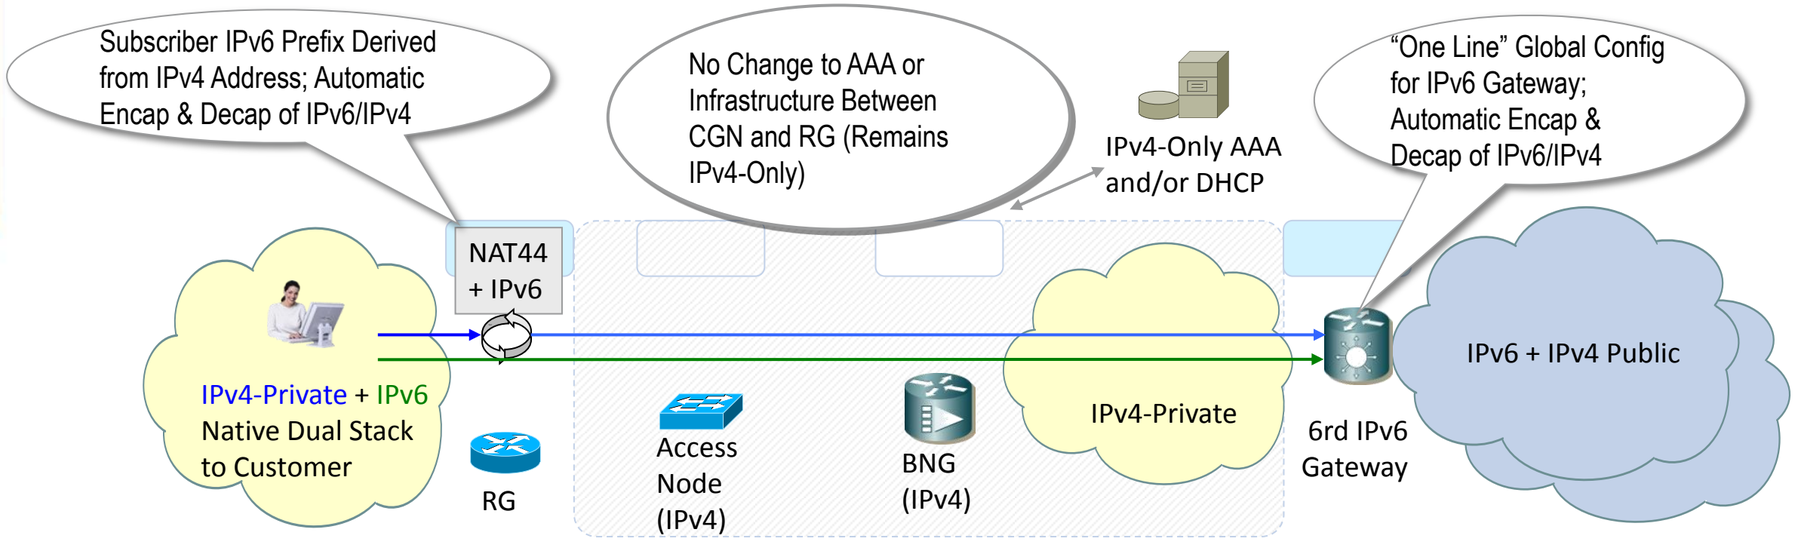

Tunnels
ISATAP (RFC 5214)

- Intra Site Automatic Tunnel Addressing Protocol
  - Tunnel from a dual stack HOST PC to an IPv6 gateway
- ISATAP hosts use a special IPv6 address format
  - Rightmost 32 bits of Interface ID contains the host IPv4 address
  - Leftmost 32 bits of Interface ID contains “0000:5EFE”
- Operates within single administrative domain
- Creates a virtual IPv6 link over an IPv4 backbone
  - IPv4 network treated as an NBMA link layer
  - Routers provide ISATAP service
  - DNS may hold potential router list or ISATAP gateways

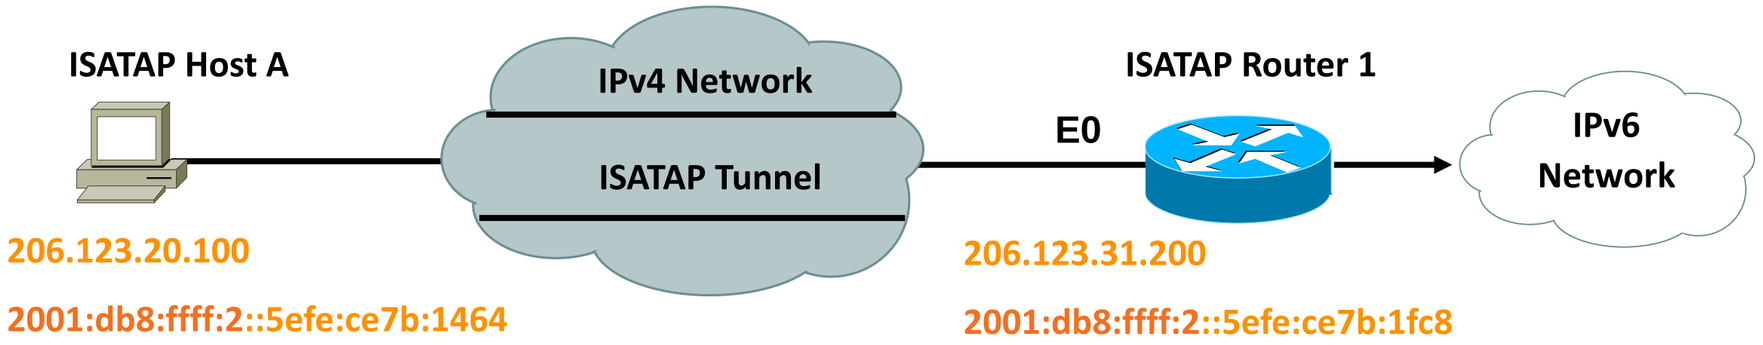

Tunnels

DMVPN for IPv6

- Connects private IPv6 islands across public IPv4 cloud
- Public IPv4 network treated as NBMA network
  - Static hub and spoke topology
  - Dynamic spoke to spoke topology
- Supports dynamic unicast and multicast routing

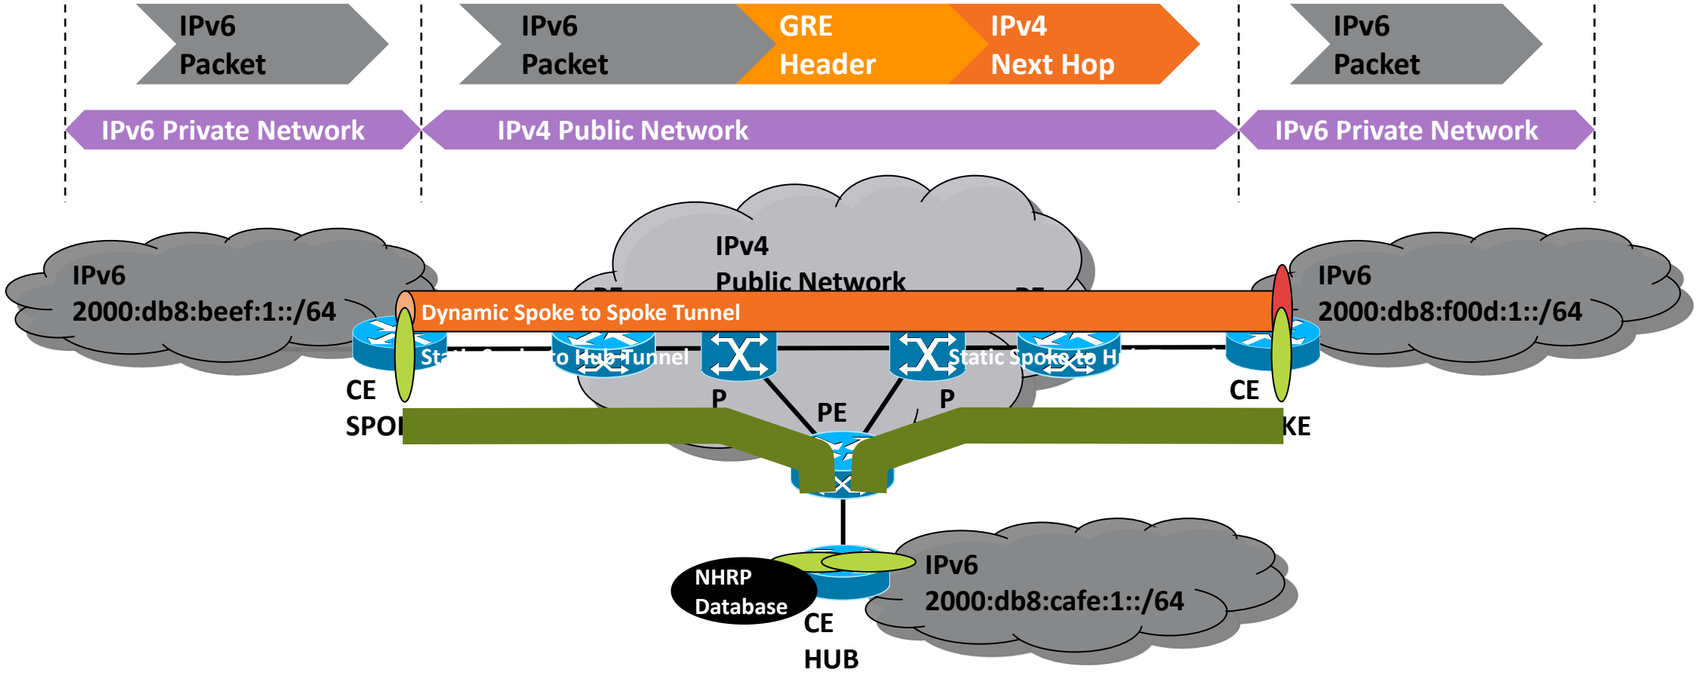

Tunnels
LISP

- Needs:
  - Rapid IPv6 Deployment
  - Minimal Infrastructure disruption
- LISP Solution:
  - LISP encapsulation is Address Family agnostic
  - IPv6 interconnected over IPv4 core
  - IPv4 interconnected over IPv6 core
- Benefits:
  - Accelerated IPv6 adoption
  - Minimal added configurations
  - No core network changes
  - Can be used as a transitional or permanent solution

