

Agenda : Day 1

| Session | Agenda |
|-------------|--|
| 1000 - 1130 | What is an IX, How to build and IXP |
| 1200 - 1330 | Internet Routing Basic |
| 1430 - 1600 | IPv6 Protocol Architecture |
| 1630 - 1730 | IGP Routing Protocol & Hands-on lab exercise |

Agenda : Day 2

| Session | Agenda |
|-------------|--|
| 1000 - 1130 | BGP Routing Protocol Operation |
| 1200 - 1330 | Attributes and Path Selection Process |
| 1430 - 1600 | Hands On Lab Exercise: iBGP Peering- What is Exchange Point? |
| 1630 - 1730 | Hands On Lab Exercise: eBGP Peering- What is Exchange Point? |

Agenda : Day 3

| Session | Agenda |
|-------------|--|
| 0930 - 1030 | IXP Design Considerations |
| 1100 - 1230 | Hands On Lab Exercise: IXP Configuration |
| 1330 - 1500 | Route Collectors & Servers |
| 1530 - 1700 | IXP BCP and What can go wrong? |

Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

Acknowledgment

- Cisco System
- Philip Smith

Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

Overview


IXP Workshop

- **What is an Internet Exchange Point (IXP)?**
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

What is an Internet Exchange Point (IXP)?


- The Internet is an interconnection of networks
 - Each controlled by separate entities
 - Generally called Internet Service Providers (ISPs)
 - Grouped by Autonomous Systems (AS) number
- Transit
 - Where ISP will pay to send/receive traffic
 - Downstream ISP will pay upstream ISP for transit service
- Peering
 - ISPs will not pay each other to interchange traffic
 - Works well if win win for both
 - Reduce cost on expensive transit link

Peering and Transit example



A and B peer for free, but need transit arrangements with D to get packets to/from C


Private Interconnect



Public Interconnect

- A location or facility where several ISPs are present and connect to each other over a common shared media
- Why?
 - To save money, reduce latency, improve performance
- IXP – Internet eXchange Point
- NAP – Network Access Point


High Level View of the Global Internet



Detailed View of the Global Internet

- Global Transit Providers
 - Connect to each other
 - Provide connectivity to Regional Transit Providers
- Regional Transit Providers
 - Connect to each other
 - Provide connectivity to Content Providers
 - Provide connectivity to Access Providers
- Access Providers
 - Connect to each other across IXPs (free peering)
 - Provide access to the end user

Categorizing ISPs



Inter-provider relationships

- Peering between equivalent sizes of service providers (e.g. Tier 2 to Tier 2)
 - Shared cost private interconnection, equal traffic flows
 - No cost peering
- Peering across exchange points
 - If convenient, of mutual benefit, technically feasible
- Fee based peering
 - Unequal traffic flows, “market position”

Internet Exchange Point- Why peer?

- Consider a region with one ISP
 - They provide internet connectivity to their customers
 - They have one or two international connections
- Internet grows, another ISP sets up in competition
 - They provide internet connectivity to their customers
 - They have one or two international connections
- How does traffic from customer of one ISP get to customer of the other ISP?
 - Via the international connections

Internet Exchange Point- Why peer?

- Yes, International Connections...
 - If satellite, RTT is around 550ms per hop
 - So local traffic takes over 1s round trip
- International bandwidth
 - Costs significantly more than domestic bandwidth
 - Congested with local traffic
 - Wastes money, harms performance

Internet Exchange Point- Why peer?

- Solution:
 - Two competing ISPs peer with each other
- Result:
 - Both save money
 - Local traffic stays local
 - Better network performance, better QoS,...
 - More international bandwidth for expensive international traffic
 - Everyone is happy

Internet Exchange Point- Why peer?

- A third ISP enters the equation
 - Becomes a significant player in the region
 - Local and international traffic goes over their international connections
- They agree to peer with the two other ISPs
 - To save money
 - To keep local traffic local
 - To improve network performance, QoS,...

Internet Exchange Point- Why peer?

- Private peering means that the three ISPs have to buy circuits between each other
 - Works for three ISPs, but adding a fourth or a fifth means this does not scale
- Solution:
 - Internet Exchange Point

Internet Exchange Point

- Every participant has to buy just one whole circuit
 - From their premises to the IXP
- Rather than $N-1$ half circuits to connect to the $N-1$ other ISPs
 - 5 ISPs have to buy 4 half circuits = 2 whole circuits → already twice the cost of the IXP connection


Internet Exchange Point

- Solution
 - Every ISP participates in the IXP
 - Cost is minimal – one local circuit covers all domestic traffic
 - International circuits are used for just international traffic – and backing up domestic links in case the IXP fails
- Result:
 - Local traffic stays local
 - QoS considerations for local traffic is not an issue
 - RTTs are typically sub 10ms
 - Customers enjoy the Internet experience
 - Local Internet economy grows rapidly


Internet Exchange Points

- Layer 2 exchange point
 - Ethernet (100Gbps/10Gbps/1Gbps/100Mbps)
 - Older technologies include ATM, Frame Relay, SRP, FDDI and SMDS
- Layer 3 exchange point
 - Router based
 - Has historical status now

Layer 2 Exchange



Layer 2 Exchange



Layer 2 Exchange

- Two switches for redundancy
- ISPs use dual routers for redundancy or loadsharing
- Offer services for the “common good”
 - Internet portals and search engines
 - DNS Root & TLDs, NTP servers
 - Routing Registry and Looking Glass

Layer 2 Exchange

- Requires neutral IXP management
 - Usually funded equally by IXP participants
 - 24x7 cover, support, value add services
- Secure and neutral location
- Configuration
 - Private address space if non-transit and no value add services
 - Otherwise public IPv4 (/24) and IPv6 (/64)
 - ISPs require AS, basic IXP does not

Layer 2 Exchange

- Network Security Considerations
 - LAN switch needs to be securely configured
 - Management routers require TACACS+ authentication, vty security
 - IXP services must be behind router(s) with strong filters

“Layer 3 IXP”

- IX will provide layer two connection/switch port to ISPs
- Each ISP will peer with a route server on the IX
- Route server will collect and distribute directly connected routes to every peers

Layer 2 versus Layer 3

- Layer 3
 - IXP team requires good BGP knowledge
 - Rely on 3rd party for BGP configuration
 - One peering will get all IXP routes
 - Less freedom on who peers with whom
 - Usually competes with IXP membership
 - Tends to be distributed over wide area
 - IXP can grow faster

Layer 2 versus Layer 3

- Layer 2
 - IXP team does not need routing knowledge
 - Easy to get started
 - More complicated to distribute over wide area
 - ISPs free to set up peering agreements with each other as they wish



Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- **What is the value of Peering?**
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

ISP Goals

- Minimise the cost of operating the business
- Transit
 - ISP has to pay for circuit (international or domestic)
 - ISP has to pay for data (usually per Mbps)
 - Repeat for each transit provider
 - Significant cost of being a service provider
- Peering
 - ISP shares circuit cost with peer (private) or runs circuit to public peering point (one off cost)
 - No need to pay for data
 - Reduces transit data volume, therefore reducing cost

Transit – How it works

- Small access provider provides Internet access for a city's population
 - Mixture of dial up, wireless and fixed broadband
 - Possibly some business customers
 - Possibly also some Internet cafes
- How do their customers get access to the rest of the Internet?
- ISP buys access from one, two or more larger ISPs who already have visibility of the rest of the Internet
 - This is transit – they pay for the physical connection to the upstream and for the traffic volume on the link

Peering – How it works

- If two ISPs are of equivalent sizes, they have:
 - Equivalent network infrastructure coverage
 - Equivalent customer size
 - Similar content volumes to be shared with the Internet
 - Potentially similar traffic flows to each other's networks
- This makes them good peering partners
- If they don't peer
 - They both have to pay an upstream provider for access to each other's network/customers/content
 - Upstream benefits from this arrangement, the two ISPs both have to fund the transit costs

Example: South Asian ISP @ LINX

- Date: October 2011
- Facts:
 - Route Server plus bilateral peering offers 81k prefixes
 - IXP traffic averages 55Mbps/15Mbps
 - Transit traffic averages 35Mbps/3Mbps
- Analysis:
 - 61% of inbound traffic comes from 81k prefixes available by peering
 - 39% of inbound traffic comes from remaining 287k prefixes from transit provider

Example: South Asian ISP @ HKIX

- Date: October 2011
- Facts:
 - Route Server plus bilateral peering offers 34k prefixes
 - IXP traffic is 130Mbps/30Mbps
 - Transit traffic is 125Mbps/40Mbps
- Analysis:
 - 51% of inbound traffic comes from 42k prefixes available by peering
 - 49% of inbound traffic comes from remaining 326k prefixes from transit provider

Example: South Asian ISP

- Summary:
 - Traffic by Peering: 185Mbps/45Mbps
 - Traffic by Transit: 160Mbps/43Mbps
 - 54% of incoming traffic is by peering
 - 52% of outbound traffic is by peering

Example: South Asian ISP

- Router at remote co-lo
 - Benefits: can select peers, easy to swap transit providers
 - Costs: co-lo space and remote hands
- Servers at remote co-lo
 - Benefits: mail filtering, content caching, etc
 - Costs: co-lo space and remote hands
- Overall advantage:
 - Can control what goes on the expensive connectivity “back to home”

Value propositions

- Peering at a local IXP
 - Reduces latency & transit costs for local traffic
 - Improves Internet quality perception
- Participating at a Regional IXP
 - A means of offsetting transit costs
- Managing connection back to home network
- Improving Internet Quality perception for customers

Summary

- Benefits of peering
 - Private
 - Internet Exchange Points
- Local versus Regional IXPs
 - Local services local traffic
 - Regional helps defray transit costs

Worked Example

Single International Transit

Versus

Local IXP + Regional IXP + Transit

Worked Example

- ISP A is local access provider
 - Some business customers (around 200 fixed links)
 - Some co-located content provision (datacentre with 100 servers)
 - Some consumers on broadband (5000 DSL/Cable/Wireless)
 - Some consumers on dial (1000 on V.34 type speeds)
- They have a single transit provider
 - Connect with a 16Mbps international leased link to their transit's PoP
 - Transit link is highly congested

