Cisco IT – Building an IPv6 Only Network

Deploying IPv6 only in SJC23

Khalid Jawaid CCIE 6765
Solutions Engineer, Global Infrastructure Services, Cisco IT
30th Oct 2017

Acknowledgements

Great Team Behind This

- **Ben Irving (Sponsor Director)**
- **Travis Norling (Manager ETE)**
- **Hitesh Panchal**
- **Charles Radke**
- **Norman Fong**
- **Tsung Chan**
- **John Banner**
- **Many More!**

Agenda list

1. Cisco IT Overview
2. IPv6 Only in Building 23 and Issues
3. IPv6 Only DC Plans and Issues
4. Q/A – Interactive Discussion

Cisco IT Overview

- 50,000+ Devices
- 300+ locations in 92 countries
- 500+ buildings
- 200,000 Sq Ft of DC space
- 1000+ labs worldwide
- 150,000+ Users
- ~ 5 Million IP Addresses (All Inclusive)
- ~ 6800 Applications

Cisco IT Overview

- 11 iPoPs advertising Cisco IPv4/IPv6 space
- EIGRP for IPv6/IPv4 + BGP
- Dual Stacked Everywhere (ExceptExtranet and CVO)
- Dual Stack DC Gateways (not server VLANs)
- Management over IPv4 (Except IPv6 Service Monitoring and SJC23)
- CNR for DHCP Services

Cisco Global Internet Presence IPv6 Advertisements (ARIN 2001:420::/32)

