IPv6 Deployment

Ali Mummar

amummar@cisco.com
Systems Engineer – SP Group

December, 2011

All information presented herein this session are considered sole proprietary of Cisco, for the purpose of providing technology awareness and update, It shall not be used in any way in production networks without appropriate testing in lab environment.

Please engage the local Cisco account team for support.

Please, don't forget to fill the evaluation sheet, there is a raffle withdraw that may gets you lucky and WIN a prize !!!

Reference Materials

- Deploying IPv6 in Campus Networks: http://www.cisco.com/en/US/docs/solutions/Enterprise/Campus/CampIPv6.html
- Deploying IPv6 in Branch Networks: http://www.cisco.com/en/US/solutions/ns340/ns414/ns742/ns816/landing_br_ipv6.html
- SRND: Deploying IPv6 in Unified Communications Networks
  - http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/ipv6/ipv6srnd.html
- IOS IPv6 VOIP implementation Guide
  - http://www.cisco.com/en/US/docs/ios/ipv6/configuration/guide/ip6_voip.pdf
- DNS and BIND, 5th Edition, by Cricket Liu and Paul Albitz, O'Reilly Media, May 2006
- RFC 3596: DNS Extensions to Support IP Version 6, by S. Thomson, C. Huitema, V. Ksinant, and M. Souissi, October 2003 (format: TXT=14093 bytes)(obsoletes RFC 3152 and RFC 1886) (status: Draft Standard)

Recommended Reading

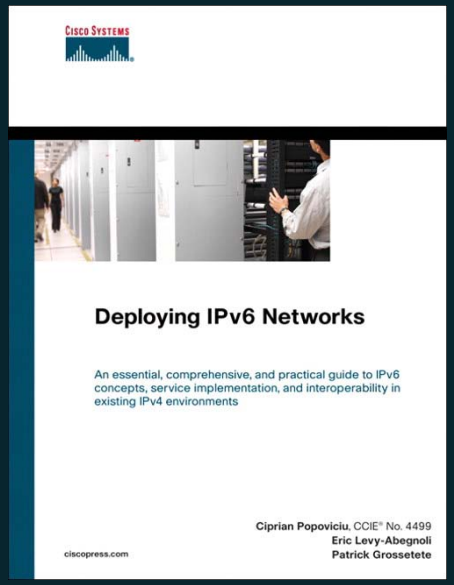

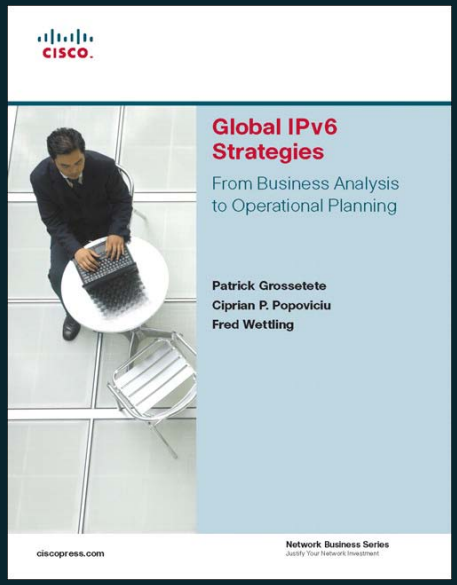

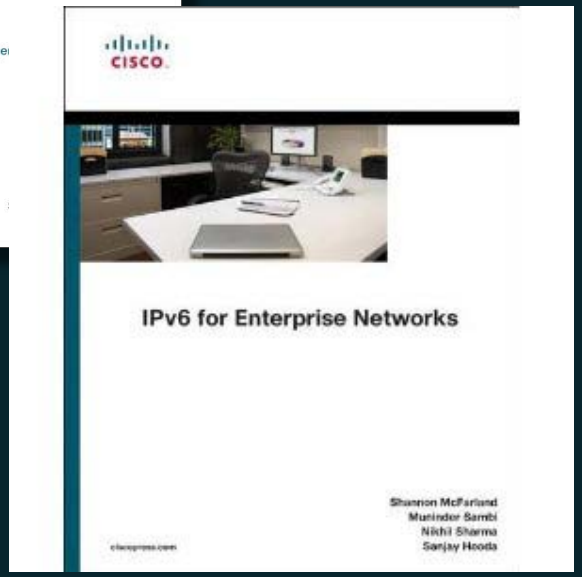

Recommended Reading

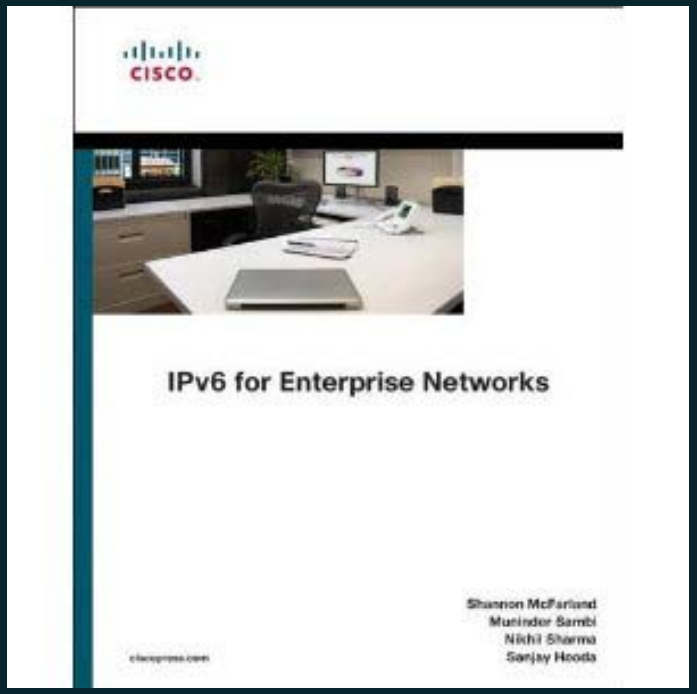

New

Deploying IPv6 in Broadband Networks
Adeel Ahmed, Salman Asadullah
ISBN0470193387,
John Wiley & Sons Publications®

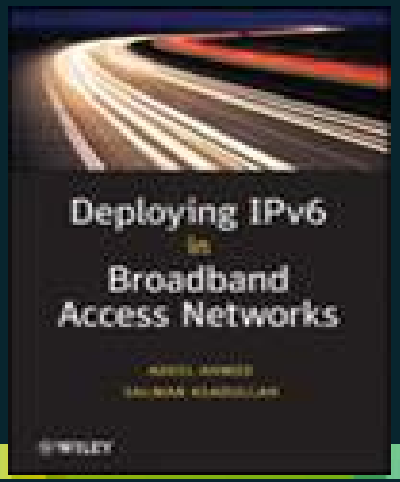

IPv6 Lab Material

- Login to my WebEx Meeting Room and get the Lab document

http://cisco.webex.com/meet/amummar

Folder Name: Files -> CiscoExpo_2011

File Name: IPv6_Labs_v1.2.p...

Password: "*Expo2011*" (Case sensitive)

File download expires on December 15, 2011

Agenda

IPv6 Top-On-Mind Questions

IPv6 Deployments

IPv6 Business Impact

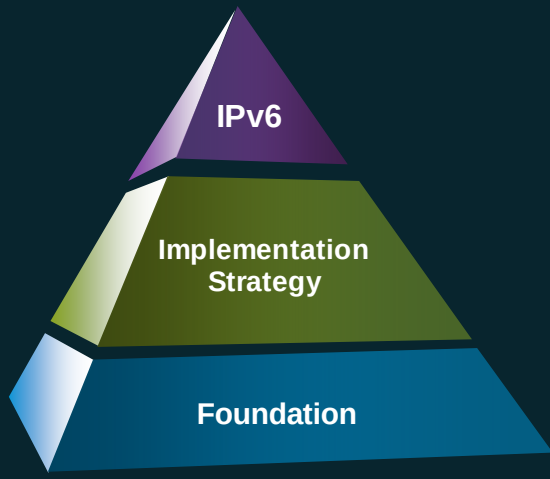

- How soon IPv6 transition will happen ?
- Where we are today ?
- What happens if we don't deploy IPv6 ?
- Can we continue using IPv4 and rely on mechanisms like NAT ?
- Are there any IPv6 internet application today ?

How soon IPv6 transition will happen ?

A Major Transition Is Coming

IPv4 Address Exhaustion Has Happened

2010 — 2011 — 2012

NOVEMBER, 2010

Globalization: 25% of the world's population using 100% of IPv4 addresses

JAN, 2011

Date the last IPv4 addresses was allocated

SEPTEMBER, 2012

Civilian US Government Agencies mandated to provide external IPv6 connectivity

Where are we today ?

We are here today - 2011

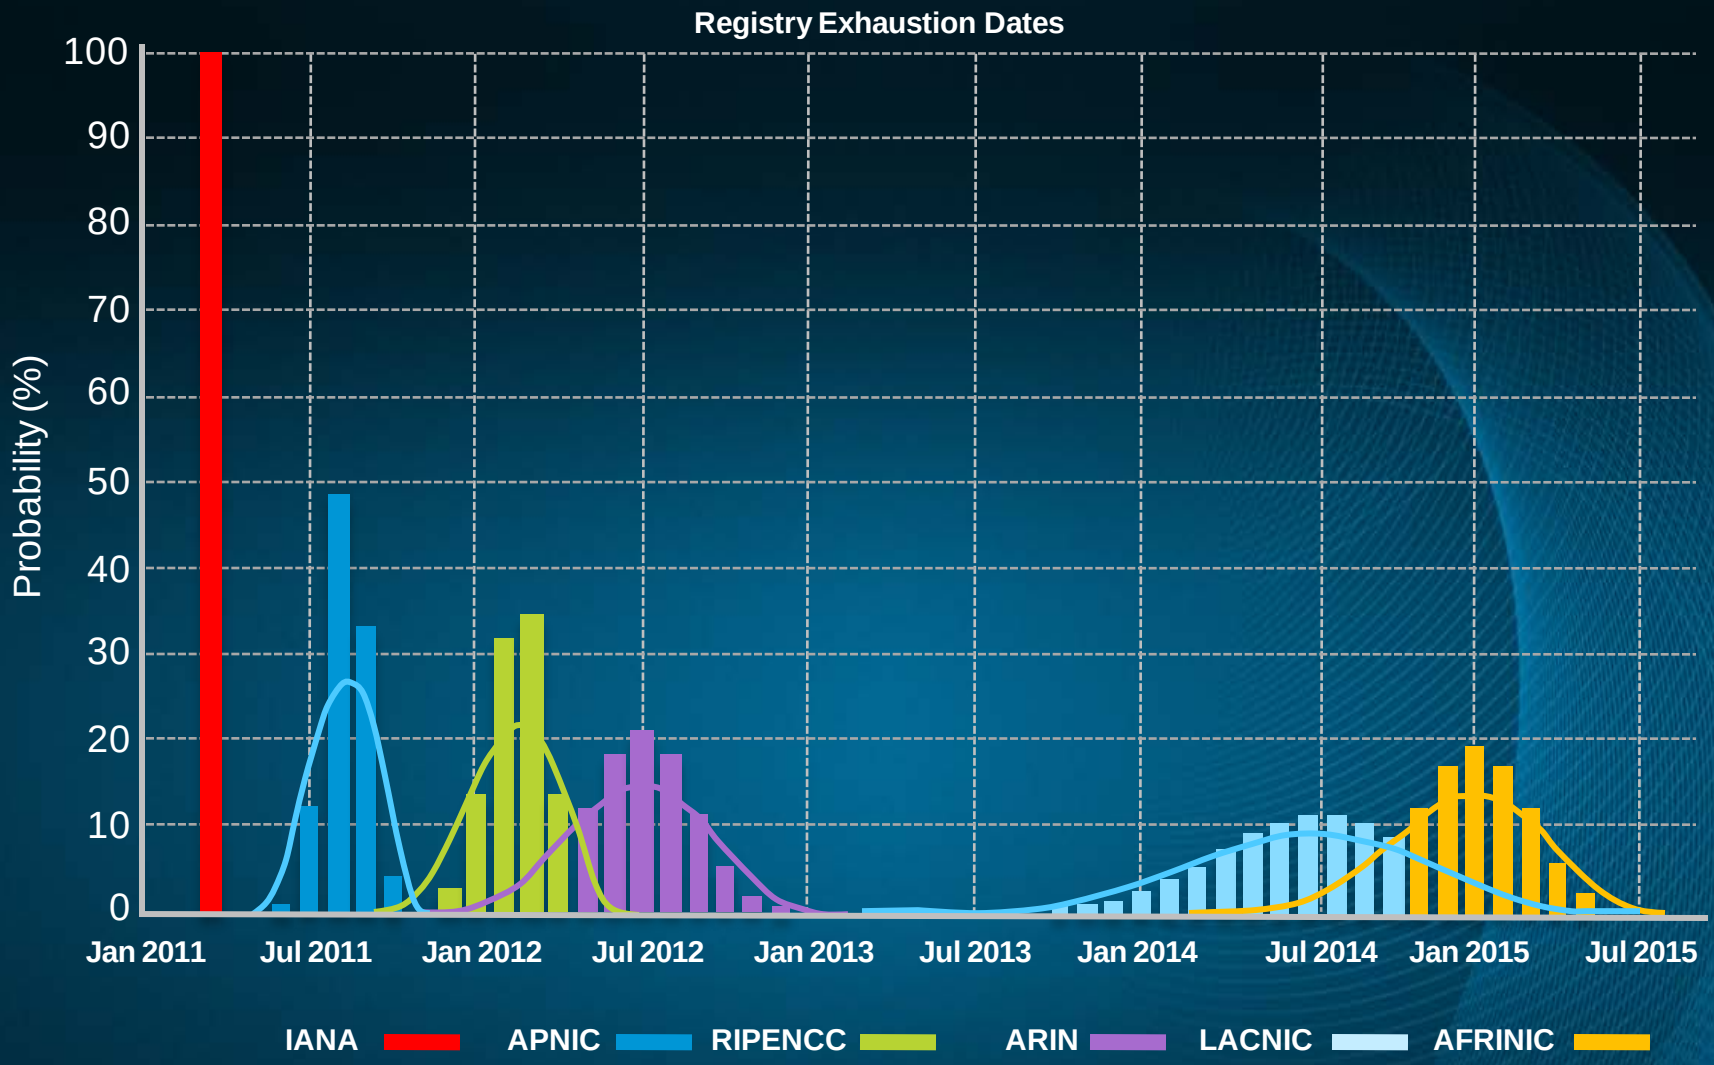

Source: Geoff Huston, APNIC

What happens if we don't deploy IPv6 ?

Failure to Act Will Impact Business

IPv6 Estimated Adoption Timeframes

- 2010: Low Impact – Buying behavior shift limited to mandated and early adopter sites

2011: Internet Evolution begins – "*…IPv6 is important to all of us (…) to everyone around the world, It is crucial to our ability to tie together everyone and every device". John Chambers*

•2012: Mandates take effect – Transition to IPv6 forces us to acquire product or managed services to sustain business and customer reach

- 2014: IPv6 is mainstream – customers without transition infrastructure experience reduced service levels, diminished customer reach, increase operational complexity

IPv6 Business Impact – The Cost of Waiting Goes Up

Low Risk | Moderate Risk | High Risk

Can we continue using IPv4 and rely on mechanisms like NAT ?

Myth: We'll Never Run Out of IP Addresses

Mobile and the Internet of Things drive growth

In 2013….There Will Be

50 BILLION

Devices Connected to the Network,

up from 35 BILLION in 2010

Source: Forrester, Cisco IBSG

Are there any IPv6 internet application today ?

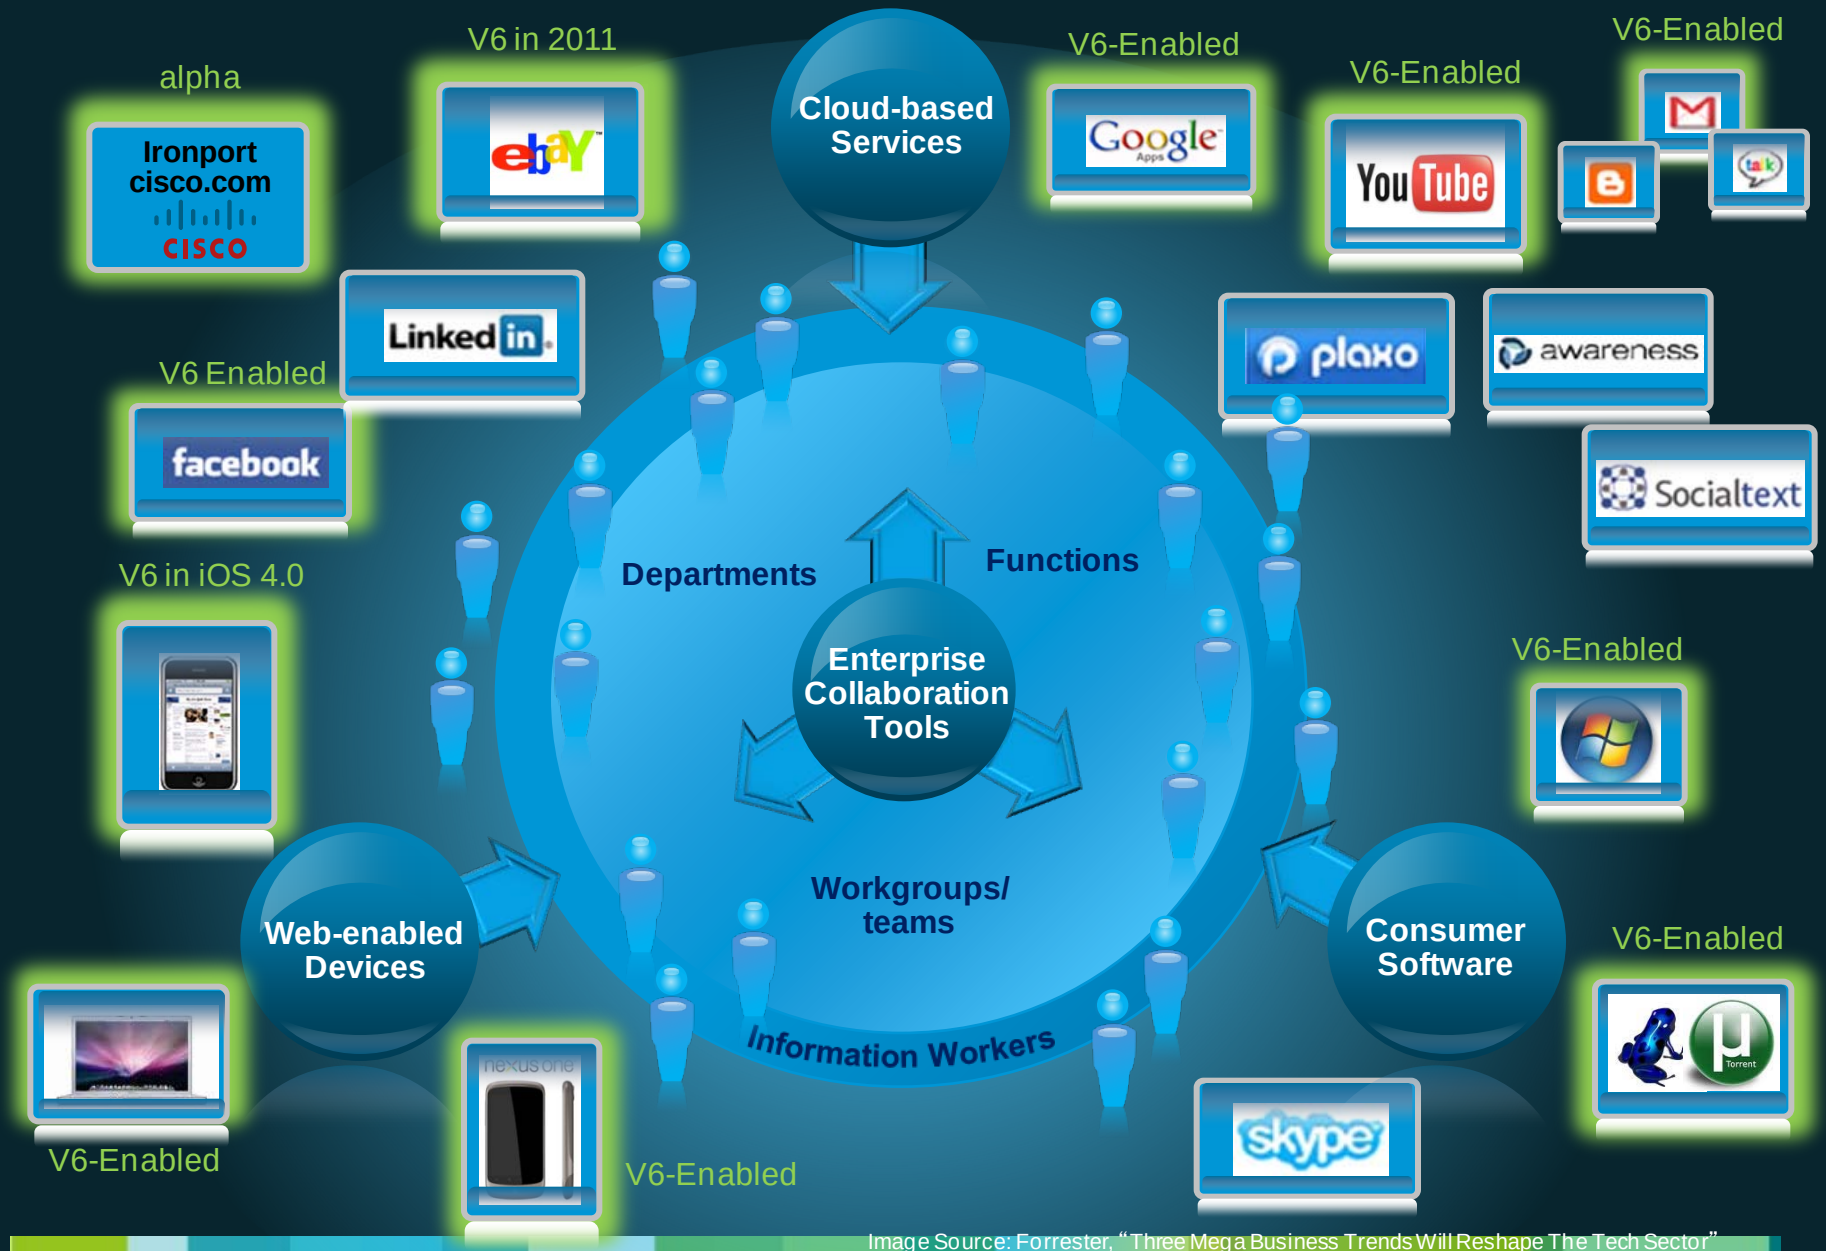

Image Source: Forrester, "Three Mega Business Trends Will Reshape The Tech Sector"

Where do I start from ?