Worked Example (2)

- There are two other ISPs serving the same locality
 - There is no interconnection between any of the three ISPs
 - Local traffic (between all 3 ISPs) is traversing International connections
- Course of action for our ISP:
 - Work to establish local IXP
 - Establish presence at overseas co-location
- First Step
 - Assess local versus international traffic ratio
 - Use NetFlow on border router connecting to transit provider

Worked Example (3)

- Local/Non-local traffic ratio
 - Local = traffic going to other two ISPs
 - Non-local = traffic going elsewhere
- Example: balance is 30:70
 - Of 16Mbps, that means 5Mbps could stay in country and not congest International circuit
 - 16Mbps transit costs \$50 per Mbps per month traffic charges = \$250 per month, or \$3000 per year for local traffic
 - Circuit costs \$100k per year: \$30k is spent on local traffic
- Total is \$33k per year for local traffic

Worked Example (4)

- IXP cost:
 - Simple 8 port 10/100 managed switch plus co-lo space over 3 years could be around US\$30k total; or \$3k per year per ISP
 - One router to handle 5Mbps (e.g. 2801) would be around \$3k (good for 3 years)
 - One local 10Mbps circuit from ISP location to IXP location would be around \$5k per year, no traffic charges
 - Per ISP total: \$9k
 - Somewhat cheaper than \$33k
 - Business case for local peering is straightforward - \$24k saving per annum

Worked Example (5)

- After IXP establishment
 - 5Mbps removed from International link
 - Leaving 5Mbps for more International traffic – and that fills the link within weeks of the local traffic being removed
- Next step is to assess transit charges and optimise costs
 - ISPs visits several major regional IXPs
 - Assess routes available
 - Compares routes available with traffic generated by those routes from its Netflow data
 - Discovers that 30% of traffic would transfer to one IXP via peering

Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- **How to build an IXP?**
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

How to Build an IXP?

- The IXP Core is an Ethernet switch
- Has superseded all other types of network devices for an IXP
 - From the cheapest and smallest 12 or 24 port 10/100 switch
 - To the largest 192 port 10GigEthernet switch

How to Build an IXP?

- Each ISP participating in the IXP brings a router to the IXP location
- Router needs:
 - One Ethernet port to connect to IXP switch
 - One WAN port to connect to the WAN media leading back to the ISP backbone
 - To be able to run BGP

How to Build an IXP?

- IXP switch located in one equipment rack dedicated to IXP
 - Also includes other IXP operational equipment
- Routers from participant ISPs located in neighbouring/adjacent rack(s)
- Copper (UTP) connections made for 10Mbps, 100Mbps or 1Gbps connections
- Fibre used for 10Gbps and 40Gbps

Peering

- Each participant needs to run BGP
 - They need their own AS number
 - **Public** ASN, **NOT** private ASN
- Each participant configures external BGP directly with the other participants in the IXP
 - Peering with all participants
 - or
 - Peering with a subset of participants

Routing

- ISP border routers at the IXP generally should NOT be configured with a default route or carry the full Internet routing table
 - Carrying default or full table means that this router and the ISP network is open to abuse by non-peering IXP members
 - Correct configuration is only to carry routes offered to IXP peers on the IXP peering router
- Note: Some ISPs offer transit across IX fabrics
 - They do so at their own risk – see above

Routing (more)

- ISP border routers at the IXP should not be configured to carry the IXP LAN network within the IGP or iBGP
 - Use next-hop-self BGP concept
- Don't generate ISP prefix aggregates on IXP peering router
 - If connection from backbone to IXP router goes down, normal BGP failover will then be successful

Address Space

- Some IXPs use private addresses for the IX LAN
 - Public address space means IXP network could be leaked to Internet which may be undesirable
 - Because most ISPs filter RFC1918 address space, this avoids the problem
- Some IXPs use public addresses for the IX LAN
 - Address space available from the RIRs
 - IXP terms of participation often forbid the IX LAN to be carried in the ISP member backbone

APNIC Policy on IXP Address Space

- The End-User Assignments policy caters for IXPs Public Address space under IXP Address Assignment
- It requires the IXP with minimum 3 ISPs connected and have clear and open policy for joining
- The minimum IXP Assignment is /24 of IPv4 and /48 for IPv6

Hardware

- Try not to mix port speeds
 - if 10Mbps and 100Mbps connections available, terminate on different switches (L2 IXP)
- Don't mix transports
 - if terminating ATM PVCs and G/F/Ethernet, terminate on different devices
- Insist that IXP participants bring their own router
 - moves buffering problem off the IXP
 - security is responsibility of the ISP, not the IXP

Services Offered

- Services offered should not compete with member ISPs (basic IXP)
 - e.g. web hosting at an IXP is a bad idea unless all members agree to it
- IXP operations should make performance and throughput statistics available to members
 - Use tools such as MRTG to produce IX throughput graphs for member (or public) information

Services to Offer

- ccTLD DNS
 - the country IXP could host the country's top level DNS
 - e.g. "SE." TLD is hosted at Netnod IXes in Sweden
 - Offer back up of other country ccTLD DNS
- Root server
 - Anycast instances of I.root-servers.net, F.root-servers.net etc are present at many IXes
- Usenet News
 - Usenet News is high volume
 - could save bandwidth to all IXP members

Services to Offer

- Route Collector
 - Route collector shows the reachability information available at the exchange
 - Technical detail covered later on
- Looking Glass
 - One way of making the Route Collector routes available for global view (e.g. www.traceroute.org)
 - Public or members only access

Services to Offer

- Content Redistribution/Caching
 - For example, Akamised update distribution service
- Network Time Protocol
 - Locate a stratum 1 time source (GPS receiver, atomic clock, etc) at IXP
- Routing Registry
 - Used to register the routing policy of the IXP membership (more later)



Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- **How Internet works & Routing Protocol Basic**
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

1: How Does Routing Work?

- Internet is made up of the ISPs who connect to each other's networks
- How does an ISP in Kenya tell an ISP in Japan what customers they have?
- And how does that ISP send data packets to the customers of the ISP in Japan, and get responses back
 - After all, as on a local ethernet, two way packet flow is needed for communication between two devices

2: How Does Routing Work?

- ISP in Kenya could buy a direct connection to the ISP in Japan
 - But this doesn't scale – thousands of ISPs, would need thousands of connections, and cost would be astronomical
- Instead, ISP in Kenya tells his neighbouring ISPs what customers he has
 - And the neighbouring ISPs pass this information on to their neighbours, and so on
 - This process repeats until the information reaches the ISP in Japan

3: How Does Routing Work?

- This process is called “Routing”
- The mechanisms used are called “Routing Protocols”
- Routing and Routing Protocols ensures that the Internet can scale, that thousands of ISPs can provide connectivity to each other, giving us the Internet we see today

4: How Does Routing Work?

- ISP in Kenya doesn't actually tell his neighbouring ISPs the names of the customers
 - (network equipment does not understand names)
- Instead, he has received an IP address block as a member of the Regional Internet Registry serving Kenya
 - His customers have received address space from this address block as part of their “Internet service”
 - And he announces this address block to his neighbouring ISPs – this is called announcing a “route”

Routing Protocols

- Routers use “routing protocols” to exchange routing information with each other
 - **IGP** is used to refer to the process running on routers inside an ISP's network
 - **EGP** is used to refer to the process running between routers bordering directly connected ISP networks

What Is an IGP?

- Interior Gateway Protocol
- Within an Autonomous System
- Carries information about internal infrastructure prefixes
- Two widely used IGPs in service provider network:
 - OSPF
 - ISIS

Why Do We Need an IGP?

- ISP backbone scaling
 - Hierarchy
 - Limiting scope of failure
 - Only used for ISP's **infrastructure** addresses, not customers or anything else
 - Design goal is to **minimise** number of prefixes in IGP to aid scalability and rapid convergence

What Is an EGP?

- Exterior Gateway Protocol
- Used to convey routing information between Autonomous Systems
- De-coupled from the IGP
- Current EGP is BGP

Why Do We Need an EGP?

- Scaling to large network
 - Hierarchy
 - Limit scope of failure
- Define Administrative Boundary
- Policy
 - Control reachability of prefixes
 - Merge separate organisations
 - Connect multiple IGPs

Interior versus Exterior Routing Protocols


- **Interior**

- Automatic neighbour discovery
- Generally trust your IGP routers
- Prefixes go to all IGP routers
- Binds routers in one AS together
- Carries ISP infrastructure addresses only
- ISPs aim to keep the IGP small for efficiency and scalability

- **Exterior**

- Specifically configured peers
- Connecting with outside networks
- Set administrative boundaries
- Binds AS's together
- Carries customer prefixes
- Carries Internet prefixes
- EGPs are independent of ISP network topology


Hierarchy of Routing Protocols



FYI: Cisco IOS Default Administrative Distances


| Route Source | Default Distance |
|-----------------------------|------------------|
| Connected Interface | 0 |
| Static Route | 1 |
| Enhanced IGRP Summary Route | 5 |
| External BGP | 20 |
| Internal Enhanced IGRP | 90 |
| IGRP | 100 |
| OSPF | 110 |
| IS-IS | 115 |
| RIP | 120 |
| EGP | 140 |
| External Enhanced IGRP | 170 |
| Internal BGP | 200 |
| Unknown | 255 |

What does a Router do?




