Cisco live!

February 15 – 19, 2016 • Berlin, Germany

We're ready. Are you?

How to write an IPv6 Addressing Plan

Veronika McKillop, CSE

Wim Verrydt, CSA

BRKRST-2667

Cisco live!

The presenters

Veronika McKillop

Consulting Systems Engineer

CCIE #23705

Wim Verrydt

Solutions Architect

CCIE #8118

Abstract reminder

An **IPv6 addressing plan** is the **cornerstone of a successful IPv6 deployment**. The huge addressing space available provides flexibility, we have never experienced with IPv4. At the same time may become a source of frustration - how should I deal with this? **How** should I **carve up the IPv6 prefix** in such a way as to meet today's needs while catering **for future requirements**?

This session will help the attendees to **learn about** the **best practices for** writing an **IPv6 addressing plan**. We will focus on Enterprise, SP and, newly in 2016, on DC/Cloud IPv6 Addressing practices.

Also, a part of the session will be dedicated to a **practical exercise**, which will give the participants an opportunity to apply the newly acquired knowledge to a real world example.

If you fail to PLAN (an address plan), you plan to FAIL.”

Core Message of this Session

Key Takeaways

1. IPv6 addressing is easy and highly flexible.

1. IPv6 addressing strategy is the key to successful deployment.

1. A good addressing plan simplifies IPv6 network operation and troubleshooting.

Disclaimer

- This session is **focused** on **Unicast IPv6 Addressing** only
- We expect you to have:
  - A basic knowledge of IPv6 address types
  - A strong familiarity with IPv4 ☺
  - A good understanding of IP design in routed networks
  - An understanding how IP addresses are allocated and assigned on the Internet today

Agenda

- IPv6 Address Types – Refresher
- Typical Requirements
  - Enterprise
  - Service Provider
  - (Virtualized) DC / Cloud
- Methodology for writing the IPv6 Addressing Plan
- Exercise
- Final Thoughts & Conclusion

IPv6 Address Types – Refresher

Types of Unicast Addresses - RFC 4291

- (Node) Loopback Address
- **Link-Local Address (LLA) – fe80::/10**
- **Unique Local Address (ULA) – fc00::/7**
  - **Site-Local Address** has been deprecated by IETF (RFC 3879, September 2004)
- **Global Unicast Address (GUA) – 2001::/3**
- Anycast – it is the same as GUA
- **NOTE:** An interface will have multiple IPv6 addresses

Link-Local Address - RFC 4291

- Range: **fe80::/10**
  - No subnet significance
- Used for communication with hosts on the same link
  - Examples: Stateless Address Autoconfiguration (SLAAC), Neighbor Discovery, Duplicate Address Detection
- For link operation purposes
  - Leveraged by routing protocols and gateways
- Never routed to other links
  - No meaning outside the link
- Typically, first 64 bits are fixed, only Interface Identifier is modified
- **Example:** fe80::0224:d7ff:fe2c:7831

Unique Local Address - RFC 4193

For Your Reference

- Range: **fc00::/7**
  - Currently used fd00::/8
- Globally unique address for local communications
- **40-bit global ID** generated using a pseudo-random algorithm
- Not designed to be aggregated
- Not expected to be routed on the Internet but routable within an **administrative domain**
- Scope needs to be managed
  - ACLs and Prefix lists
  - Your upstream ISP will filter it anyway
- **Example**: fd68:df3d:80ee::/48 (LACNIC)

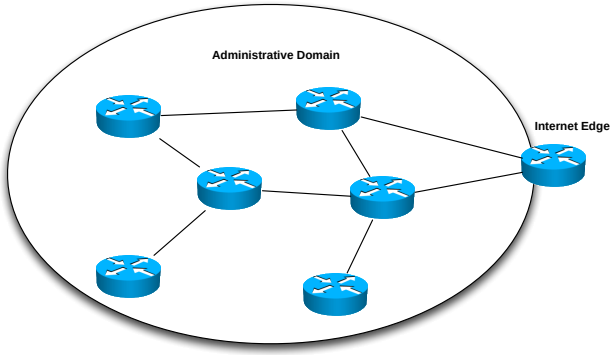

fc00:gggg:gggg:	ssss:	xxxx:xxxx:xxxx:xxxx

Unique-Local (ULA) – fc00::/7

Cisco live!

Global Unicast Address - RFC 3587

- Globally unique and routable
  - Defined for use across the IPv6 Internet
- Primary goal is to provide plenty of globally accessible addresses
- Reserved and identified by **high-level 3 bits** set to "**001**"
  - Range: **2000::/3**
- Global IPv6 Prefix received from an LIR or RIR
- Presence in Global Routing Table
  - Aggregation is critical
  - Hierarchical assignment enforced through IANA allocation policy

- **Example:** 2001:420:0:1::1

Have NO Fear! ☺

Source: http://www.sensiblehealth.com/Blog/wp-content/uploads/2014/08/rollercoasters-in-cities-venice-frozen-over-nois7-surreal-photos-images-manipulations-R.jpg

Agenda

- IPv6 Address Types – Refresher
- Typical Requirements
  - Enterprise
  - Service Provider
  - (Virtualized) DC / Cloud
- Methodology for writing the IPv6 Addressing Plan
- Exercise
- Final Thoughts & Conclusion

Enterprise

Where To Get an IPv6 Prefix?

- IPv6 Prefix assigned from:
  - an **Local Internet Registry (LIR)** - typical enterprise scenario or
  - directly from an **Regional Internet Registry (RIR)** – RIPE in Europe
    - typical for ISPs and large enterprises which span multiple countries, have dual-homing requirements (have AS number)
- Assigned prefix has fixed length, work is done with the bits between the assigned prefix and the /64