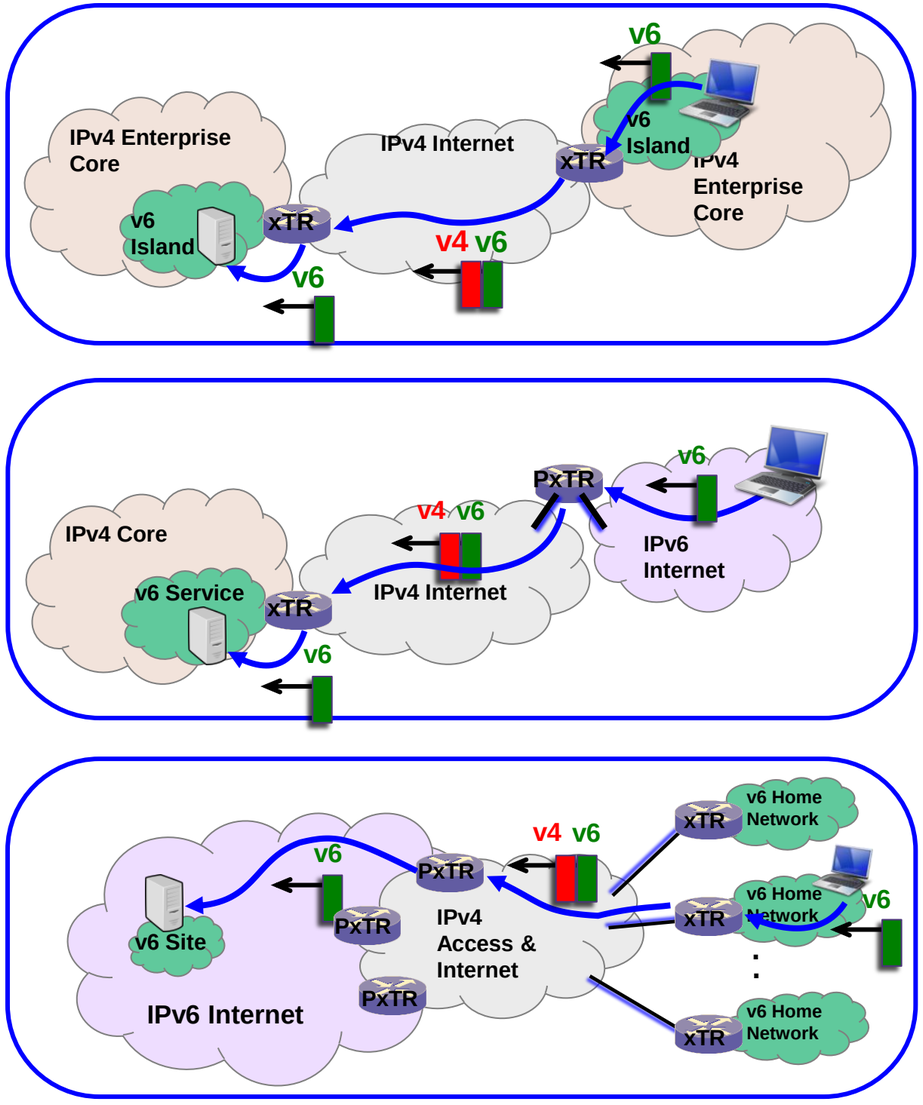

MPLS

IPv6 Over Pseudowires

- No impact on existing MPLS Core (IPv6 unaware)
- CE routers need only be single stack IPv6 aware
- P2P PWs or P2MP PWs (VPLS) supported at PE
- Good for interconnecting discrete customer islands
- Sub-Optimal Multicast support, Full Mesh required
- Many CE interface typed supported – Ethernet, ATM, Frame Relay

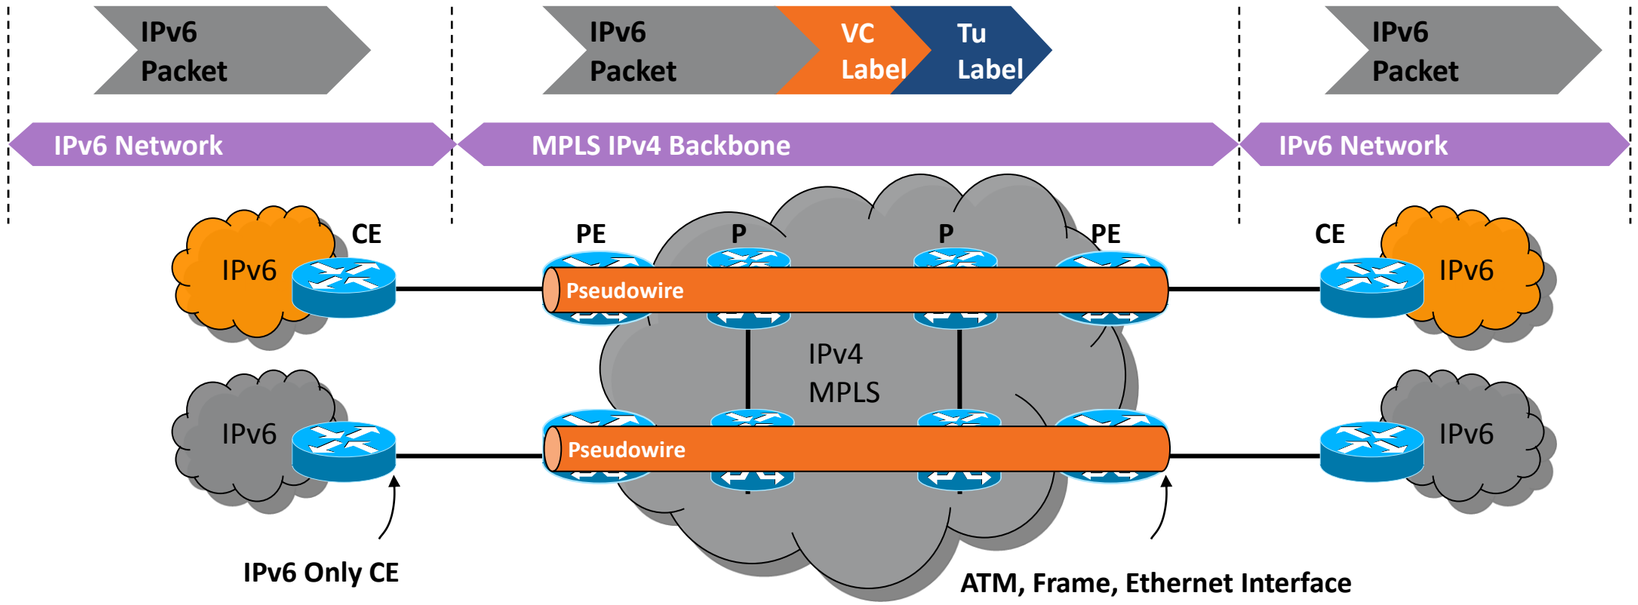

MPLS

IPv6 Transit Using 6PE (RFC 4798)

- 6PEs must support dual stack IPv4+IPv6 (acts as normal IPv4 PE)
- IPv6 packets transported from 6PE to 6PE over Label Switch Path
- IPv6 addresses exist in global table of PE routers only
  - IPv6 addresses exchanged between 6PE using MP-BGP session
- Core uses IPv4 control plane (LDPv4, TEv4, IGPv4, MP-BGP)
- Benefits from MPLS features such as FRR, TE

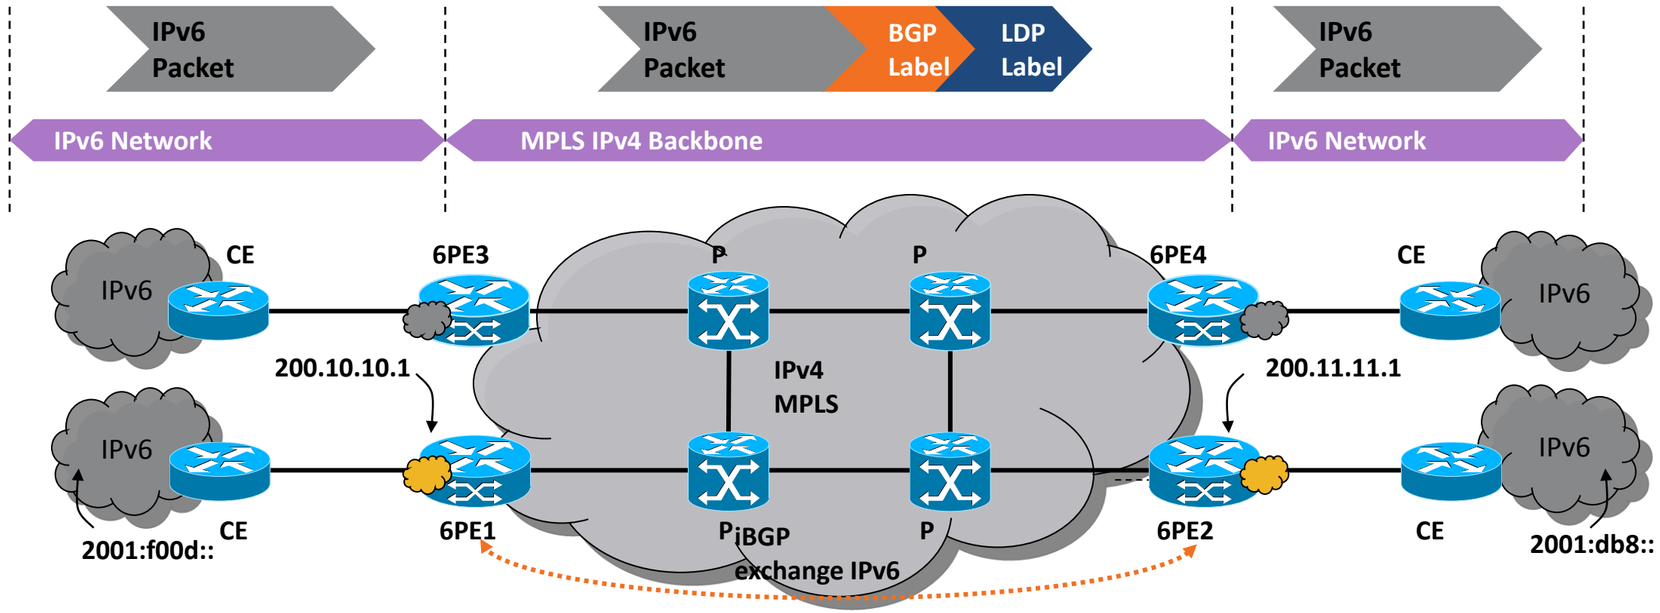

MPLS

IPv6 VPN 6VPE (RFC 4659)

- 6VPE uses existing IPv4 MPLS infrastructure to provide IPv6 VPN
  - Core uses IPv4 control plane (LDPv4, TEv4, IGPv4)
- PEs must support dual stack IPv4+IPv6
- Offers same architectural features as MPLS-VPN for IPv4
  - RTs, VRFs, RDs are appended to IPv6 to form VPNv6 address
  - MP-BGP distributed both VPN address families
  - BGP NH uses IPv4 to IPv6 mapped address format ::ffff:A.B.C.D
- VRF can contain both VPNv4 and VPNv6 routes

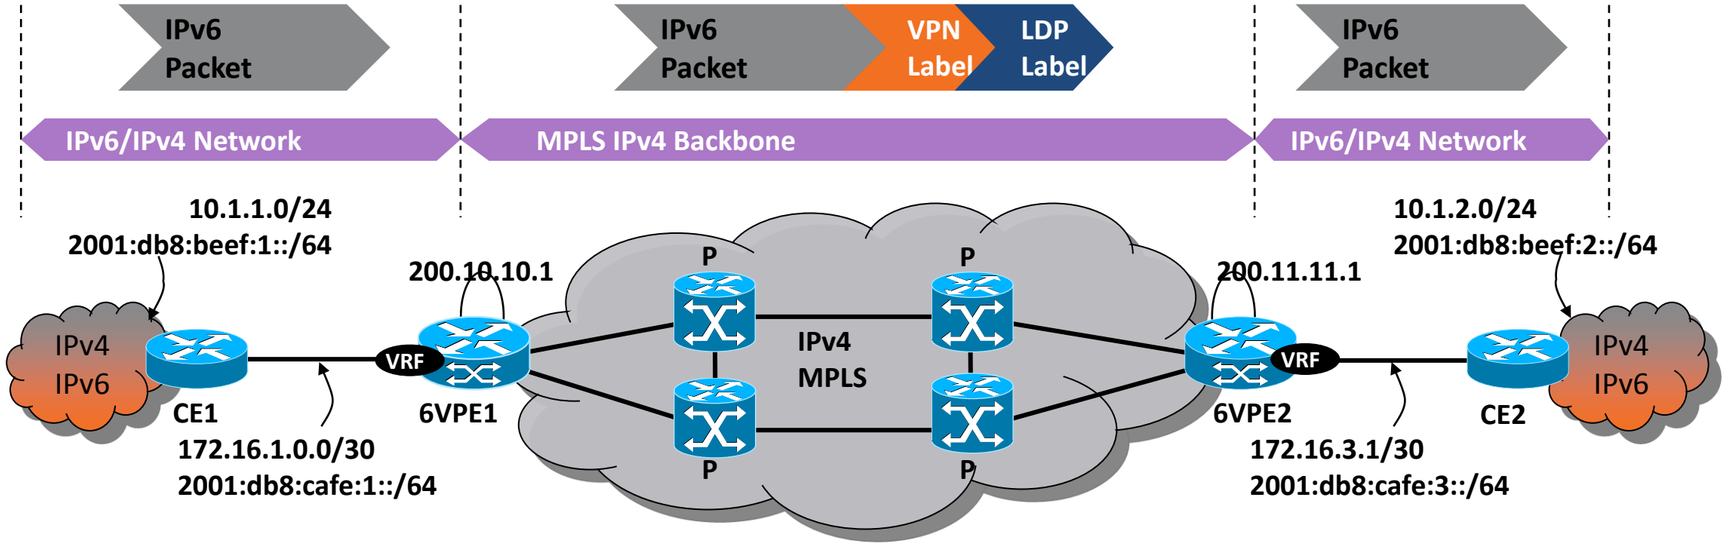

Translation

Address Family Translation (AFT)