A day in a life of a Router

- Find path
- Forward packet, forward packet, forward packet, forward packet...
- Find alternate path
- Forward packet, forward packet, forward packet, forward packet...
- Repeat until powered off



Routing versus Forwarding

- Routing = building maps and giving directions
- Forwarding = moving packets between interfaces according to the “directions”

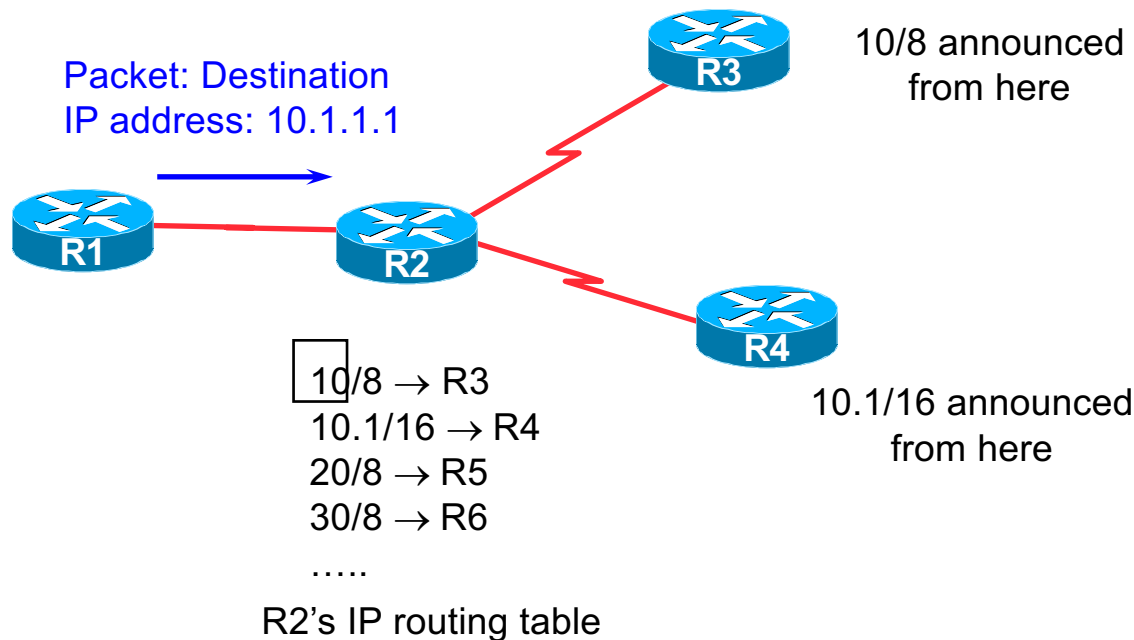


IP Route Lookup

- Based on destination IP address
- “longest match” routing
 - More specific prefix preferred over less specific prefix
 - **Example:** packet with destination of 10.1.1.1/32 is sent to the router announcing 10.1/16 rather than the router announcing 10/8.

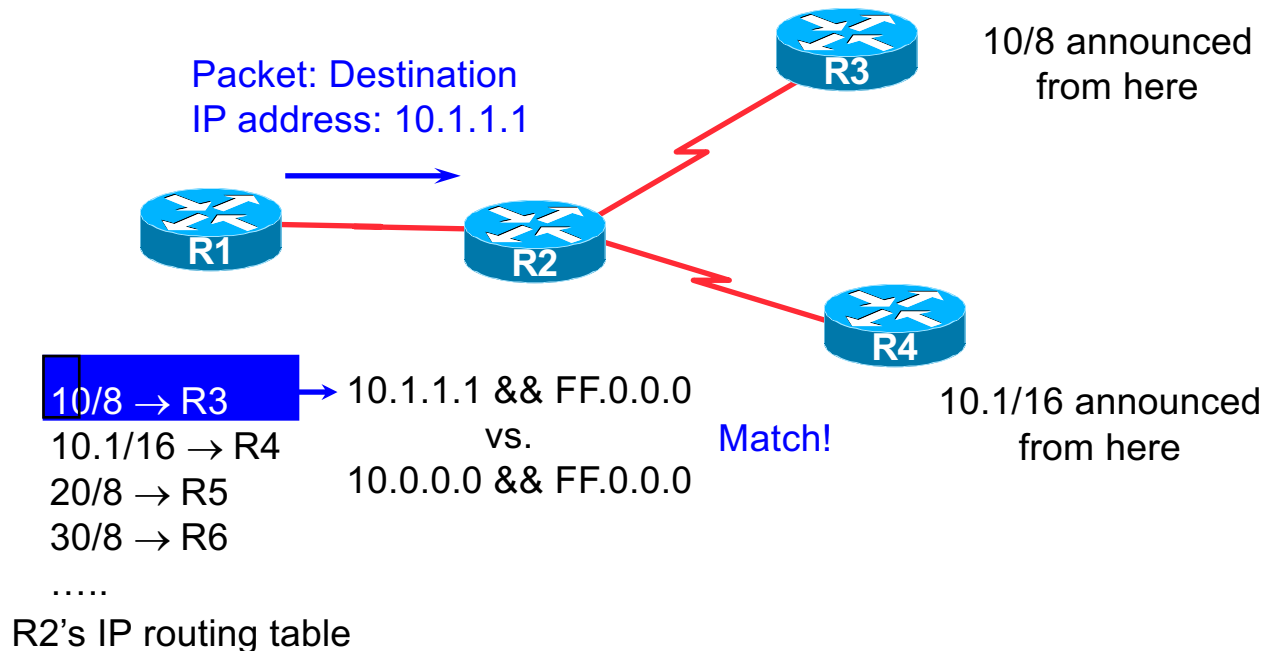
IP route lookup

- Based on destination IP address




IP route lookup: Longest match routing

- Based on destination IP address



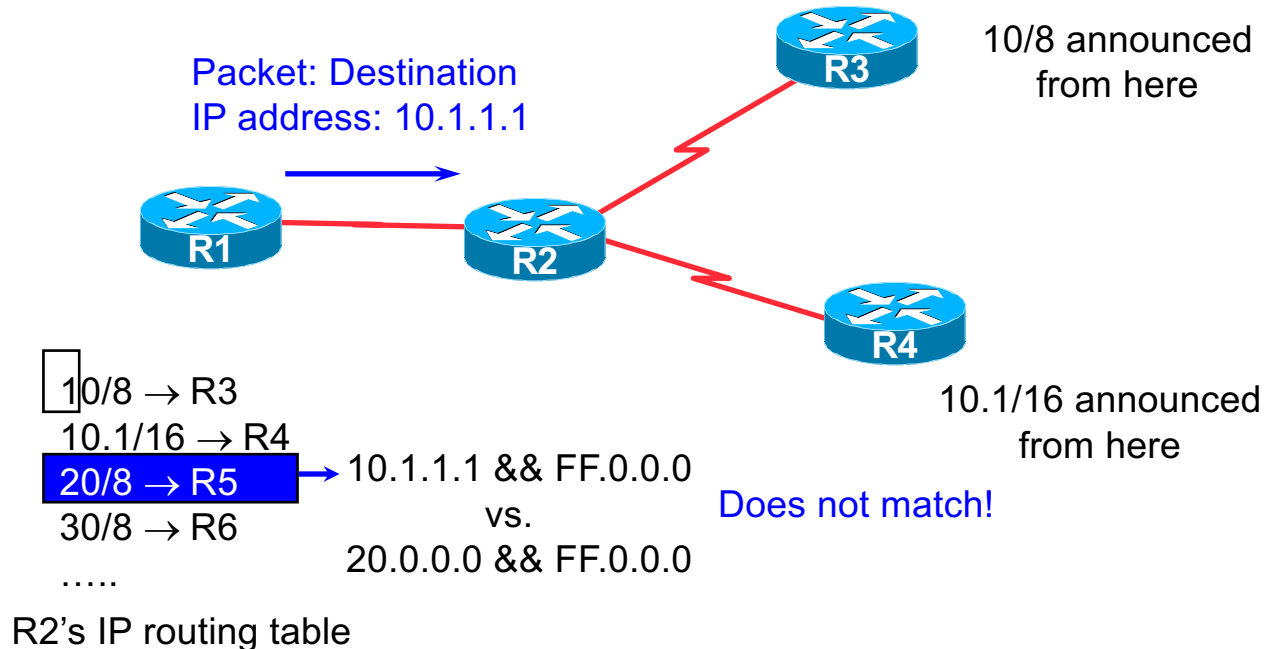
IP route lookup: Longest match routing

- Based on destination IP address



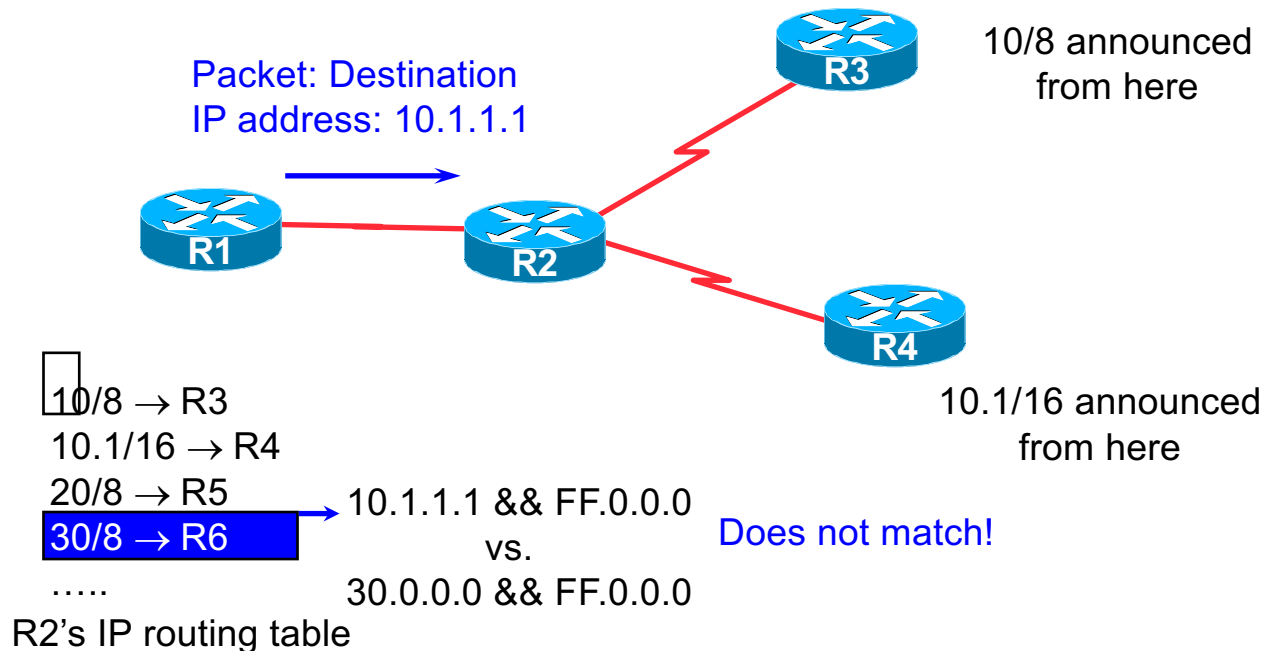
IP route lookup: Longest match routing

- Based on destination IP address




IP route lookup: Longest match routing

- Based on destination IP address



IP route lookup: Longest match routing

- Based on destination IP address



RIBs and FIBs

- RIB is the Routing Table
 - It contains a list of all the destinations and the various next hops used to get to those destinations – and lots of other information too!
 - One destination can have lots of possible next-hops – only the best next-hop goes into the FIB
- FIB is the Forwarding Table
 - It contains destinations and the interfaces to get to those destinations
 - Used by the router to figure out where to send the packet
 - Careful! Some people still call this a route!