Global Unicast Address Example

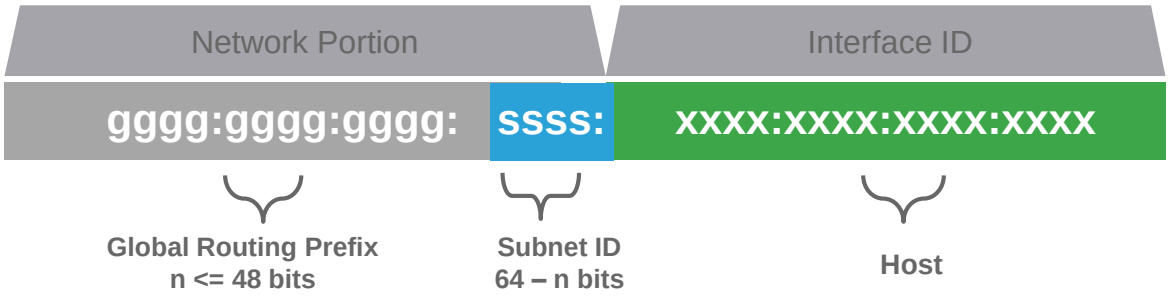

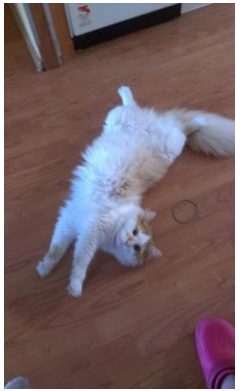

© K. Bednarova, 2015

- **Get cozy! The IPv6 addressing space is HUGE….**

Enterprise Requirements

- Different **locations**
- Different **places in the network** (PINs)
  - Campus vs. Branch vs. Internet Edge vs. DC
- Services **centralized** (one DC) or **distributed** (in branches/multiple DCs)
- **Encoding of information** within the IPv6 address (e.g. locations, PINs, services)
  - Use for accounting, administrative reasons, troubleshooting etc.
- **Transition mechanisms** deployed:
  - ISATAP – to bypass legacy equipment
  - NAT64 / SLB64 / NPTv6 – for the Internet Edge
- Who are the "**consumers**" of IPv6?
  - **Internal users** and their access to internal systems and IPv6 Internet
  - **External customers** and **IT systems** that enable interaction
  - **Partners** and **suppliers** (extranet or more often public Internet)
- **Security**
  - Easily manageable ACL
  - Exposing information about the network (e.g. VLAN number)
- **Rate of change and growth** – mergers & acquisitions

Considerations

- Usually splitting up **/48 per location**
  - Definition of a location?
    - E.g. single building vs. campus
    - Narrower than a city
- Multiple /48 prefixes are more likely
- **16 bits for the per-location addressing plan**
  - Important because it can help with identifying buildings within a location/subnets etc.
- **Aggregation** is very important
  - **# of required prefixes** at each level, **# of levels** required
  - NOT important: # hosts within subnet (**/64 = 2^{64}**)
- **Simplicity**
- Larger IPv6 prefix if the enterprise is big (possibly it is an PA directly form the RIR)
- **Have a well defined process and guidelines for IPv6 address allocations**

Service Provider

(Internet) Service Provider Requirements

Source: http://tinyurl.com/telco-map

- **Clear addressing** for different parts of network
  - Core business vs. internal enterprise vs. infrastructure
  - Example: **PIES** – private, internal, external, subscribers
- **Customers** – broadband subscribers / business customers
  - Reserved vs. assigned IPv6 prefixes (e.g. broadband: /56 allocated, /60 assigned)
  - Identification of services within the IPv6 address (aka **Prefix Coloring**)
- **Customer facing systems / services**
  - CDN
  - Cloud
  - Hosting offerings
  - Subscriber Access Types
- Aggregation
  - **At least /32** to work with, typically much more (**the ISPs can get /29 and they should obtain it!**)
- Scalability
- Stability

Considerations

- **Aggregation** of prefixes assigned to the customers
  - Per region, per PoP level, per BNG, etc.
  - Multihoming of business customers
- **Capabilities of network devices**
  - How many IPv6 prefixes (IPv6 customers) can your BNG/PE handle?
  - Leave room for growth
- Don’t hand out IPv6 prefixes on a wrong **bit boundary**
  - Multiples of 4 bits (nibble) to align on hexadecimal boundary
  - Check the **“readability” of the prefix = supportability**!
- **Ease of operation and troubleshooting** is absolutely necessary
- **Aggregation** to the upstream SPs
- **Have a well defined process and guidelines for IPv6 address allocations**
- **All the rules can be “bent” if required. You must know what you are doing!**

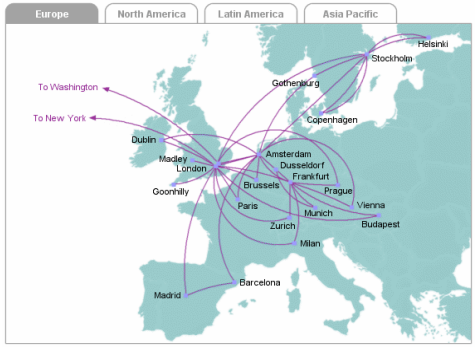

Source: http://www.bt.net/info/

Other Influences in the Industry

- **IEFT Homenet & IPSs' IPv6 Addressing**
  - This WG focuses on supporting next-generation services on unmanaged home networks
    - In the center of their work is **IPv6**
  - Multiple ISP connections to the home
    - Example: broadband, VPN router, smart meters, home security etc.
    - Terminated at a CPE (6rd, DHCPv6-PD, MAP, static IPv6 etc.)
  - **Not really impacting the ISP IPv6 addressing** as it's behind the CPE
    - Rather realize the potential **of prefix coloring & IPv6 Segment Routing**
  - Homenet presentation by Mark Townsley @ UKNOF 27 (January 2014)

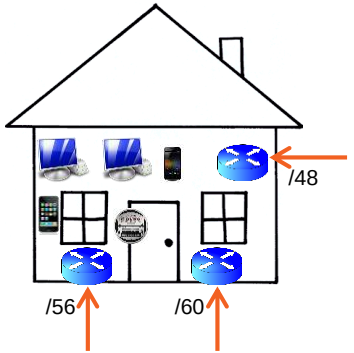

© V.McKillop, 2015

- **IETF v6 Ops & Unique IPv6 Prefix Per Host**
  - Large scale environments with the need to assign IPv6-prefix per host (E.g. SP Wi-Fi)
  - Advantages:
    - Monitoring the prefix instead of IPv6 address
    - Host isolation (prefix has an Off-link flag set), limitation in ND communication
  - Think about it from the perspective of the IPv6 prefix allocation from your RIR/LIR
  - **How many /64 are you going to need?**
    - This will impact the required allocation size
  - presentation at UKNOF 33 (January 2016)

(Virtualized) DC / Cloud

The biggest concern in the DC / Cloud?