Managing an Orderly IPv6 Transition

IPv6 Is Not a Rip-and-Replace Proposition

Preserve

Preserve the customer's existing investment
- Audit and leverage existing IPv6 capabilities

Prepare

Prepare a migration and deployment plan
- Identify and enable critical IPv6 functional areas

Prosper

Prosper through the transition to IPv6 Internet
- Enable all systems with dual-stack capabilities
- Grow seamlessly as customers transition to IPv6

IPv6 is the foundation of a lifecycle management discussion

Preserve

Audit the Customer's Existing IPv6 Capabilities

Over a Decade of Cisco Investment - Shipping Since 1996

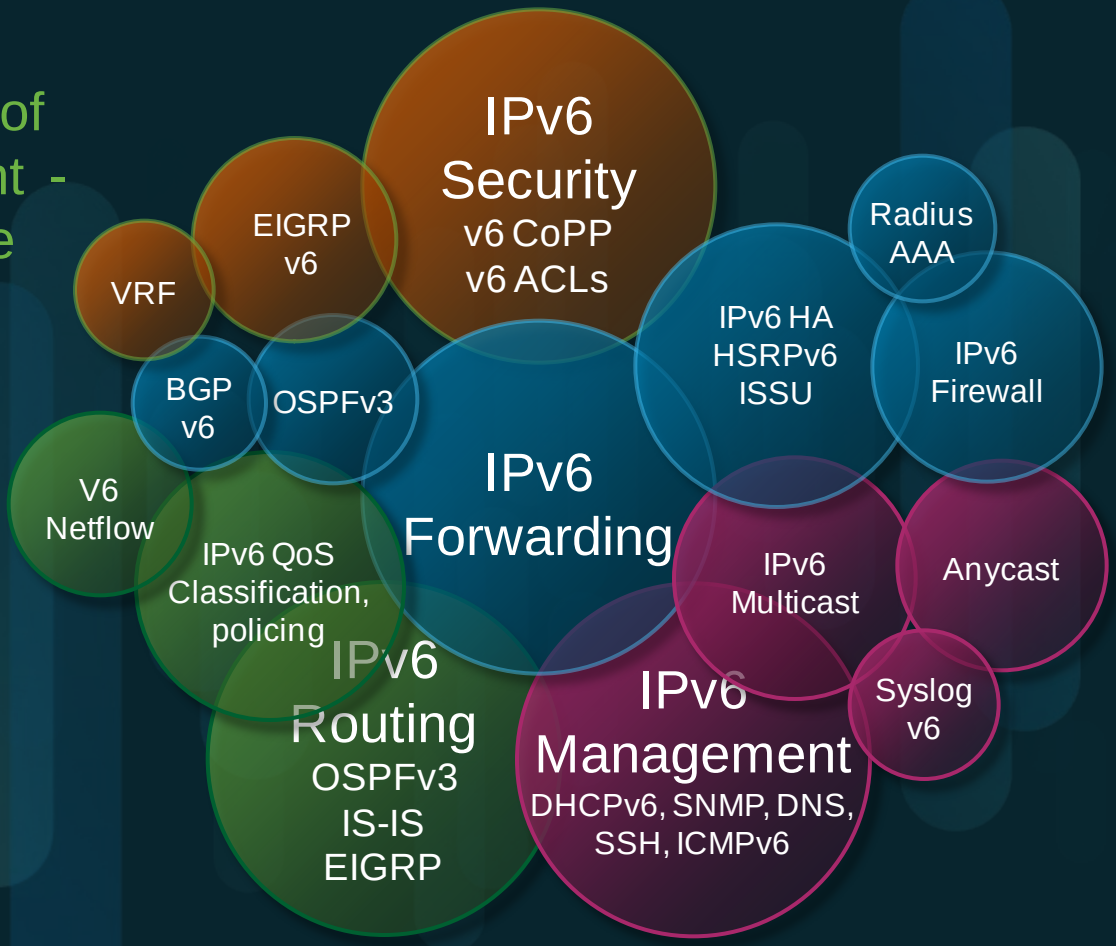

These capabilities and more are already part of your customer's investment

Prepare

Start with a Phased Plan Aligned with Your Business Strategy

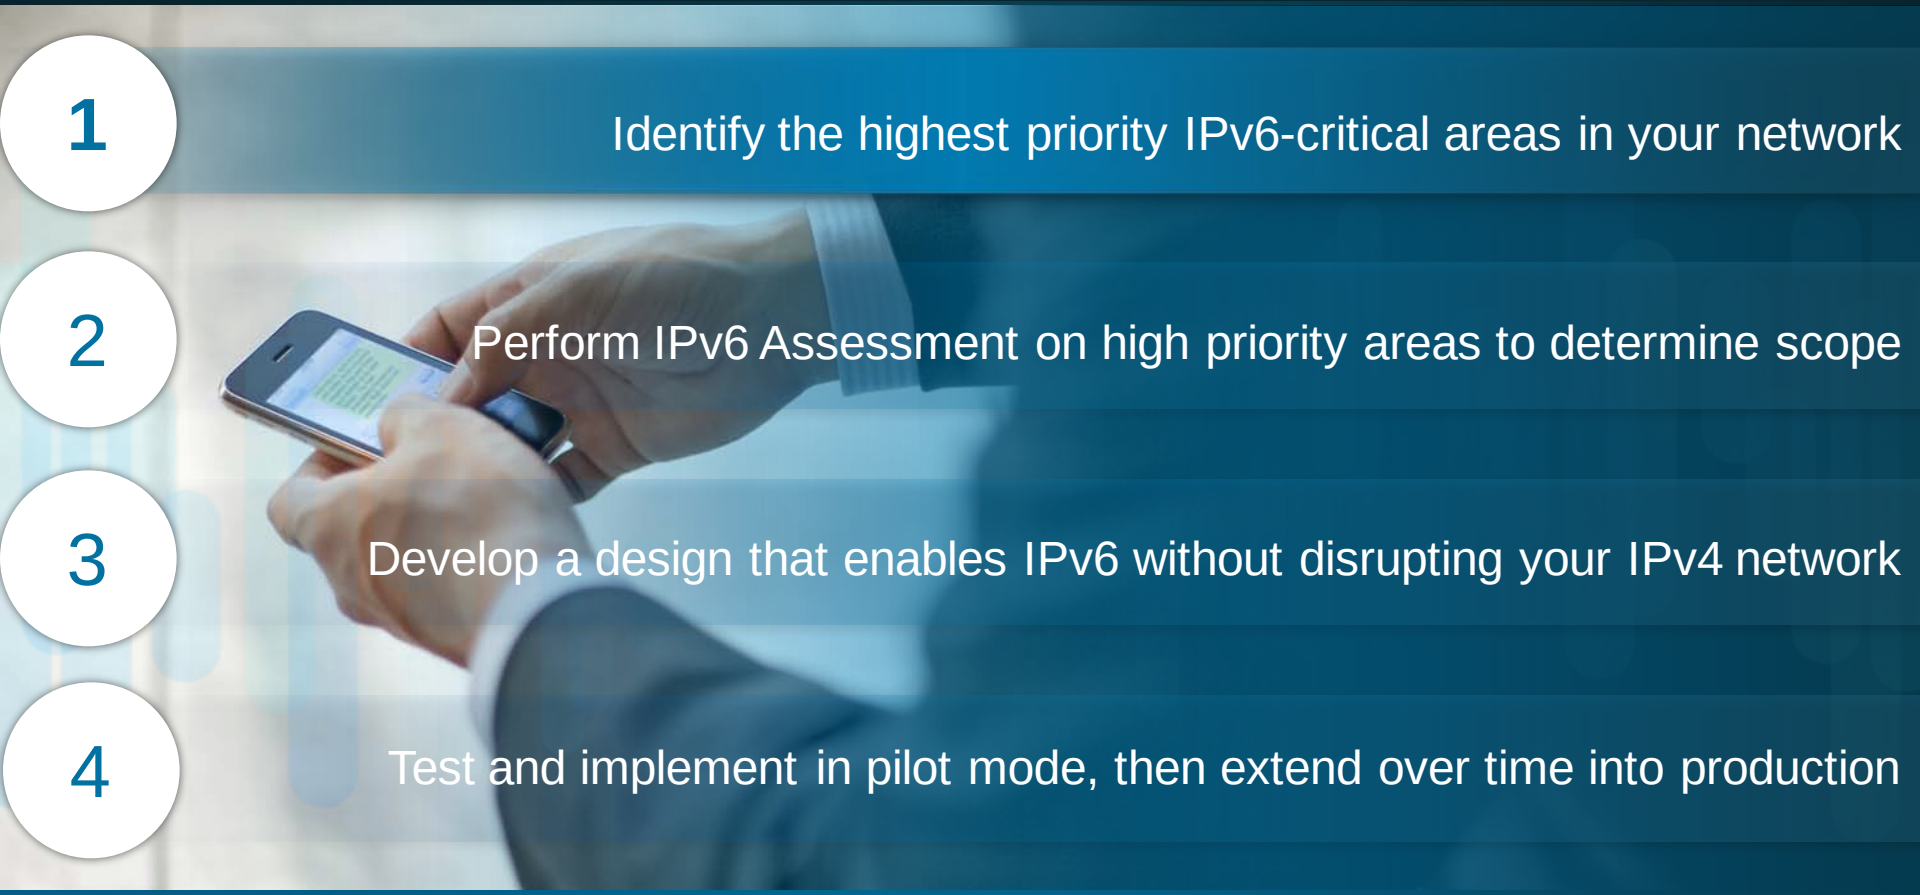

1. Identify the highest priority IPv6-critical areas in your network
2. Perform IPv6 Assessment on high priority areas to determine scope
3. Develop a design that enables IPv6 without disrupting your IPv4 network
4. Test and implement in pilot mode, then extend over time into production

Repeat for the Next IPv6-Critical Area in Your Network

Prosper

Grow Seamlessly as Customers Transition to the IPv6 Internet

A well-structured migration plan provides insurance against unexpected costs as customers, partners, and suppliers move to IPv4 and IPv6 coexistence

Enterprise IPv6 Adoption and Delivery

1. Sales Certs (USGv6, JITC UCR2008)
2. IPv6 Pilot and Basic Infrastructure
3. IPv6 Internet Presence (websites, remote users, B2B …)
4. IPv6 Islands (Wireless/Consumer devices, Labs …)
5. Internal Data Center, Enterprise Apps
6. Ubiquitous Dual-Stack
7. IPv4 EOL

"Mandated"
1, 2, 3

Who?

- Government Agencies
- Customers who sell to government agencies

"Motivated"
2 3 4

Who?

- Customers with IPv4 address exhaustion
- Global Enterprises with consumer or business interaction on the public internet
- Customers with user-provided devices on their networks

"Early Adopter"
2 4 3 5 6 7

Who?

- Companies looking for competitive advantage
- Companies using IPv6 to solve business problems
- Early adopters preparing for coexistence

"Mainstream"
2

Who?

- Large US/European Enterprises
- Small-Medium Enterprises

IPv6 Adoption Spectrum

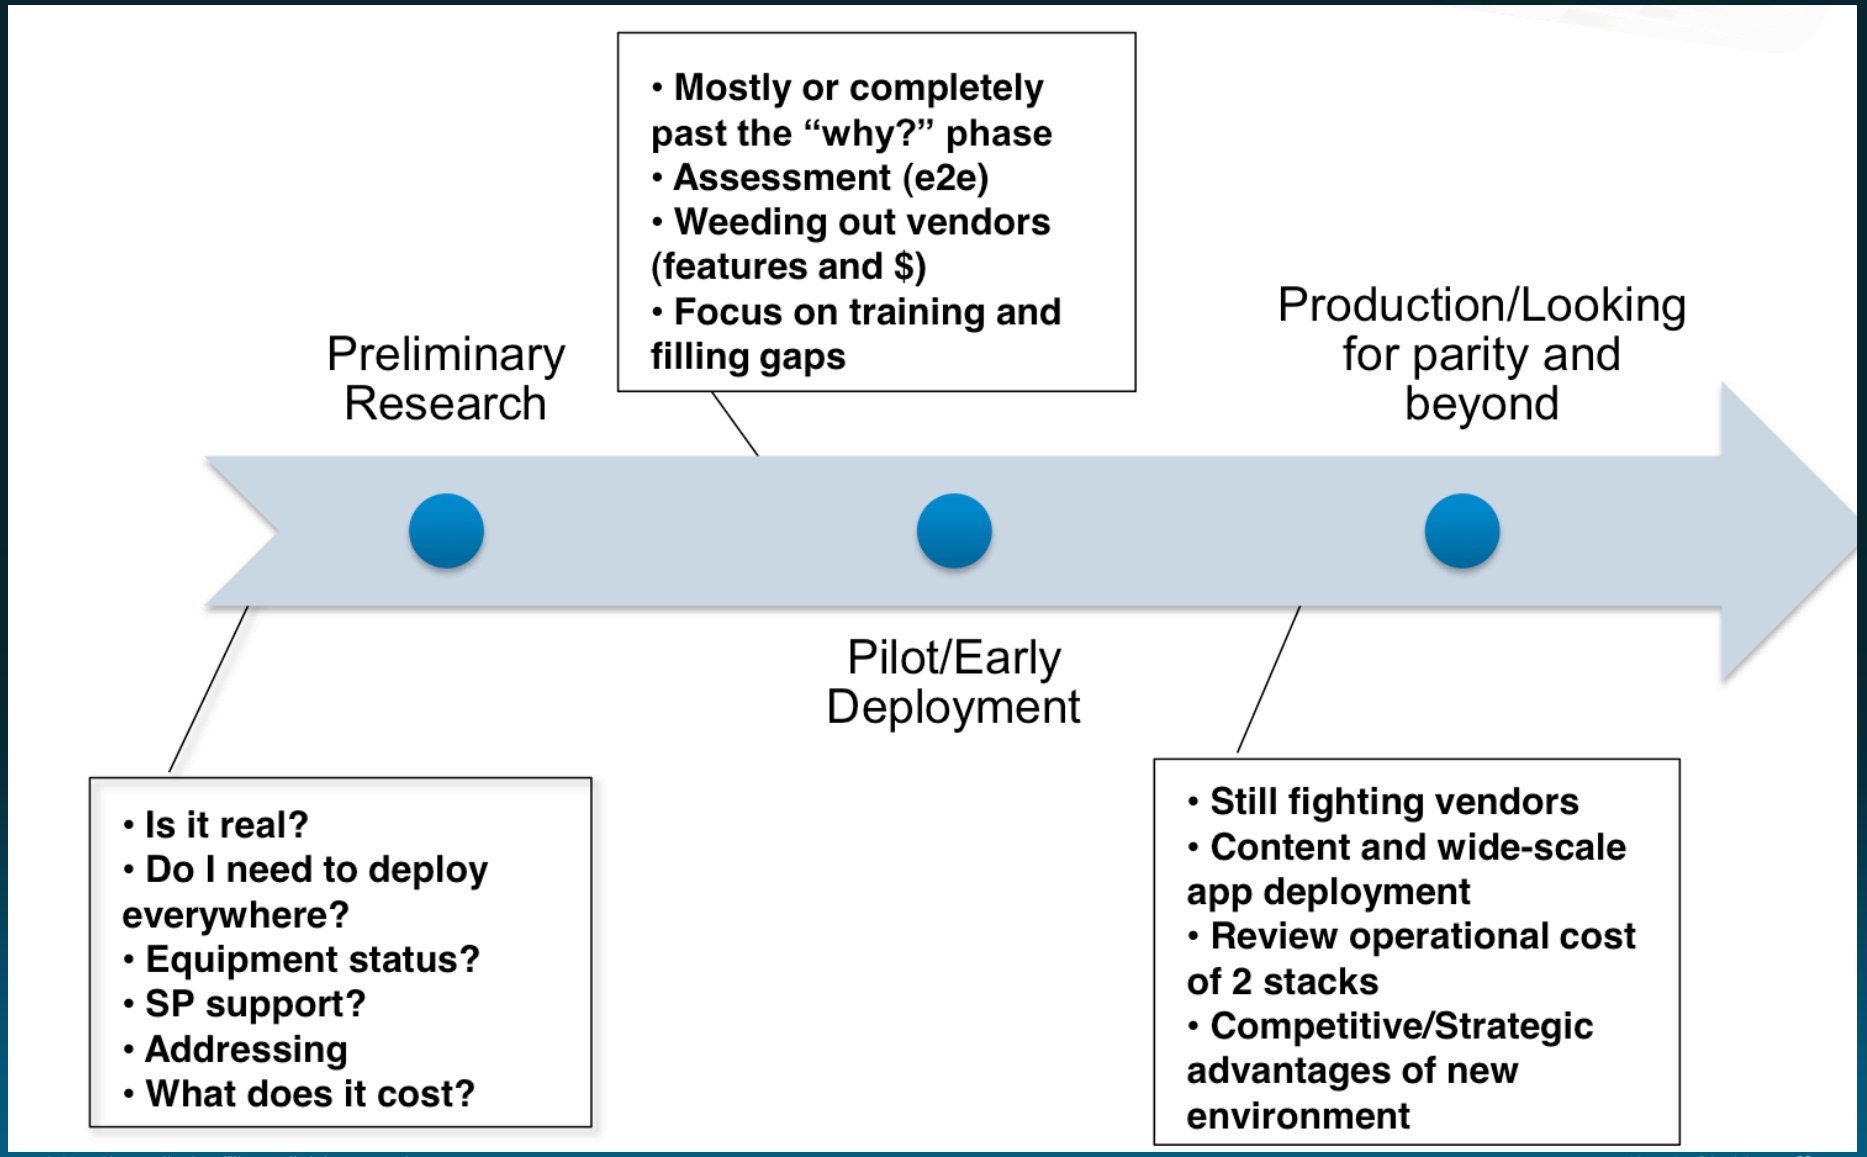

Who Owns the Translation?