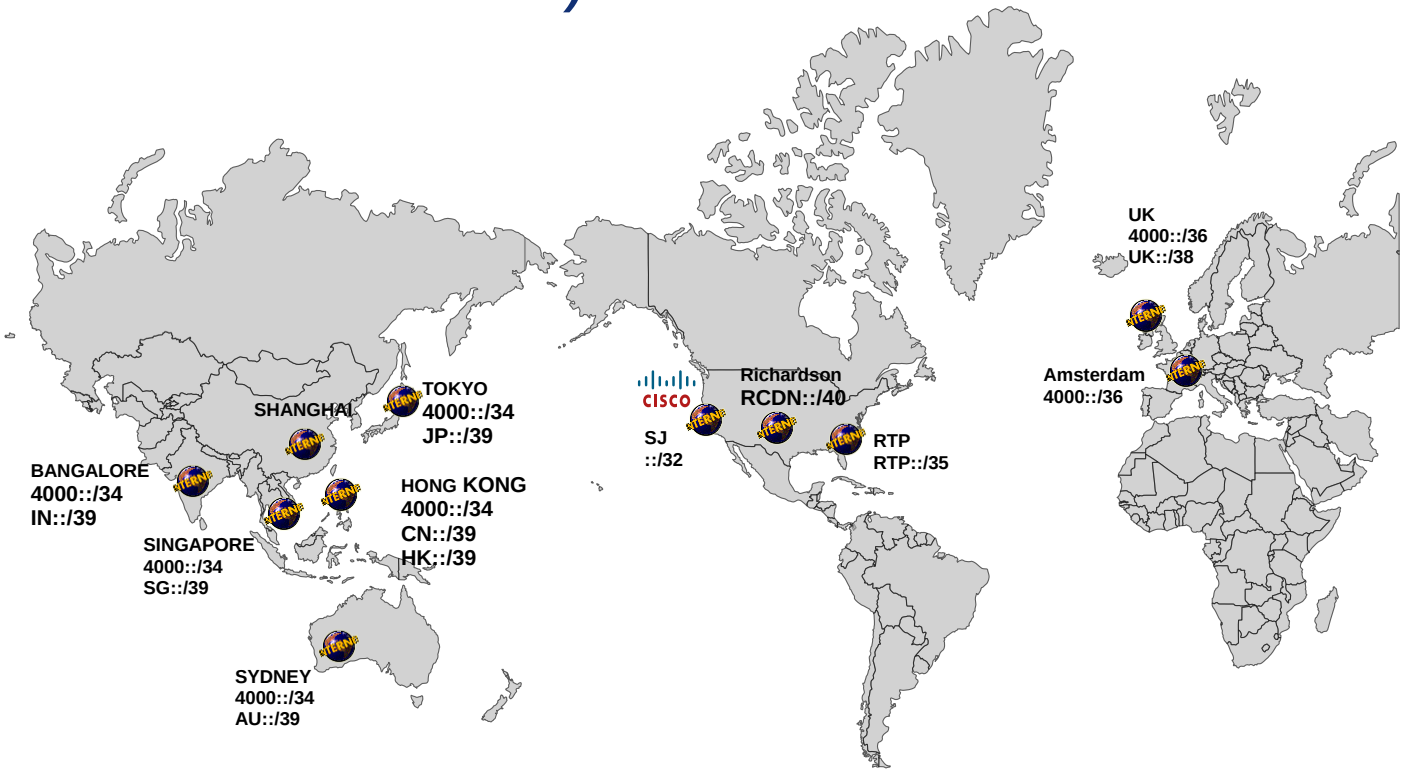

Our IPv6 Timeline

2010 – 2016 – Dual Stack

2016 – 2018 – Dual Stack + IPv6 Only

SJC23 – IPv6 Only

RTP IPv6 Only DC POD

2018/19 – 20??

Training / Development

IPv6 Only Mandate (New Apps)

An IPv6 Only Experience

Goals

User Experience

Product Gaps

Operational Gaps

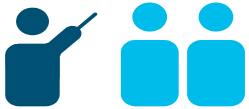

Knowledge/
Awareness

SJC23 – IPv6 Only Access

Target

- Single Campus Building
- Wired and Wireless
- Android and iOS
- NAT64/DNS64
- Management + Data
- UC / Collaboration