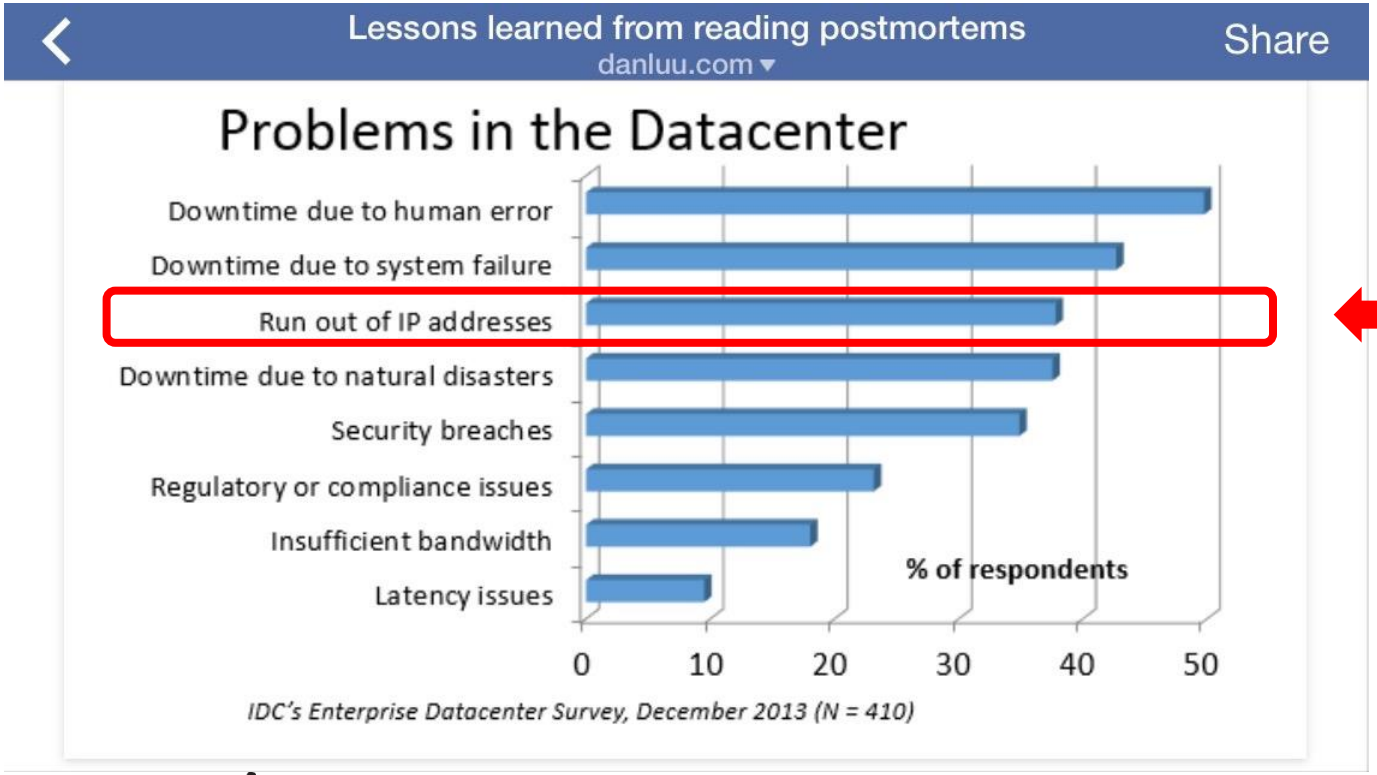

Multiple layers of NAT44 in the Cloud … More stuff that the automation system has to deal with

What needs IP in the (v)DC / Cloud?

- **Many components** that require IP:
  - Physical servers, Virtual Machines, Containers, API Endpoints
  - Provisioning, Orchestration, Management services
  - Virtual networking services and Physical networking
  - Highly-Available services, Anycast etc.

- **Example 1** – a virtual machine has:
  - Access to multiple networks depending on the application types the VM hosts
  - There can be numerous IPv6 prefixes required to support the application use case
  - Use **/64 prefix per network/VLAN**
- **Example 2** – the subnet for Docker containers should at least have a size of /80
  - Container can leverage its MAC address
  - Use standard **/64**

Cisco live!

(Virtualised) DC / Cloud Considerations

- **/48 per DC / Cloud** deployment is typical
- In a traditional DC (legacy mainframe, bare metal servers) things don't differ from IPv4
- **Cloud** services **add a layer of hierarchy** into the DC architecture
  - Moving deeper, Top of Rack (ToR) switch becomes the L3 boundary
  - As can a virtual Router on a physical host
- **Tenant addressing** – "BYO Prefix" or do you allocate?
  - If a tenant requires a **virtual router** within the tenant domain:
    - Use a **/64 or /127** on ToR-facing link - run IGP or static route towards ToR
    - The vRouter delegates /64 to the VMs
  - If a tenant uses "provider networks":
    - where the upstream switches/routers (i.e. DC Aggregation layer) provide the L3 services then use **/64 per network/VLAN** for tenant-facing networks
- **/64** is the least common denominator
  - For L2 south of ToR, use /64 per VLAN
    - Avoid large L2 domains – chatty IPv6 NDP!

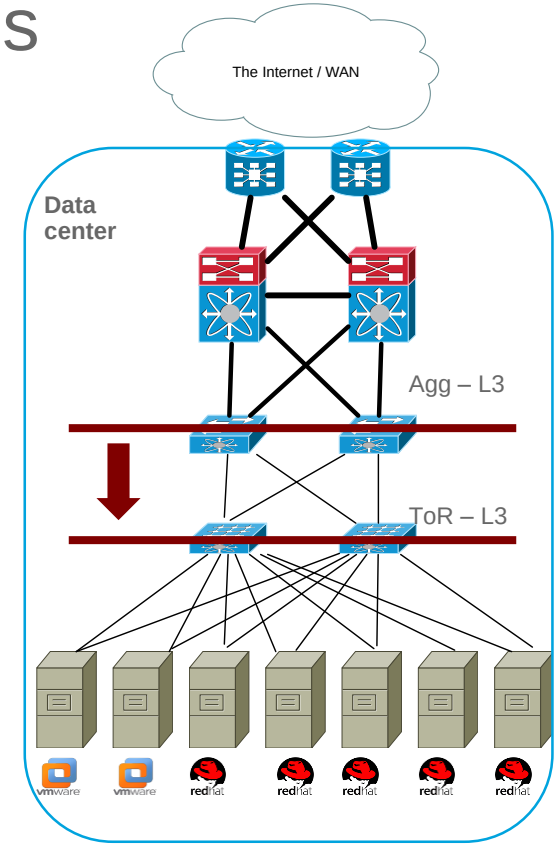

Agenda

- IPv6 Address Types – Refresher
- Typical Requirements
  - Enterprise
  - Service Provider
  - (Virtualized) DC / Cloud
- Methodology for writing the IPv6 Addressing Plan
- Exercise
- Final Thoughts & Conclusion