Explicit versus Default Routing

- Default:
 - simple, cheap (cycles, memory, bandwidth)
 - low granularity (metric games)
- Explicit (default free zone)
 - high overhead, complex, high cost, high granularity
- Hybrid
 - minimise overhead
 - provide useful granularity
 - requires some filtering knowledge

Egress Traffic


- How packets leave your network
- Egress traffic depends on:
 - route availability (what others send you)
 - route acceptance (what you accept from others)
 - policy and tuning (what you do with routes from others)
 - Peering and transit agreements

Ingress Traffic

- How packets get to your network and your customers' networks
- Ingress traffic depends on:
 - what information you send and to whom
 - based on your addressing and AS's
 - based on others' policy (what they accept from you and what they do with it)

Autonomous System (AS)


- Collection of networks with same routing policy
- Single routing protocol
- Usually under single ownership, trust and administrative control



Definition of terms

- **Neighbours**
 - AS's which directly exchange routing information
 - Routers which exchange routing information
- **Announce**
 - send routing information to a neighbour
- **Accept**
 - receive and use routing information sent by a neighbour
- **Originate**
 - insert routing information into external announcements (usually as a result of the IGP)
- **Peers**
 - routers in neighbouring AS's or within one AS which exchange routing and policy information

Routing flow and packet flow



For networks in AS1 and AS2 to communicate:

AS1 must announce to AS2

AS2 must accept from AS1


AS2 must announce to AS1

AS1 must accept from AS2

Routing flow and Traffic flow


- Traffic flow is always in the opposite direction of the flow of Routing information
 - Filtering outgoing routing information inhibits traffic flow inbound
 - Filtering inbound routing information inhibits traffic flow outbound

Routing Flow/Packet Flow: With multiple ASes



- For net N1 in AS1 to send traffic to net N16 in AS16:
 - AS16 must originate and announce N16 to AS8.
 - AS8 must accept N16 from AS16.
 - AS8 must forward announcement of N16 to AS1 or AS34.
 - AS1 must accept N16 from AS8 or AS34.
- For two-way packet flow, similar policies must exist for N1

Routing Flow/Packet Flow: With multiple ASes




- As multiple paths between sites are implemented it is easy to see how policies can become quite complex.

Routing Policy


- Used to control traffic flow in and out of an ISP network
- ISP makes decisions on what routing information to accept and discard from its neighbours
 - Individual routes
 - Routes originated by specific ASes
 - Routes traversing specific ASes
 - Routes belonging to other groupings
 - Groupings which you define as you see fit

Routing Policy Limitations



- AS99 uses red link for traffic to the red AS and the green link for remaining traffic
- To implement this policy, AS99 has to:
 - Accept routes originating from the red AS on the red link
 - Accept all other routes on the green link

Routing Policy Limitations



- AS99 would like packets coming from the green AS to use the green link.
- But unless AS22 cooperates in pushing traffic from the green AS down the green link, there is very little that AS99 can do to achieve this aim


IPv6 Addressing

- An IPv6 address is 128 bits long
- So the number of addresses are 2^{128}
=340282366920938463463374607431768211455
(39 decimal digits)
=0xffffffffffffffffffffffffffffffff (32 hexadecimal digits)
- In hex 4 bit (nibble) is represented by a hex digit
- So 128 bit is reduced down to 32 hex digit


IPv6 Address Representation

- Hexadecimal values of eight 16 bit fields
 - X:X:X:X:X:X:X:X (X=16 bit number, ex: A2FE)
 - 16 bit number is converted to a 4 digit hexadecimal number
- Example:
 - FE38:DCE3:124C:C1A2:BA03:6735:EF1C:683D
 - Abbreviated form of address
 - 4EED:0023:0000:0000:0000:036E:1250:2B00
 - →4EED:23:0:0:0:36E:1250:2B00
 - →4EED:23::36E:1250:2B00
 - (Null value can be used only once)