Physical Topology – IPv6 Only @ SJC23

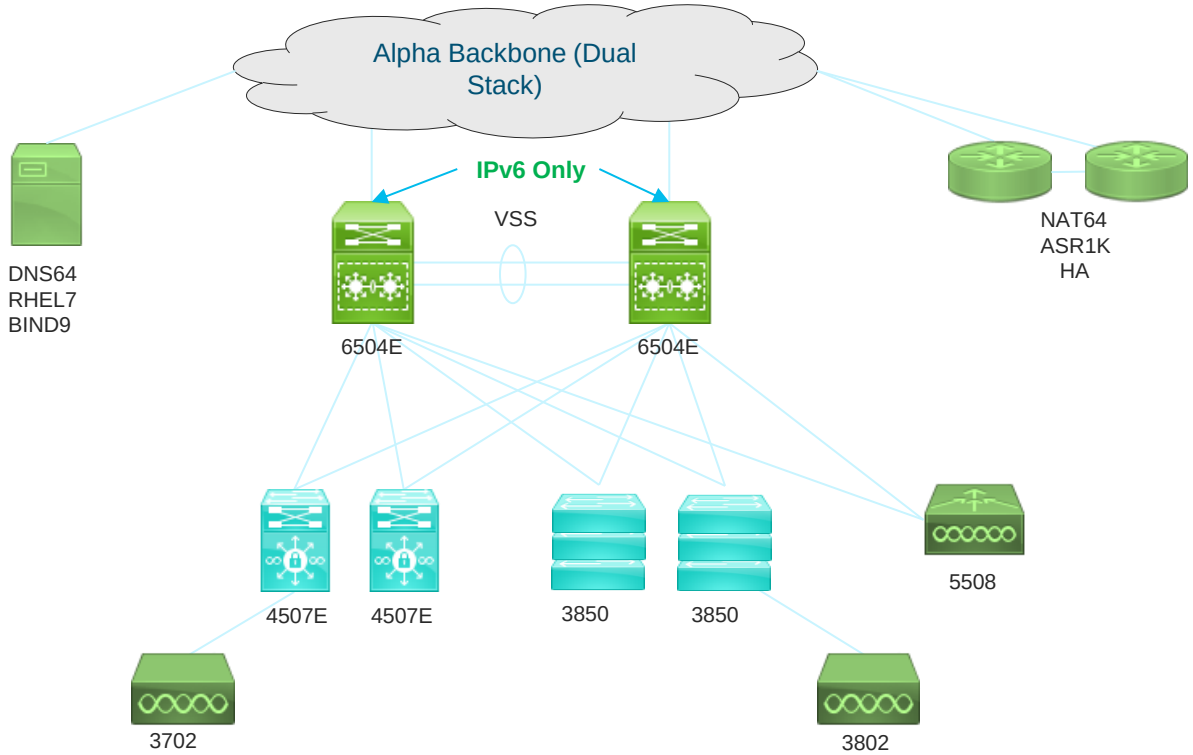

NAT64 Topology – IPv6 Only @ SJC23

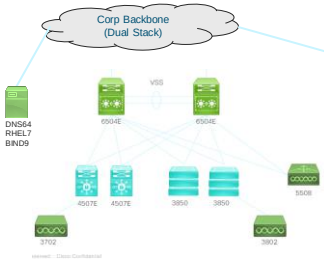

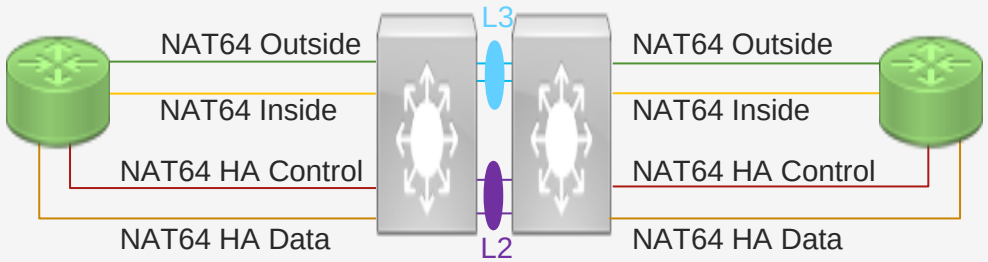

L2 Adjacency for NAT64 HA

Products Used SJC23

ASR1K

6504E

WLC 5508

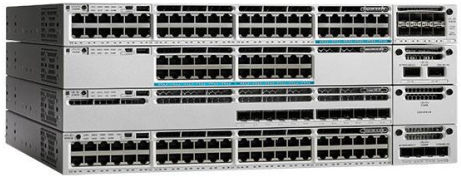

3850

AP 3702/3802

4507R+E

IPv6 Features Deployed

- HSRPv2 for IPv6 (First Hop Routing)
- First Hop Security
  - IPv6 Snooping (DHCPv6 Guard, Destination Guard, DHCPv6 Binding)
  - ND Inspection
  - RA Guard
  - uRPF
- DHCPv6 Stateful (Default and Preferred)
- SLAAC (Special case)
- EIGRP for IPv6
- NAT64/DNS64

Statistics

- Average 300 Users, peak 500
- 3 Months (start to finish)
- Approx. 7 – 8 engineers
- Average Traffic 250 Mbps (v6 Only Links)
- Average 32K NAT64 Xlate Entries

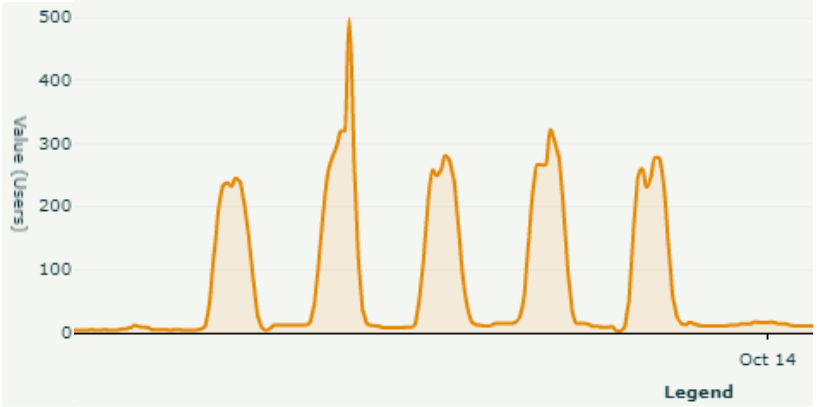

Charts (NAT64 Xlates)