Methodology of IPv6 Address Planning

Think globally, act locally! ☺

Core message of the Methodology section

4 Rules

1. **Keep it SIMPLE**
   - You don't want to spend weeks explaining it!
2. **Embed information to help operations**
   - To help troubleshooting and operation of the network
   - Examples: location, country, PIN, VLAN, IPv4 addresses in Link Local and/or Global Addresses (consider this carefully!)
     - (not over the top though – remember Rule #1)
3. **Plan for expansion (build in reserve)**
   - Cater for future growth, mergers & acquisitions, new locations
   - Reserved vs. assigned
4. **Take advantage of hierarchy / aggregation**
   - Good aggregation is essential, just one address block (per location), we can take advantage of this (unlike in IPv4!)
   - Ensures scalability and stability

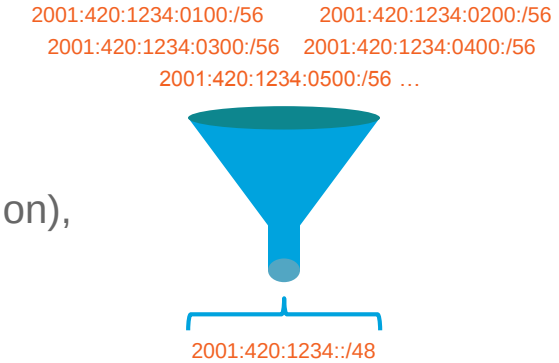

Methodology

- Structure
- Prefix sizing
- Information encoding
- Infrastructure addressing

Methodology (1) – Structure

- Analyze, where will IPv6 be? **EVERYWHERE**
  - Addressing plan needs to be **designed globally**
- Identify the structure of the addressing plan
  - Based on requirements and considerations discussed earlier
  - Top-down approach (different from the IPv4 days when #hosts/subnet was important) or middle-up
- Where and how many locations
  - Countries, regions, locations, buildings, etc…
  - Needs to map onto the physical / logical network topology
- Which **services, applications** and **systems** connected in each location
  - E.g. Fixed networks, mobile networks, end-users, ERP, CRM, R&D, etc…
- ULA recommendations
  - Don't deploy for end-point addressing
  - Unless in completely closed system – example: CPEs – management address
  - Could be considered for infrastructure addressing (e.g loopbacks, links)

Methodology (2) – Number of Prefixes per Level

- How many prefixes will you need at each level of the addressing plan
  - Example: a BNG can handle 64000 subscribers = 64000 IPv6 prefixes
  - Example: the number of interconnects (P2P) in your network
  - As always, put aside a reserve!
- How many /64 prefixes (subnets) you need to deploy at a location
  - Example: desktops, WIFI, guestnet, sensors, CCTV, network infrastructure, etc…
  - As always, put aside a reserve!
- **Don’t worry about the number of hosts**
  - **We have 2^{64} of IPv6 addresses for hosts!**

© K. Bednarova, 2015

Example - How Many Subnets in a Location?