Allows Access Between IPv6 and IPv4 Networks (IETF BEHAVE)

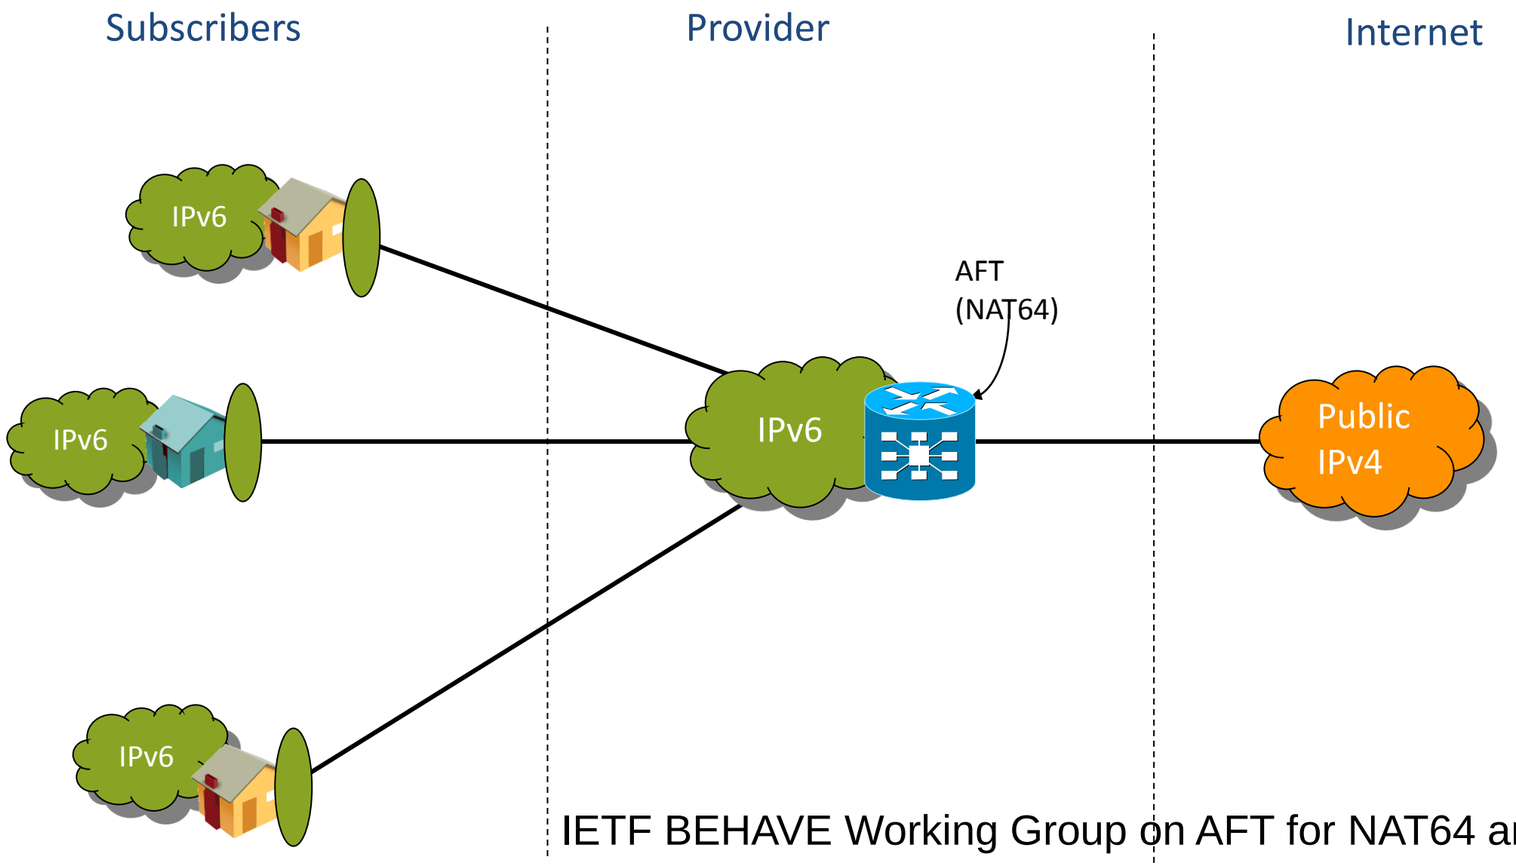

IETF BEHAVE Working Group on AFT for NAT64 and NAT46

Translation: Stateful/Stateless

- Stateless
- NAT is more scalable
- NATs can be located anywhere
- IPv4- or IPv6-initiated connections
- 1:1 mapping
  - one IPv4 address is consumed for every participating IPv6 address

- Stateful
- NAPT (Network Address and Port Translator)
- NAT is less scalable
- NAT placement related to network topology
- IPv6-initiated connections
- 1:N mapping
  - Many IPv6 hosts consume 1 IPv4 address
- Limited TCP ports

Stateful NAT64 + DNS64 Operation

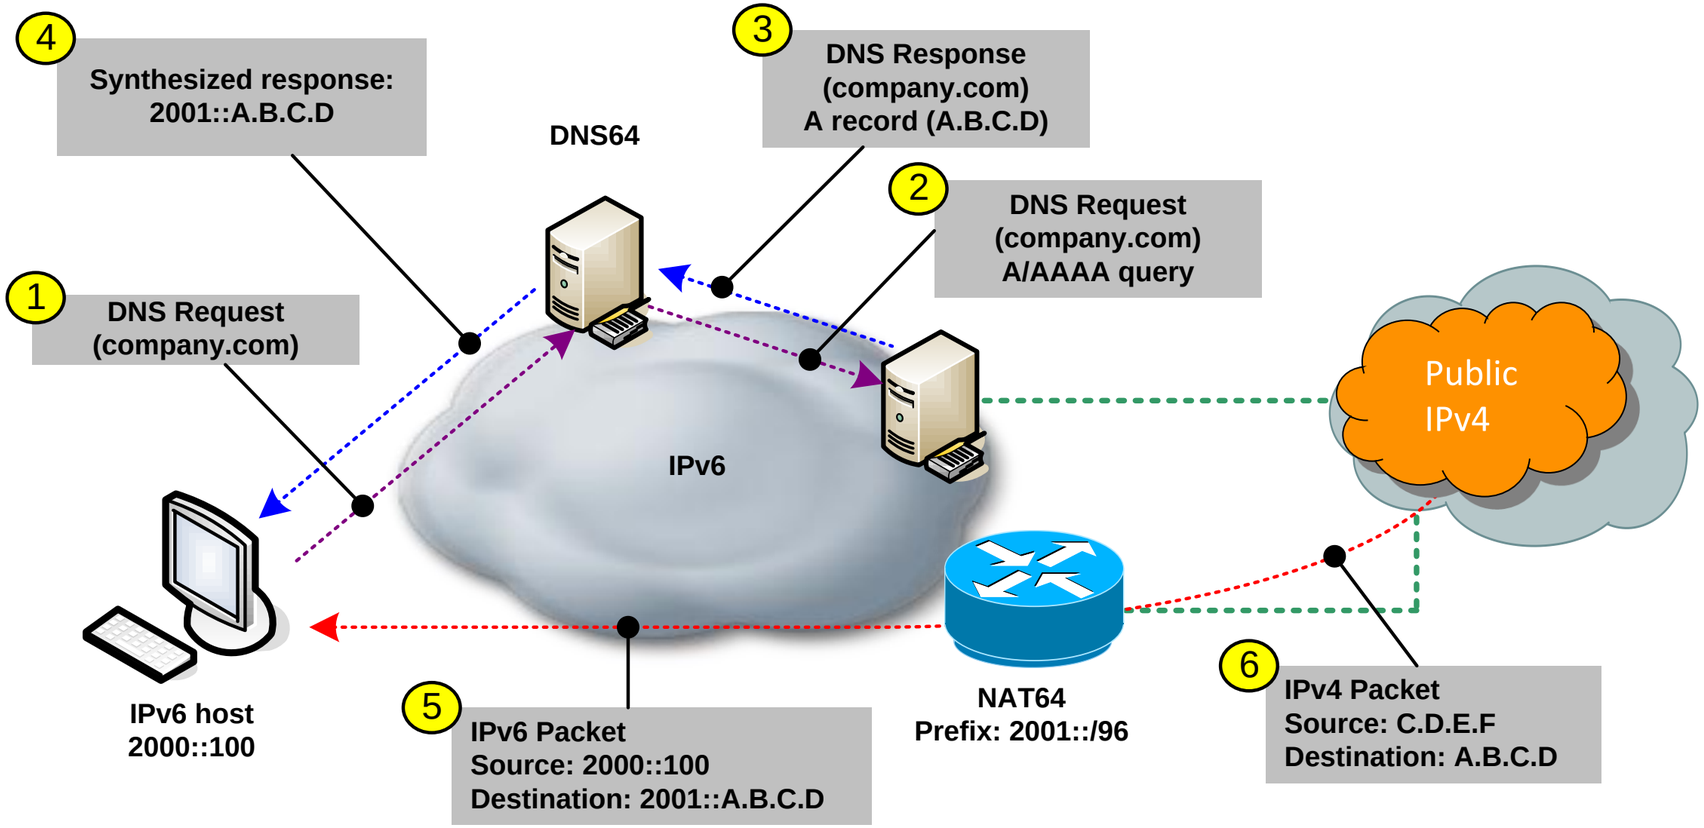

Who Owns the Translation?