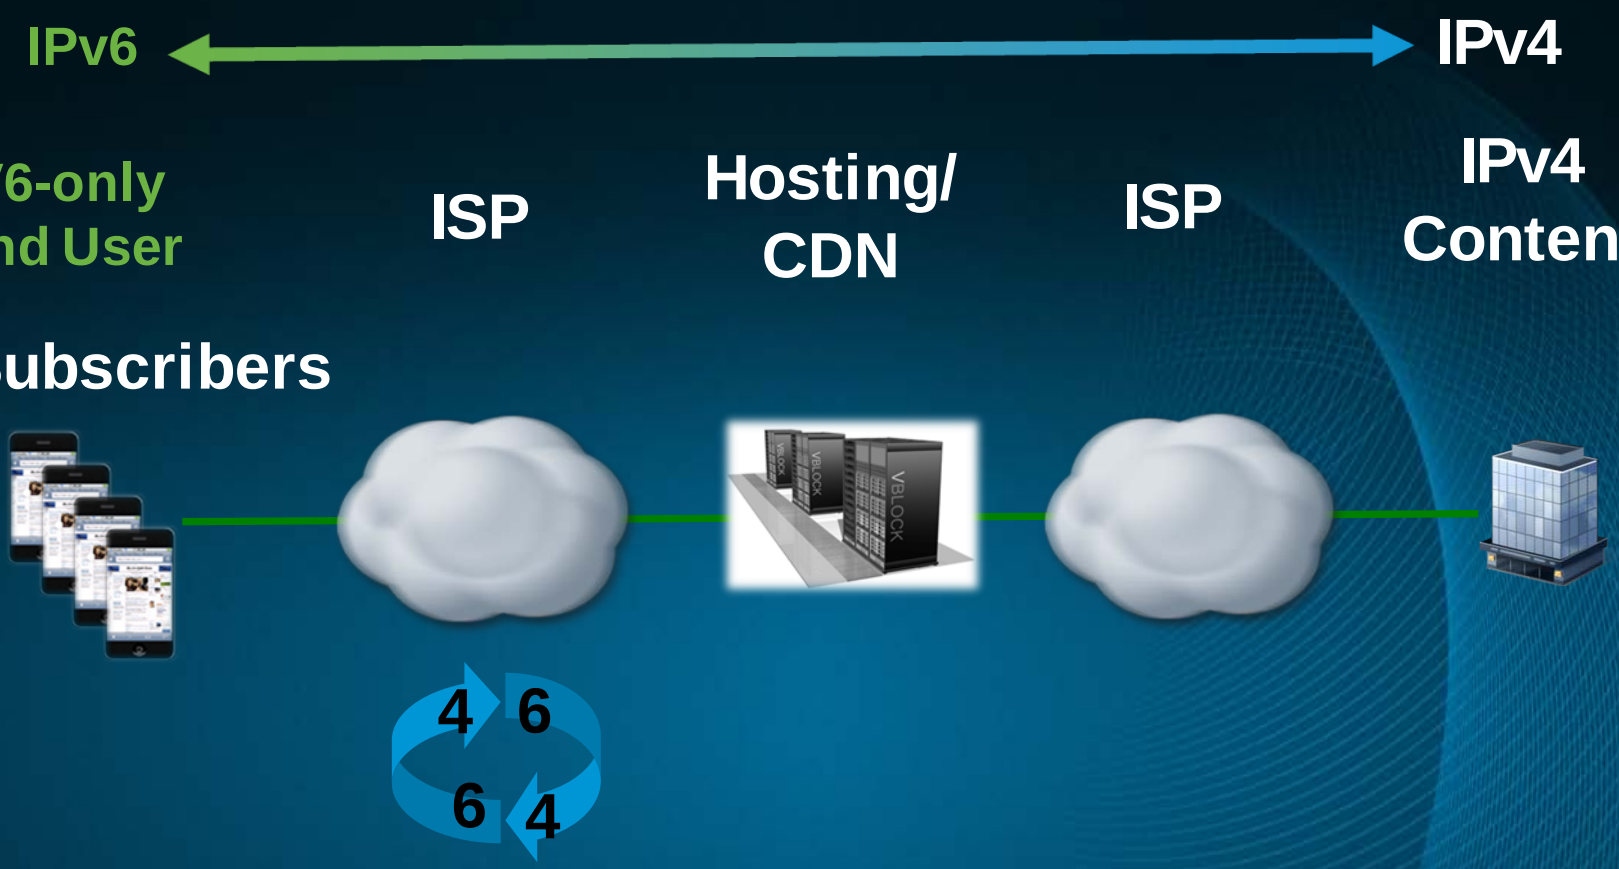

Considerations:
Experience, Scale, Cost, Operations, Technology…

IPv6 Integration Outline

Pre-Deployment Phases	Deployment Phases
• Establish the network starting point • Importance of a network assessment and available tools • Build a pilot or lab environment • Obtain addressing or use ULA or documentation prefix (in lab) • Learn the basics (DNS, routing changes, address assignment)	• Transport considerations for integration • Internet Edge (ISP, Apps) • Campus IPv6 integration options • Data Center integration options • WAN IPv6 integration options • Execute on gaps found in assessment

Where do I start in my IPv6?

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

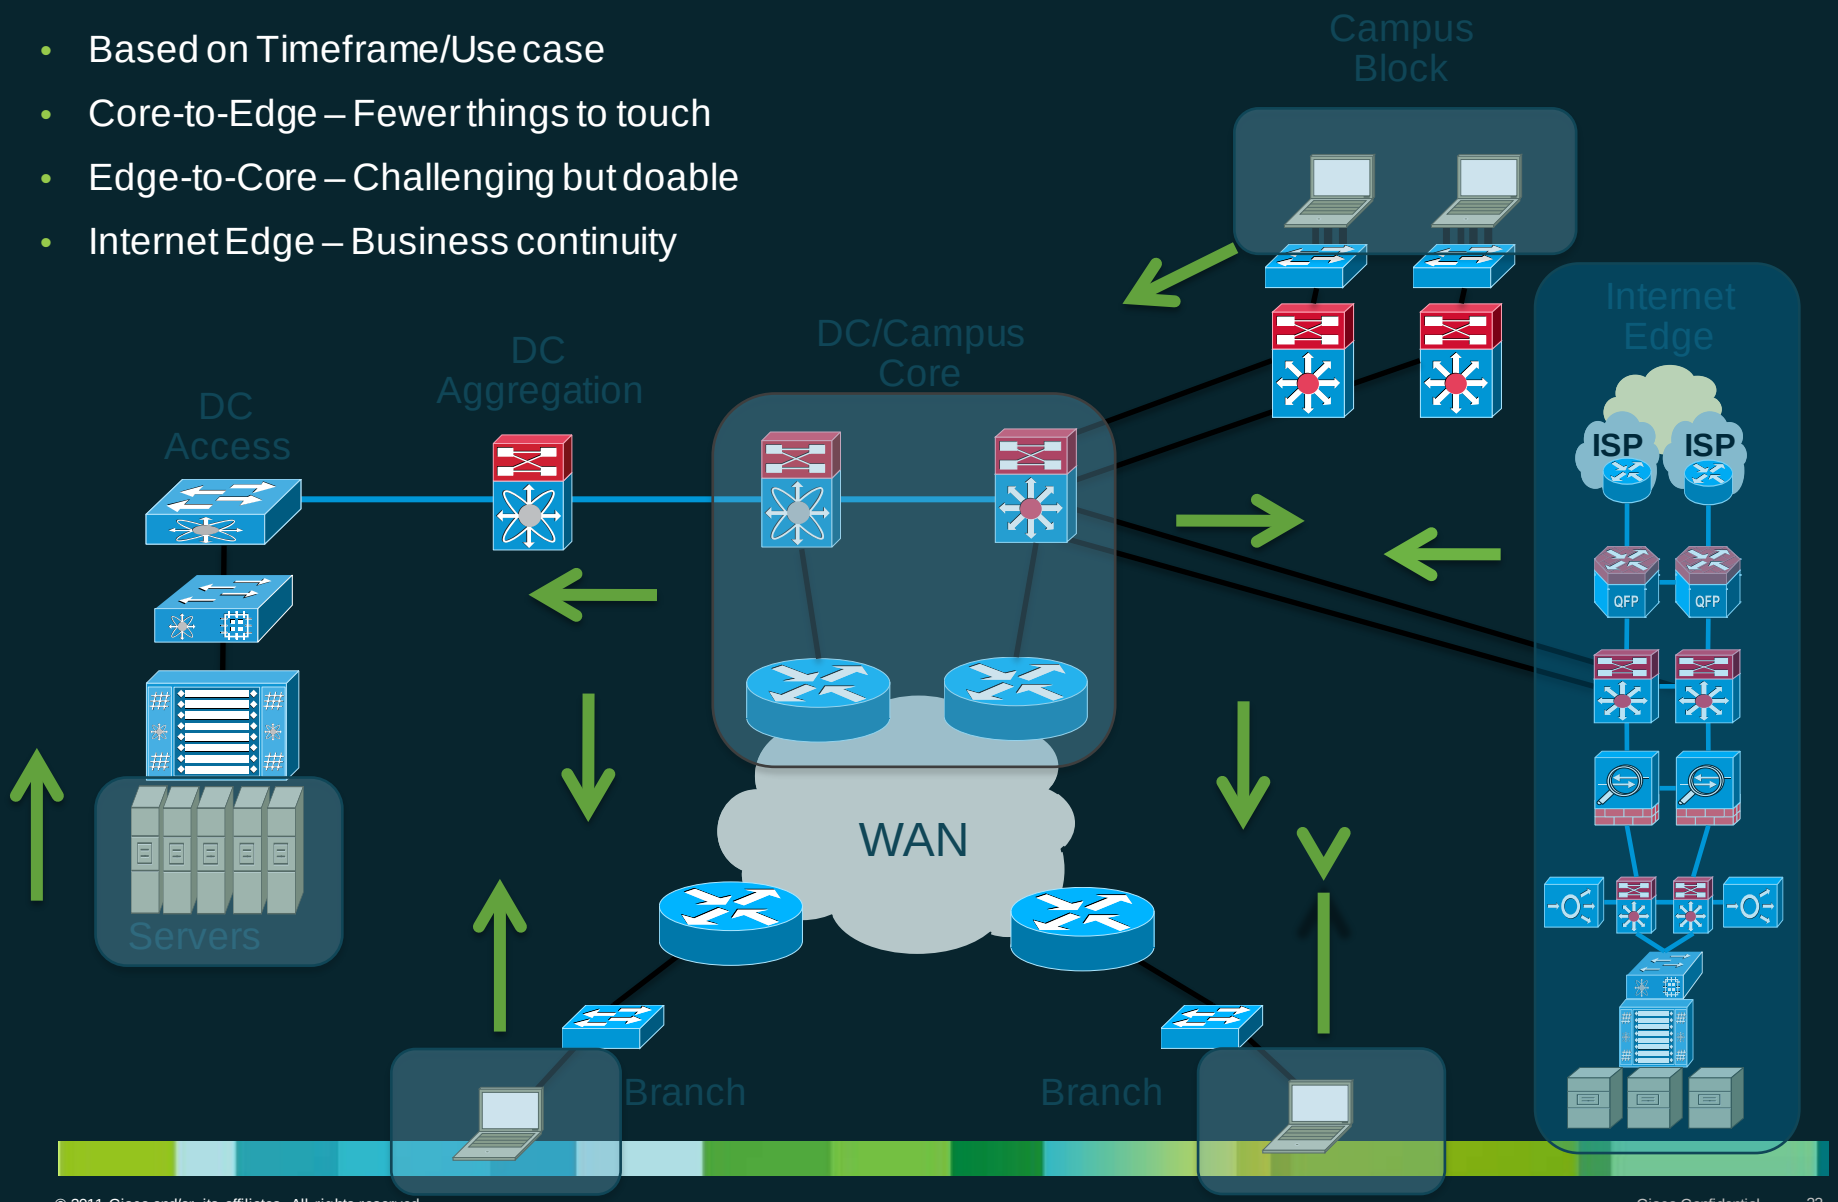

Agenda

IPv6 Top-On-Mind Questions

IPv6 Deployments

IPv6 Deployment

- IPv6 Enterprise Activities
- IPv6 Address Considerations
- General Network Considerations
- Infrastructure Deployment
  - Campus
  - Data Center/Internet Edge
  - WAN/Branch

IPv6 Enterprise Activities

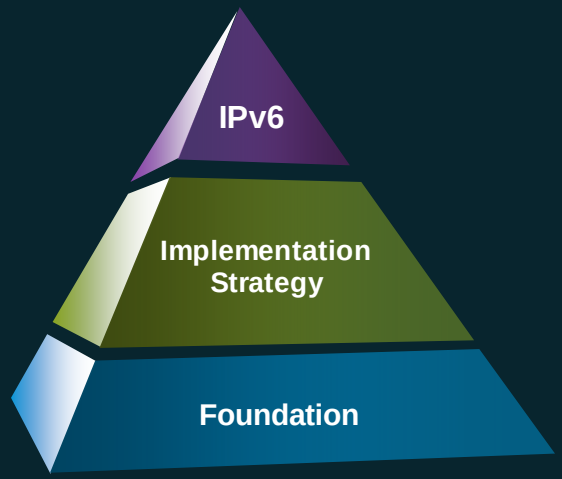

Dynamic Increase in Enterprise Activity Why?

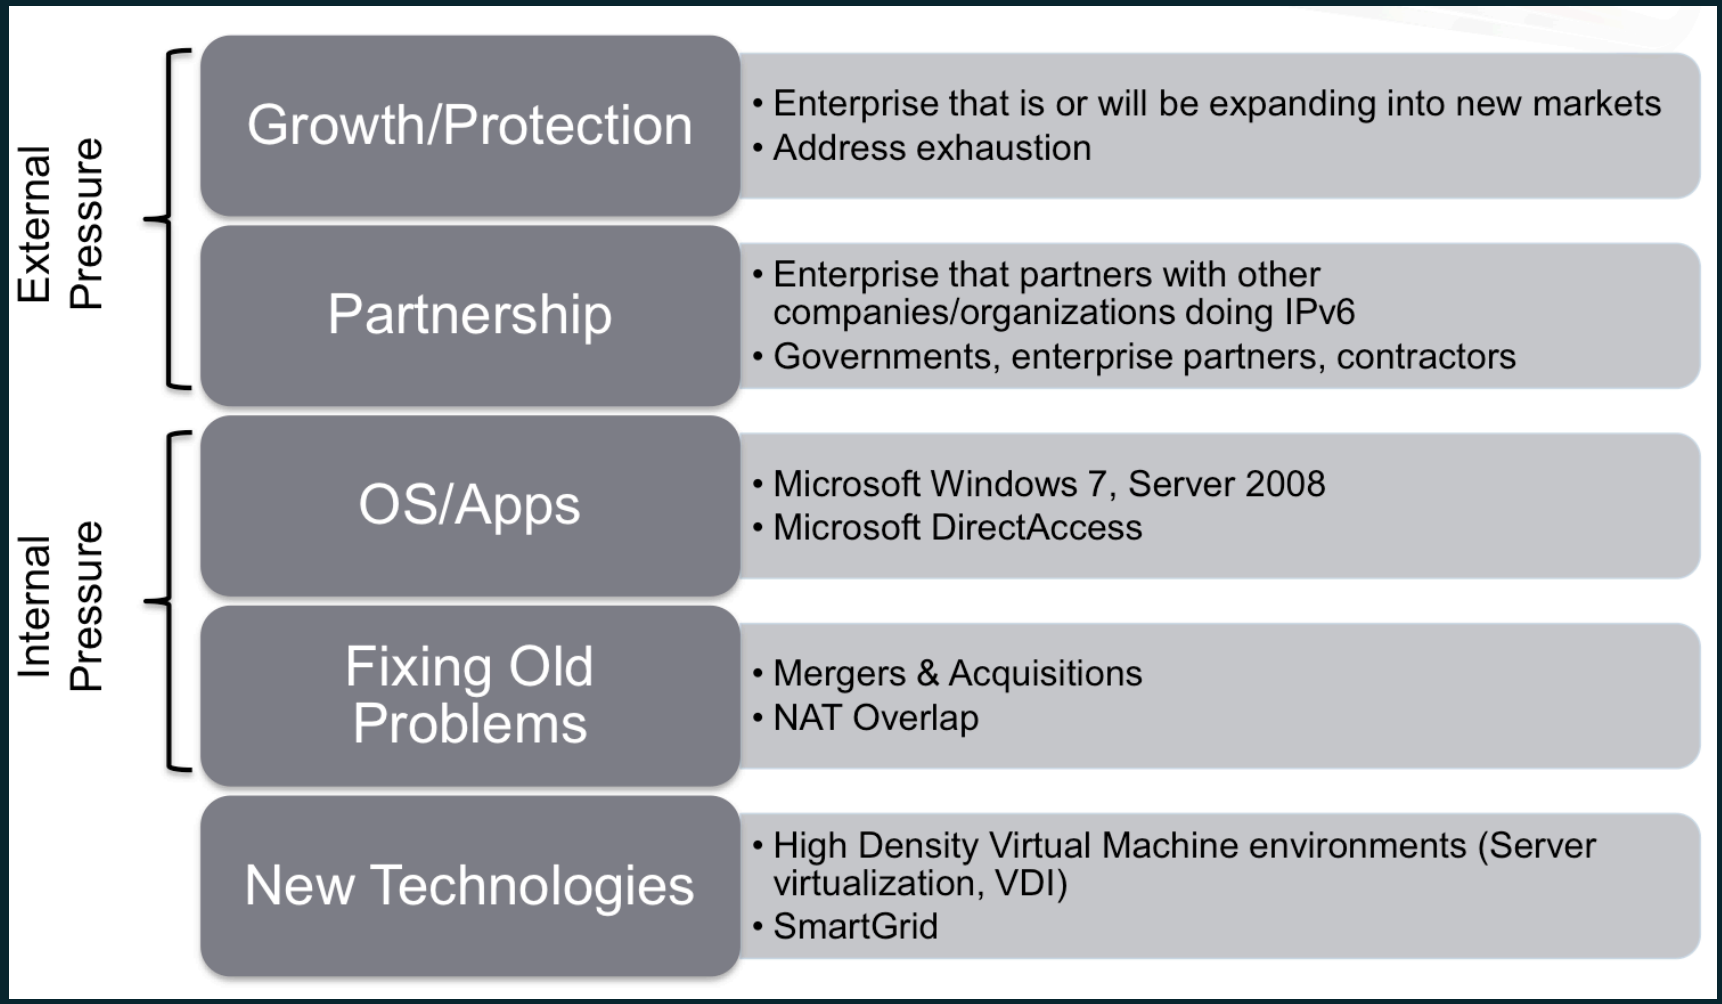

Mergers & Acquisitions

- Unique blend of technical and business problems
- Colliding RFC1918 space
- Common options
  - If you don't collide then leave as-is until renumbering is complete
  - NAT overlap pools (into non-colliding space) until renumbering is complete
  - IPv6 as an overlay network
  - IPv6 added as a native protocol (dual stack)
- This is a growing issue and IPv6 ends up being a perfect tool for resolving the technical issues

NAT Overlap + IPv6 Overlay

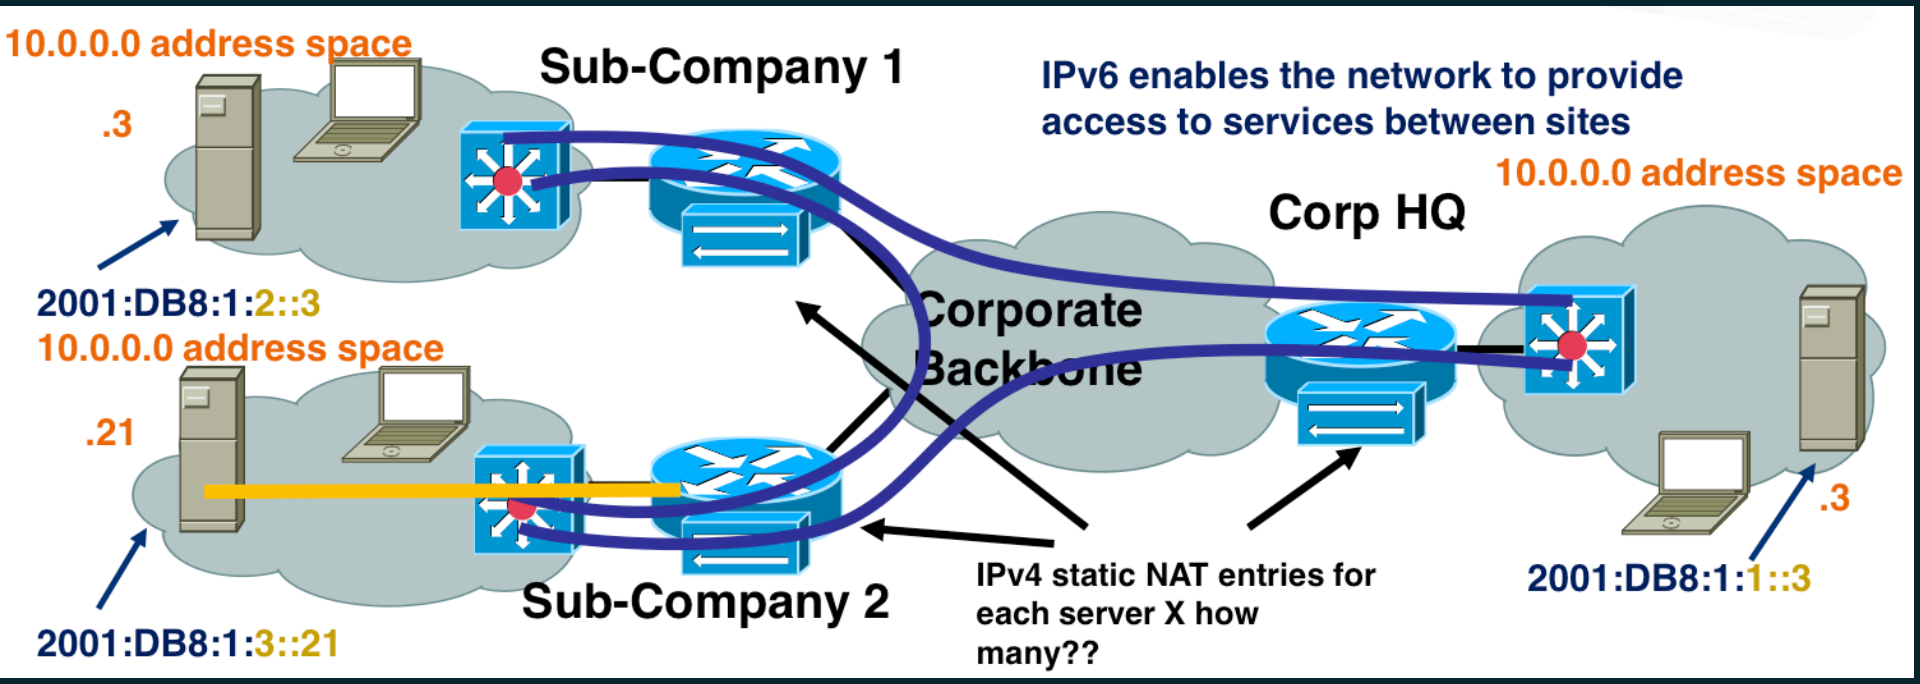

- Build an overlay network to encapsulate IPv6 over IPv4
- IPv6 is deployed only at those sites and for specific hosts that need end-to-end routability between entities
- Can be very operationally difficult to maintain in large environments
- May be a show stopper if you have to get a lot of tunnels past a bunch of IPv4 NAT