For Your Reference

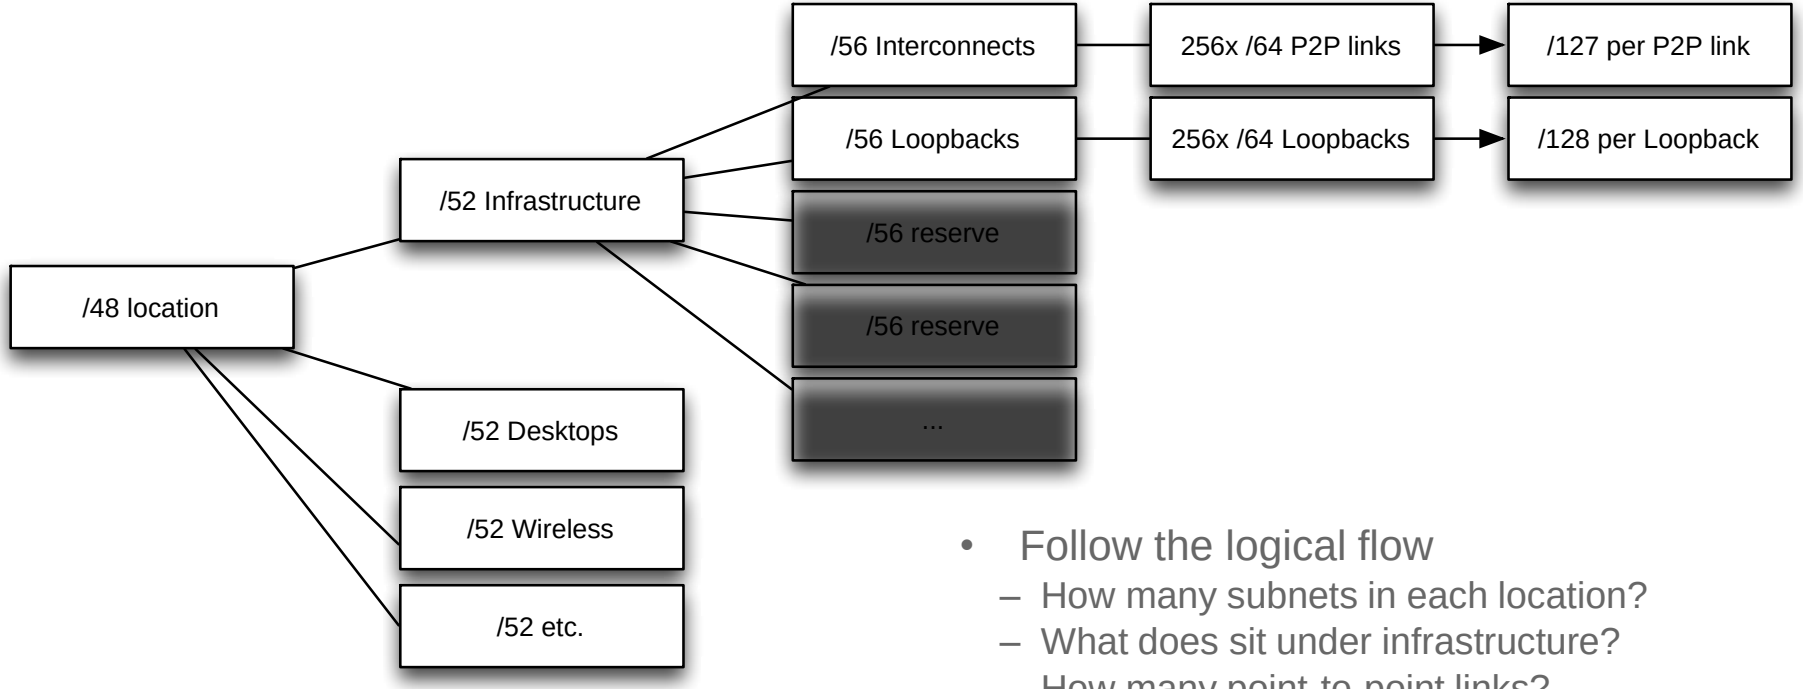

- Follow the logical flow
  - How many subnets in each location?
  - What does sit under infrastructure?
  - How many point-to-point links?
  - Where is the reserve?

Methodology (3) – Information Encoding

- Remember transition mechanisms – these will have specific address format requirements
  - ISATAP
  - NAT64 (/96)
  - 6rd, MAP
- Possible encoding of information in particular portions of the IPv6 prefix – examples:
  - Places In the Network (PINs)
  - VLANs in the prefix (or as part of the LLA)
    - VLAN **4096** → 2001:db8:1234:**4096**::/64 (alternatively in hex ☺)
  - The whole IPv4 address or just a portion
    - consider this carefully – trade-off between linkage vs. independence
    - IPv4 address **10.0.13.1** → 2001:db8:1234:100::**10:0:13:1**
  - Router IDs or IPv4 address in Link-Local
    - Router ID 1.1.1.1 → LLA: FE80::1.1.1.1
    - All interfaces on specific router can have identical LLAs
  - **Consider security implications!**

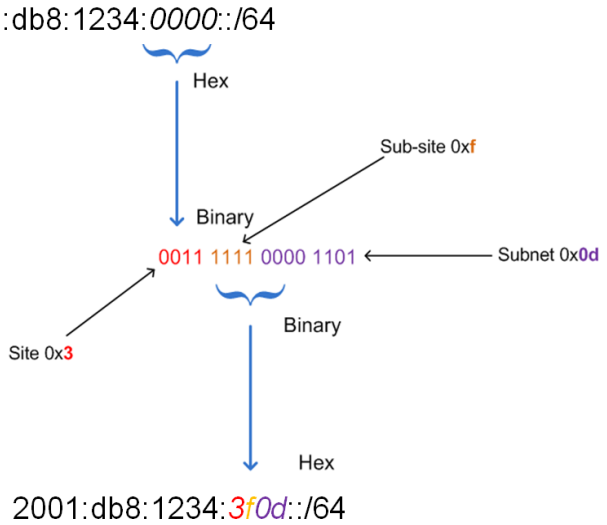

Choice of LLA can help NetOps/Troubleshooting

For Your Reference

```
R1#show ipv6 route
IPv6 Routing Table - default - 6 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user
Static route
       […output omitted…]
O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1, OE2 -
OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2, ls
- LISP site, ld - LISP dyn-EID, a - Application

LC   2001:DB8::1/128 [0/0]
     via Loopback0, receive
O    2001:DB8::2/128 [110/1]
     via FE80::10:0:12:1, GigabitEthernet0/1
O    2001:DB8::3/128 [110/1]
     via FE80::10:0:13:1, GigabitEthernet0/2
O    2001:DB8::4/128 [110/1]
     via FE80::10:0:14:1, GigabitEthernet0/3
OI   2001:DB8::5/128 [110/2]
     via FE80::10:0:12:1, GigabitEthernet0/1
L    FF00::/8 [0/0]
     via Null0, receive
R1#
```

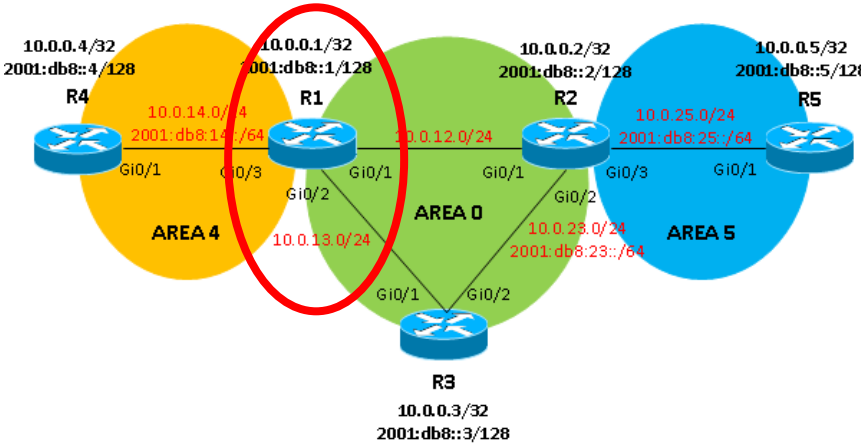

Methodology (4) – Infrastructure Addressing 1.

Point-to-Point Links

- First recommendations: configure /64, /112 or /126
  - RFC 3627, September 2003 – /127 considered harmful
  - moved to historic by RFC 6547 (Feb. 2012)
- Since April 2011, RFC 6164 recommends **/127** on inter-router (P2P) links
- Current recommendation /64, ~~/126~~ or /127
  - /127 mitigates ND exhaustion attacks
- Allocate /64 from a block (e.g. /56) for infrastructure links but configure /127
  - ~~Example: 2001:db8:1234:1::**0**/127 and 2001:db8:1234:1::**1**/127~~
  - What about **offsetting the suffix**?
  - 2001:db8:1234:1::**A**/127
  - 2001:db8:1234:1::**B**/127

- **You must follow the /127 subnet rule!!!**

Methodology (4) – Infrastructure Addressing 2.