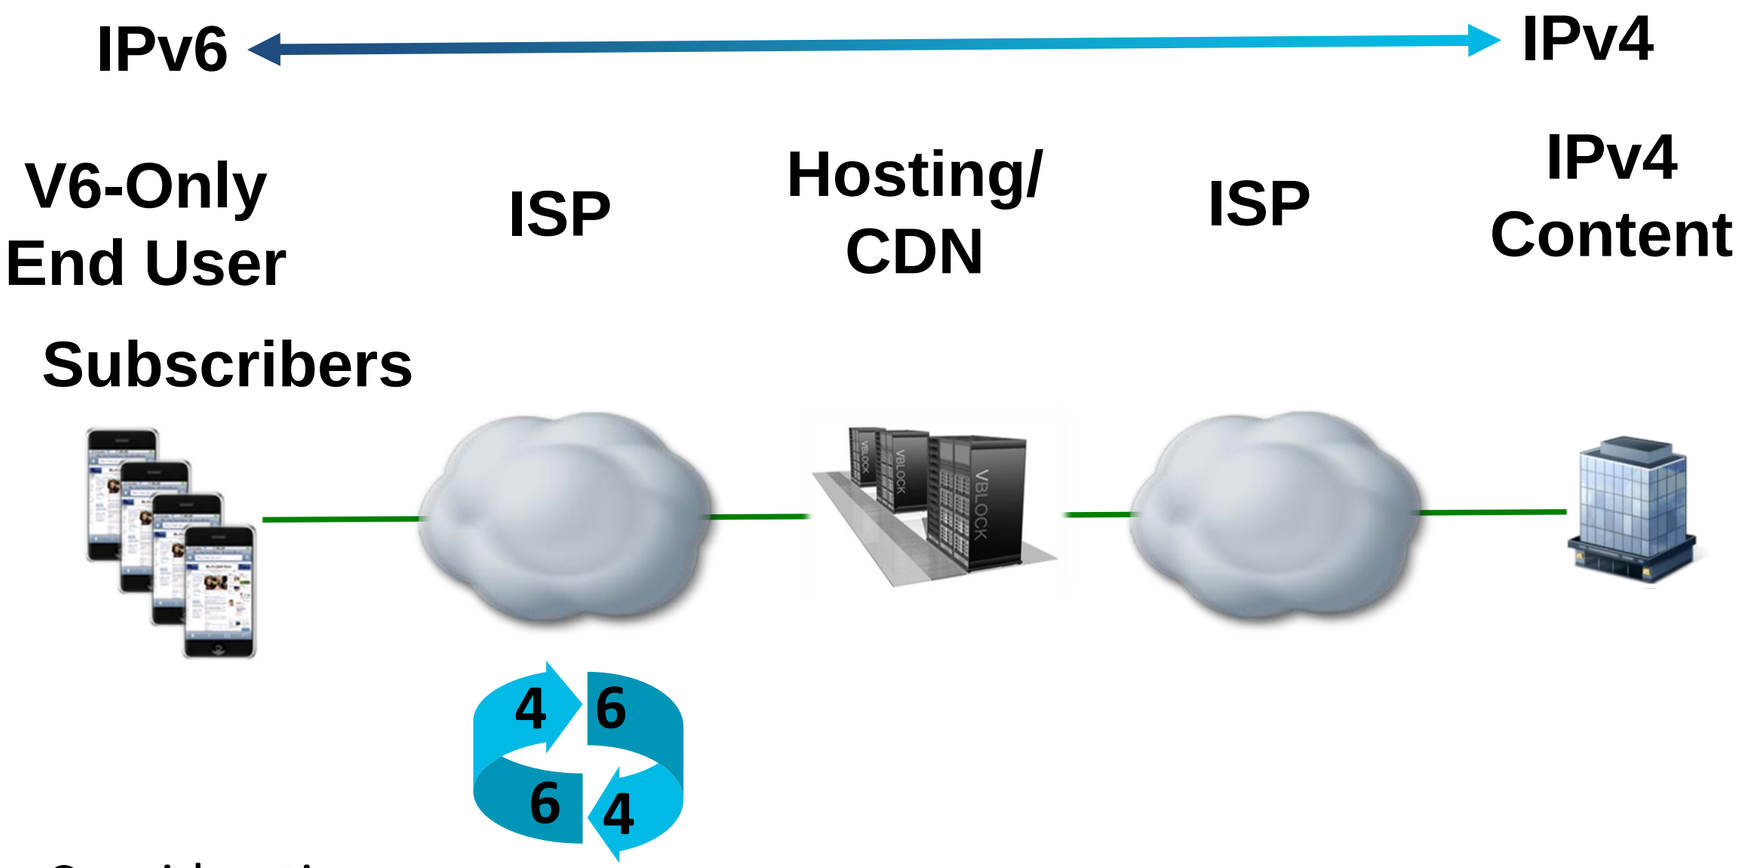

Considerations:

Experience, Scale, Cost, Operations, Technology…

The Enterprise Problem and Solution

- The Internet challenge
  - Internet access providers deploy IPv6-only, growing number of IPv6-only endpoints
  - Providers introduce network address translation (NAT) for IPv6-only endpoints to reach IPv4-only content and services
- The enterprise problem
  - Business-critical Internet-facing content and services on IPv4
  - IPv6 endpoints MUST go through NATs
  - Provider NATs introduce complexity, mask identity, and may impact performance
  - Enterprise will lose some control over user experience...this is UNACCEPTABLE
- The enterprise solution
  - Maintain the “end to end” model and existing control over user experience
  - Content and services accessible over both IPv4 and IPv6 (dual stack)