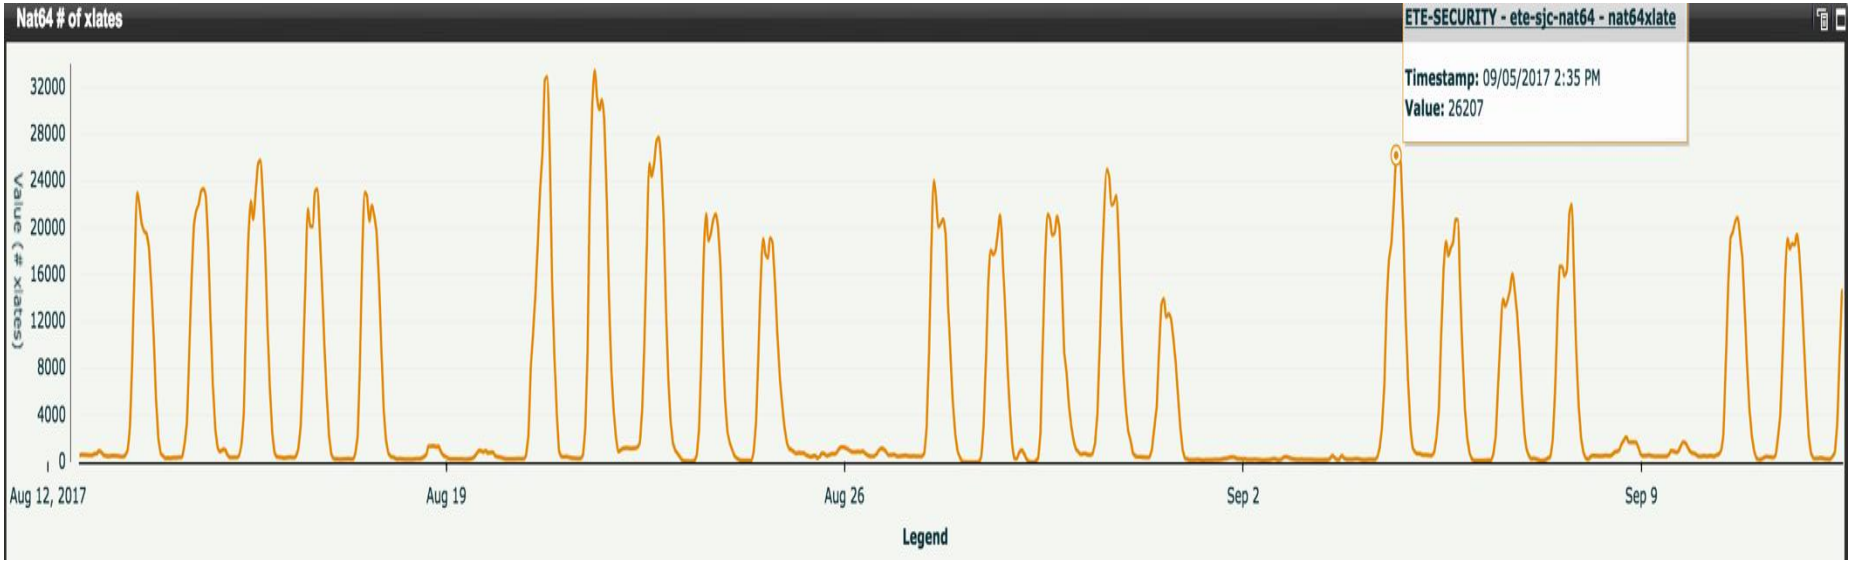

Not available via SNMP, gather with a script

Issues and Challenges

Issue – NAT64 fallback fails with SERVFAIL

Problem Statement: No fall back to NAT64 upon receiving a SERVFAIL or FORMERR.

Symptoms: No Connectivity to website

Diagnosis: NAT64 does not create a synth AAAA back to client if it gets a SERVFAIL

Workaround: Create a master zone on Cisco DNS64 for destination and get manually synth AAAA (Problems when the destination fails over)

- **LTF**: Webex upgrade of GSS

Issue – AnyConnect client fails on MAC

Problem Statement: AnyConnect client keeps reconnecting on MAC

Symptoms: No Connectivity – Client Reconnecting

Diagnosis: AnyConnect client software issue with NAT64 headend causing fragmentation. Client dropping TCP Fragments due to implicit filtering breaking TLS connection causing reconnecting loops. Also impacts IPSec/DTLS Tunnels

Workaround: No Workaround

- **LTF**: Fixed in AnyConnect Client ver 4.4MR3+

Issue – Spark Web Clients not IPv6 Ready

Problem Statement: Web based Spark Clients not working. Client apps working across all platforms for all services