Loopbacks

- E.g. Dedicate /56 for Loopback addresses per location
- Allocate /64 per Loopback but configure /128
  - Example: 2001:420:1234:**100**:**1**::1/128 and 2001:420:1234:**101**:**1**::1/128
  - Avoid a potential overlap with Embedded RP addresses
- Remember to check how many Longest Prefix Matches (LPM) [/128] your network devices can carry
  - Does not always equal the total number of supported IPv6 prefixes

Example of an IPv6 Prefix Allocation (Cisco)

For Your Reference

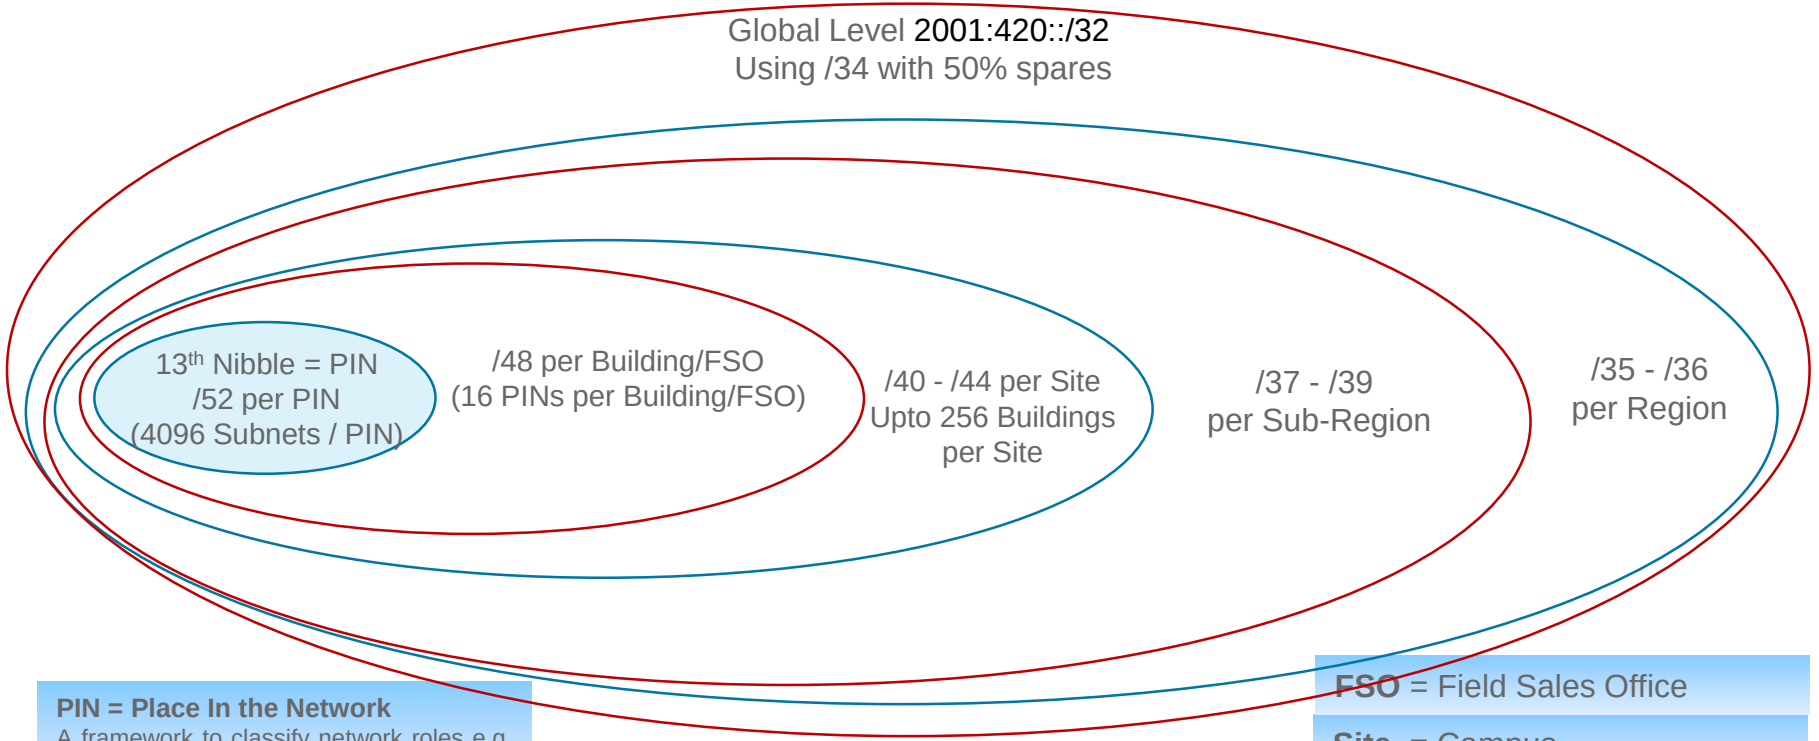

PIN = Place In the Network
A framework to classify network roles e.g. Lab, Desktop, DC, DMZ etc

FSO = Field Sales Office

Site = Campus

Example of an IPv6 Prefix Allocation (ISP)