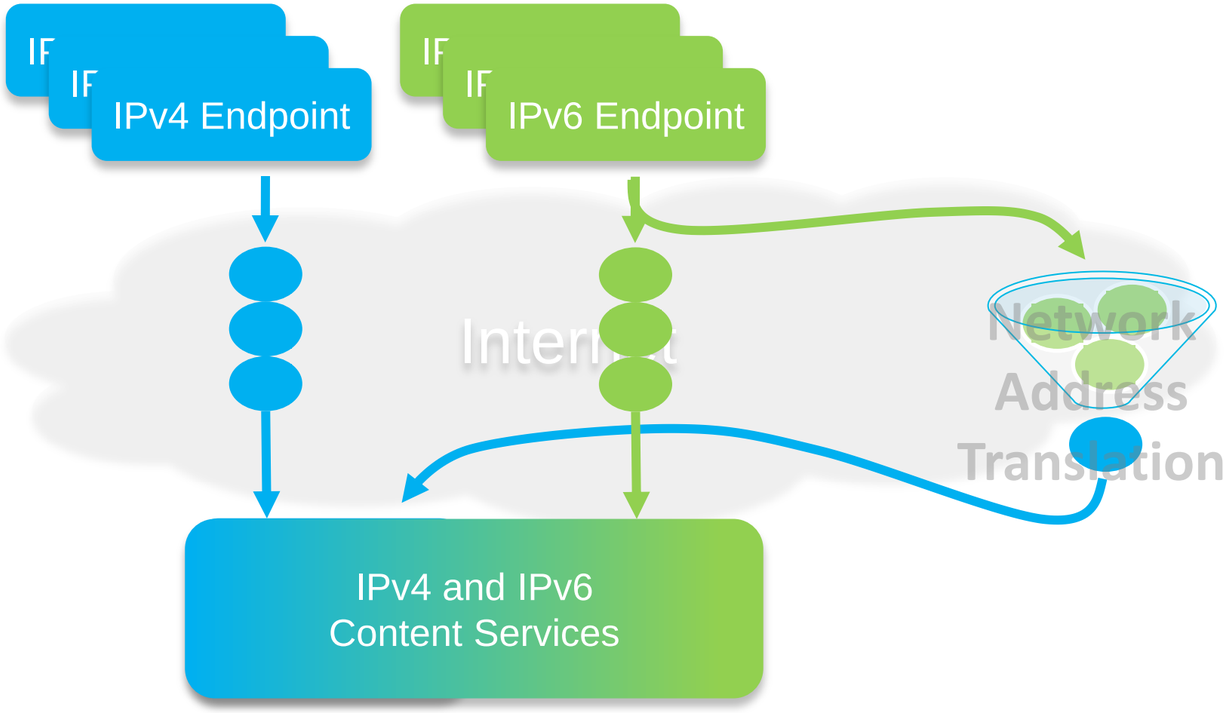

Common Deployment Models for Internet

Pure Dual Stack

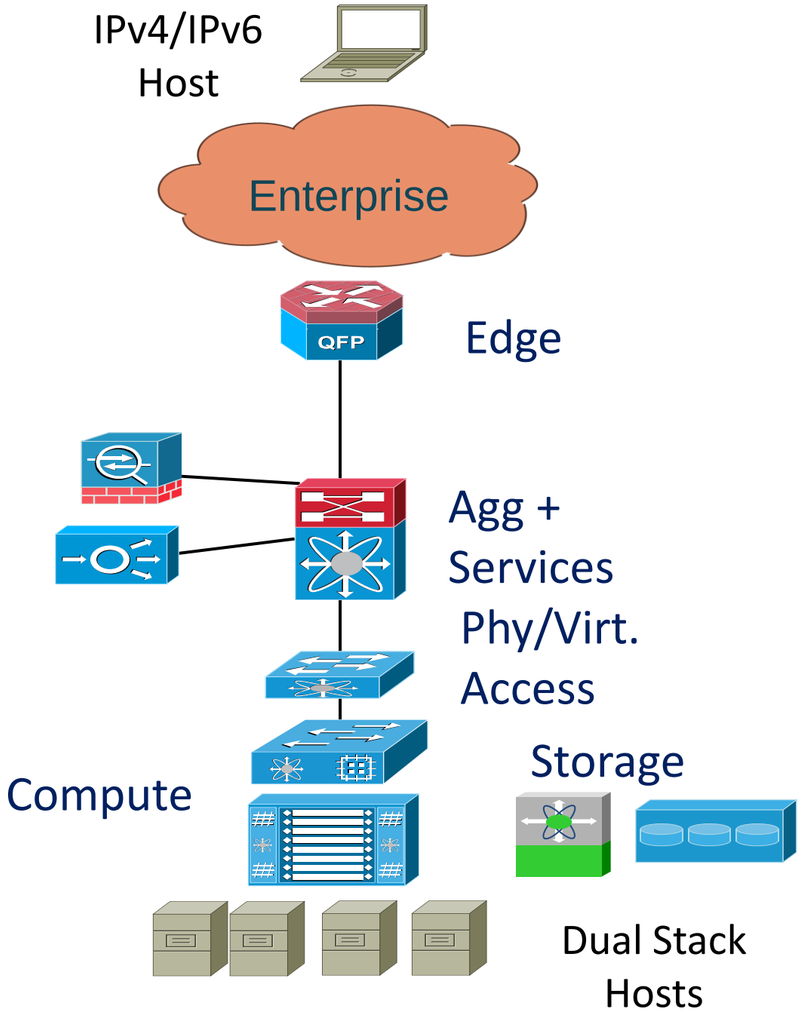

Conditional Dual Stack

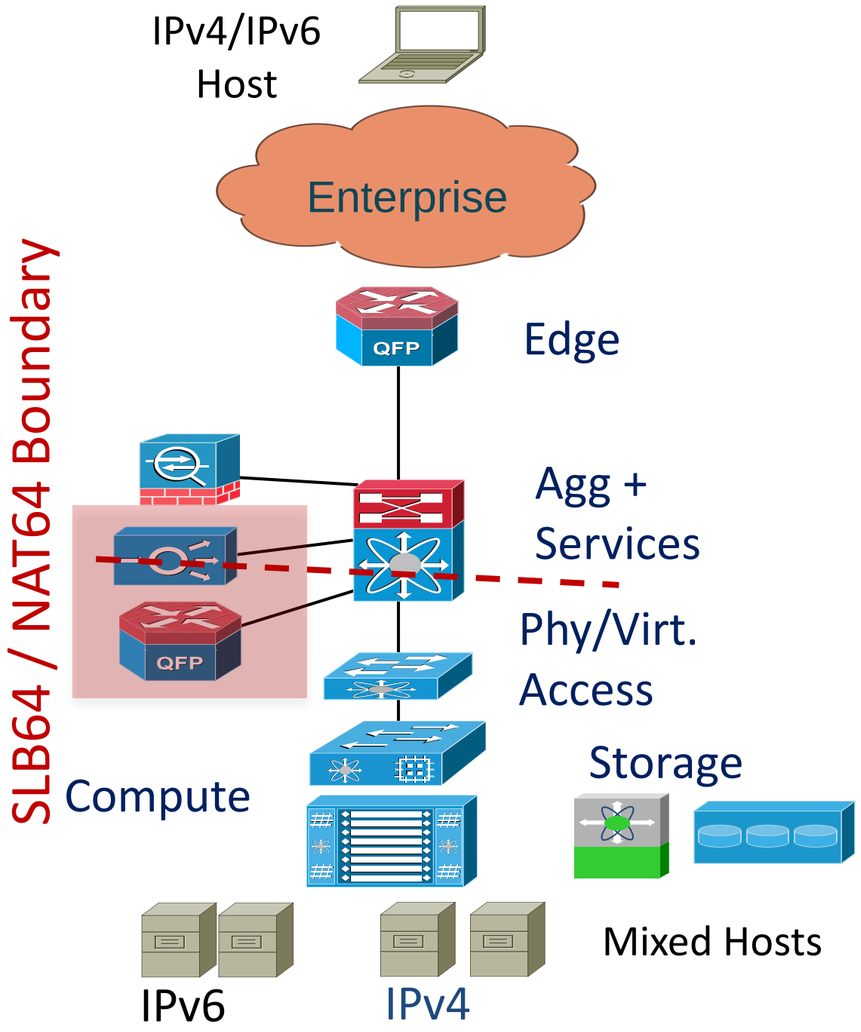

Translation as a Service

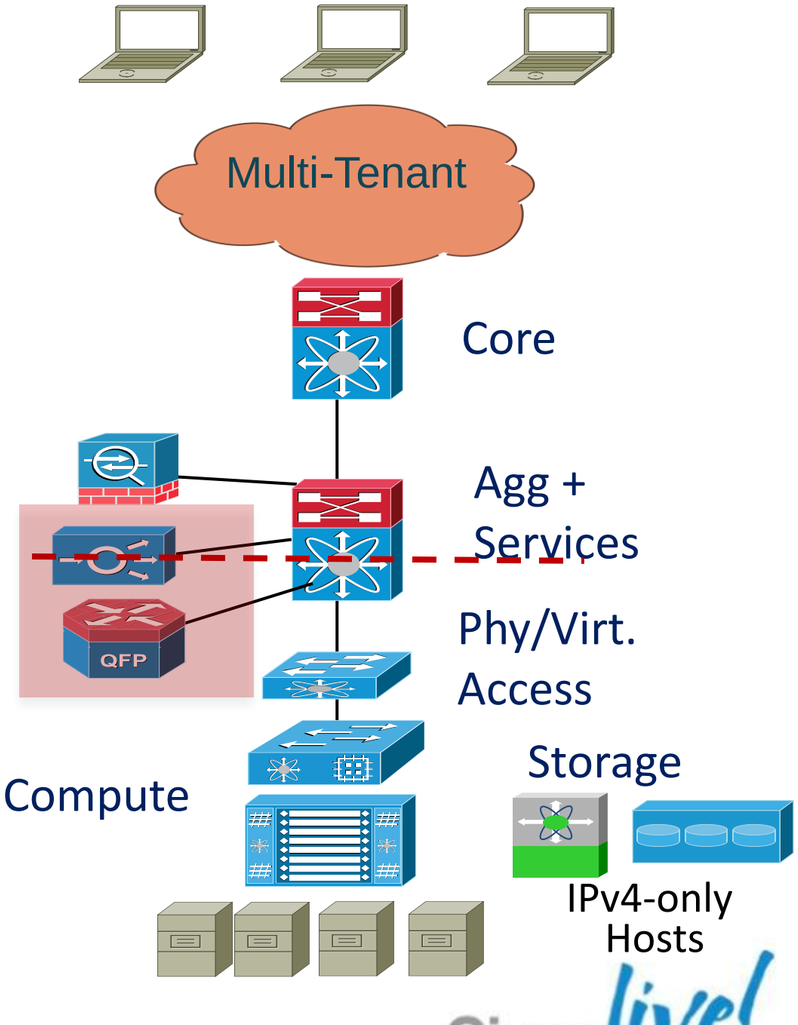

Co-Existence Considerations

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 Neighbor Cache entry:

```
2001:DB8:CAFE:2:2891:1C0C:F52A:9DF1       4       000d.6084.2c7a  STALE Vl2
2001:DB8:CAFE:2:7DE5:E2B0:D4DF:97EC       16      000d.6084.2c7a  STALE Vl2
FE80::7DE5:E2B0:D4DF:97EC                 16      000d.6084.2c7a  STALE Vl2
```

- 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!**

Scalability and Performance

- Full internet route table
  - 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 tunnels on platforms that use HW switching when attempting large scale deployments (hundreds/thousands of tunnels)

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

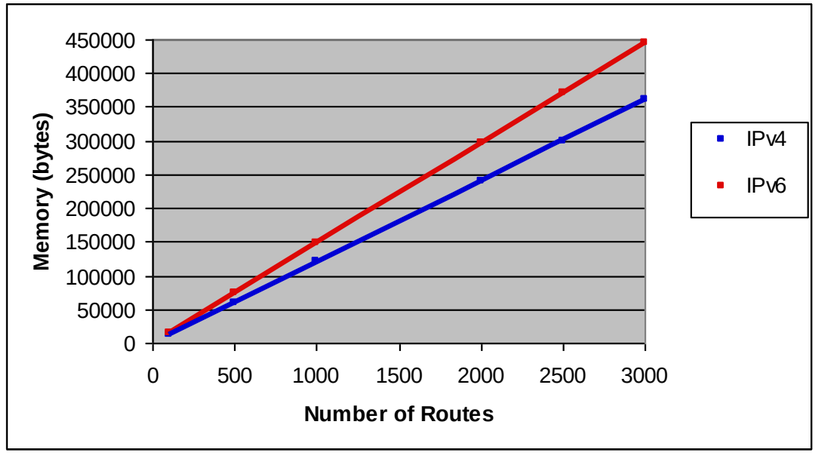

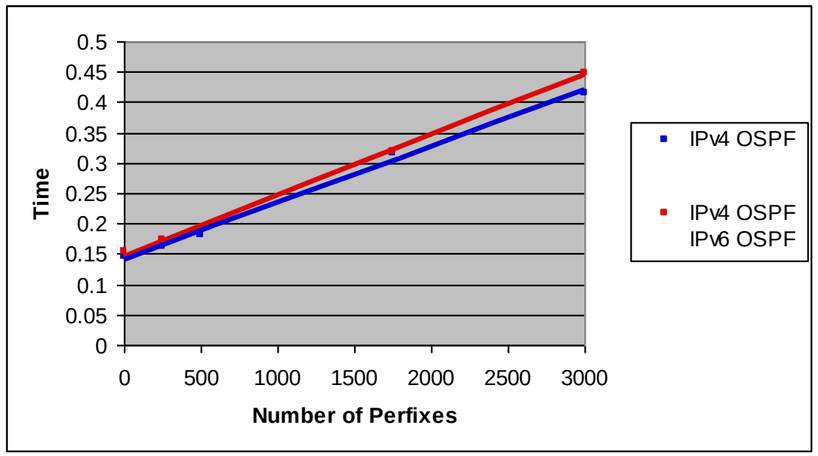

The Coexistence Twist

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

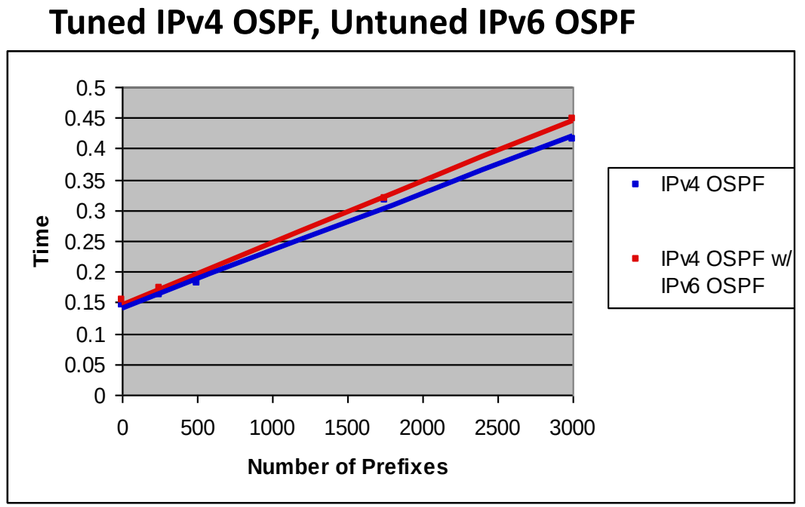

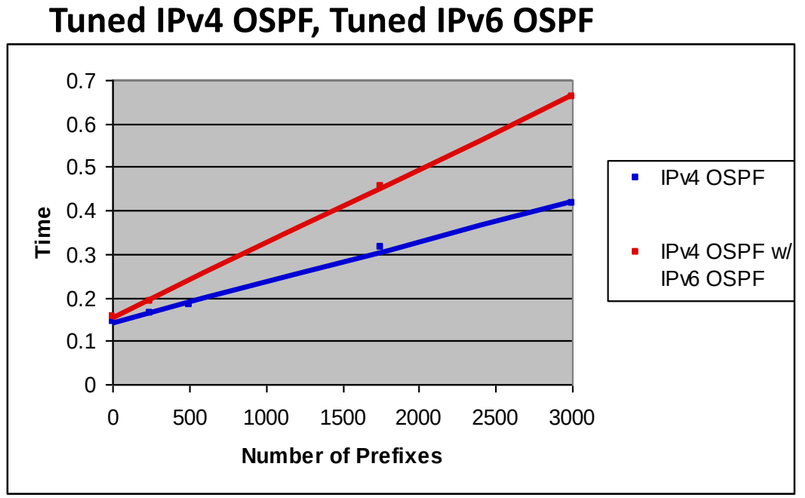

QoS Considerations

- IPv4 and IPv6 QoS features are mostly compatible (RFC 2460/3697)
- Both Transport uses DSCP (aka Traffic Class)
- Queues have to be properly sized
- Capacity has to be re-validated
- Control plane Queues need to now take into account IPv6 Overhead too
- Future use of Flow label will be a viable options if application differentiation required or source specific traffic characteristics

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
- **Metreing:** Flexible NetFlow
  - –Advantages: cache and export content flexibility
    - User selection of flow keys
    - User definition of the records

NBAR2

What can NBAR2 do for IPv6 Traffic ?