Symptoms: No Connectivity/Calling/Services

Diagnosis: Web Client connectivity Infrastructure is not IPv6 enabled

Workaround: No Workaround

- **LTF**: IPv6 Enable Web client infrastructure

Issue – NAT64 – No SNMP MIB for Xlates

Problem Statement: Can’t poll IOS-XE for Xlate data using SNMP

Symptoms: No SNMP data

Diagnosis: Not supported

Workaround: Use a script to collect Xlate output via SSH/CLI

LTF: CSCvc13935 bug filed as Enhancement Request

Issue – Jabber/Phones Fail to Register

Problem Statement: Jabber clients failed to register with CUCM

Symptoms: No Registration

Diagnosis: IPv6 support is not available for Jabber clients below CUCM Ver 12.0

Workaround: No Workaround

- **LTF**: Upgrade to CUCM 12.0 – After upgrade, all features / services working

Issue – NAT64 – No SNMP MIB for Xlates

Problem Statement: Can’t poll IOS-XE for Xlate data using SNMP

Symptoms: No SNMP data

Diagnosis: Not supported

Workaround: Use a script to collect Xlate output via SSH/CLI

LTF: CSCvc13935 bug filed as Enhancement Request

Misc Issues

- Not all apps/ drivers in standard Cisco Desktop image ipv6 ready. Needed latest updates
- IPv4 Literals – Can’t do DNS64 and therefore no NAT64
- 802.1X – Need platform support across 4500 (In Development) – use SGT/SGACLs as workaround

Web Tools with interesting stats

- AAAA and IPv6 connectivity statistics of top websites according to Alexa - http://www.employees.org/~dwing/aaaa-stats/
- NAT64Check
  - https://nat64check.go6lab.si/
- Google’s DNS64 Service
  - https://developers.google.com/speed/public-dns/docs/dns64

IPv6 Only DC (PoC Stages)