For Your Reference

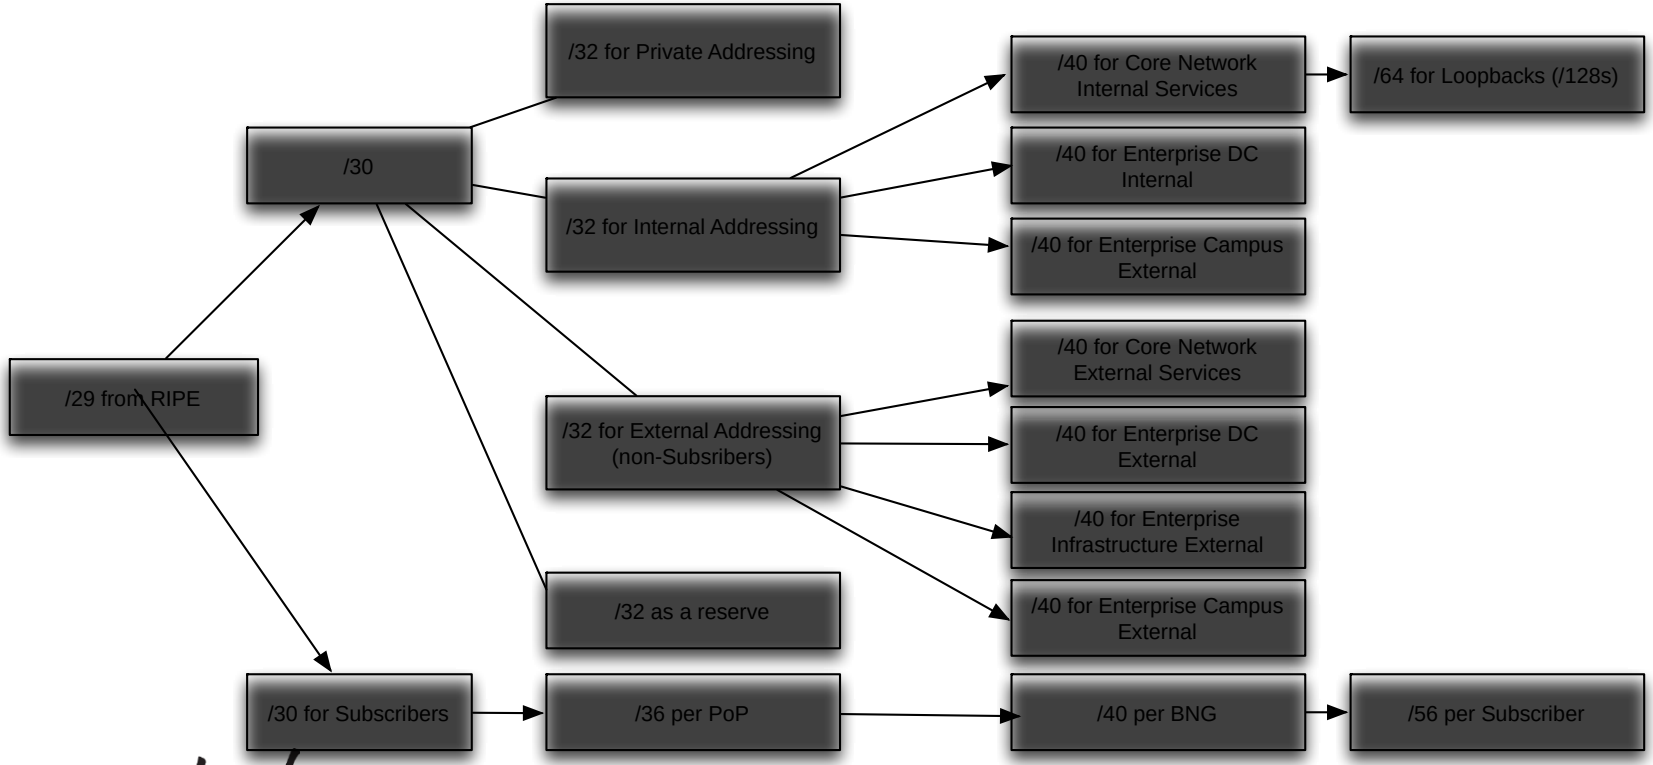

Exercise

Exercise – the Rules!

- Choose 1 from 2 options
  - Enterprise
  - ISP
- Work alone or with your colleagues
- 25 minutes
- Q&A during the exercise
  - Both speakers will help you
- The solution will be presented step by step

- **Remember: Host** portion is **ALWAYS 64 bits**!

Note: Details of the Exercise will be included on the day of the breakout session.

Agenda

- IPv6 Address Types – Refresher
- Typical Requirements
  - Enterprise
  - Service Provider
  - (Virtualized) DC / Cloud
- Methodology for writing the IPv6 Addressing Plan
- Exercise
- Final Thoughts & Conclusion

Tools for Managing IPv6 Addressing Plan

- Not just a spreadsheet, please!
- Very prone to error ☹

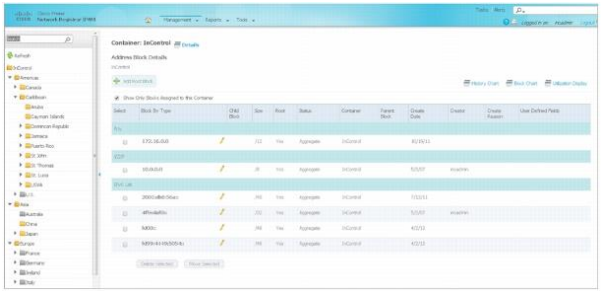

- There are many **IP Address Management** tools on the market
  - See our exhibitors in **the World of Solutions**
- **Cisco Prime Network Registrar**
  - http://www.cisco.com/en/US/products/ps11808/index.html
  - Visit CiscoLive! NOC
- Work with an **IPv6 prefix calculator**
  - Example: http://www.gestioip.net/cgi-bin/subnet_calculator.cgi

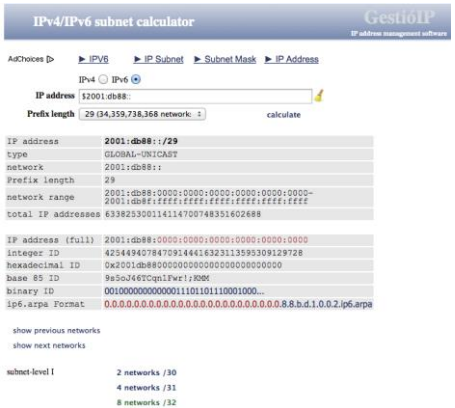

Recommended Reading

- **Create an IPv6 Addressing Plan**
  - RIPE NCC
  - https://www.ripe.net/publications/ipv6-info-centre/deployment-planning/create-an-addressing-plan
  - RFC 6177 **IPv6 Address Assignment to End Sites**
- Tom Coffeen: **IPv6 Address Planning** (O'Reilly, 2015)
  - http://shop.oreilly.com/product/0636920033622.do
- Register for **RIPE NCC IPv6 Addressing Plan webinar** – 8th March 2016
  - https://www.ripe.net/support/training/learn-online/webinars/ipv6-addressing-plan
- **ULA voluntary registry**
  - https://www.sixxs.net/tools/grh/ula/list/

Conclusion

Key Takeaways

1. IPv6 addressing is easy and highly flexible.

1. IPv6 addressing strategy is the key to successful deployment.

1. A good addressing plan simplifies IPv6 network operation and troubleshooting.

"If you fail to PLAN (an address plan), you plan to FAIL"