- NBAR2 allows network managers to detect native IPv6 traffic as well as IPv6 traffic encapsulated in IPv4 in their network, in order to apply QOS policies and to enable advanced IPv6 reporting.
- NBAR2 can detect IPv6 in IPv4 traffic
  - Support of ISATAP, 6to4, Teredo, Generic IPv6 in IPv4
  - Supported on ISR-G2 (15.1(4)M) and ASR1K (IOS XE 3.3.0S)
- Statefull Application classification for native IPv6 traffic
  - Supported on ISR-G2 (15.2(2)T) and ASR1K (IOS XE 3.5.0S)
- Statefull Application classification IPv6 in IPv4 Traffic
  - Supported on ISR-G2 (15.2(2)T) and ASR1K (IOS XE 3.5.0S)
- Advanced Integration with Flexible NetFlow IPv6
  - Supported on ISR-G2 (15.2(2)T) and ASR1K (IOS XE 3.5.0S)

NBAR2

Application Classification over ISATAP

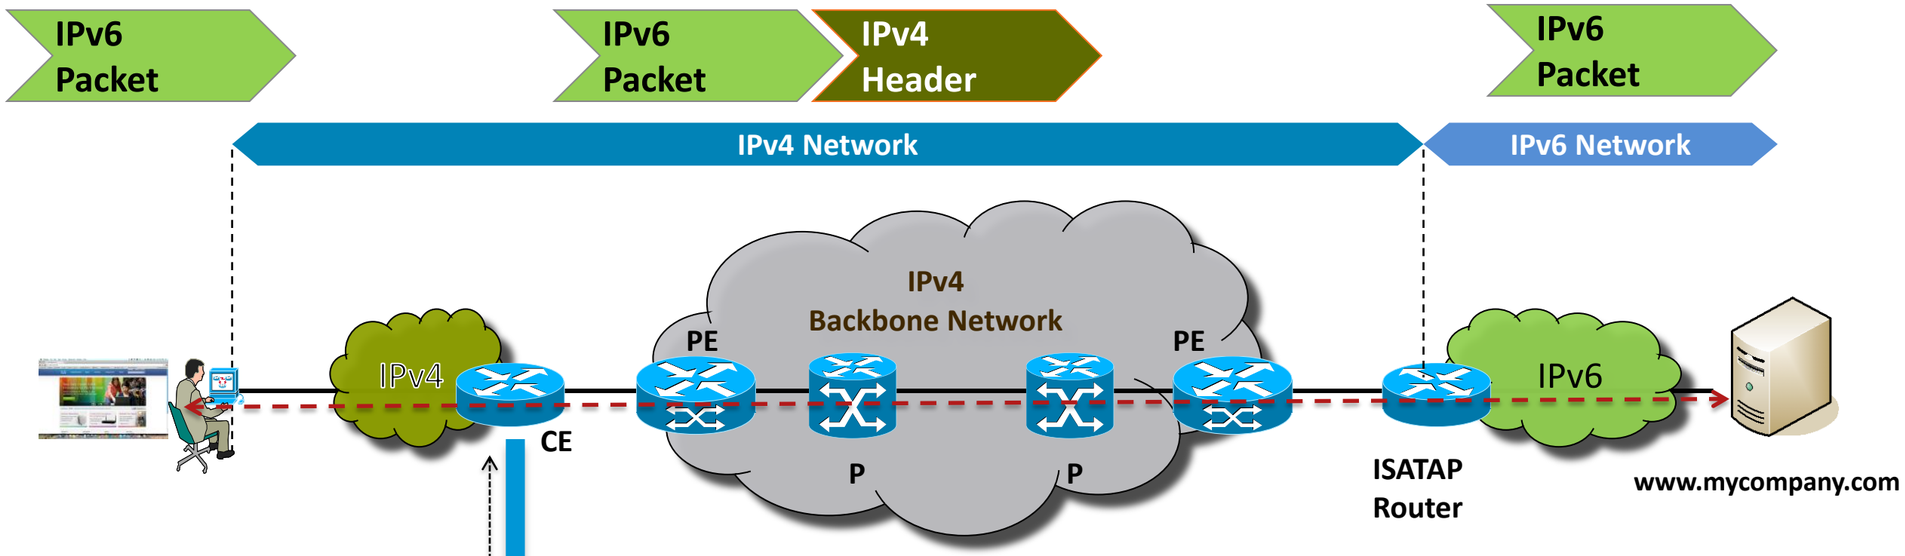

Router configured to classify within ISATAP tunnel

```
JCLabs06(config)# ip nbar classification tunneled-traffic ?
  6rd      Tunnel type 6RD
  6to4     Tunnel type 6TO4
  isatap   Tunnel type ISATAP
  teredo   Tunnel type TEREDO

JCLabs06(config)#ip nbar classification tunneled-traffic isatap
JCLabs06(config)#
```

```
JCLabs06#sh ip nbar protocol-discovery

 Last clearing of "show ip nbar protocol-discovery" counters 02:22:58

                                 Input                     Output
                                 -----                     ------
   Protocol                      Packet Count              Packet Count
                                 Byte Count                Byte Count
                                 5min Bit Rate (bps)       5min Bit Rate (bps)
                                 5min Max Bit Rate (bps)   5min Max Bit Rate (bps)
   ----------------------------- ------------------------- -------------------------
   http                           7184                      7302
                                 513776                    474146
                                 1000                      1000
                                 1000                      4000
```

IPSLA Operations Supported for IPv6

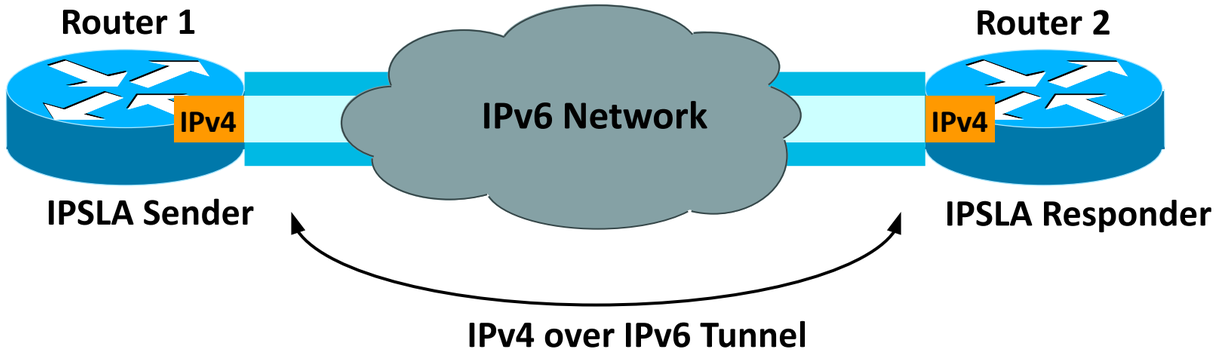

- Operations supported for IPv6:
  UDP-Jitter, UDP-Echo, ICMP Echo, TCP-Connect
- On:
  - 12.2(33)SB — C10K,C7200,C7300 Series
  - 12.2(33)XNA — ASR 1000 Series
  - 12.2(33)SRC — C7600 Series
  - 12.4(20)T — ISR Series
- For all other operations, use IPv4 tunnelling

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 DNS

Introduction to DNS and IPv6

- When you introduce IPv6, you will use both IPv4 and IPv6 addresses in your network
- Therefore, you 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)

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 TCP/53 and UDP/53 can be accessed through IPv6.
- Enable IPv6 connectivity to the external full-service resolvers that send DNS queries to authoritative servers in the world.

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.