Partial Overlay + Partial Dual Stack

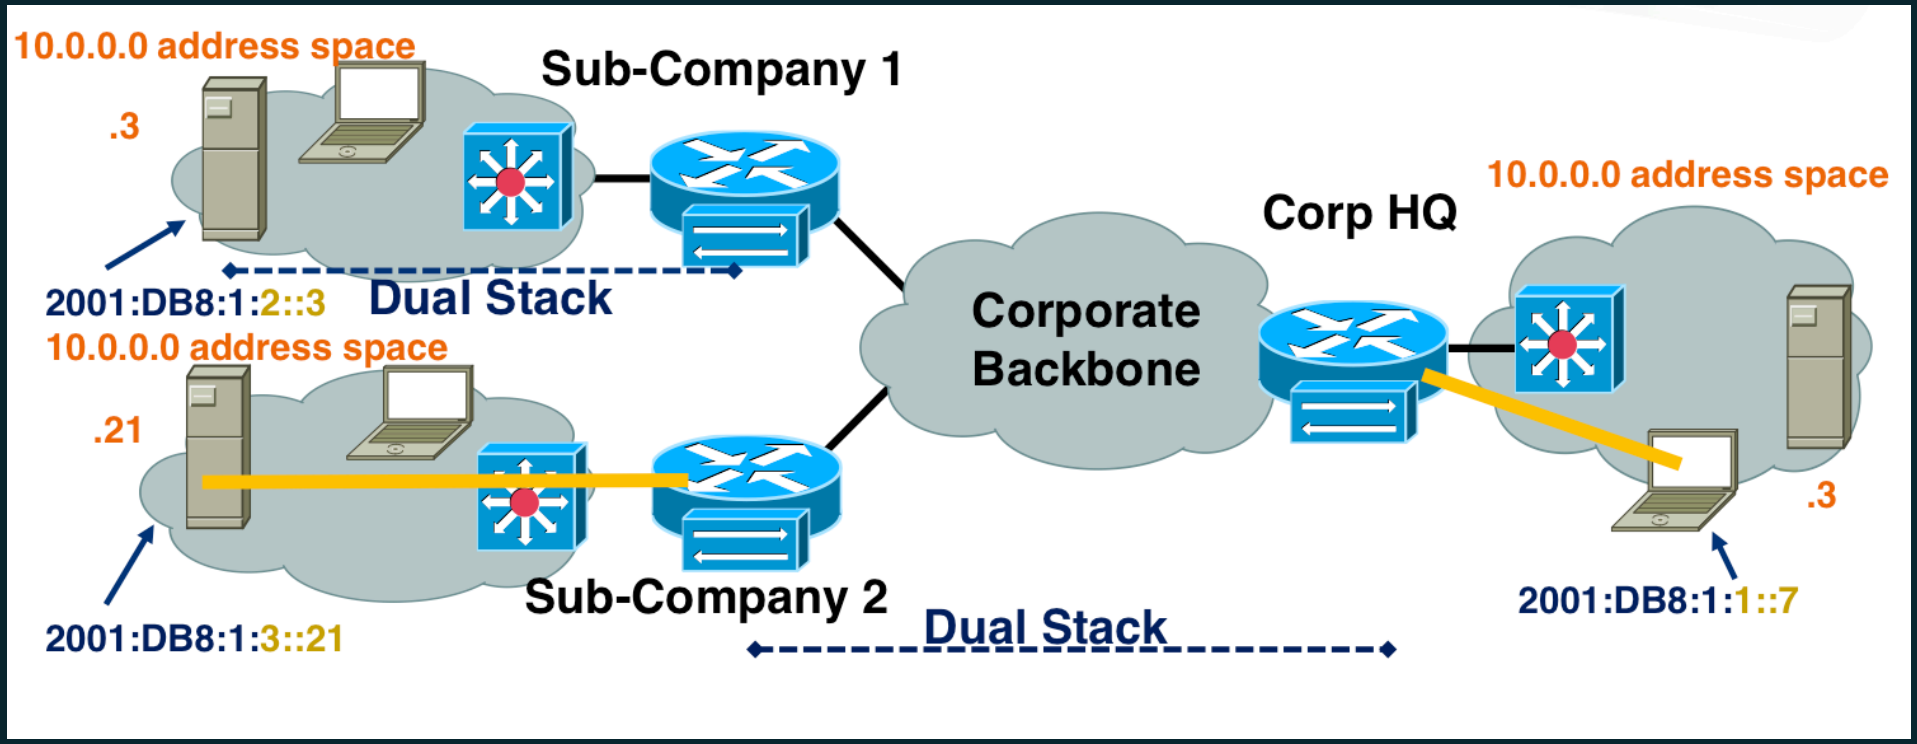

- Combine overlay network with dual stack
- Build as much dual stack as you can - tunnel only when you have to
- You don't want to keep this forever - goal is dual stack to all places that need end-to-end connectivity between sites/orgs

Dual Stack Everywhere

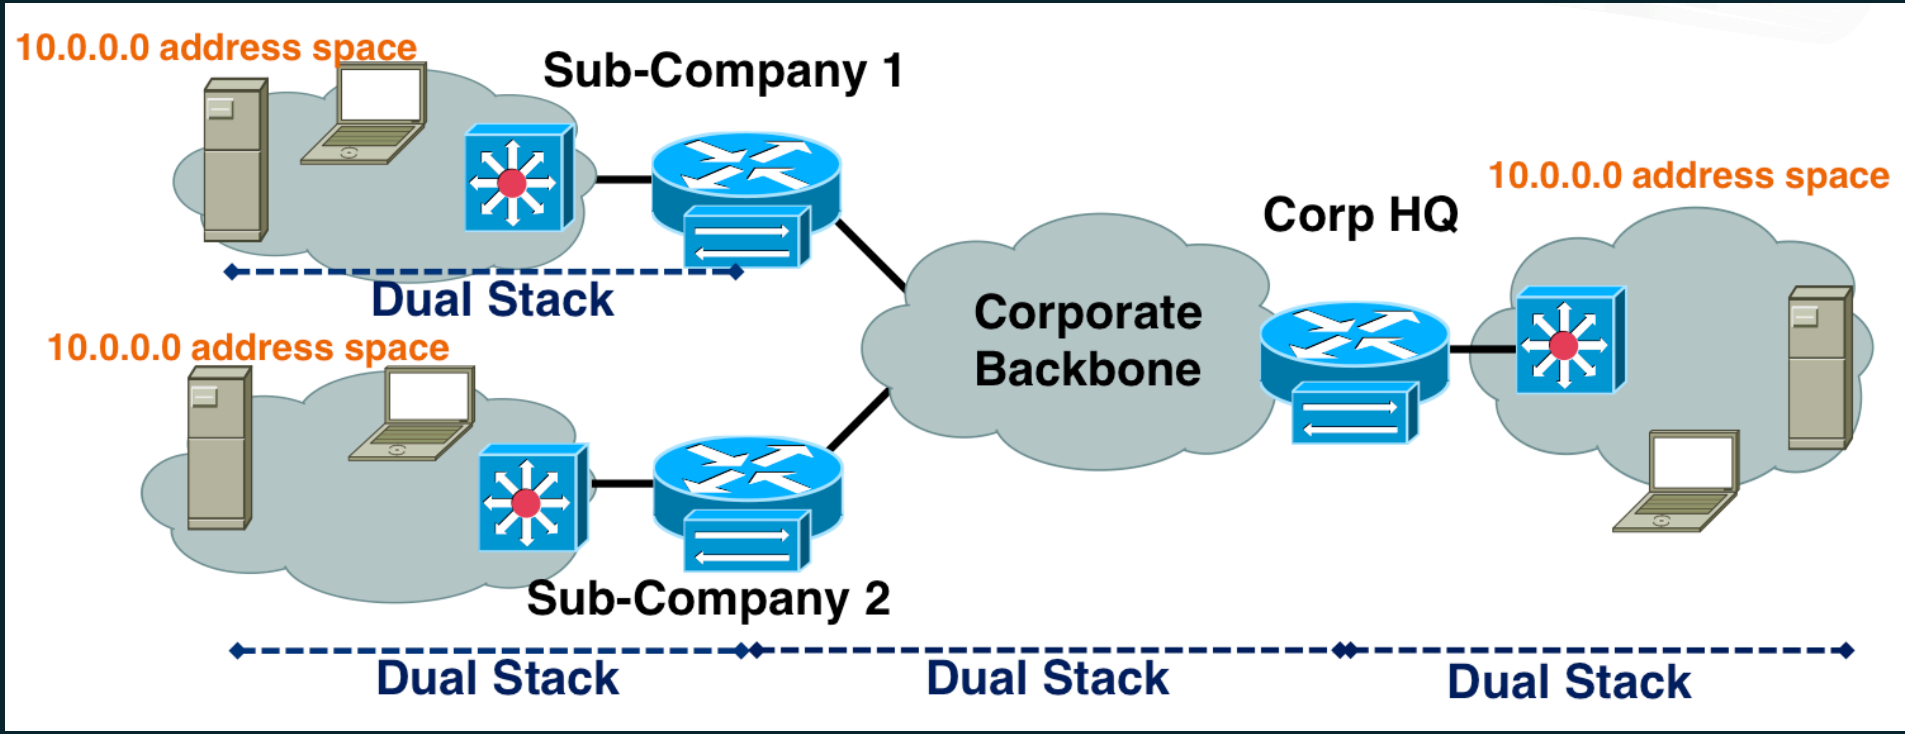

- Dual stack everywhere - there is nothing else to say ;-)
- We will discuss the deployment of dual stack and other end-to-end considerations for the rest of this talk

IPv6 Address Considerations

http://bit.ly/IPv6addrplan

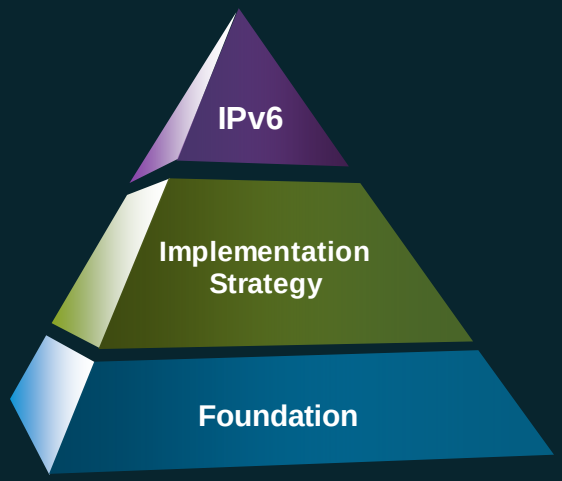

IPv6 Addresses

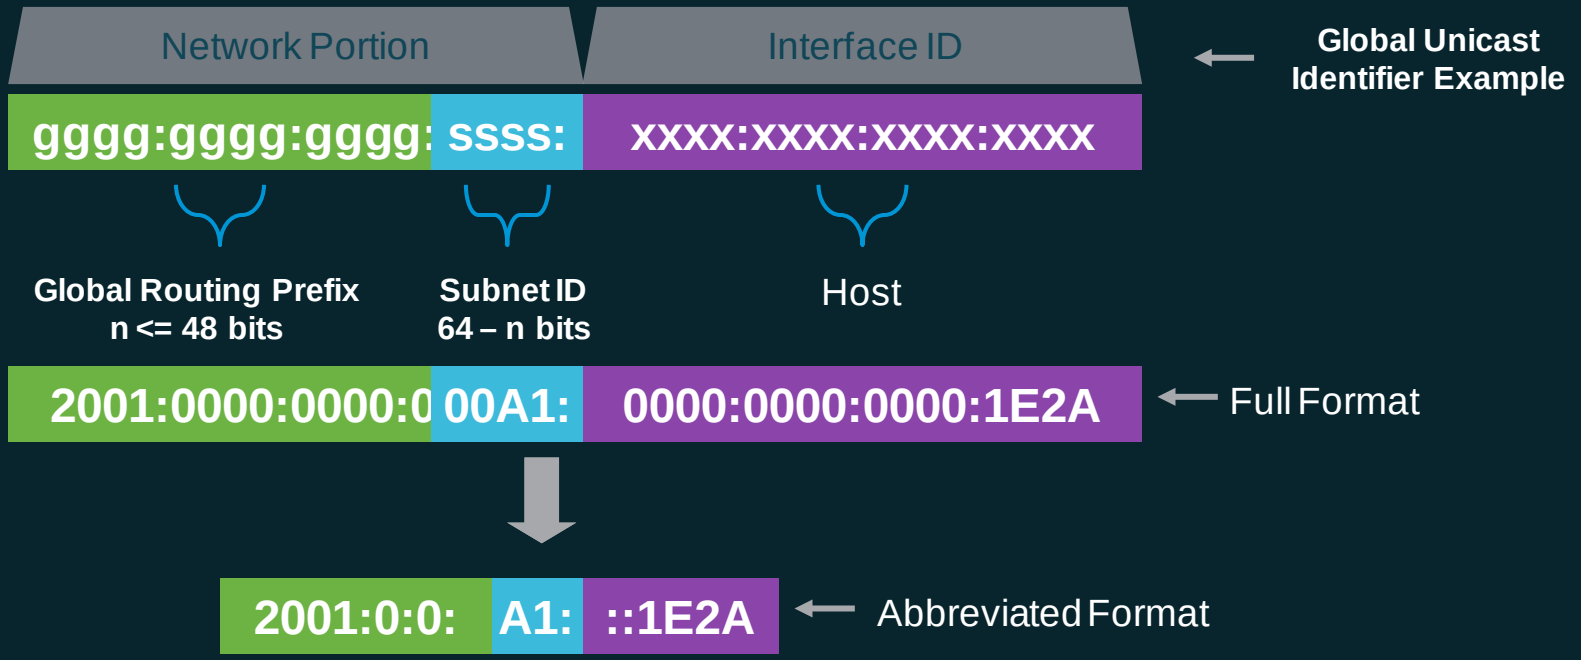

- IPv6 addresses are 128 bits long
  - Segmented into 8 groups of four HEX characters
  - Separated by a colon (:)
  - 50% for network ID, 50% for interface ID
  - Network portion is allocated by Internet registries 2^64 (1.8×10^{19})
  - Still leaves us with ~ 3 billion network prefixes for each person on earth