- Single Pod (ACI)
- Data plane only
- NAT64/DNS64
- Stateless and Stateful NAT

Products Used DC Pod

Nexus 7K

Nexus 9K
C9396PX

ASR1K

High Level Topology

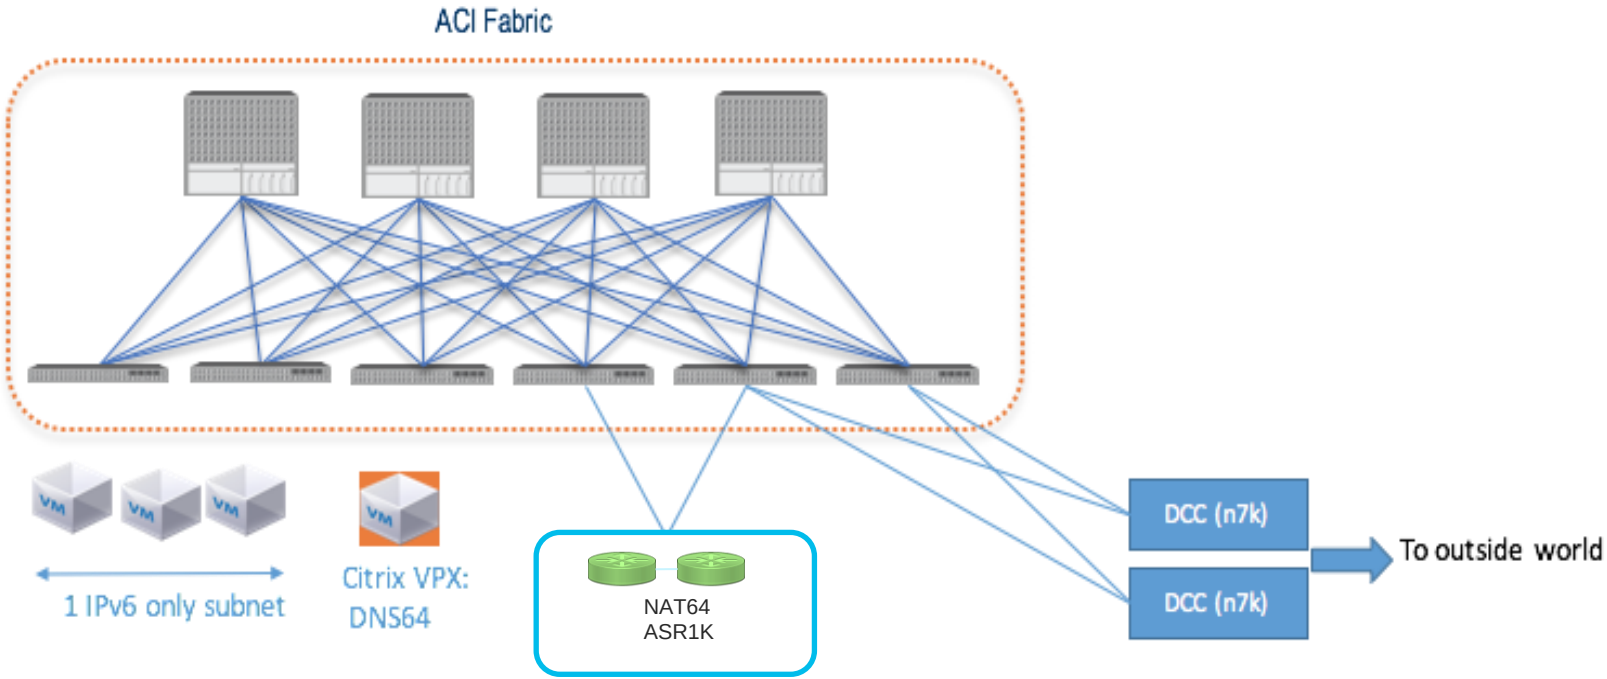

Gaps and Issues

- WAAS – Does not support IPv6 yet
- Kubernetes – IPv6 Not Supported / In Dev
- PXE Boot – Not supported over IPv6
- Storage – IPv6 only not tested – IPv4 must be served as long as it exists or storage pools will be fragmented (cost and operational impact)
- More as we further develop the design / get into deployment.

Questions?

1. How do you measure User Experience?
2. Where should IPv6 go first? DC, non-DC?
3. How do you handle privacy extensions?

Key Takeaways

- Measure User Experience = Metric for success
- Some websites required internal zone creation (We did that for high impact sites that failed)
- InfoSec and related tooling is critical. Ensure that the necessary compliance is still there and working (Privacy Extensions for eg.)
- It does work – failure scenarios will mostly be specific
- Finally, there is a price to pay. For some time, IPv6 development will trail behind latest tech/features which may be IPv4 only.


CISCO

IPv6 Address Plan (Top Level)

- Global 2001:420::/32
  - Americas 2001:0420::/34
  - EMEA and Asia Pacific 2001:0420:4000::/34
  - Global Spare1 2001:0420:8000::/34
  - Global Spare2 2001:0420:C000::/34
  - Global Infrastructure 2001:0420:C000::/42
  - Global Mobility 2001:0420:C040::/42