IPv6 addressing structure



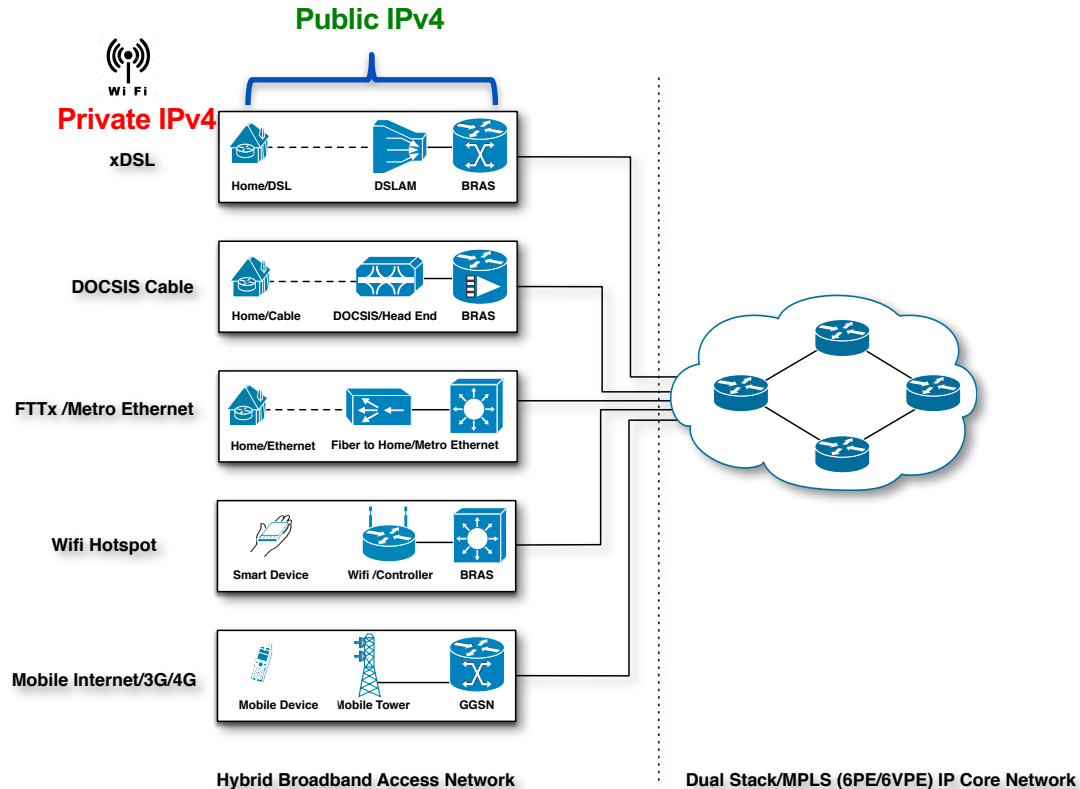
IPv4 for Broadband Access Network



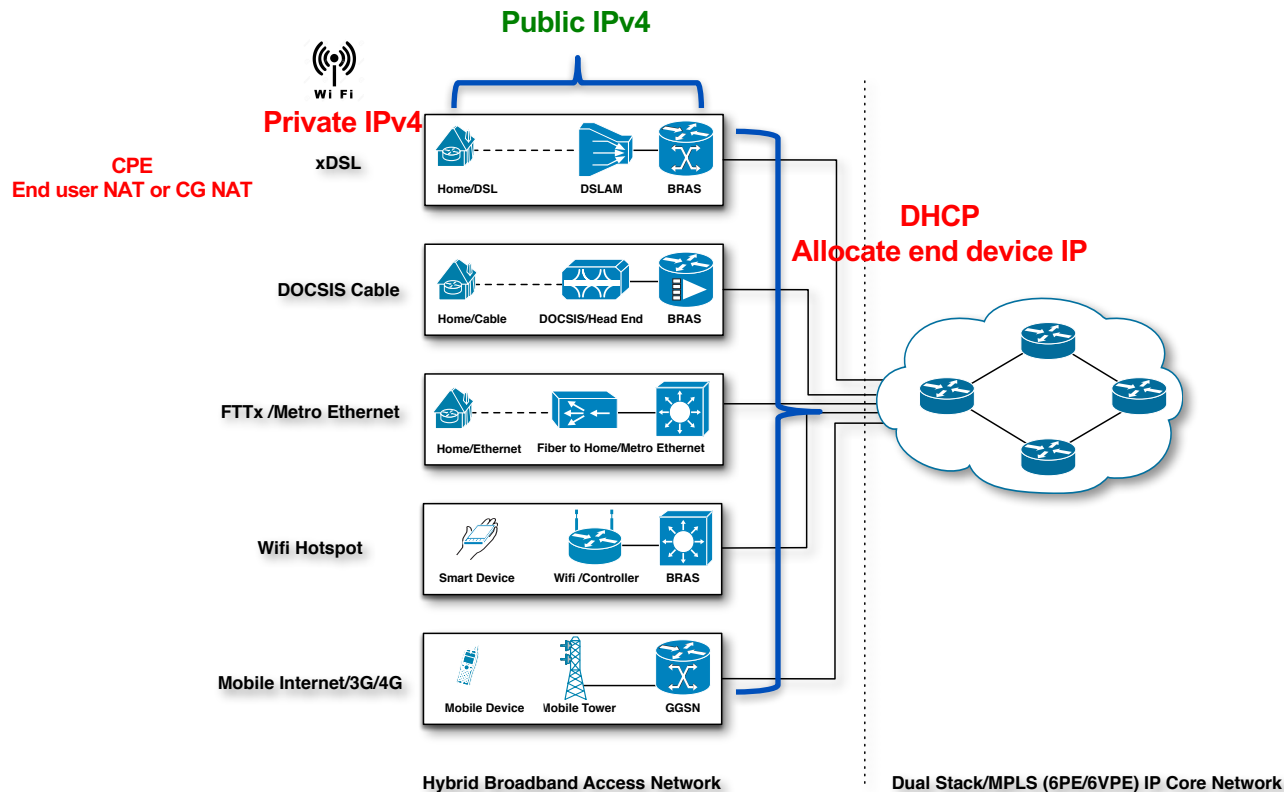
IPv4 for Broadband Access Network




IPv4 for Broadband Access Network




IPv4 for Broadband Access Network




IPv6 Broadband Access Network



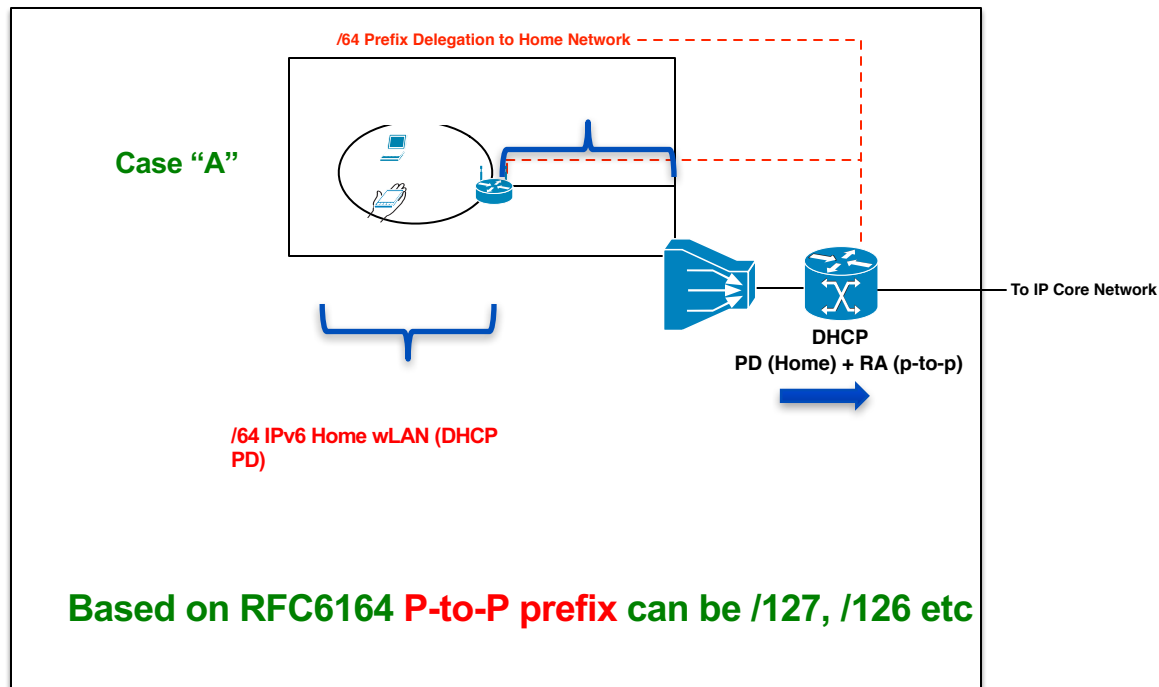
IPv6 Broadband Access Network



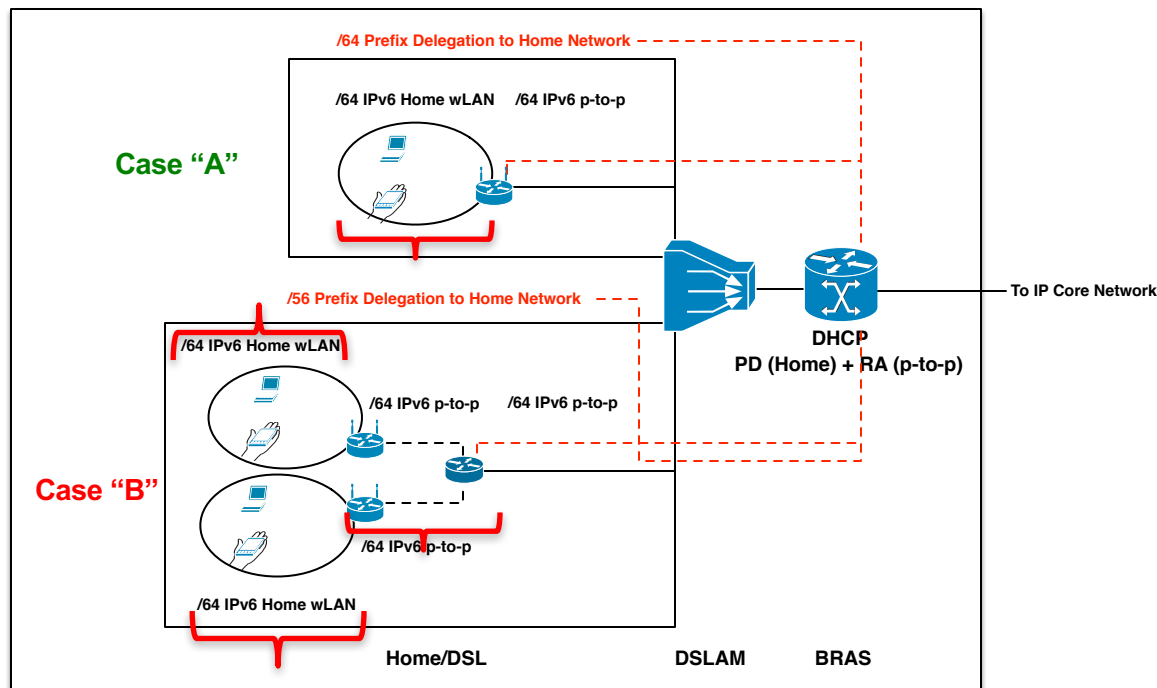
IPv6 Broadband Access Network




IPv6 Broadband Access Network




IPv6 Broadband Access Network



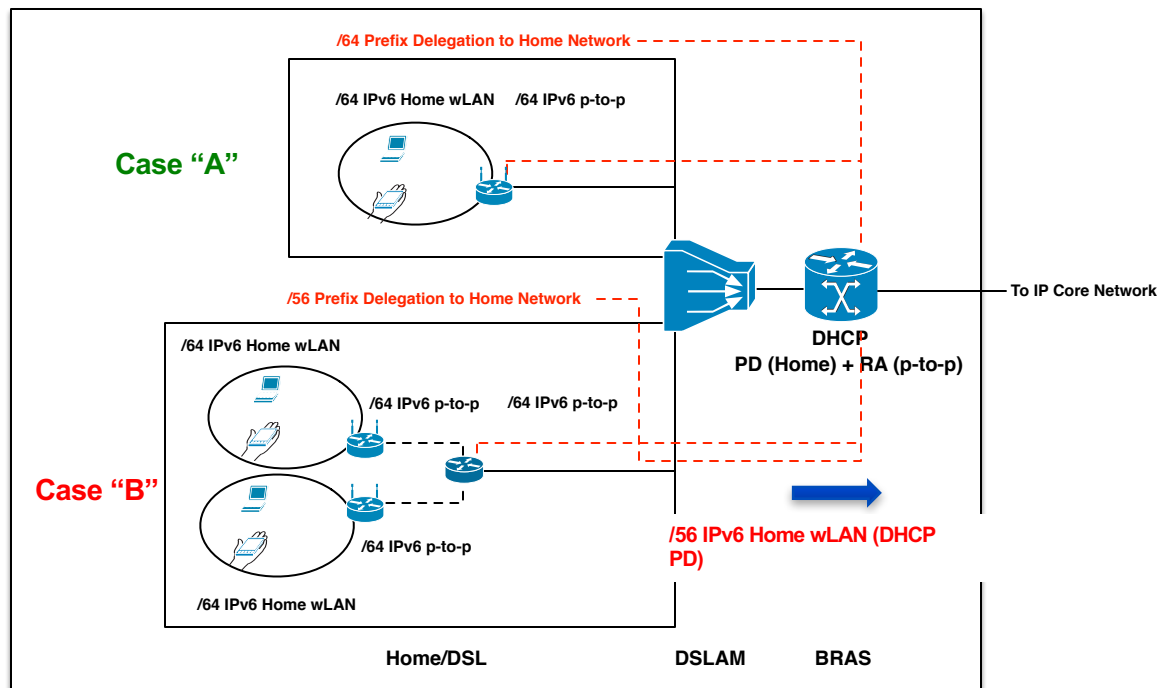
IPv6 Broadband Access Network



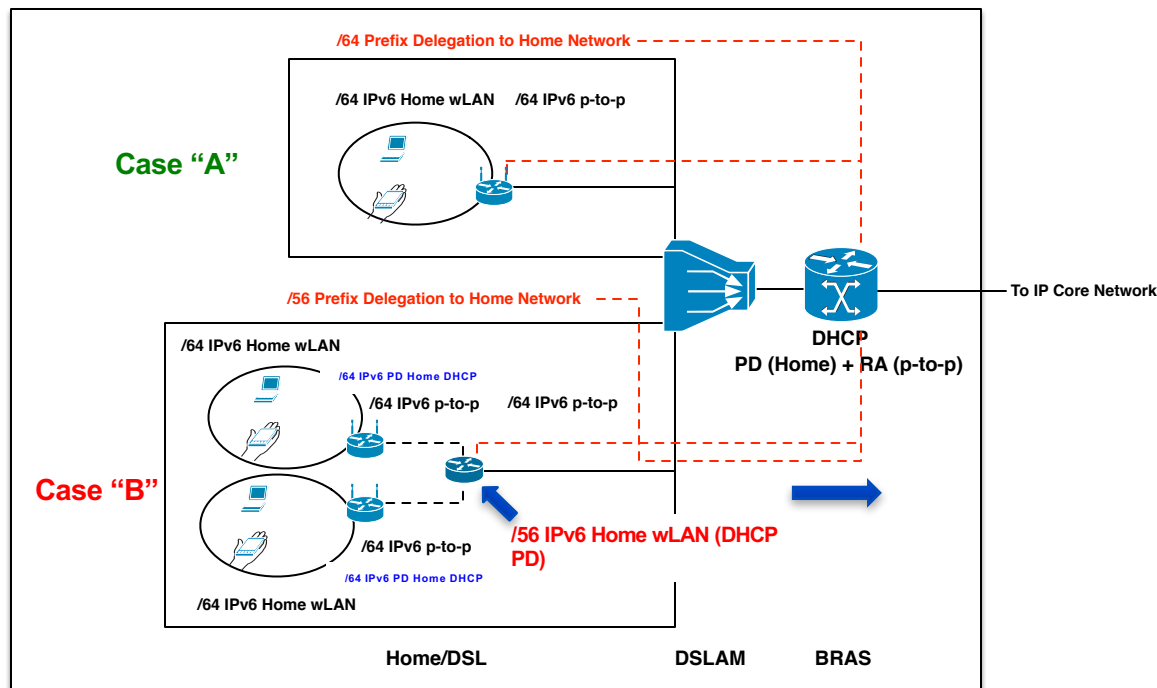
IPv6 Broadband Access Network



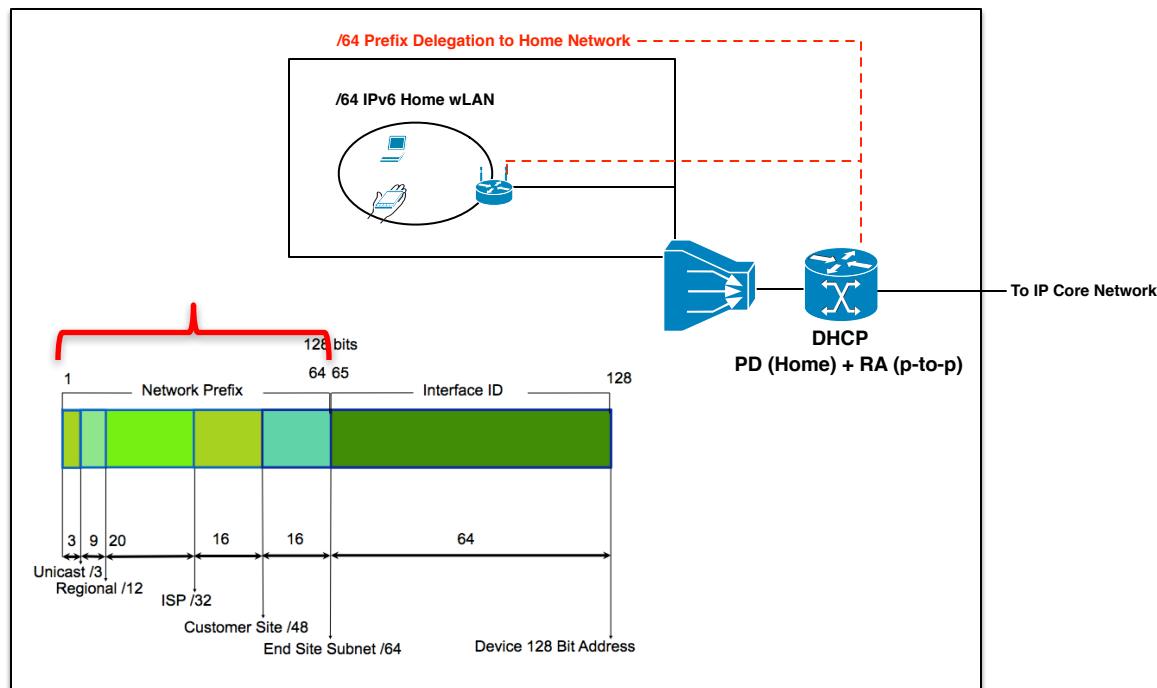
IPv6 Broadband Access Network




IPv6 Broadband Access Network



IPv6 Broadband Access Network



IPv6 Broadband Access Network



Policy Guideline on IPv6 Delegation

APNIC IPv6 Address Delegation Guideline

10. Delegations by LIRs

10.1. LIR assignments to end sites

An LIR can assign a /64 to /48 to an end site customer network based on their requirements.

The following guidelines may be useful:



- /64 where it is known that only one subnet is required.
- /56 for small sites where it is expected only a few subnets will be required within the next two years.
Subscribers can receive a /56 when connecting through on-demand or always-on connections such as small office and home office enterprises.
- /48 for larger sites, or if an end site is expected to grow into a large network.

An LIR must submit a [second opinion request](#) to APNIC if it plans to assign more than a /48 to a single end site (see Section 10.1.2 below).

Policy Guideline on IPv6 Delegation


APNIC IPv6 Address Delegation Guideline

10. Delegations by LIRs

10.1. LIR assignments to end sites

An LIR can assign a /64 to /48 to an end site customer network based on their requirements.

The following guidelines may be useful:

- 
- /64 where it is known that only one subnet is required.
 - /56 for small sites where it is expected only a few subnets will be required within the next two years.
Subscribers can receive a /56 when connecting through on-demand or always-on connections such as small office and home office enterprises.
 - /48 for larger sites, or if an end site is expected to grow into a large network.

An LIR must submit a [second opinion request](#) to APNIC if it plans to assign more than a /48 to a single end site (see Section 10.1.2 below).

Policy Guideline on IPv6 Delegation

APNIC IPv6 Address Delegation Guideline

10. Delegations by LIRs

10.1. LIR assignments to end sites

An LIR can assign a /64 to /48 to an end site customer network based on their requirements.

The following guidelines may be useful:

- /64 where it is known that only one subnet is required.
- /56 for small sites where it is expected only a few subnets will be required within the next two years.
Subscribers can receive a /56 when connecting through on-demand or always-on connections such as small office and home office enterprises.
- /48 for larger sites, or if an end site is expected to grow into a large network.

An LIR must submit a [second opinion request](#) to APNIC if it plans to assign more than a /48 to a single end site (see Section 10.1.2 below).

Overview


IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise: Basic Routing, Interface & OSPF**
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise: BGP Peering**
- IXP Design Considerations
- **Hands On Lab Exercise: IXP Configuration**
- Route Collectors & Servers
- IXP BCP and What can go wrong?

Training ISP Network Topology

- Scenario:
 - Training ISP has 4 main operating area or region
 - Each region has 2 small POP
 - Each region will have one datacenter to host content
 - Regional network are inter-connected with multiple link

Training ISP Network Topology



Training ISP Topology Diagram

Training ISP Network Topology

- Regional Network:
 - Each regional network will have 3 routers
 - 1 Core & 2 Edge Routers
 - 2 Point of Presence (POP) for every region
 - POP will use a router to terminate customer network i.e Edge Router
 - Each POP is an aggregation point of ISP customer

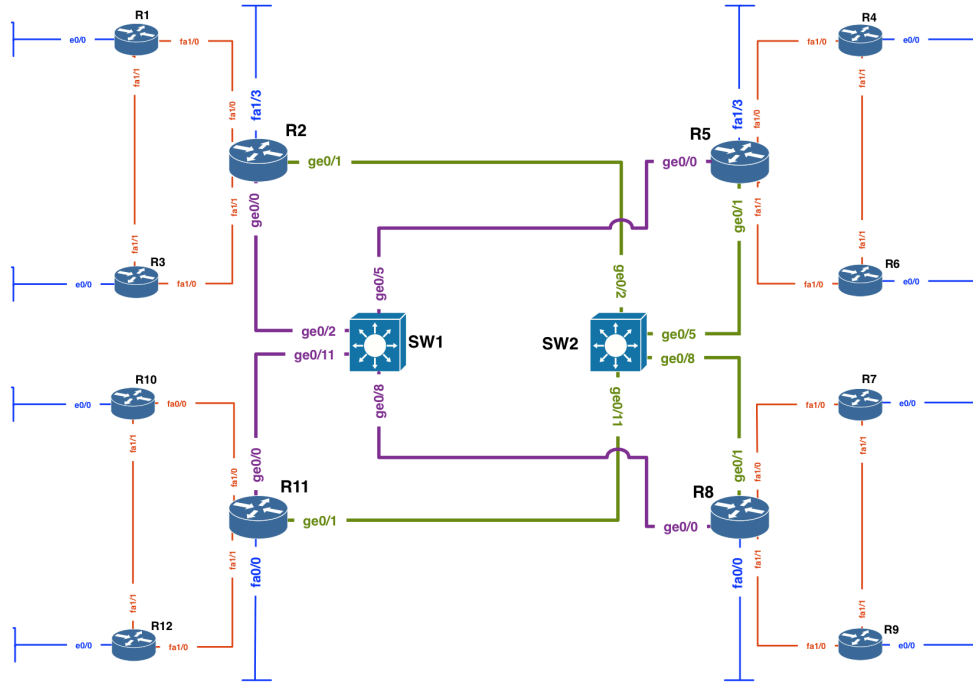
Training ISP Network Topology

- Access Network:
 - Connection between customer network & Edge router
 - Usually 10 to 100 MBPS link
 - Separate routing policy from most of ISP
 - Training ISP will connect them on edge router with separate customer IP prefix

Training ISP Network Topology

- Transport Link:
 - Inter-connection between regional core router
 - Higher data transmission capacity than access link
 - Training ISP has 2 transport link for link redundancy
 - 2 Transport link i.e Purple link & Green link are connected to two career grade switch

Training ISP Network Topology



Training ISP Core IP Backbone

Training ISP Network Topology

- Design Consideration:
 - Each regional network should have address summarization capability for customer block and CS link WAN.
 - Prefix planning should have scalability option for next couple of years for both customer block and infrastructure
 - No Summarization require for infrastructure WAN and loopback address

Training ISP Network Topology

- Design Consideration:
 - All WAN link should be ICMP reachable for link monitoring purpose (At least from designated host)
 - Conservation will get high preference for IPv4 address planning and aggregation will get high preference for IPv6 address planning.