PA & PI Allocation Process

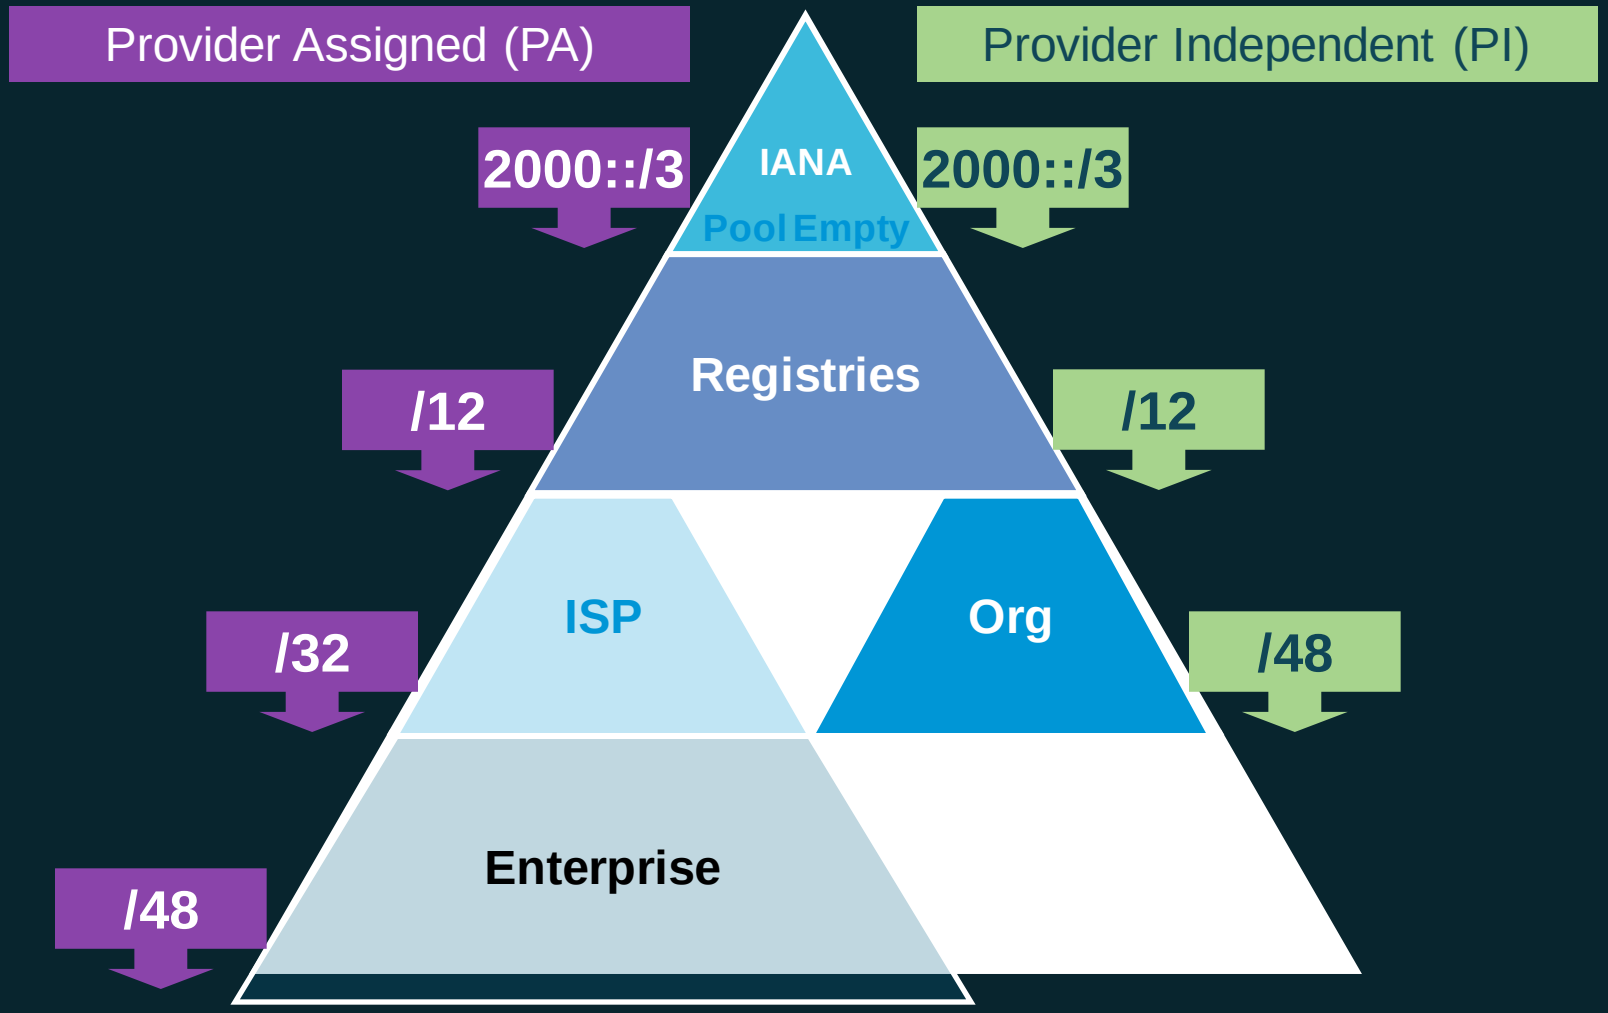

Hierarchical Addressing and Aggregation

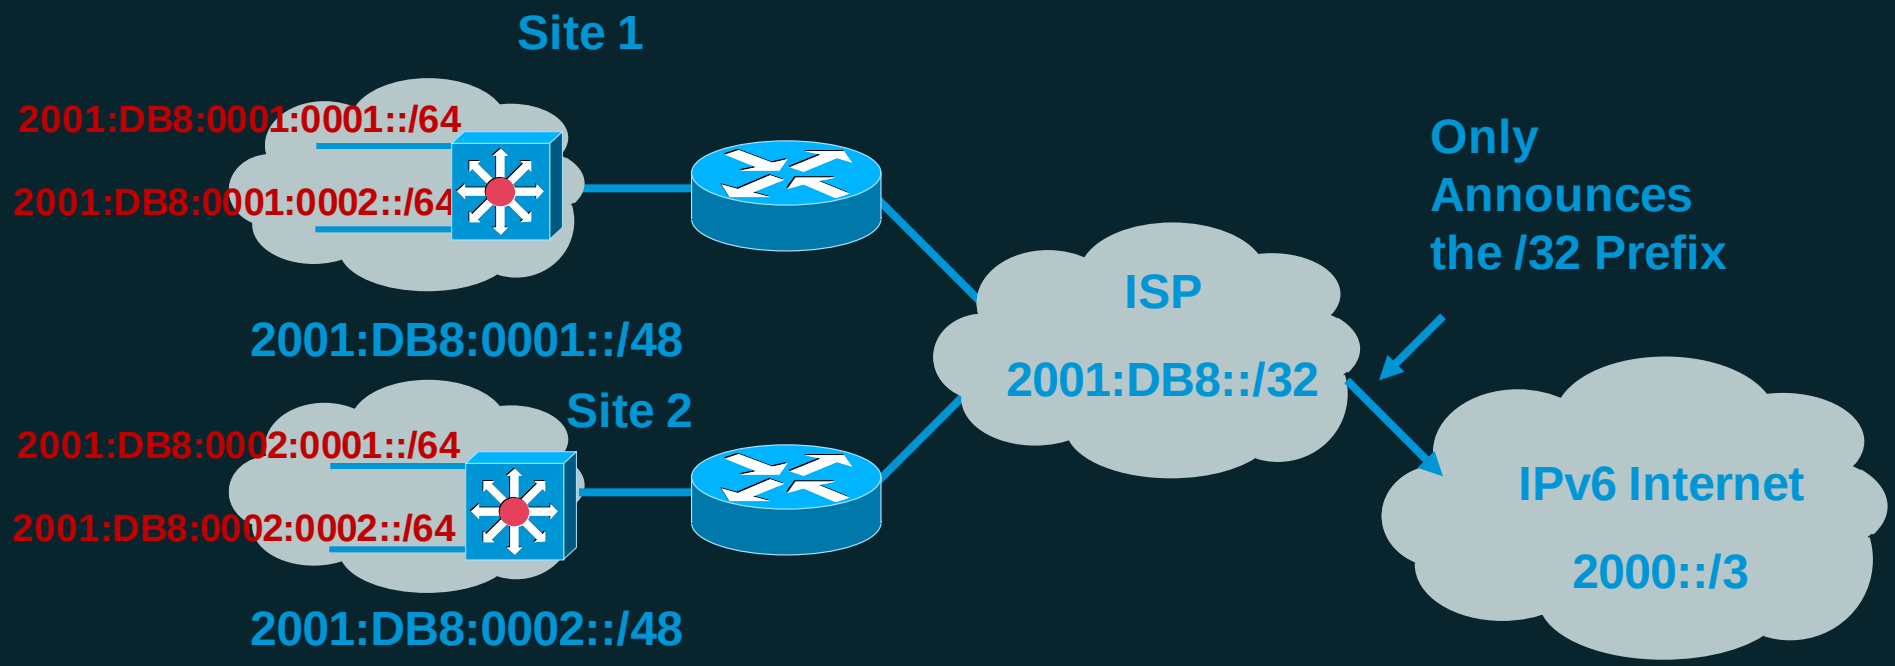

- Default is /48 – can be larger – “End-user Additional Assignment” https://www.arin.net/resources/request/ipv6_add_assign.html
- Provider independent – See Number Resource Policy Manual (NRPM) - https://www.arin.net/policy/nrpm.html

ULA, ULA + Global or Global

- What type of addressing should I deploy internal to my network? It depends:
  - o ULA-only - Today, no IPv6 NAT is useable in production so using ULA-only will not work externally to your network
  - o ULA + Global allows for the best of both worlds but at a price - much more address management with DHCP, DNS, routing and security - SAS does not always work as it should
  - o Global-only - Recommended approach but the old-school security folks that believe topology hiding is essential in security will bark at this option
- Let's explore these options ….

Not Recommended

ULA + Global

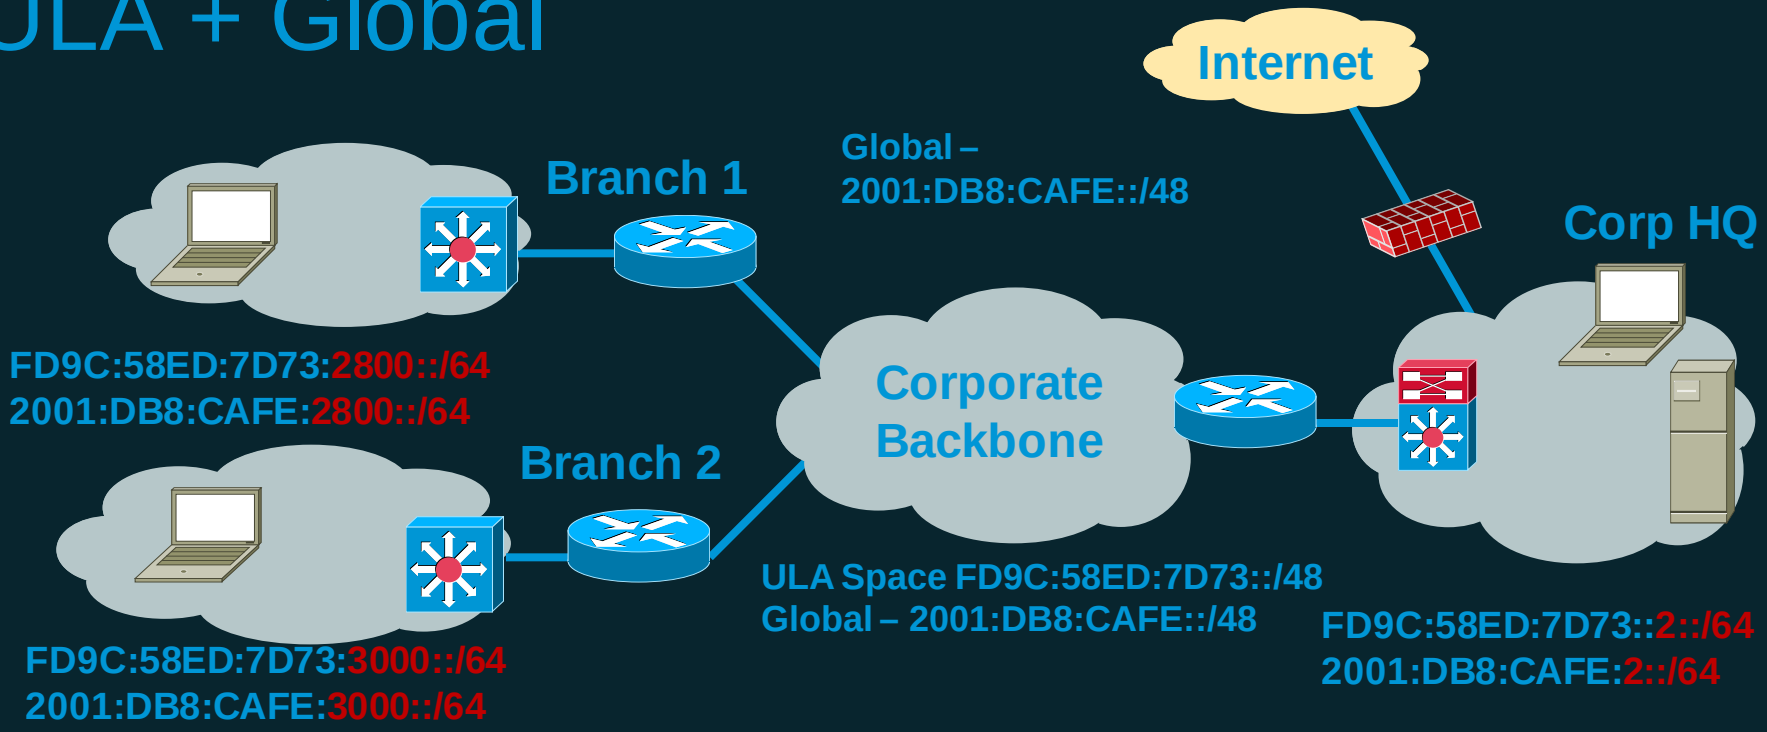

- Both ULA and Global are used internally except for internal-only hosts
- Source Address Selection (SAS) is used to determine which address to use when communicating with other nodes internally or externally
- In theory, ULA talks to ULA and Global talks to Global—SAS ‘should’ work this out
- ULA-only and Global-only hosts can talk to one another internal to the network
- Define a filter/policy that ensures your ULA prefix does not ‘leak’ out onto the Internet and ensure that no traffic can come in or out that has a ULA prefix in the SA/DA fields
- Management NIGHTMARE for DHCP, DNS, routing, security, etc…

Not Recommended

ULA + Global

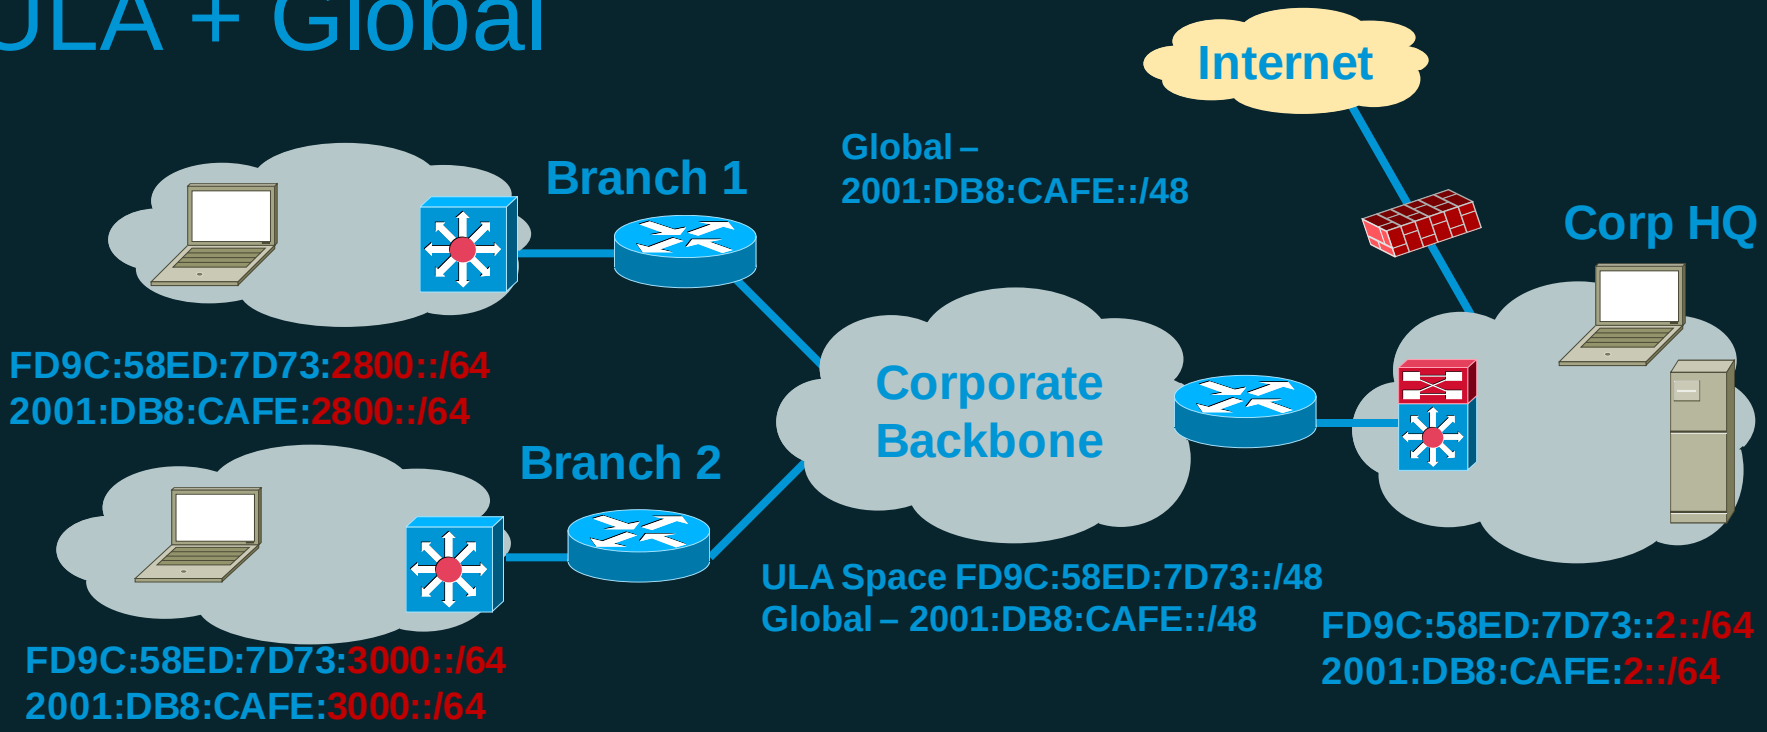

- Both ULA and Global are used internally except for internal-only hosts
- Source Address Selection (SAS) is used to determine which address to use when communicating with other nodes internally or externally
- In theory, ULA talks to ULA and Global talks to Global—SAS ‘should’ work this out
- ULA-only and Global-only hosts can talk to one another internal to the network
- Define a filter/policy that ensures your ULA prefix does not ‘leak’ out onto the Internet and ensure that no traffic can come in or out that has a ULA prefix in the SA/DA fields
- Management NIGHTMARE for DHCP, DNS, routing, security, etc…

Recommended

Global-Only

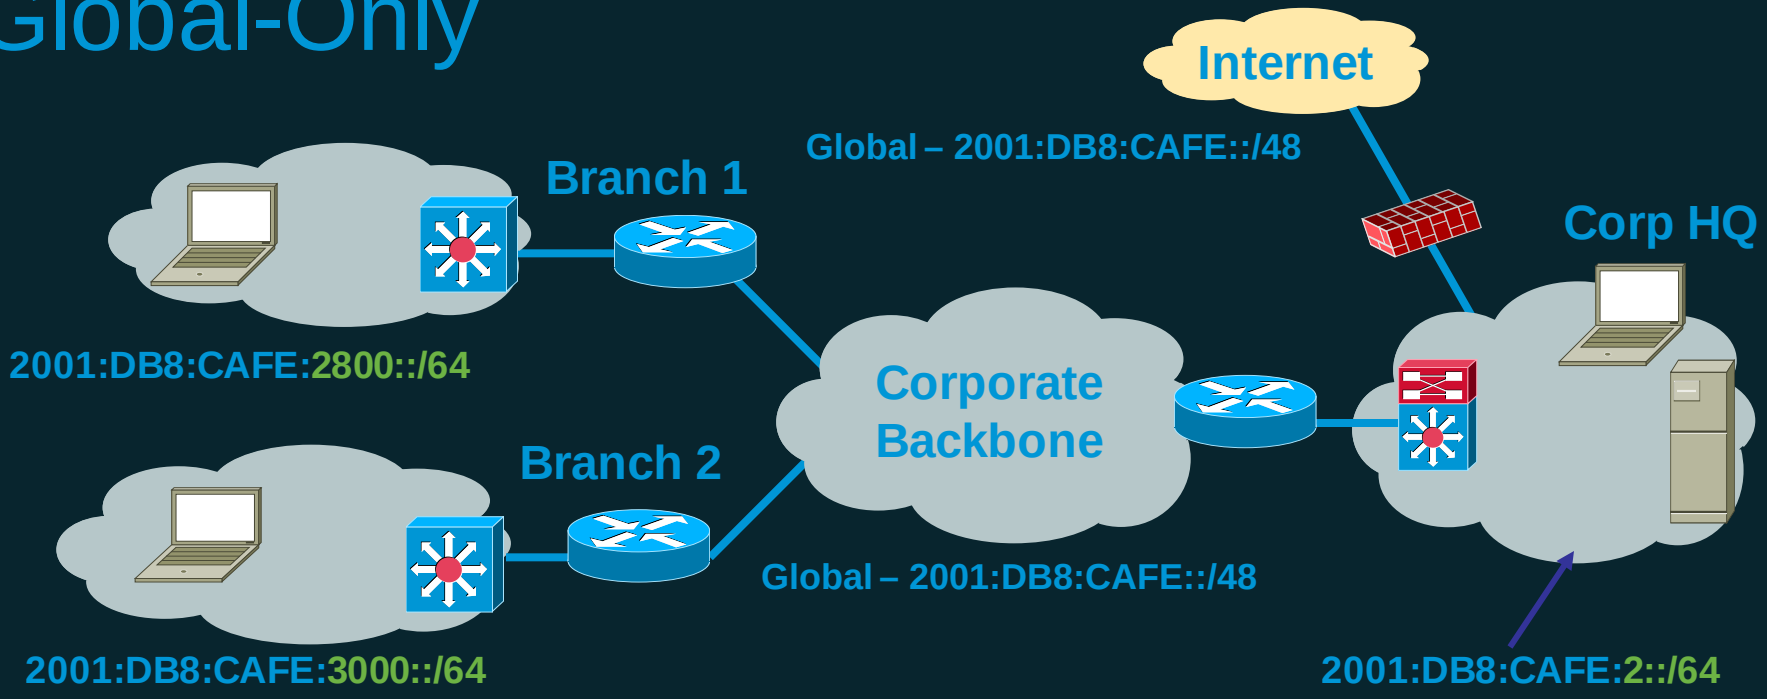

- Global is used everywhere
- No issues with SAS
- No requirements to have NAT for ULA-to-Global translation—but, NAT may be used for other purposes
- Easier management of DHCP, DNS, security, etc.
- Your heartburn comes from the security team – topology hiding

Link Level—Prefix Length Considerations

64 bits

- Recommended by RFC3177 and IAB/IESG
- Consistency makes management easy
- MUST for SLAAC (MSFT DHCPv6 also)
- Significant address space loss (18.466 Quintillion)

> 64 bits

- Address space conservation
- Special cases:
  /126—valid for p2p
  /127—valid for p2p if you are careful (draft-kohno-ipv6-prefixlen-p2p-xx/(RFC3627))
  /128—loopback
- Must avoid overlap with specific addresses:
  Router Anycast (RFC3513)
  Embedded RP (RFC3956)
  ISATAP addresses

- /64 everywhere
- /64 + /126
  - –64 on host networks
  - –126 on P2P
- /64 + /127
  - –64 on host networks
  - –127 on P2P
- Always use /128 on loop

SLAAC & Stateful/Stateless DHCPv6

- StateLess Address AutoConfiguration (SLAAC) RA- based assignment (a MUST for Mac)
- Stateful and stateless DHCPv6 server

Cisco Network Registrar:http://www.cisco.com/en/US/products/sw/netmgtsw/ps1982/

Microsoft Windows Server 2008:http://technet.microsoft.com/en-us/library/bb218968.aspx?missingurl=%2fen-US%2flibrary%2fbab0f%25201a1-54aa-4cef-9164-139e8bcc4475.aspx

- DHCPv6 Relay - supported on routers and switches

SLAAC & Stateful/Stateless DHCPv6

```
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 dhcp relay destination 2001:DB8:CAFE:10::2
```

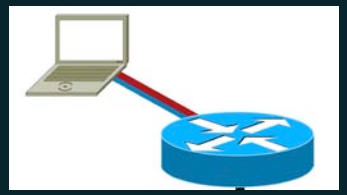

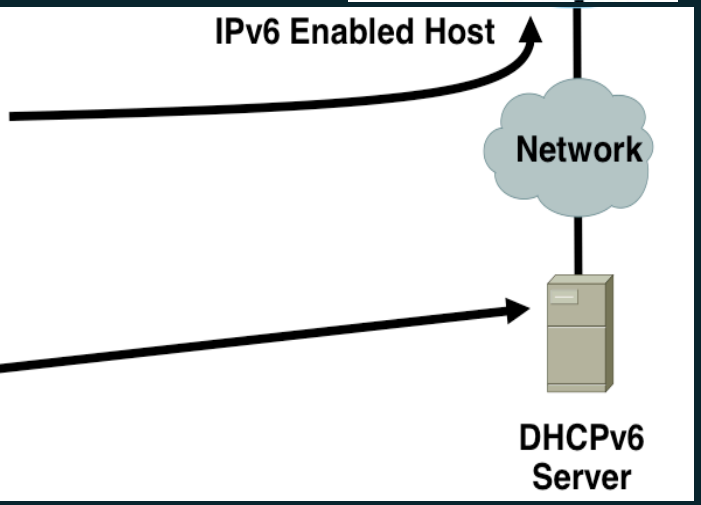

General Network Considerations

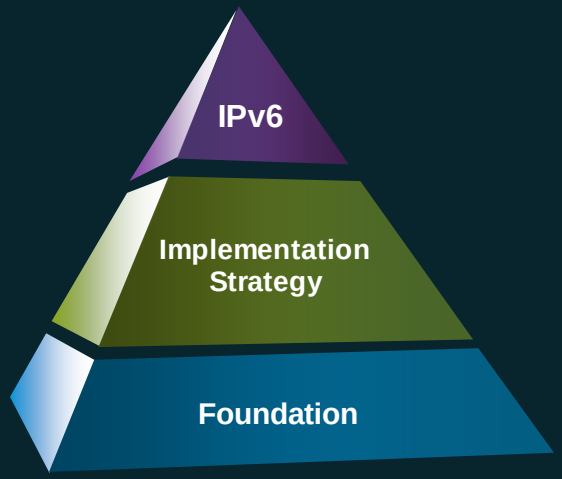

HSRP for IPv6

- 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)