Address Overview

Breaking down the /32

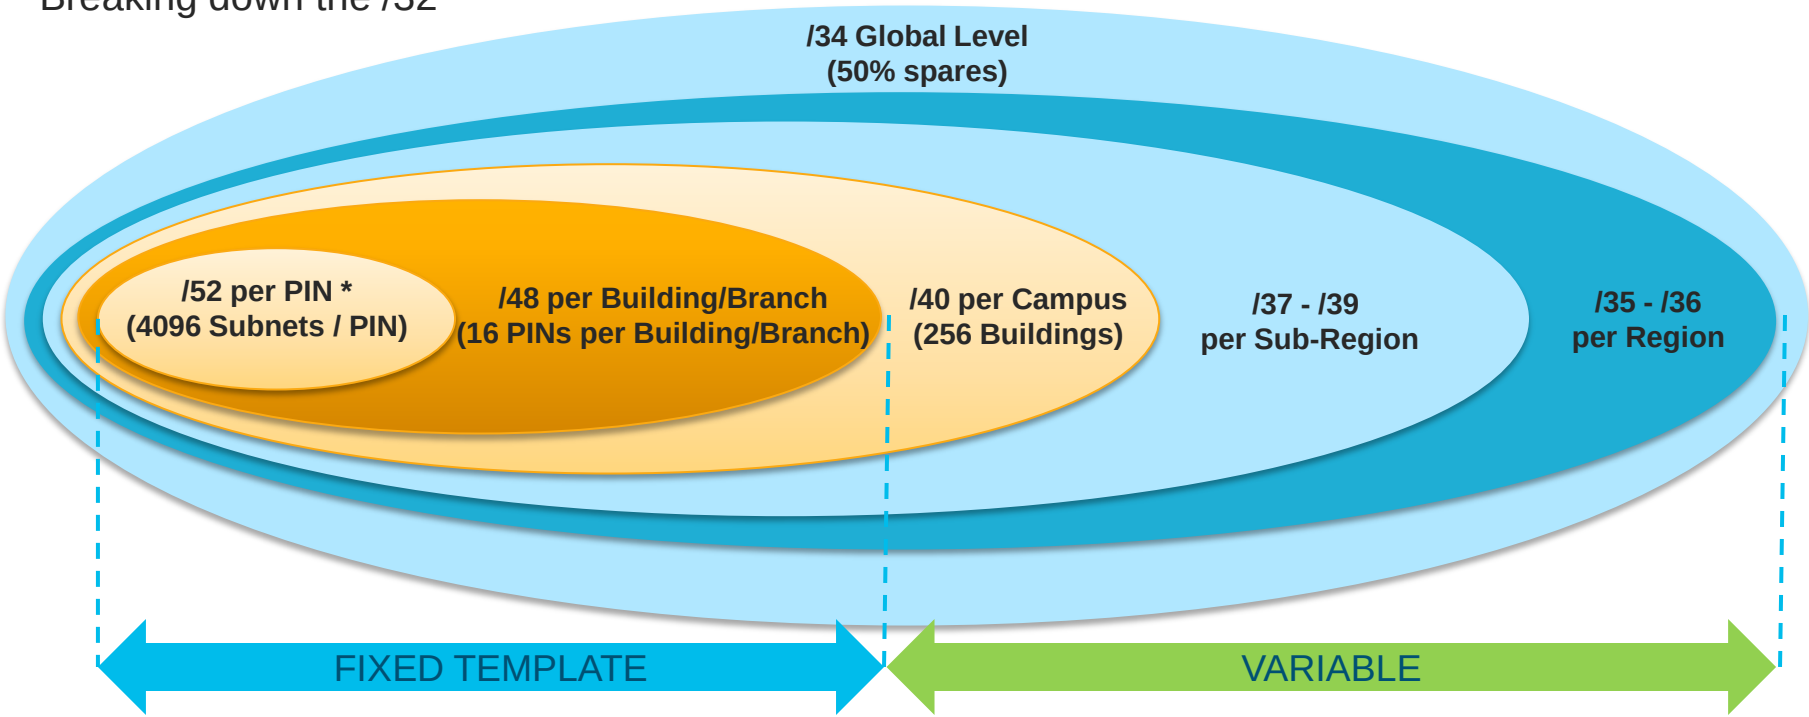

***** PIN = Place In the Network - A framework to classify functional areas of the network
eg, Lab, Desktop, DC, DMZ etc

Address Planning

Template Addressing

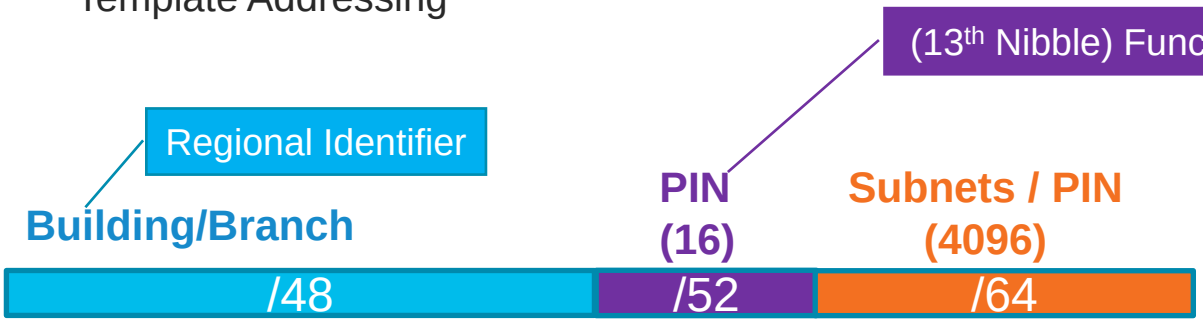

0 = Infra
1 = Desktop / Wireless
2 = Lab
3 = Guest
4 = DMZ
D = Building DC
... etc