Training ISP Network Topology

- Design Consideration:
 - OSPF is running in ISP network to carry infrastructure IP prefix
 - Each region is a separate OSPF area
 - Transport core is in OSPF area 0
 - Customer will connect on either static or eBGP (Not OSPF)
 - iBGP will carry external prefix within ISP core IP network

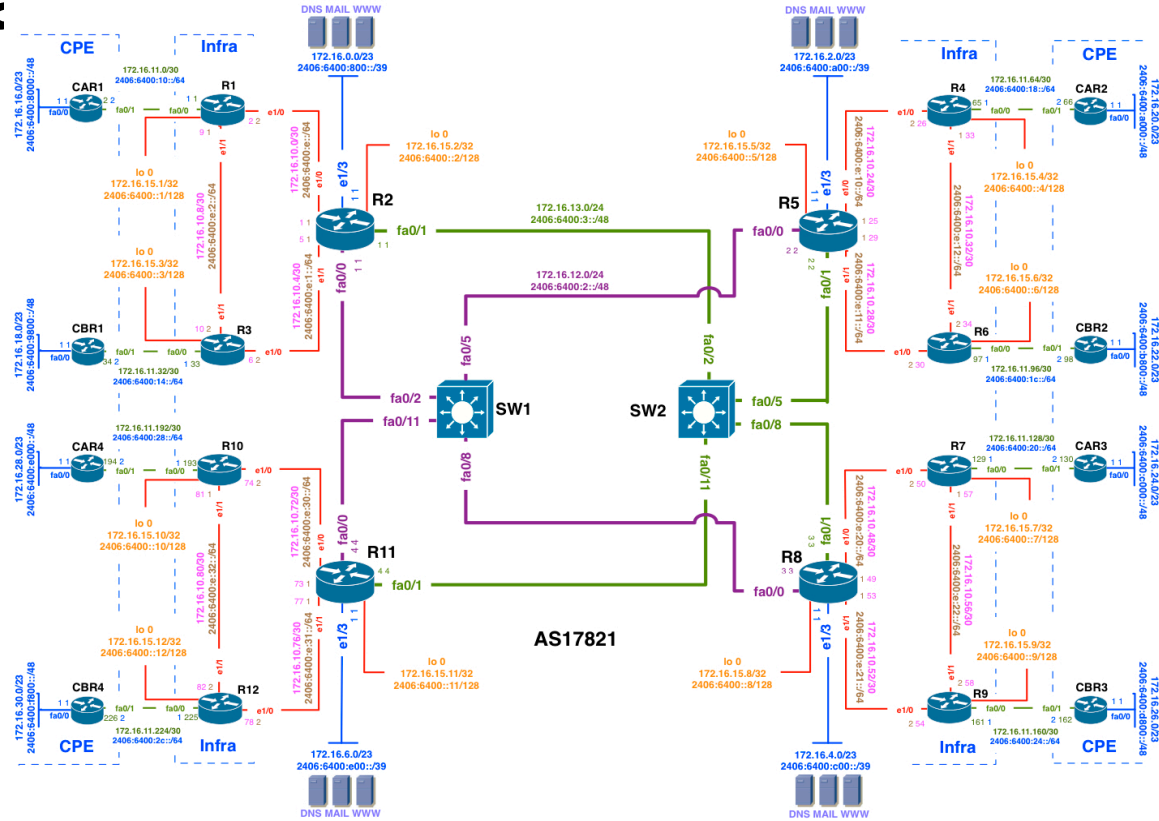
Training ISP IPV6 Addressing Plan

- IPv6 address plan consideration:
 - Big IPv6 address space can cause very very large routing table size
 - Most transit service provider apply IPv6 aggregation prefix filter (i.e. anything other than /48 & \leq /32 prefix size
 - Prefix announcement need to send to Internet should be either /32 or /48 bit boundary

Training ISP IPV6 Addressing Plan

- IPv6 address plan consideration (RFC3177):
 - WAN link can be used on /64 bit boundary
 - End site/Customer sub allocation can be made between /48~/64 bit boundary
 - APNIC Utilization/HD ratio will be calculated based on /56 end site assignment/sub-allocation

Training ISP IPV6 Addressing Plan



Addressing Plans – ISP Infrastructure

- What about LANs?
 - /64 per LAN
- What about Point-to-Point links?
 - Protocol design expectation is that /64 is used
 - /127 now recommended/standardised
 - <http://www.rfc-editor.org/rfc/rfc6164.txt>
 - (reserve /64 for the link, but address it as a /127)
 - Other options:
 - /126s are being used (mirrors IPv4 /30)
 - /112s are being used
 - Leaves final 16 bits free for node IDs
 - Some discussion about /80s, /96s and /120s too

Addressing Plans – ISP Infrastructure

- ISPs should receive /32 from their RIR
- Address block for router loop-back interfaces
 - Generally number all loopbacks out of one /48
 - /128 per loopback
- Address block for infrastructure
 - /48 allows 65k subnets
 - /48 per region (for the largest international networks)
 - /48 for whole backbone (for the majority of networks)
 - Summarise between sites if it makes sense

Addressing Plans – Customer

- Customers get one /48
 - Unless they have more than 65k subnets in which case they get a second /48 (and so on)
- In typical deployments today:
 - Several ISPs give small customers a /56 or single LAN end-sites a /64, e.g.:
 - /64 if end-site will only ever be a LAN
 - /56 for medium end-sites (e.g. small business)
 - /48 for large end-sites
 - (This is another very active discussion area)

Addressing Plans – Advice

- Customer address assignments should not be reserved or assigned on a per PoP basis
 - Same principle as for IPv4
 - ISP iBGP carries customer nets
 - Aggregation within the iBGP not required and usually not desirable
 - Aggregation in eBGP is very necessary
- Backbone infrastructure assignments:
 - Number out of a single /48
 - Operational simplicity and security
 - Aggregate to minimise size of the IGP

Addressing Plans Planning

- Registries will usually allocate the next block to be contiguous with the first allocation
 - Minimum allocation is /32
 - Very likely that subsequent allocation will make this up to a /31
 - So plan accordingly

Example Address Plan

- IPv6 Allocation From Registry is
 - 2406:6400::/32
- IPv4 Allocation From Registry is
 - 172.16.0.0/19

Training ISP IPV6 Addressing Plan

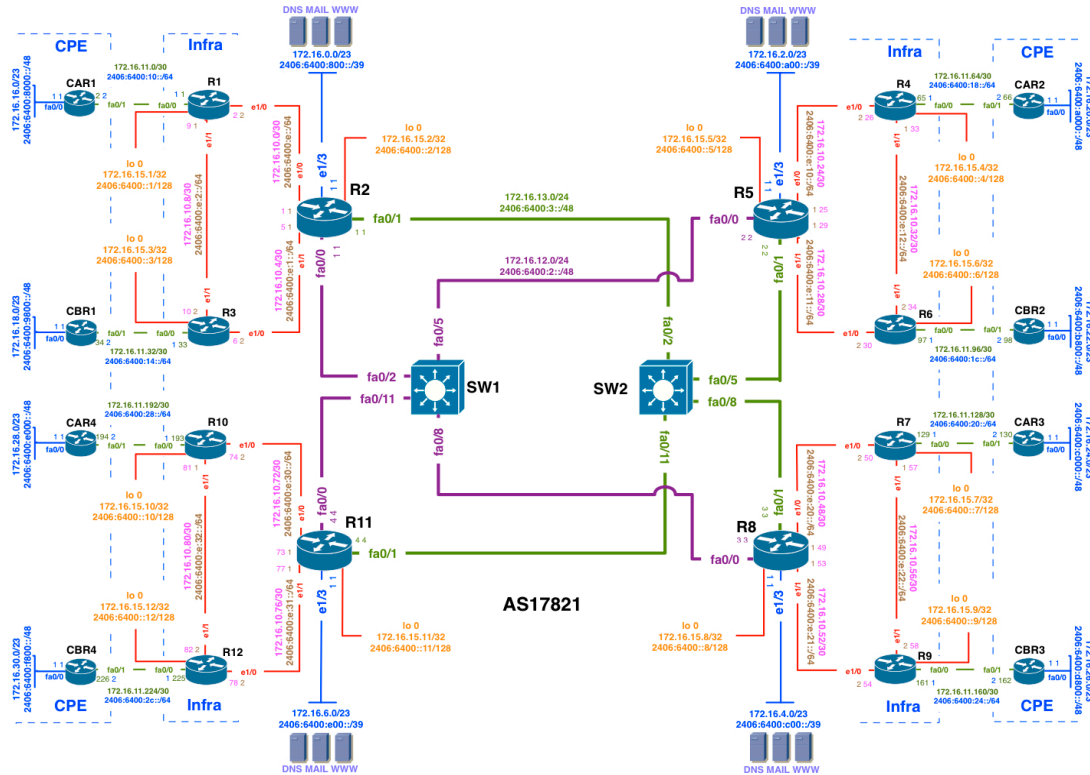
| Table 1: Top level distribution infrastructure & customer | | | | | |
|---|--------------------------|---------------------------|-----------------------------|---------|--------------|
| Block# | Prefix | Description | Reverse Domain | SOR | Registration |
| 1 | 2406:6400::/32 | <i>Parent Block</i> | 0.0.4.6.6.0.4.2.ip6.arpa. | N/A | APNIC |
| 2 | 2406:6400:0000:0000::/36 | Infrastructure | 0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Optional |
| | 2406:6400:1000:0000::/36 | | | | |
| | 2406:6400:2000:0000::/36 | | | | |
| | 2406:6400:3000:0000::/36 | | | | |
| | 2406:6400:4000:0000::/36 | | | | |
| | 2406:6400:5000:0000::/36 | | | | |
| | 2406:6400:6000:0000::/36 | | | | |
| | 2406:6400:7000:0000::/36 | | | | |
| 3 | 2406:6400:8000:0000::/36 | Customer network Region 1 | 8.0.0.4.6.6.0.4.2.ip6.arpa. | Not yet | Optional |
| | 2406:6400:9000:0000::/36 | | | | |
| 4 | 2406:6400:a000:0000::/36 | Customer network Region 2 | a.0.0.4.6.6.0.4.2.ip6.arpa. | Not yet | Optional |
| | 2406:6400:b000:0000::/36 | | | | |
| 5 | 2406:6400:c000:0000::/36 | Customer network Region 3 | c.0.0.4.6.6.0.4.2.ip6.arpa. | Not yet | Optional |
| | 2406:6400:d000:0000::/36 | | | | |
| 6 | 2406:6400:e000:0000::/36 | Customer network Region 4 | e.0.0.4.6.6.0.4.2.ip6.arpa. | Not yet | Optional |
| | 2406:6400:f000:0000::/36 | | | | |