AAAA Records on the Wire

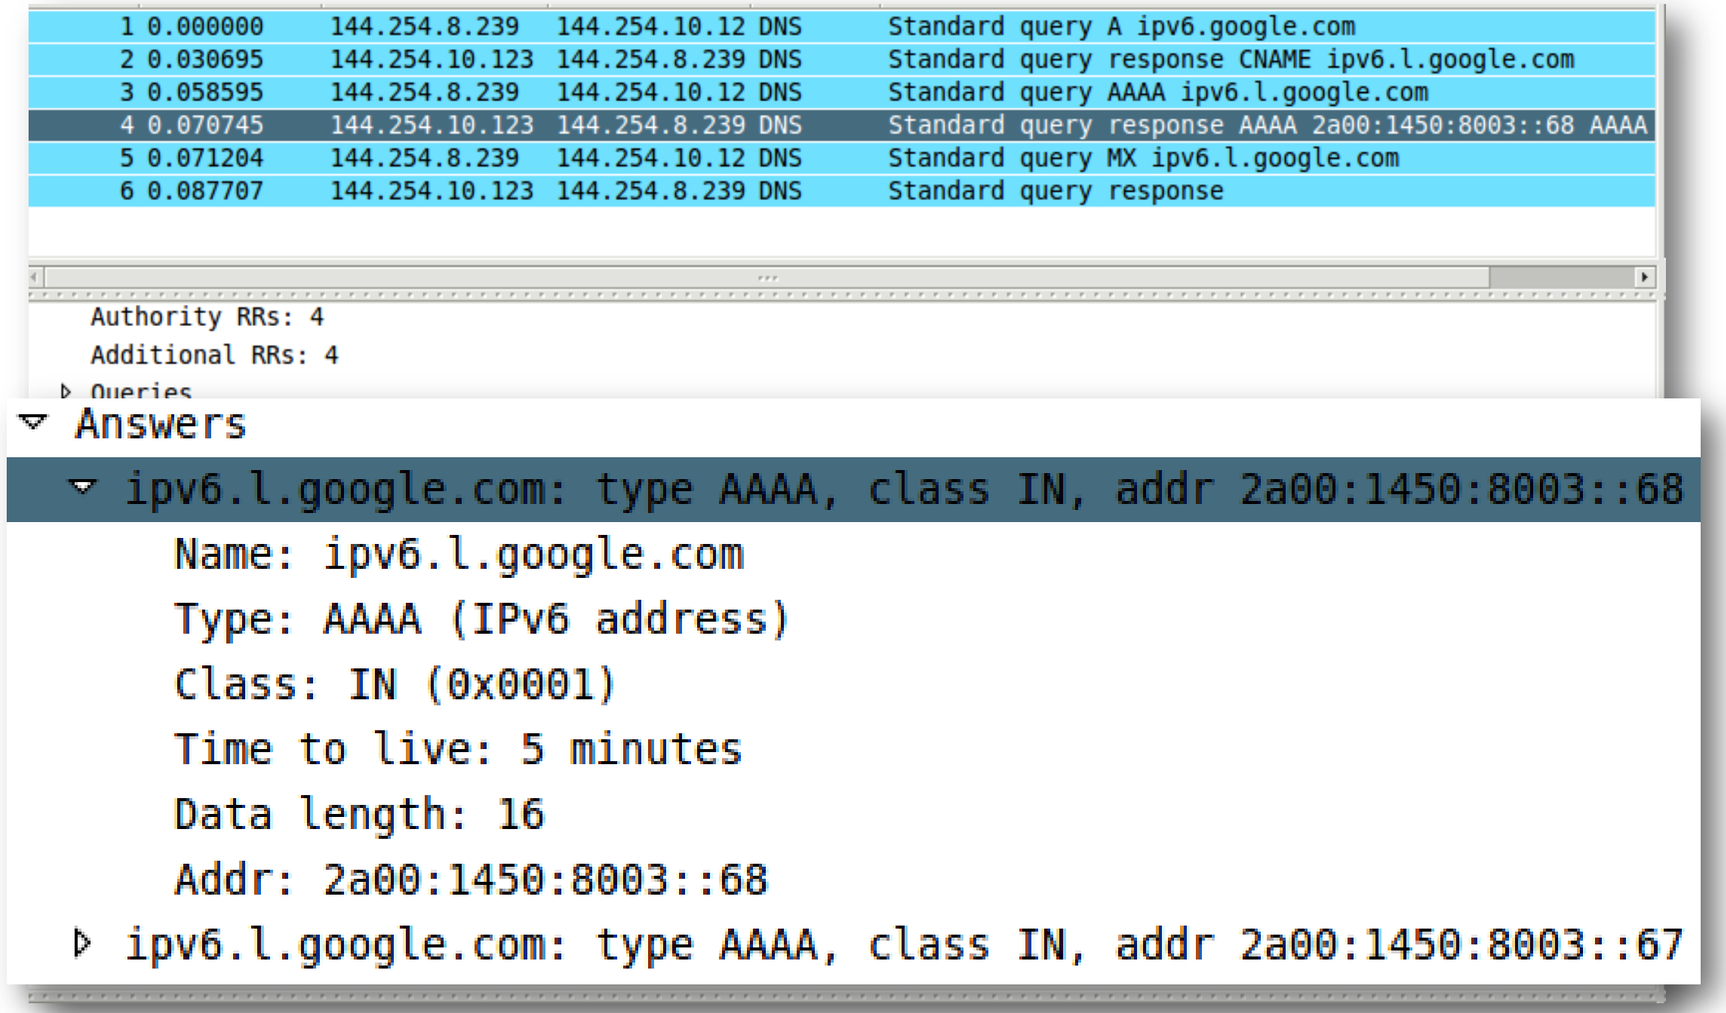

Test-IPv6.com – Verify Your Connectivity

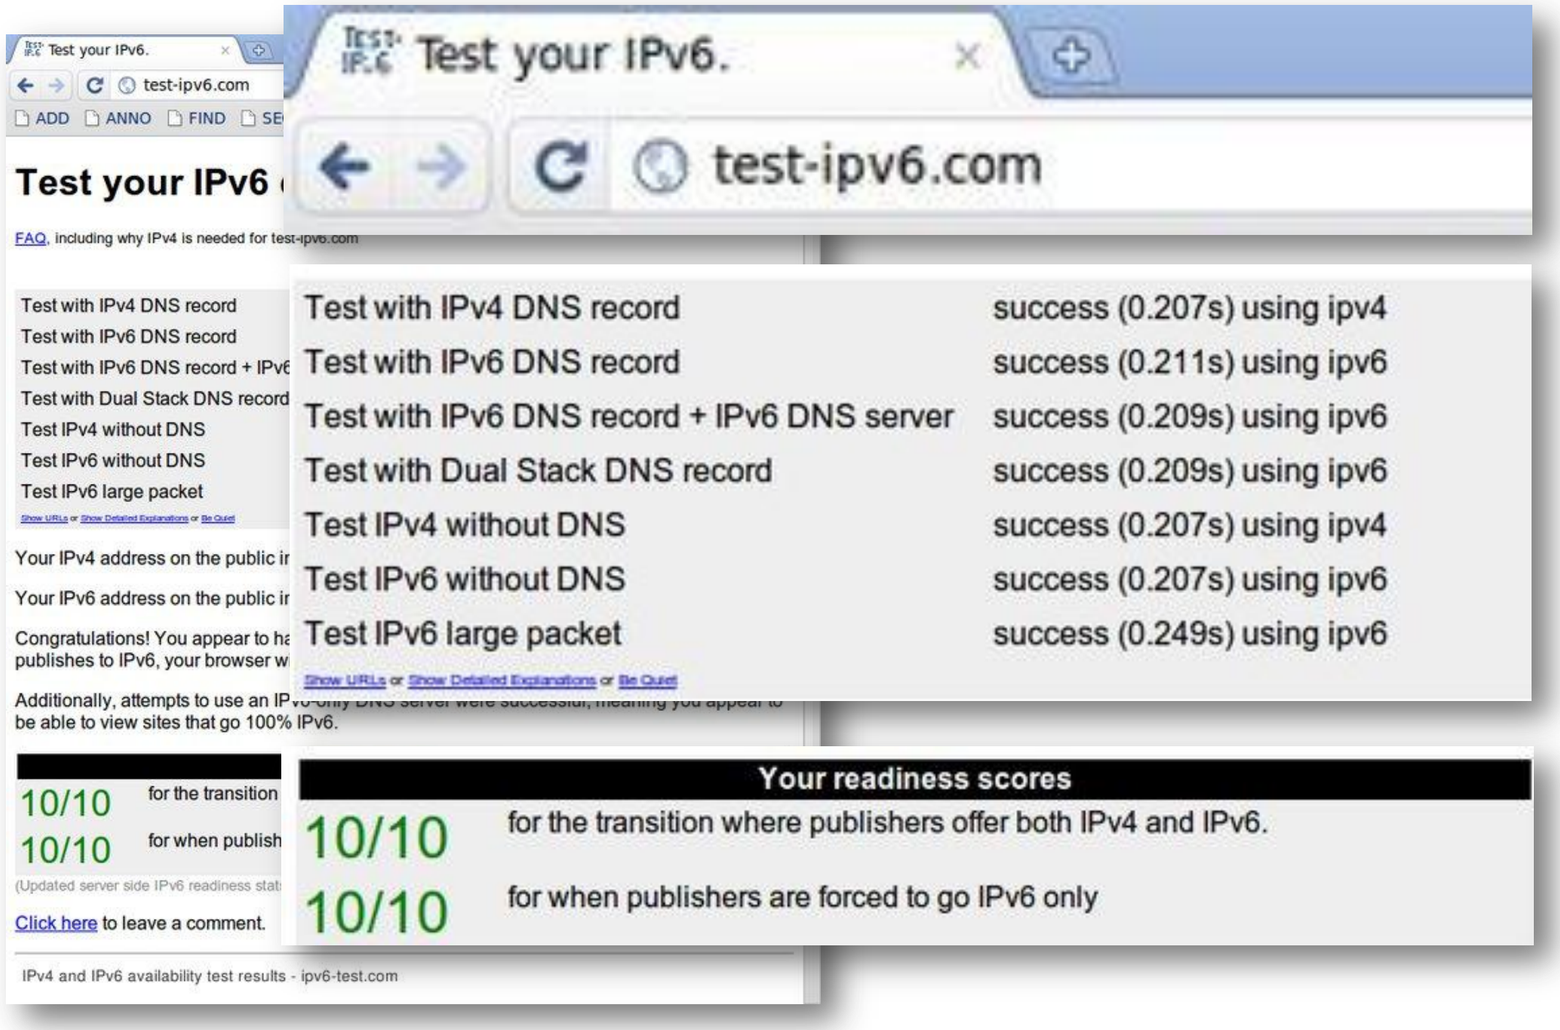

IPv6 Security

Security Considerations

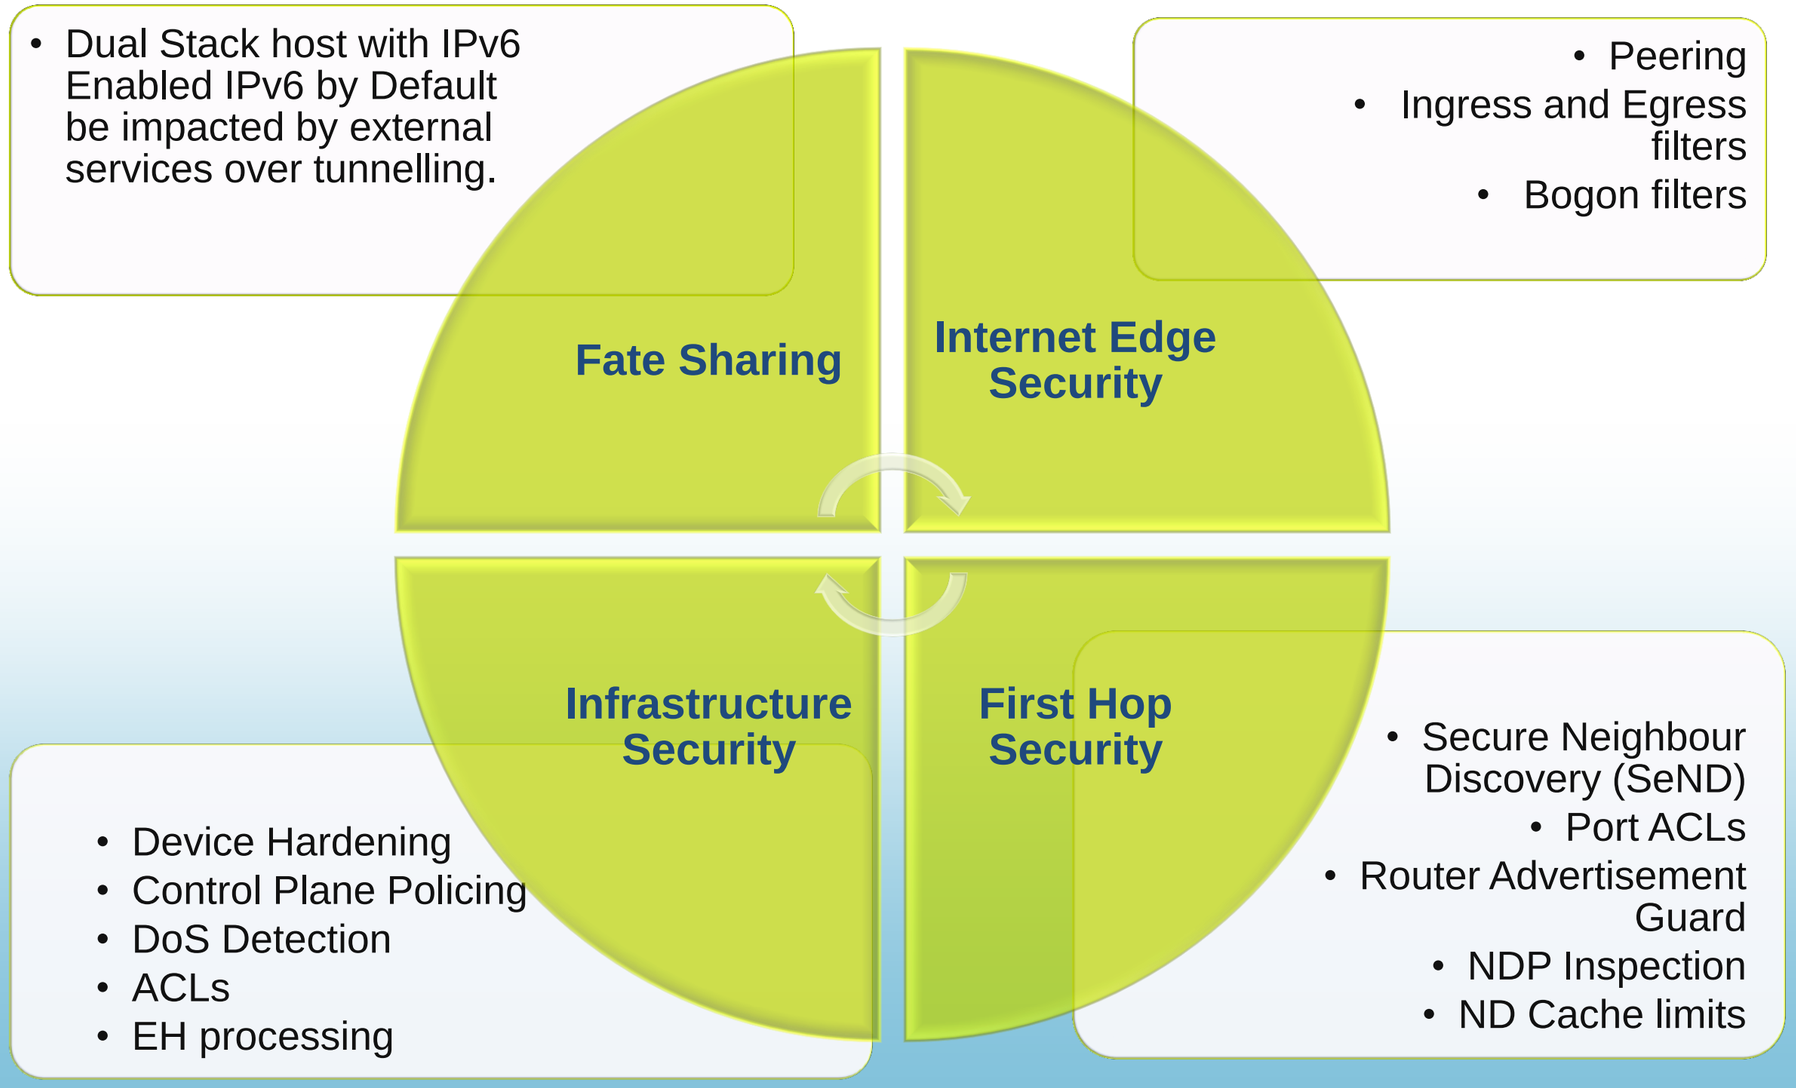

Dual Stack Host Considerations

- 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 IP versions
  - Host intrusion prevention, personal firewalls, VPN clients, etc.