2001:0420:028C:1000::/52 - Desktop PIN

2001:0420:028C:1300::/64 – Desktop VLAN 300

2001:0420:028C:1301::/64 – Desktop VLAN 301

2001:0420:028C:2000::/52 - Lab PIN

2001:0420:028C:2001::/64 – Lab Subnet 1

2001:0420:028C:2002::/64 – Lab Subnet 2

Charts

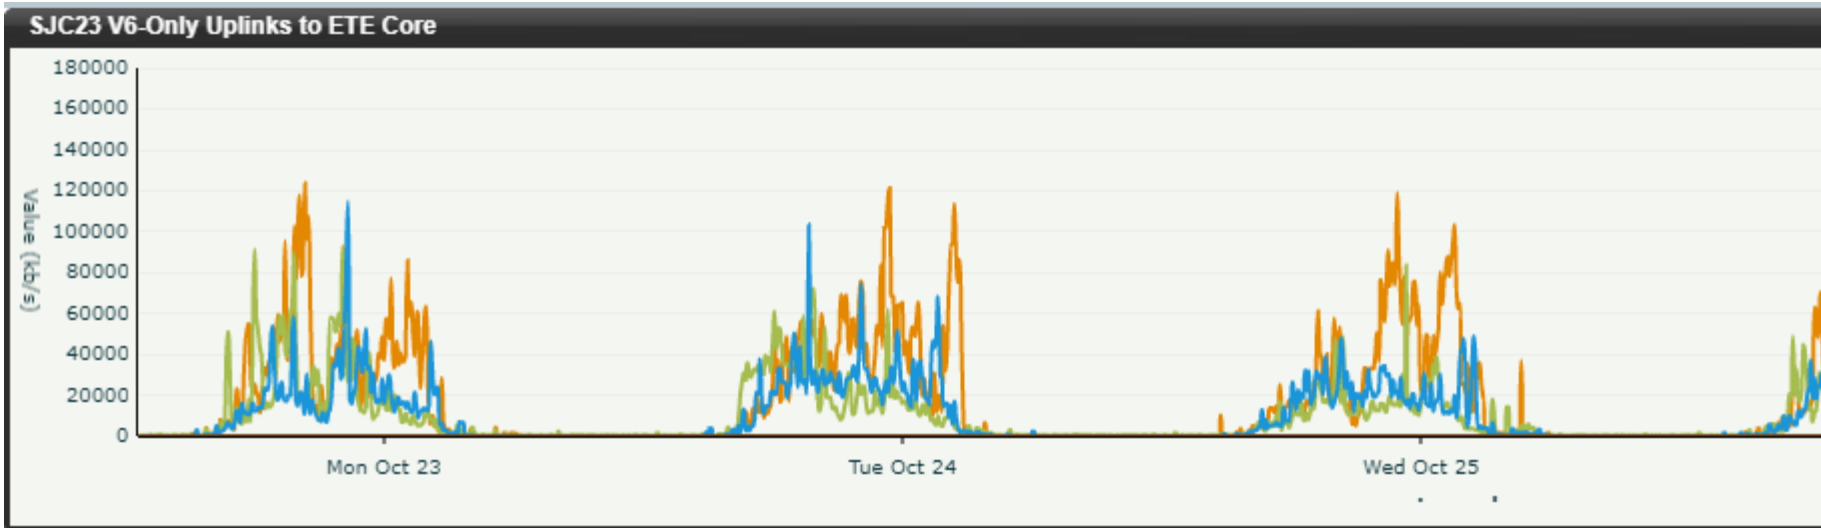

Questions

- How are we handling legacy v4 embedded in apps. Do we use 464XLAT? If yes, how do we plan on retiring it?
  - Ans – 2 situations :
    - 1) Embedded IPv4 literal will fail in ipv6 only
    - 2) Host resolving to v4 will use NAT64 and leave a foot print
- Did we include BMS, HVAC, etc? Where is the NAT64 gateway?
  - Ans: Only users and Network Infra. NAT64 is at site/network core of campus