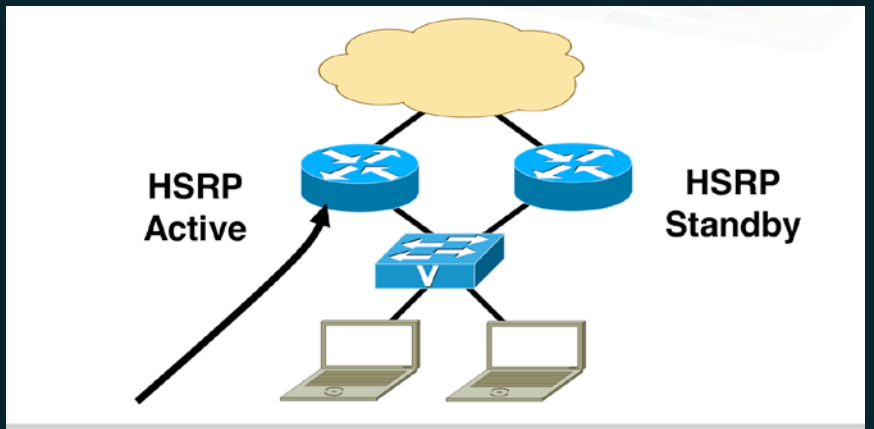


```
track 2 interface FastEthernet0/0 line-protocol
interface FastEthernet0/1
 ipv6 address 2001:DB8:66:67::2/64
 standby version 2
 standby 2 ipv6 autoconfig
 standby 2 timers msec 250 msec 800
 standby 2 preempt
 standby 2 preempt delay minimum 180
 standby 2 authentication cisco
 standby 2 track 2 decrement 10
```

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

IPv6 QoS & Syntax

- Unified QoS Policy (v4/v6 in same policy) or separate?
- 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

IPv6 Co-existence Solutions

Dual Stack

Recommended Enterprise Co-existence strategy

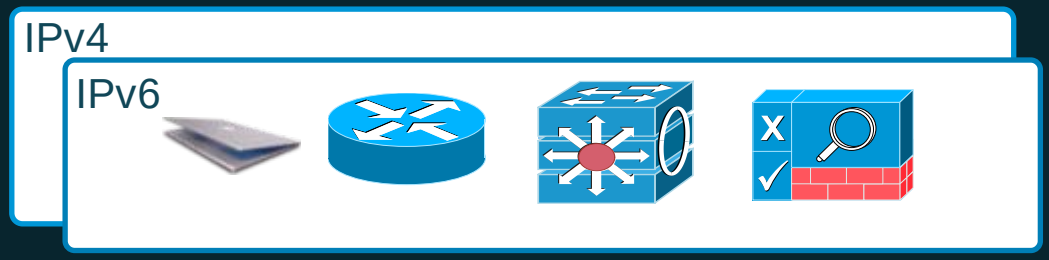

Tunneling Services

Connect Islands of IPv6 or IPv4

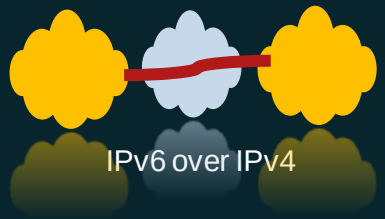

Translation Services

Connect to the IPv6 community

Business Partners
Government Agencies
International Sites
Remote Workers
Internet consumers

Infrastructure deployment

Campus Network

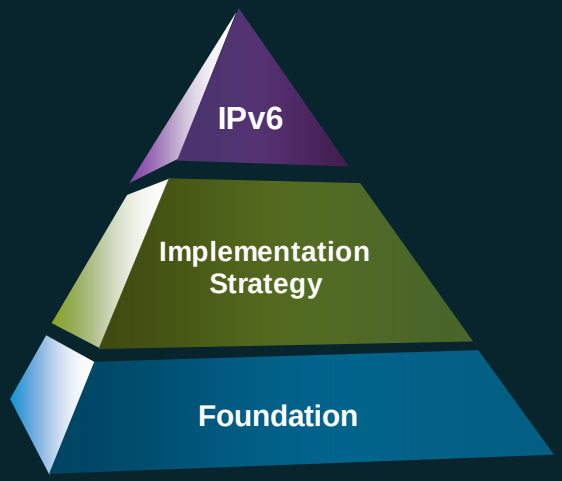

Campus IPv6 Deployment Options Dual-Stack IPv4/IPv6

- Dual Stack = Two protocols running at the same time (IPv4/IPv6)
- #1 requirement—switching/ routing platforms must support hardware based forwarding for IPv6
  - –3560/3750 +
  - –4500 Sup6E +
  - –6500 Sup32/720 +
- IPv6 is transparent on L2 switches but consider:
  - –L2 multicast—MLD snooping
  - –IPv6 management—Telnet/SSH/HTTP/SNMP
  - –Intelligent IP services on WLAN
- Expect to run the same IGPs as with IPv4

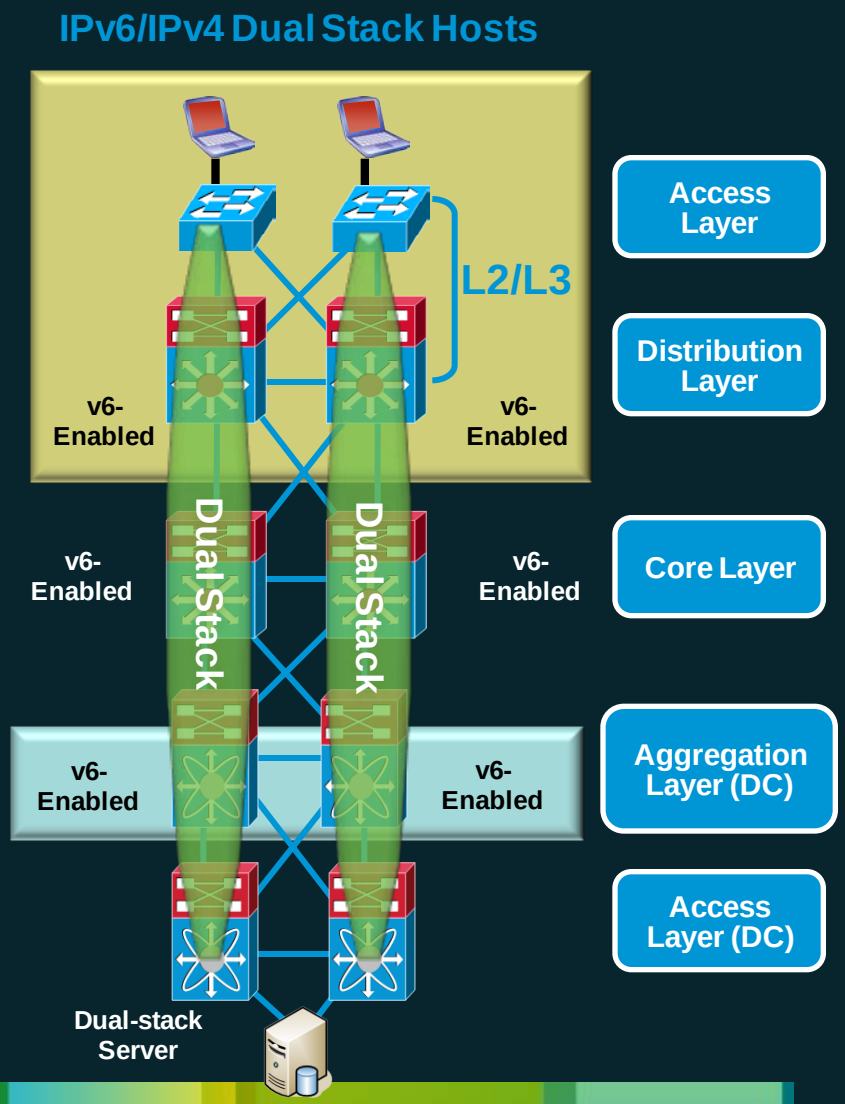

Campus IPv6 Deployment Options Hybrid Model

- Plan “B” if Layer 3 device can’t support IPv6 but you have to get IPv6 over it
- Offers IPv6 connectivity via multiple options
  - Dual-stack
  - Configured tunnels—L3-to-L3
  - ISATAP—Host-to-L3
- Leverages existing network
- Offers natural progression to full dual-stack design
- May require tunneling to less-than-optimal layers (i.e. core layer)
- Any sizable deployment will be an operational management challenge
- ISATAP creates a flat network (all hosts on same tunnel are peers)
- Provides basic HA of ISATAP tunnels via old Anycast-RP idea

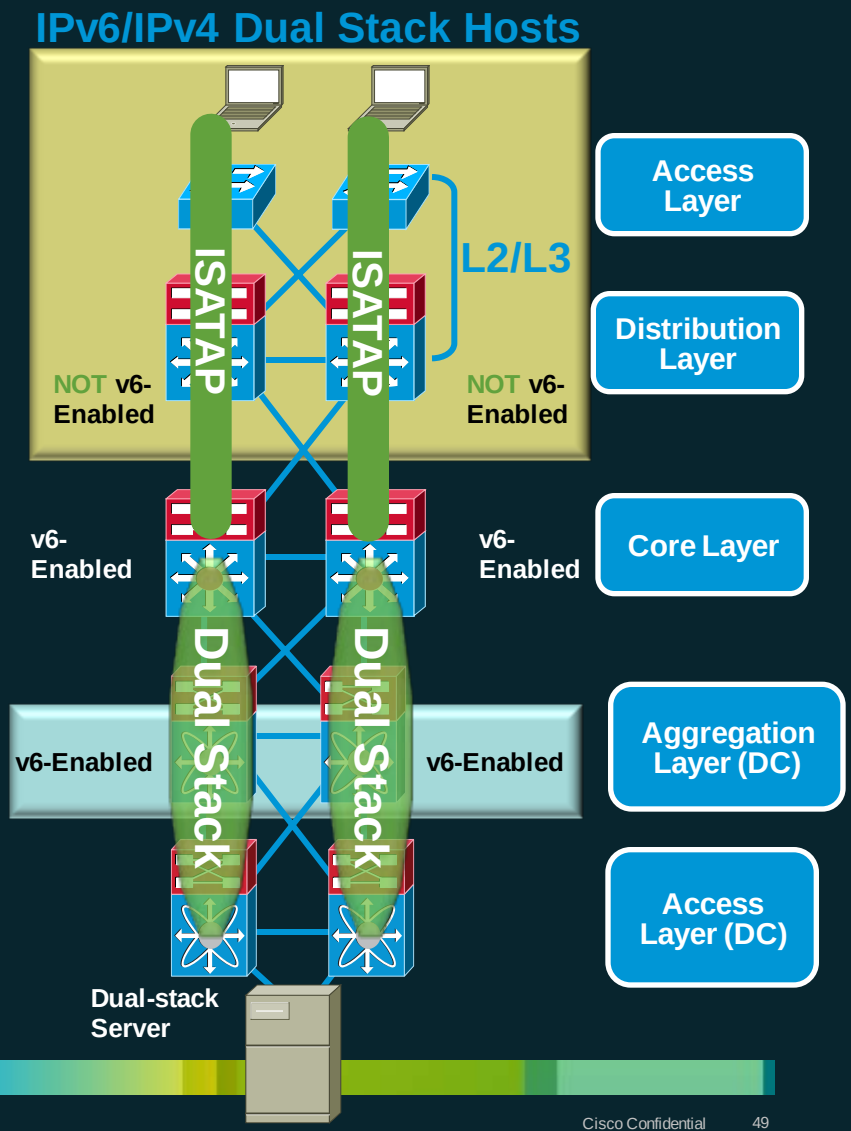

Highly Available ISATAP Design

- ISATAP tunnels from PCs in access layer to core switches
- Redundant tunnels to core or service block
- Use IGP to prefer one core switch over another (both v4 and v6 routes – deterministic
- Preference is important due to the requirement to have traffic (IPv4/IPv6) route to the same interface (tunnel)
- Works like Anycast-RP with Ipmc ☺

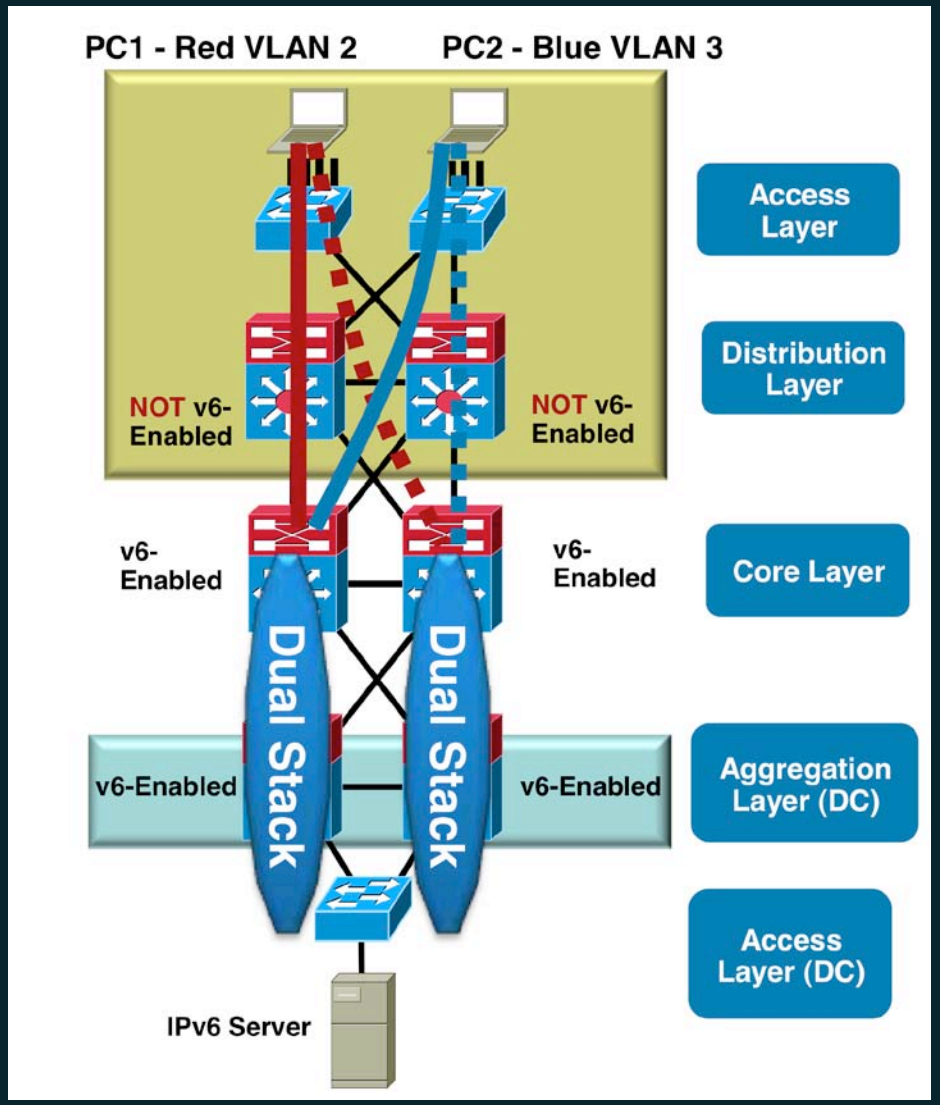

Campus IPv6 Deployment Options
IPv6 Service Block—Rapid Deployment/Pilot

- Provides ability to rapidly deploy IPv6 services without touching existing network
- Provides tight control of where IPv6 is deployed and where the traffic flows (maintain separation of groups/locations)
- Get lots of operational experience with limited impact to existing environment – Ideal for Pilot
- Similar challenges as Hybrid Model – Lots of tunneling
- Configurations are very similar to the Hybrid Model
  - –ISATAP tunnels from PCs in access layer to service block switches (instead of core layer—Hybrid)
- 1) Leverage existing ISP block for both IPv4 and IPv6 access
- 2) Use dedicated ISP connection just for IPv6—Can use IOS FW or PIX/ASA appliance

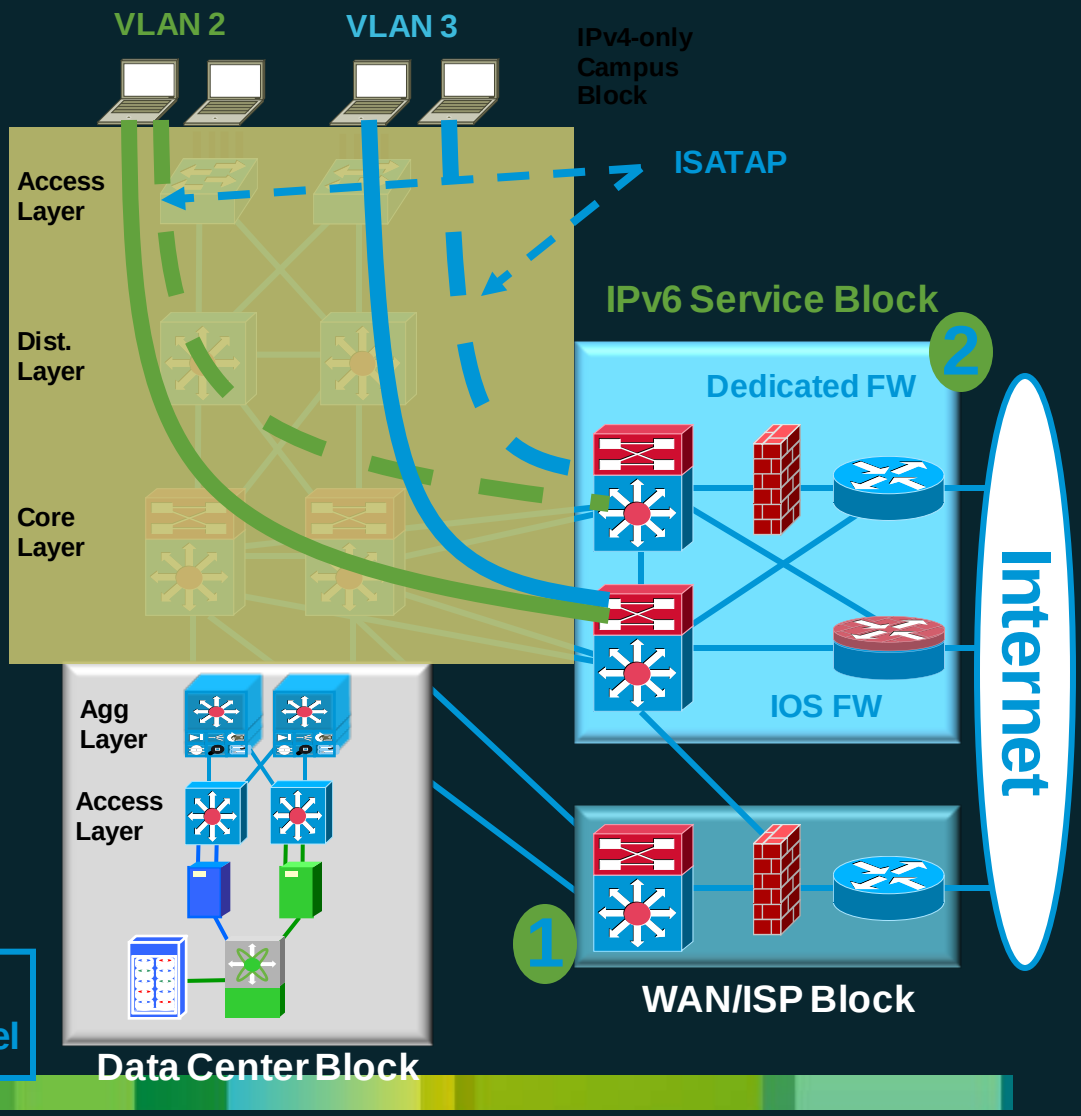

Infrastructure deployment

Data Center / Internet Edge

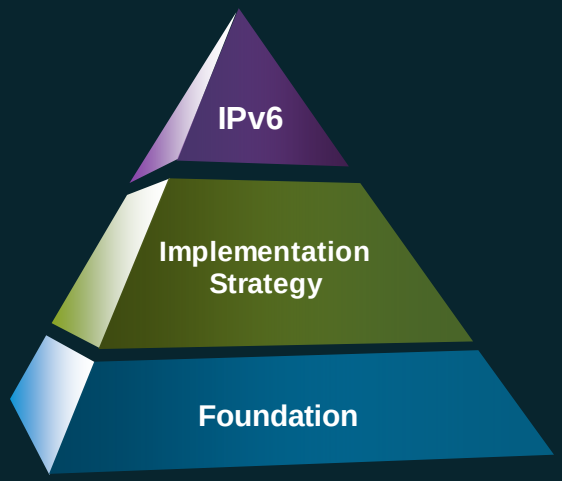

IPv6 Data Center Integration

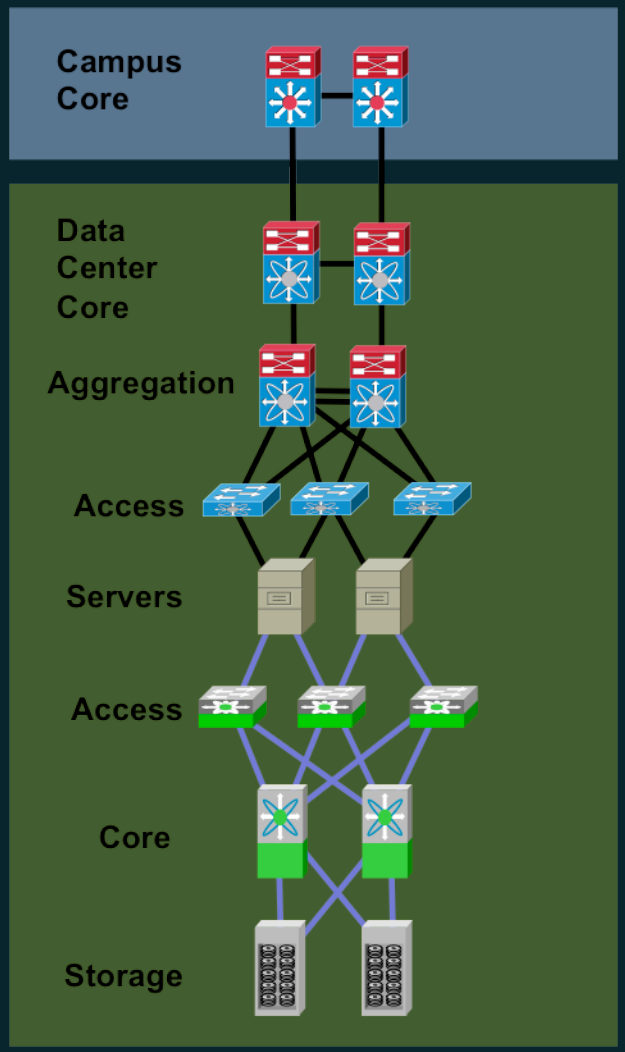

- Route/Switch design will be similar to campus based on feature, platform and connectivity similarities – Nexus, 6500 4900M
- The single most overlooked and potentially complicated area of IPv6 deployment
- Stuff people don't think about:
  - –NIC Teaming, iLO, DRAC, IP KVM, Clusters
  - –Innocent looking Server OS upgrades – Windows Server 2008 - Impact on clusters – Microsoft Server 2008 Failover clusters full support IPv6 (and L3)
- Internet-facing Data Center
- Most of the internal and Internet DC considerations are the same