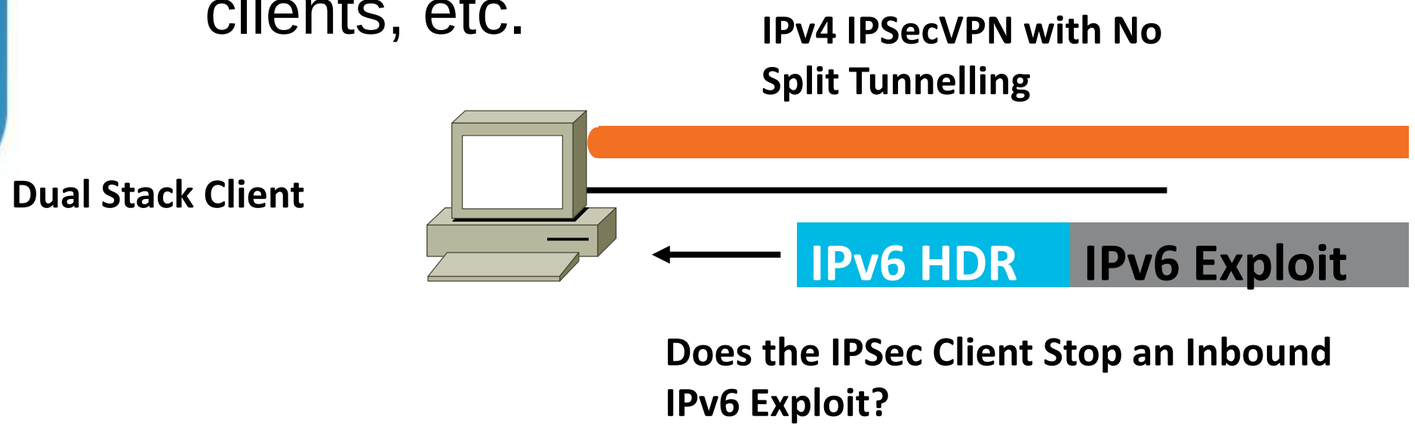

IPv6 Statistics Cisco Live London

IPv6 Enabled End-points: **75%!**

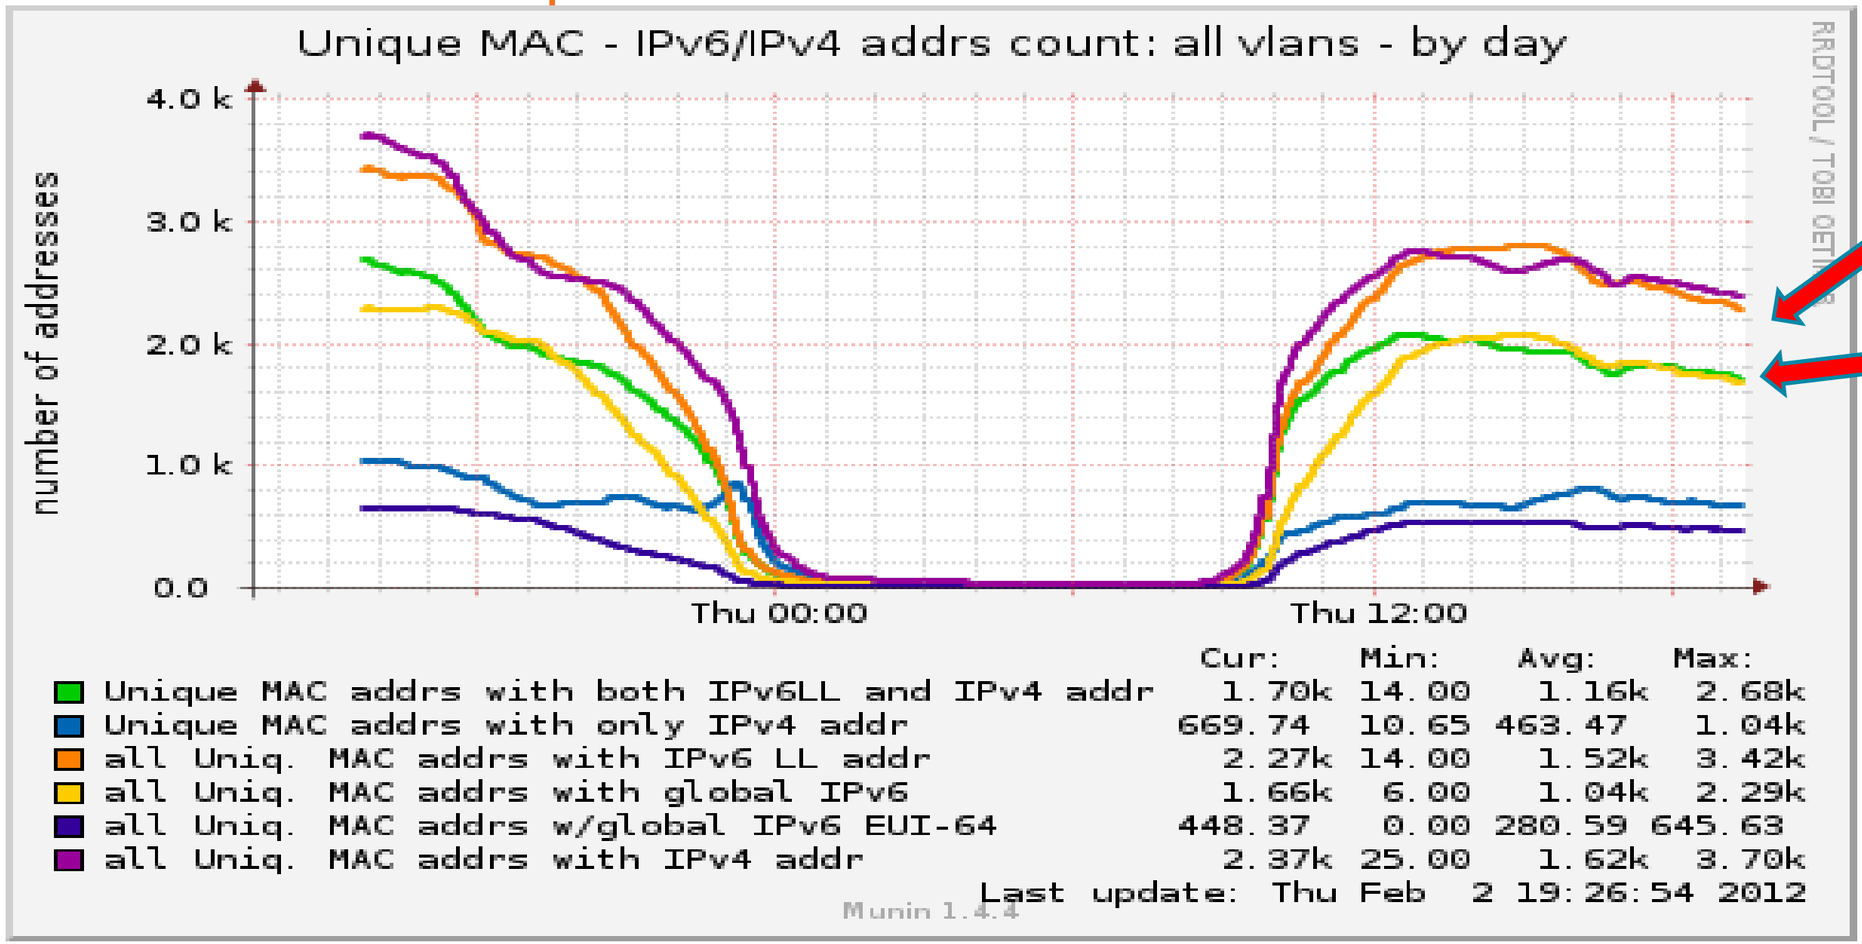

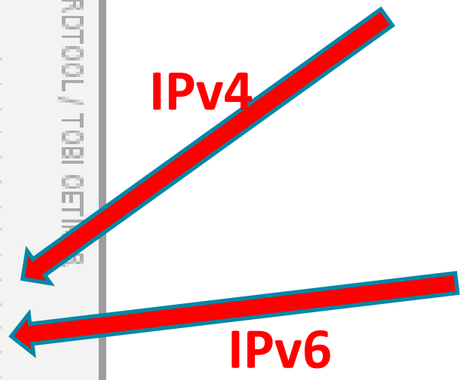

How many folks do you think new that had IPv6 ?

First Hop Security

- End-nodes are exposed to many threats:
  - Address configuration parameters : Trickery on configuration parameters
  - Address initialisation: Denial of address insertion
  - Address resolution: Address stealing
  - Default gateway discovery: Rogue routers
  - Local network configuration: Trickery on configuration parameters
  - Neighbour reachability tracking: Trickery on neighbour status

First Hop Security

Protecting Against Rogue RA

- Port ACL block all ICMPv6 Router Advertisements from hosts

```
interface FastEthernet3/13
  switchport mode access
  ipv6 traffic-filter DROP_RA in
  access-group mode prefer port
```

- RA-guard feature in host mode (12.2(33)SXI4 & 12.2(54)SG): also dropping all RA received on this port

```
interface FastEthernet3/13
  switchport mode access
  ipv6 nd raguard
  access-group mode prefer port
```

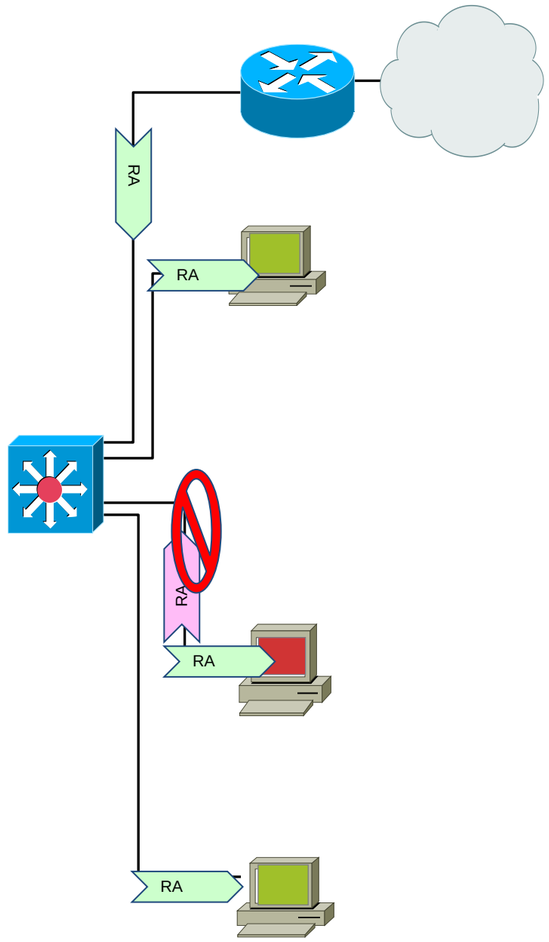

Perimeter Security

- ACLs at the edge
  - Packet Spoofing
  - ICMPv6
  - Handling extension headers
- Deep Packet Inspection
  - How to detect IPv6 packets - Native/Tunnelled
  - IPS/IDS
  - NBAR2
- Instrumentation
  - Netflow
  - Syslog
  - EEM

85

General Considerations

- 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
  - Host ACL entry may not be appropriate – multiple addresses per interface, privacy
- Selective ICMPv6 pass through via firewalls and access-lists
- Education is key!
- Security best practice links

  http://www.cisco.com/web/about/security/security_services/ciag/documents/v6-v4-threats.pdf

Action Plan

- 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
- Education

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

Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 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 8 March 12:00pm-2:00pm

Cisco live! 365

Don’t forget to activate your Cisco Live 365 account for access to all session material, communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.
www.ciscoliveaustralia.com/portal/login.ww

Cisco live!