Core message of this session

Call to Action

- Attend the following related sessions
  - **Service Provider IPv6 Deployment** [BRKSPG-2300] – Wednesday 17th Feb, 4:30 PM
  - **Don't Be Left Behind: Consumer Internet Traffic is Shifting to IPv6, Will your Organization Follow?** [PNLCRS-2307] – Wednesday 17th Feb, 4:30 PM
  - **Segment Routing for IPv6 Networks** [BRKRST-3123] – Thursday, 18th Feb, 2:30 PM
  - **Enterprise Multi-Homed Internet Edge Architectures** [BRKRST-2044] – Thursday 19th Feb, 2:30 PM
  - **Enterprise IPv6 Deployment** [BRKRST-2301] – Friday 19th Feb, 9:00 AM
- Visit the **World of Solutions** for
  - Cisco Campus
  - Walk in Labs
  - Technical Solution Clinics
- Meet the Engineer – book a meeting with us!
- Lunch and Learn Topics
  - **IPv6 in Enterprise** – Thursday 18th Feb, 1 – 2:15 PM, Hall 4.1

Learning more about IPv6

Walk-in Self-Paced Lab:

LABCRS-1000 Intro IPv6 Addressing and Routing Lab

Experiment with IPv6-only WiFi:

SSID: **CL-NAT64**

WPA passphrase: **cl-nat64**

SLAAC + stateless DHCP

NAT64 included to access legacy

Ask **all** **World of Solutions** **exhibitors** for their IPv6 support ☺

Complete Your **Online Session Evaluation**

- Please complete your online session evaluations after each session. Complete 4 session evaluations & the Overall Conference Evaluation (available from Thursday) to receive your Cisco Live T-shirt.

- All surveys can be completed via the Cisco Live Mobile App or the Communication Stations

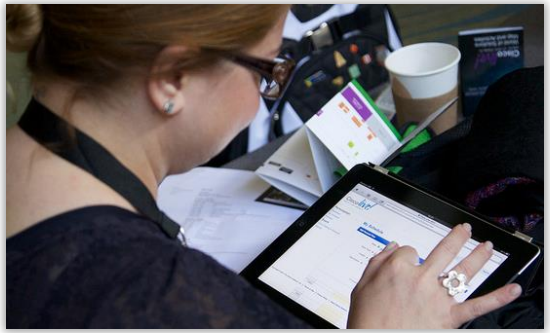

Ciscolive!

Participate in the “My Favorite Speaker” Contest

Promote Your Favorite Speaker and You Could Be a Winner

- Promote your favorite speaker through Twitter and you could win $200 of Cisco Press products (@CiscoPress)
- Send a tweet and include
  - Your favorite speaker’s Twitter handle: @vqmckillop and @wverrydt
  - Two hashtags: #CLUS #MyFavoriteSpeaker
- You can submit an entry for more than one of your “favorite” speakers
- Don’t forget to follow @CiscoLive and @CiscoPress
- View the official rules at http://bit.ly/CLUSwin

“If you fail to PLAN (an address plan), you plan to FAIL”

Core message of this session

Thank you

Appendix

Enterprise: How to Get an IPv6 Prefix

Aka PI vs PA

- Enterprise **becomes a RIPE member** and acquires Provider-aggregatable (PA) address space
- Enterprise acquires Provider-independent (PI) address space through a "**sponsoring LIR**"
  - /48 per location, **not aggregatable**
- Enterprise acquires Provider-independent (PI) address space through a "**sponsoring LIR**"
  - /48 per location + **super-block, aggregatable**
- Enterprise acquires Provider-aggregatable (PA) address space **from its ISP/LIR** allocation

Enterprise becomes a RIPE member and acquires Provider-aggregatable (PA) address space

- +++
  - /29 PA address space assigned
  - Globally routable
  - Solution for all RIR regions, can be de-aggregated in /32s per region (Europe, US/Canada, AsiaPac, Latam, Africa).
  - Within region, /32 can be further de-aggregated in /48s per site. /32 in central site/DC can be used for backup.
- ---
  - RIPE fees (2000 EUR one-off, 1500 EUR yearly). See: https://www.ripe.net/ripe/docs/ripe-620. However, fees are coming down year-after-year as RIPE is non-profit.
  - Potentially, PA /48s could be filtered by ISPs in the future. Currently, this is not an issue. If this becomes an issue, backup through the /32 aggregate is a solution.

Enterprise acquires PI address space through "sponsoring LIR" – /48 per location, **not** aggregatable

- For example: 40 locations in Europe → 40x /48s
- +++
  - /48 blocks are globally routable and ISP independent
- ---
  - RIPE will charge 50 EUR per /48. See: https://www.ripe.net/ripe/docs/ripe-620. ISP (sponsoring LIR) will mark-up to cover admin cost → few 100s EUR per /48 (?).
  - Individual /48s likely not aggregatable → no backup solution using an aggregate
  - This solution only applies to a specific RIR region. This needs to be replicated for RIPE, ARIN, APNIC, LACNIC, AFRINIC.
- Unknown
  - Fee that ISP will charge for the PI blocks?

Enterprise acquires PI address space through "sponsoring LIR" - /48 per location + super-block

- For example: 40 locations in Europe → 40x /48s + /42 super-block/aggregate
- +++
  - /48 blocks are globally routable and ISP independent
  - Super-block/aggregate available → backup solution possible
- ---
  - RIPE will charge 50 EUR per super-block. See: https://www.ripe.net/ripe/docs/ripe-620. ISP (sponsoring LIR) will mark-up to cover admin cost → few 100s EUR per /48 (?)
  - This solution only applies to a specific RIR region. This needs to be replicated for RIPE, ARIN, APNIC, LACNIC, AFRINIC.
- Unknown
  - Fee that ISP will charge for the PI blocks?

Enterprise acquires PA address space from its ISP/LIR allocation

- +++
  - Globally routable
  - Solution for all RIR regions
- ---
  - Dependence on ISP
  - Need to re-number network in case of ISP change
- Unknown
  - Fee that ISP will charge for its PA block?

References

- **RIPE NCC Charging Scheme 2015**
  - https://www.ripe.net/ripe/docs/ripe-620
- **RIPE IPv6 Address Allocation and Assignment Policy**
  - https://www.ripe.net/publications/docs/ripe-641

We’re ready. Are you?