IPv6 in the Enterprise Data Center
Biggest Challenges Today

- Application support for IPv6 – Know what you don't know
  - If an application is protocol centric (IPv4):
    - Needs to be rewritten
    - Needs to be translated until it is replaced
    - Wait and pressure vendors to move to protocol agnostic framework
- Deployment of translation
  - NAT64 (Stateful for most enterprises)
  - Apache Reverse Proxy
  - Windows Port Proxy
  - 3rd party proxy solutions
- Network services above L3 (A short-term challenge)
  - SLB, SSL-Offload, application monitoring (probes)
  - Application Optimization
  - High-speed security inspection/perimeter protection

ACE + IPv6 / ASR + NAT64

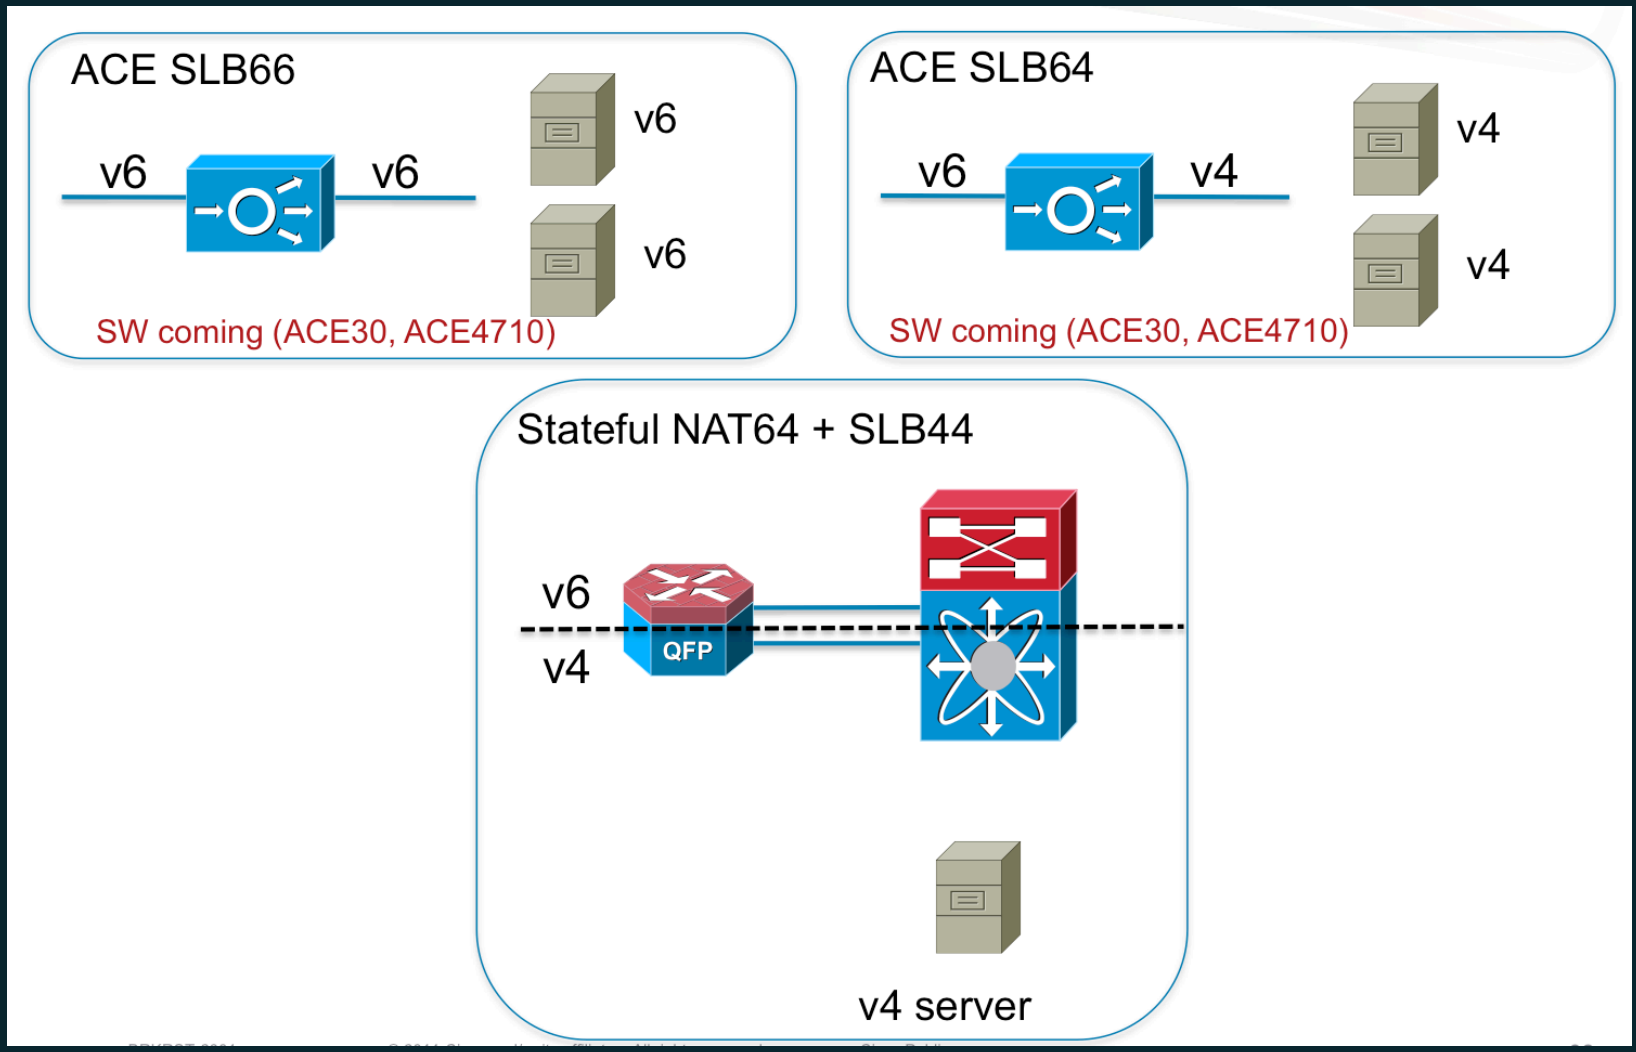

NAT64

- Two flavors – Stateless and Stateful
  - draft-ietf-behave-v6v4-xlate-xx (and others associated with that draft)
  - draft-ietf-behave-v6v4-xlate-stateful-xx
- Stateless – Not your friend in the enterprise (corner case deployment)
  - 1:1 mapping between IPv6 and IPv4 addresses (i.e. 254 IPv6 hosts-to-254 IPv4 hosts)
  - Requires the IPv6-only hosts to use an "IPv4 translatable" address format
- Stateful – What we are after for translating IPv6-only hosts to IPv4-only host(s)
  - It is what it sounds like – keeps state between translated hosts
  - Several deployment models (PAT/Overload, Dynamic 1:1, Static, etc…)
  - This is what you will use to translate from IPv6 hosts (internal or Internet) to IPv4-only servers (internal DC or Internet Edge)

Dual Stack the Internet Edge

- Dual stack the same network you have
- If not, do just enough IPv6-only to get you going
- Most design elements should be the same as with IPv4 (minus pure NAT/PAT)
- You may have to embrace SLB64/Proxy/NAT64 for IPv4- only apps

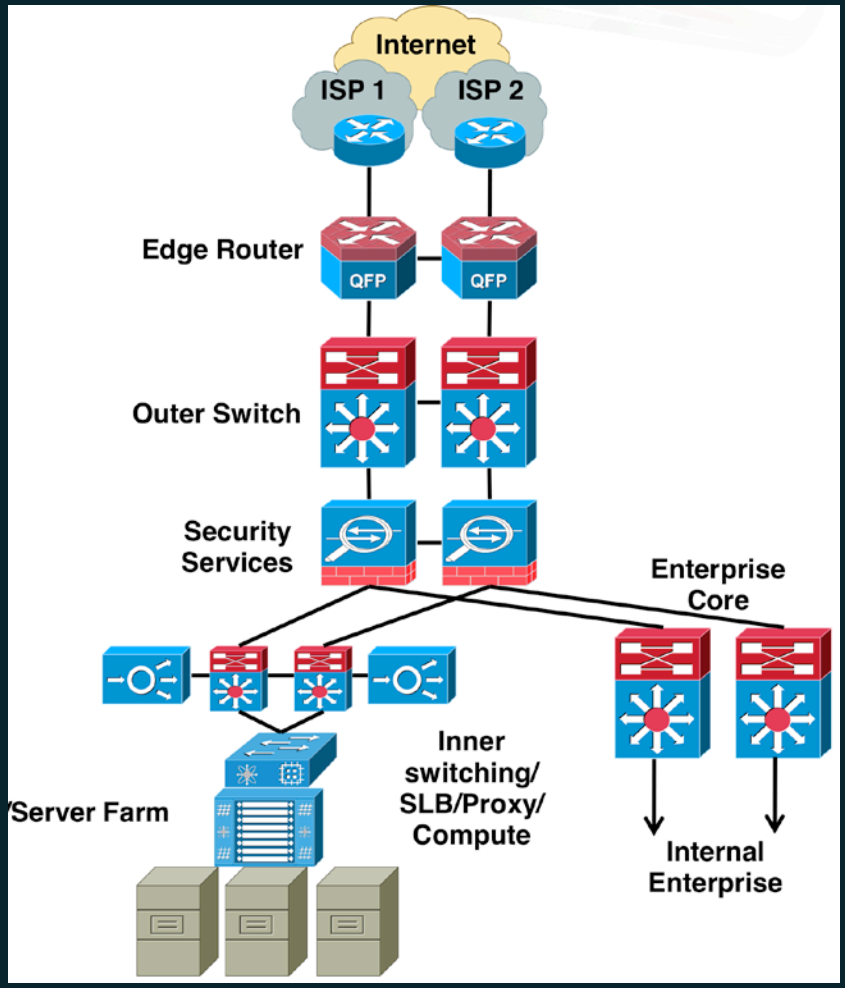

What if I Can't Dual Stack My Edge?

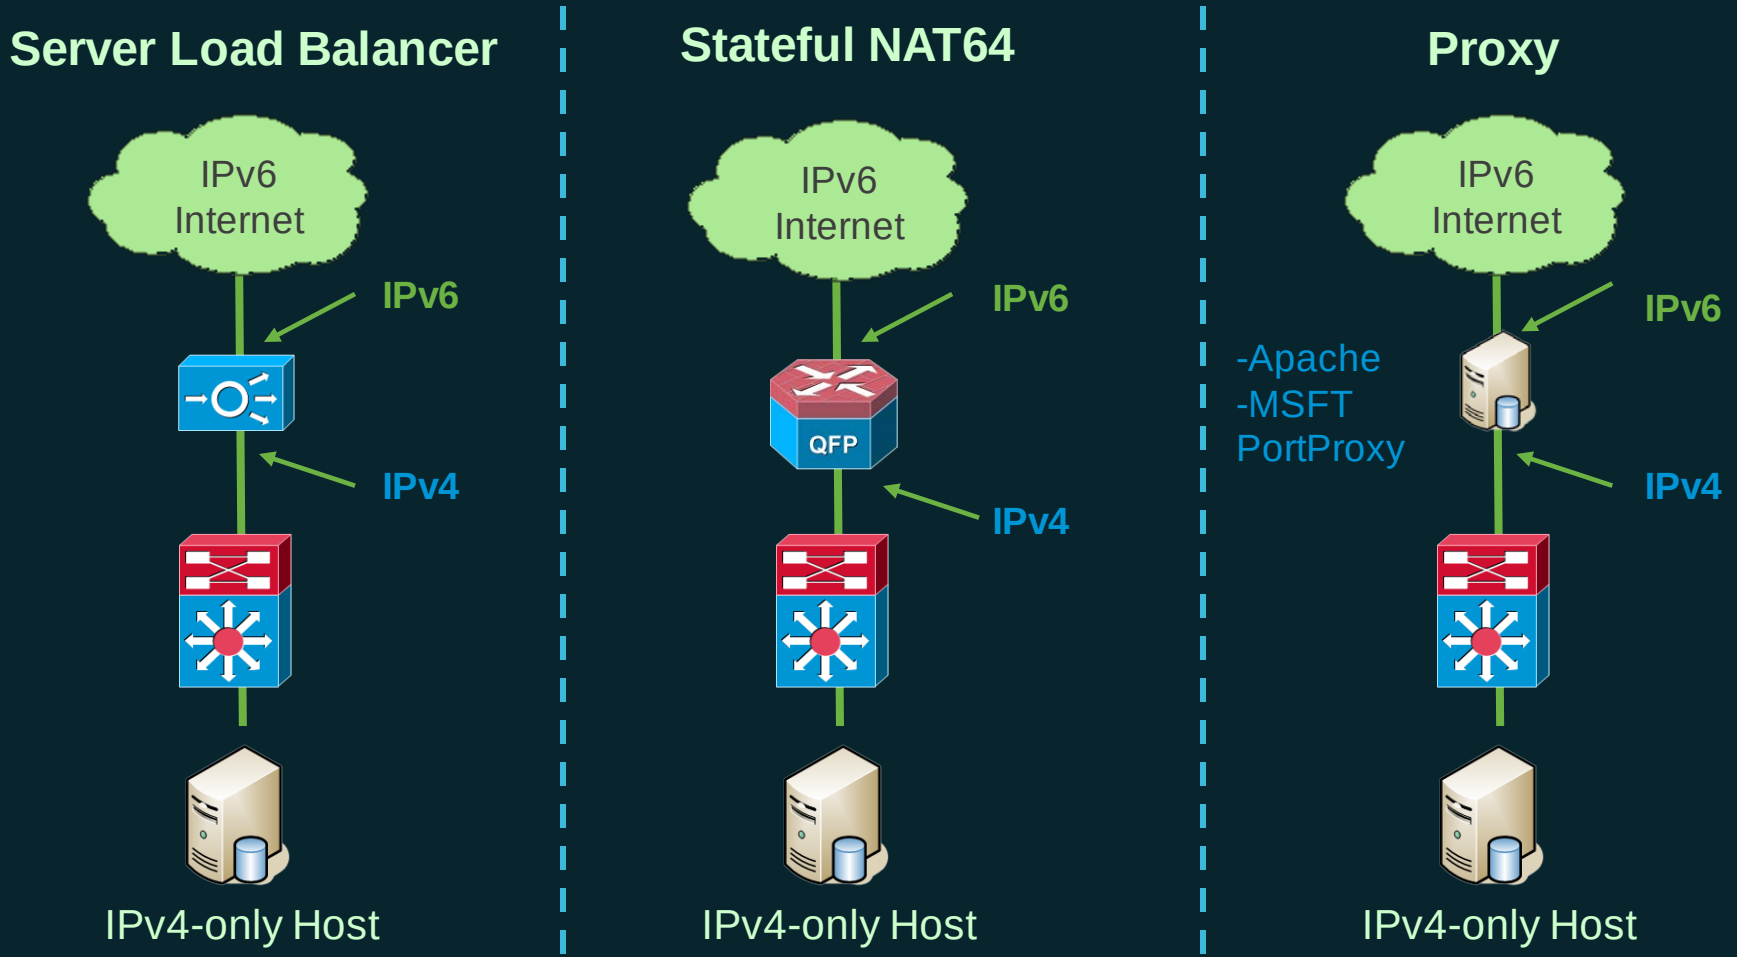

Internet Edge - to – ISP
Boatloads of options

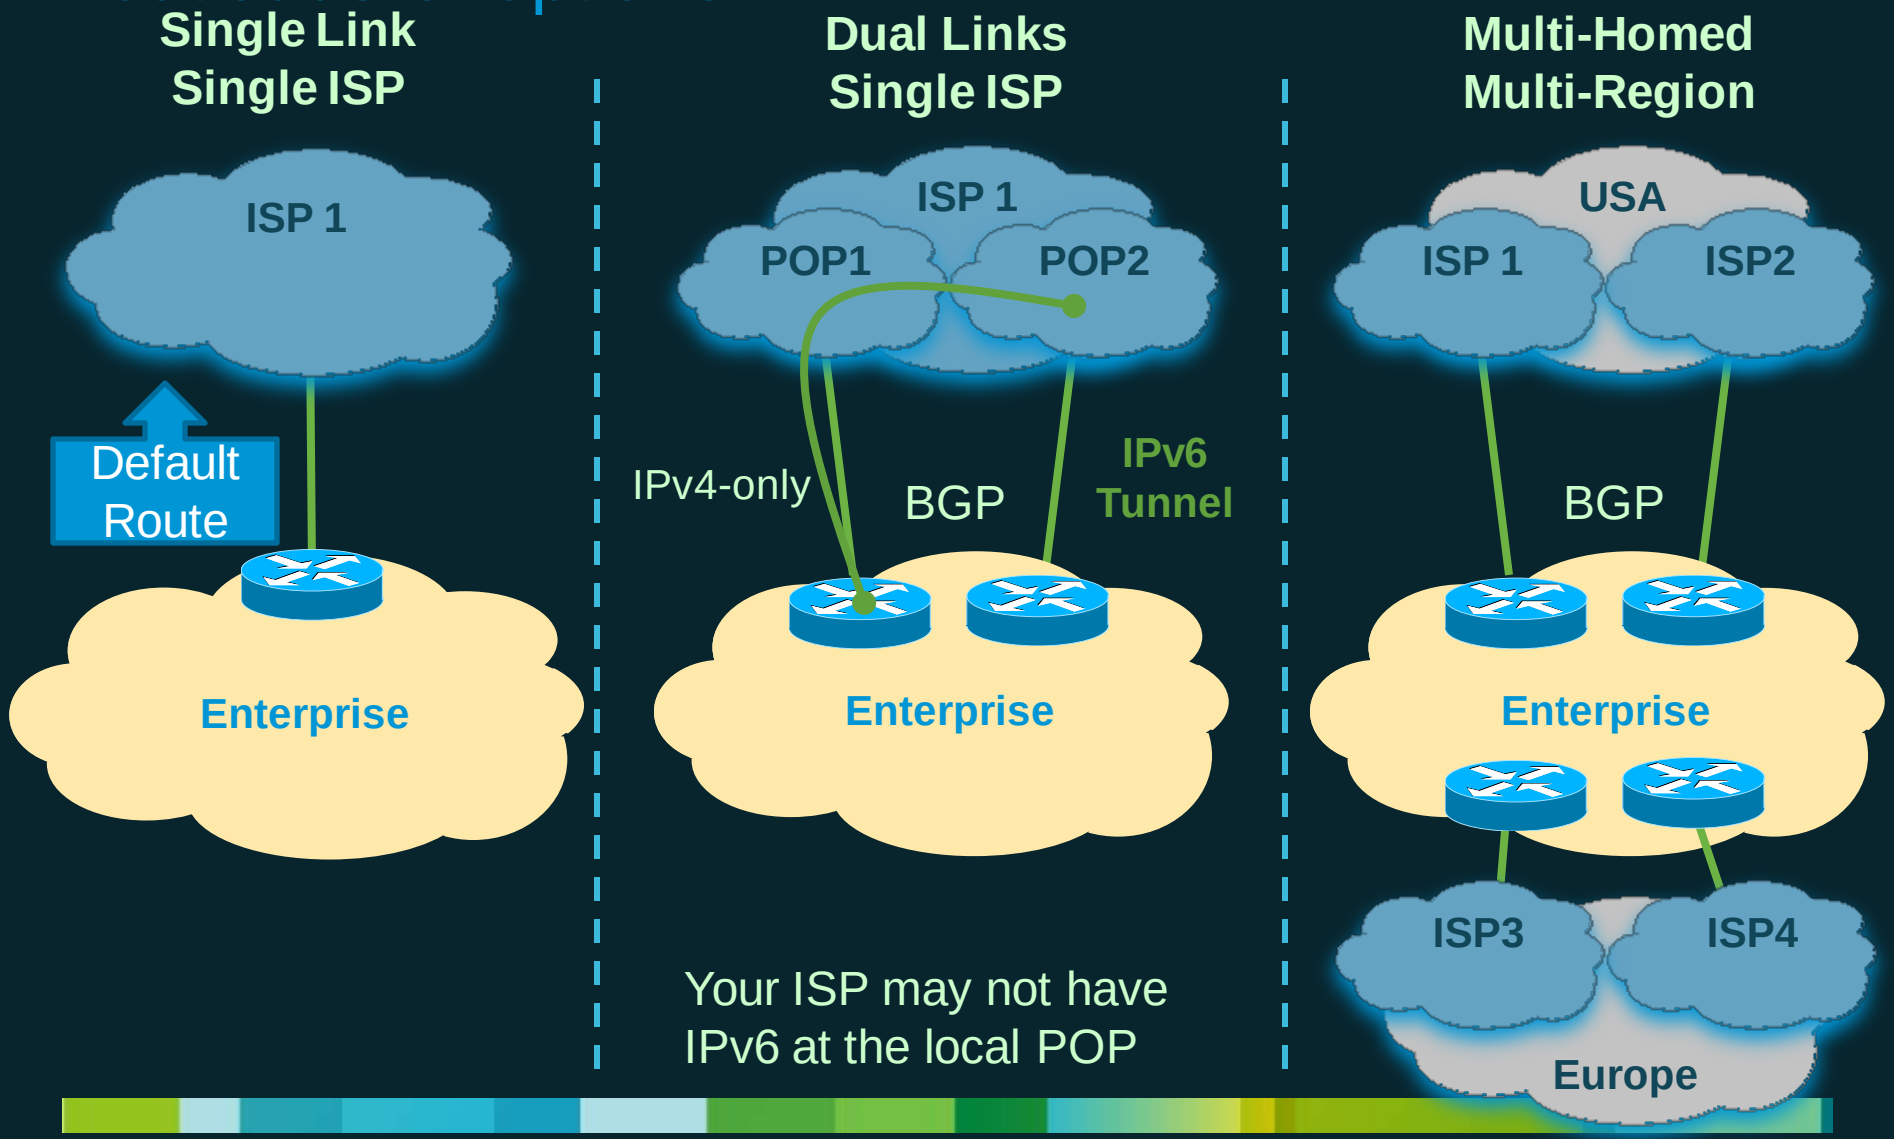

Infrastructure deployment

WAN/ Branch

Deploying IPv6 in Branch Networks:
http://www.cisco.com/univercd/cc/td/doc/solution/brchipv6.pdf

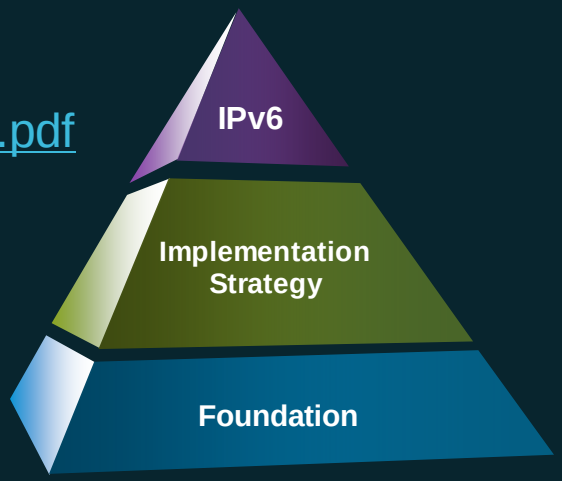

WAN/Branch Deployment

- Cisco routers have supported IPv6 for a long time
- Dual-stack should be the focus of your implementation—but, some situations still call for tunneling
- Support for every media/WAN type you want to use (Frame Relay, leased-line, broadband, MPLS, etc.)
- Don't assume all features for every technology are IPv6-enabled
- Better feature support in WAN/branch than in campus/DC

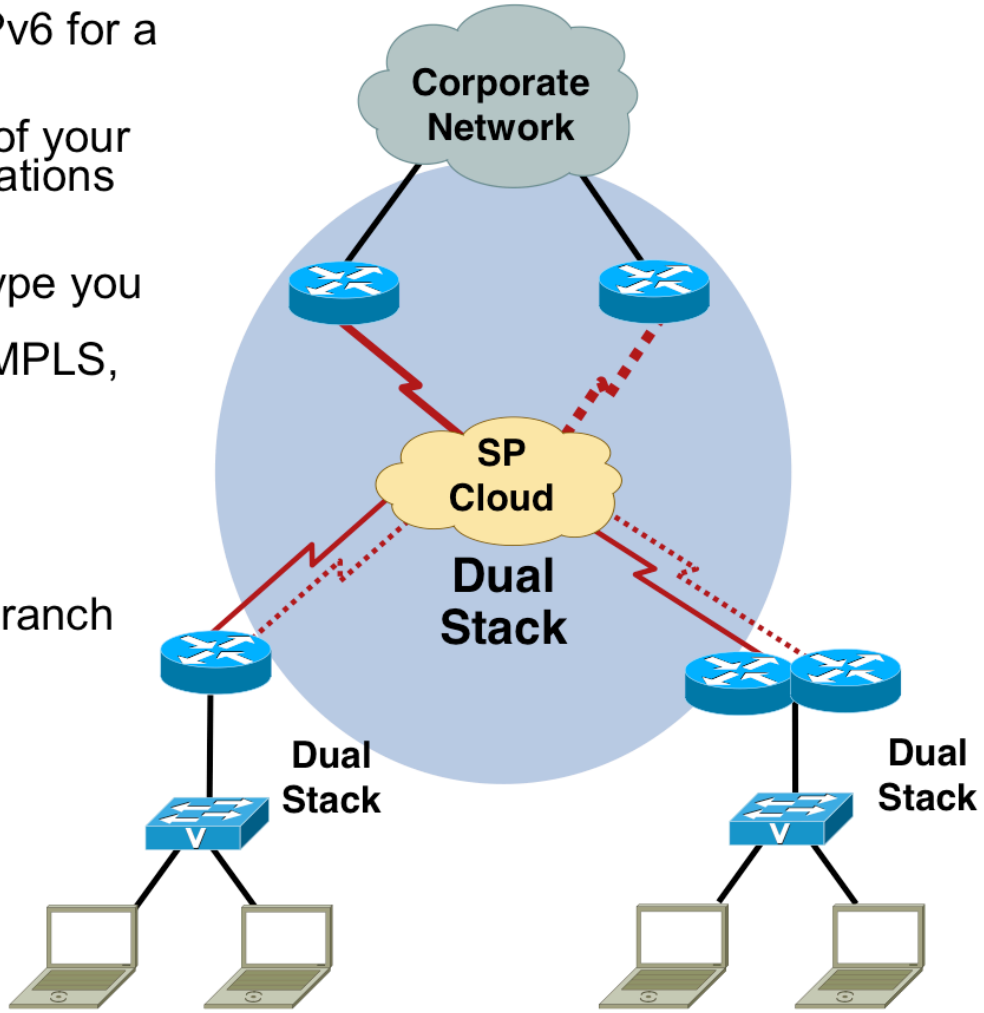

IPv6 Enabled Branch

Focus more on the provider and less on the gear

Branch Single Tier

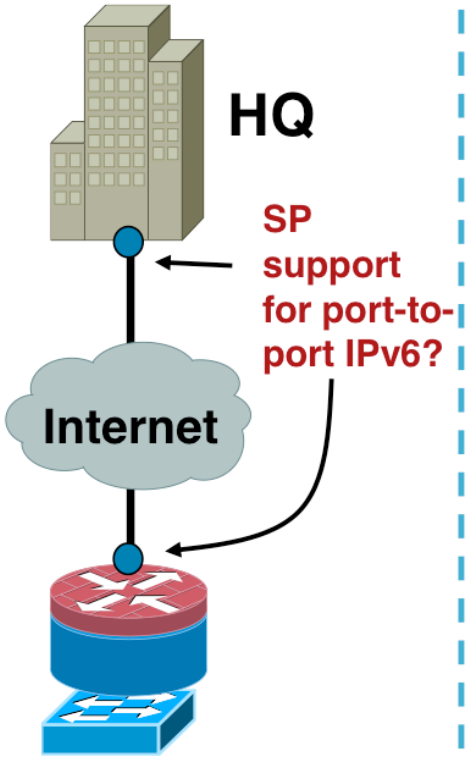

Dual-Stack
IPSec VPN (IPv4/IPv6)
Firewall (IPv4/IPv6)
Integrated Switch (MLD-snooping)

Branch Dual Tier

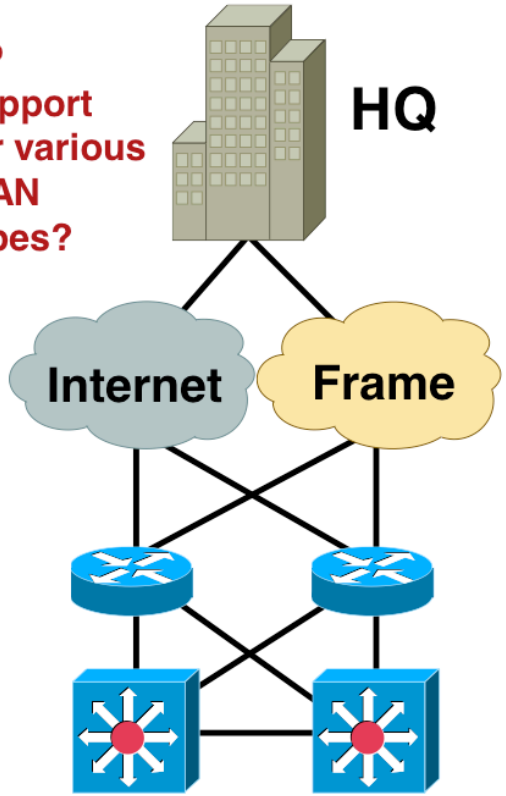

Dual-Stack
IPSec VPN or Frame Relay
Firewall (IPv4/IPv6)
Switches (MLD-snooping)

Branch Multi-Tier

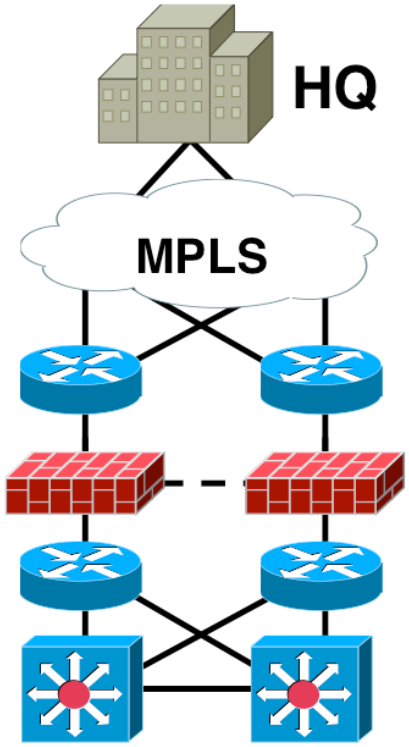

Dual-Stack
IPSec VPN or
MPLS (6PE/6VPE)
Firewall (IPv4/IPv6)
Switches (MLD-snooping)

IPv6 Enabled Branch

Focus more on the provider and less on the gear

- Mixture of attributes from each profile
- An example to show configuration for different tiers
- Basic HA in critical roles is the goal

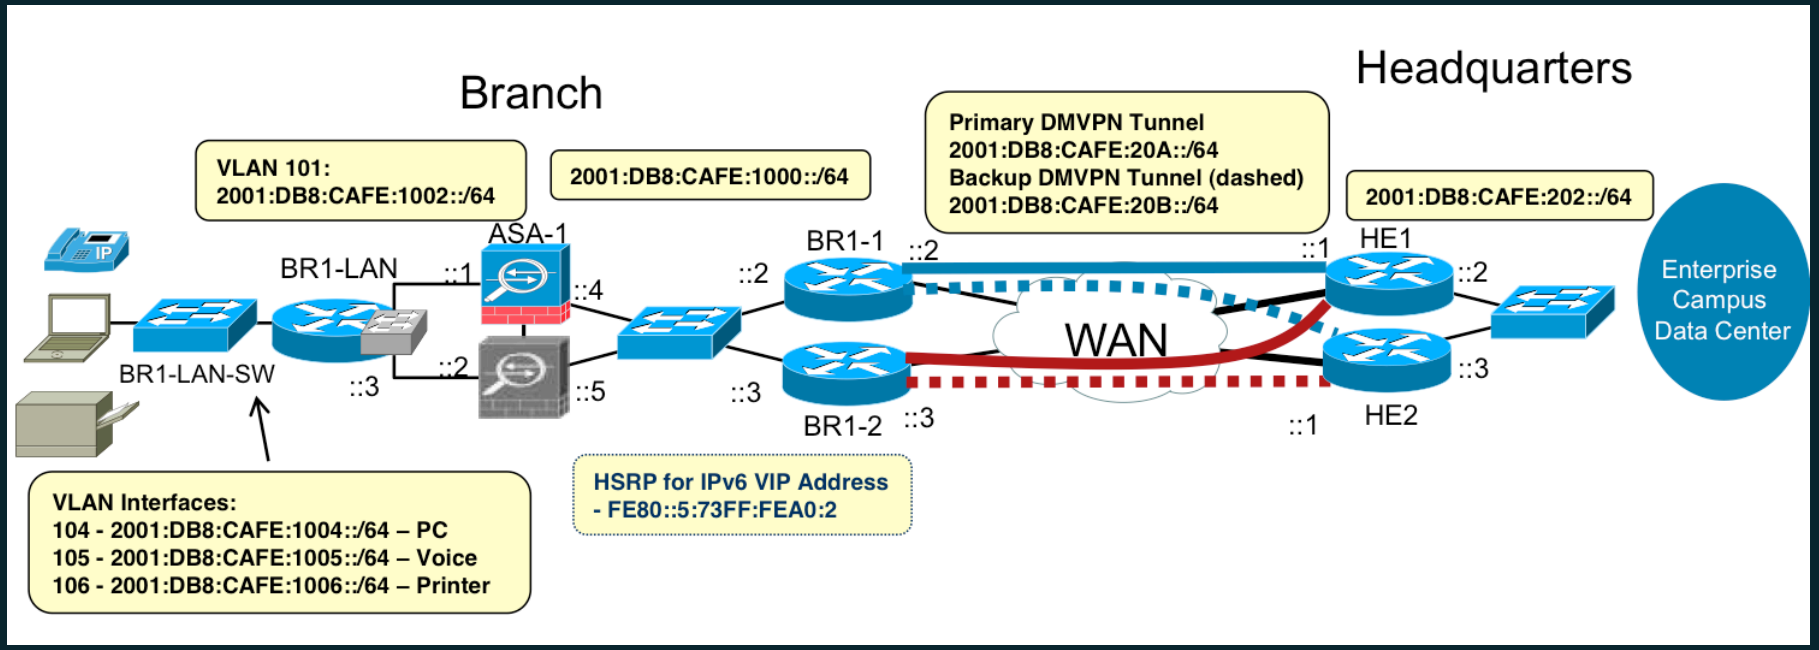

IPv6 Enabled Branch

Focus more on the provider and less on the gear

- Mixture of attributes from each profile
- An example to show configuration for different tiers
- Basic HA in critical roles is the goal

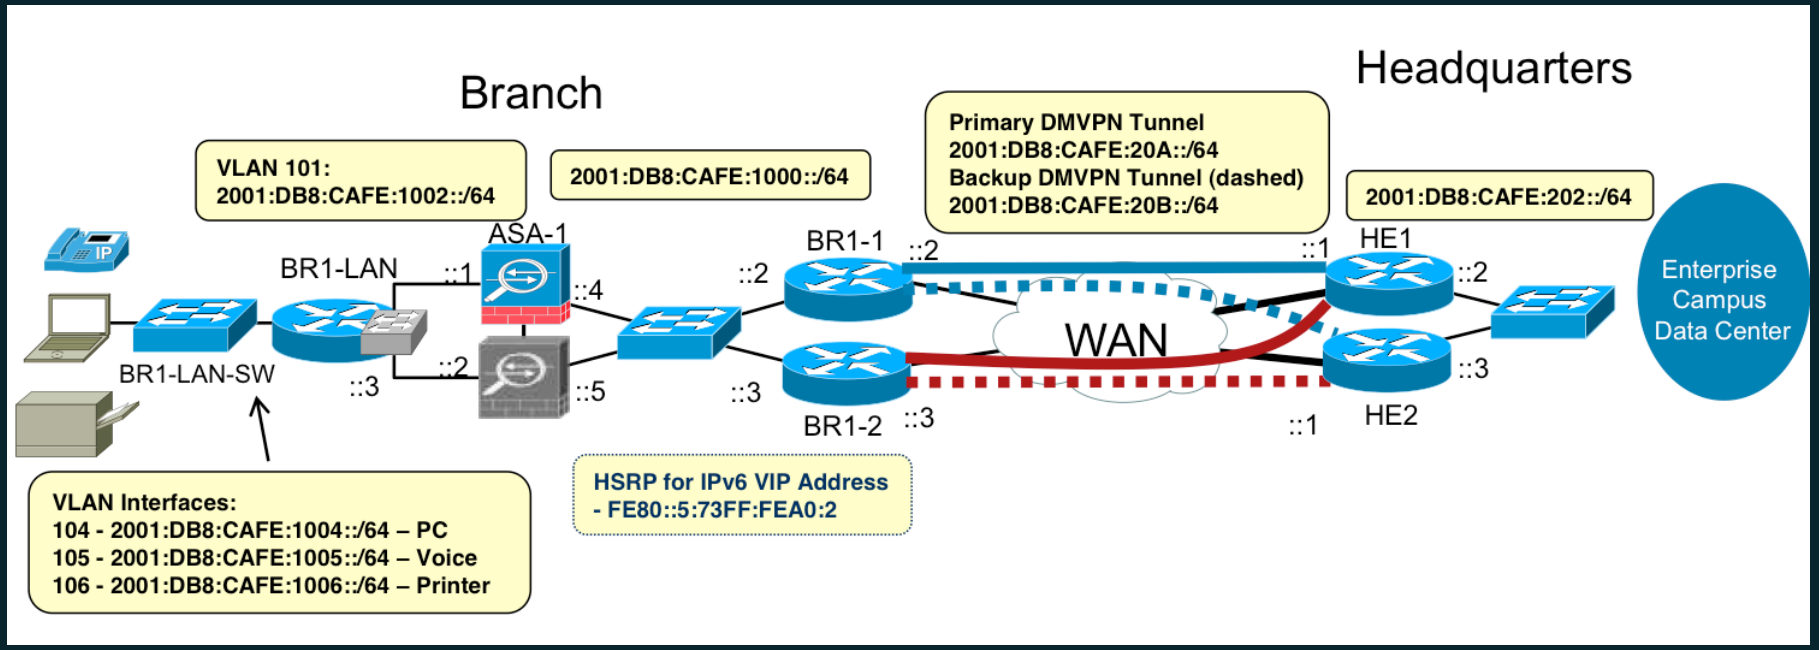

Thank you.