Training ISP IPV6 Addressing Plan

Table 2: Top level summarization option infrastructure & customer

| Block# | Prefix | Description | Reverse Domain |
|--------|--------------------------|------------------------------|-------------------|
| 7 | 2406:6400:8000:0000::/35 | CS net summary region1 [R2] | 2x/36 arpa domain |
| 8 | 2406:6400:a000:0000::/35 | CS net summary region2 [R5] | 2x/36 arpa domain |
| 9 | 2406:6400:c000:0000::/35 | CS net summary region3 [R8] | 2x/36 arpa domain |
| 10 | 2406:6400:e000:0000::/35 | CS net summary region4 [R11] | 2x/36 arpa domain |

Training ISP IPV6 Addressing Plan



Training ISP IPV6 Addressing Plan

| Table 3: Detail distribution infrastructure | | | | | |
|---|--------------------------|--------------------------------------|-------------------------------|-----|--------------|
| Block# | Prefix | Description | Reverse Domain | SOR | Registration |
| 2 | 2406:6400:0000:0000::/36 | Infrastructure | 0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Optional |
| 11 | 2406:6400:0000:0000::/40 | Loopback, Transport & WAN [Infra+CS] | 0.0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Optional |
| | 2406:6400:0100:0000::/40 | | | | |
| | 2406:6400:0200:0000::/40 | | | | |
| | 2406:6400:0300:0000::/40 | | | | |
| | 2406:6400:0400:0000::/40 | | | | |
| | 2406:6400:0500:0000::/40 | | | | |
| | 2406:6400:0600:0000::/40 | | | | |
| | 2406:6400:0700:0000::/40 | | | | |
| 16 | 2406:6400:0800:0000::/40 | R2 DC | 8.0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Recommended |
| | 2406:6400:0900:0000::/40 | | | | |
| 17 | 2406:6400:0a00:0000::/40 | R5 DC | a.0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Recommended |
| | 2406:6400:0b00:0000::/40 | | | | |
| 18 | 2406:6400:0c00:0000::/40 | R8 DC | c.0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Recommended |
| | 2406:6400:0d00:0000::/40 | | | | |
| 19 | 2406:6400:0e00:0000::/40 | R11 DC | e.0.0.0.4.6.6.0.4.2.ip6.arpa. | No | Recommended |
| | 2406:6400:0f00:0000::/40 | | | | |

Training ISP IPV6 Addressing Plan

Table 4: Datacenter prefix summarization options

| Block# | Prefix | Description | Reverse Domain |
|--------|--------------------------|---------------------------|----------------|
| 12 | 2406:6400:0800:0000::/39 | Region 1 DC Summary [R2] | |
| 13 | 2406:6400:0a00:0000::/39 | Region 2 DC Summary [R5] | |
| 14 | 2406:6400:0c00:0000::/39 | Region 3 DC Summary [R8] | |
| 15 | 2406:6400:0e00:0000::/39 | Region 4 DC Summary [R11] | |

APNIC



Training ISP IPV6 Addressing Plan

| Table 5: Further detail loopback, transport & infrastructure WAN | | | | | |
|--|--------------------------|---------------------------------|-------------------------------|-----|--------------|
| Block# | Prefix | Description | Reverse Domain | SOR | Registration |
| 11 | 2406:6400:0000:0000::/40 | Loopback, Transport & Infra WAN | 0.0.0.0.4.6.6.0.4.2.ip6.arpa. | | |
| 20 | 2406:6400:0000:0000::/48 | Loopback | | No | Recommended |
| | 2406:6400:0001:0000::/48 | | | | |
| 21 | 2406:6400:0002:0000::/48 | Purple Transport | | No | Recommended |
| 22 | 2406:6400:0003:0000::/48 | Green Transport | | No | Recommended |
| | 2406:6400:0004:0000::/48 | | | | |
| | 2406:6400:0005:0000::/48 | | | | |
| | 2406:6400:0006:0000::/48 | | | | |
| | 2406:6400:0007:0000::/48 | | | | |
| | 2406:6400:0008:0000::/48 | | | | |
| | 2406:6400:0009:0000::/48 | | | | |
| | 2406:6400:000A:0000::/48 | | | | |
| | 2406:6400:000B:0000::/48 | | | | |
| | 2406:6400:000C:0000::/48 | | | | |
| | 2406:6400:000D:0000::/48 | | | | |
| 23 | 2406:6400:000E:0000::/48 | WAN Prefix Infra Link | | No | Recommended |
| | 2406:6400:000F:0000::/48 | | | | |

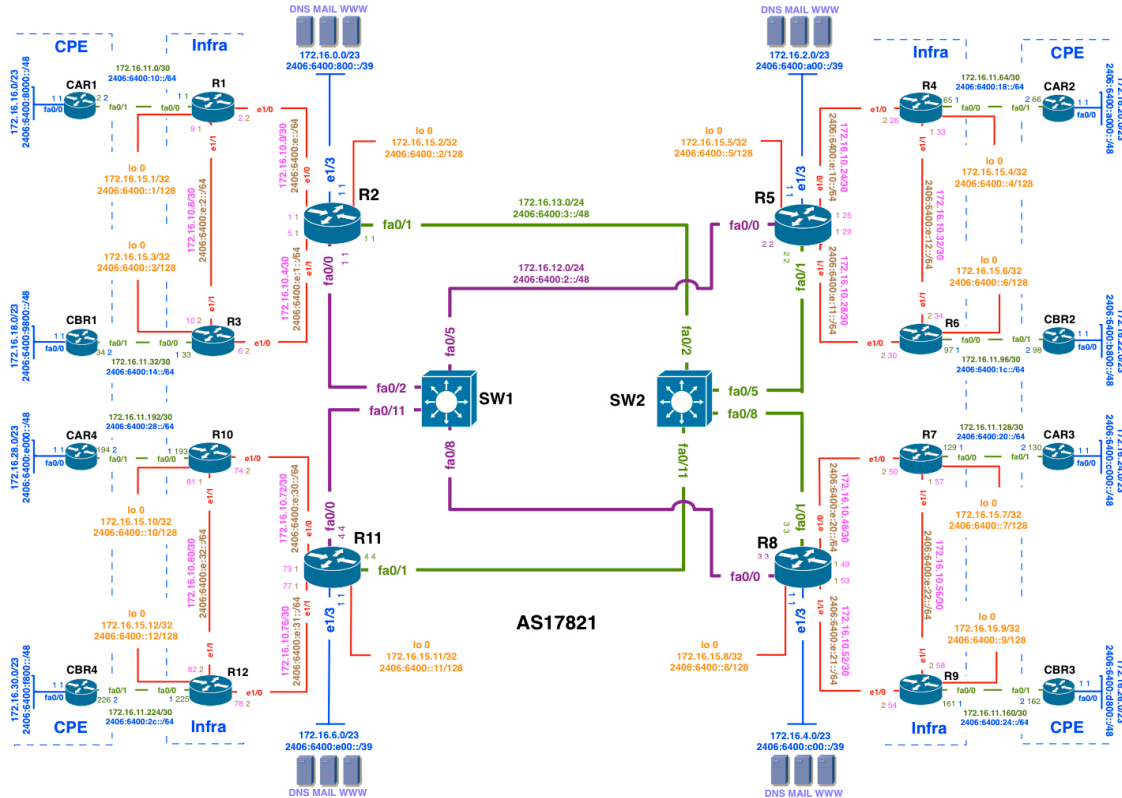
Training ISP IPV6 Addressing Plan

| Table 6: Further detail CS link WAN | | | | | |
|-------------------------------------|--------------------------|--------------------------------|----------------|-----|--------------|
| Block# | Prefix | Description | Reverse Domain | SOR | Registration |
| 27 | 2406:6400:0010:0000::/48 | WAN Prefix CS Link R1 Region1 | | No | Recommended |
| | 2406:6400:0011:0000::/48 | | | | |
| | 2406:6400:0012:0000::/48 | | | | |
| | 2406:6400:0013:0000::/48 | | | | |
| 28 | 2406:6400:0014:0000::/48 | WAN Prefix CS Link R3 Region1 | | No | Recommended |
| | 2406:6400:0015:0000::/48 | | | | |
| | 2406:6400:0016:0000::/48 | | | | |
| | 2406:6400:0017:0000::/48 | | | | |
| 32 | 2406:6400:0018:0000::/48 | WAN Prefix CS Link R4 Region2 | | No | Recommended |
| | 2406:6400:0019:0000::/48 | | | | |
| | 2406:6400:001A:0000::/48 | | | | |
| | 2406:6400:001B:0000::/48 | | | | |
| 33 | 2406:6400:001C:0000::/48 | WAN Prefix CS Link R6 Region2 | | No | Recommended |
| | 2406:6400:001D:0000::/48 | | | | |
| | 2406:6400:001E:0000::/48 | | | | |
| | 2406:6400:001F:0000::/48 | | | | |
| 37 | 2406:6400:0020:0000::/48 | WAN Prefix CS Link R7 Region3 | | No | Recommended |
| | 2406:6400:0021:0000::/48 | | | | |
| | 2406:6400:0022:0000::/48 | | | | |
| | 2406:6400:0023:0000::/48 | | | | |
| 38 | 2406:6400:0024:0000::/48 | WAN Prefix CS Link R9 Region3 | | No | Recommended |
| | 2406:6400:0025:0000::/48 | | | | |
| | 2406:6400:0026:0000::/48 | | | | |
| | 2406:6400:0027:0000::/48 | | | | |
| 42 | 2406:6400:0028:0000::/48 | WAN Prefix CS Link R10 Region4 | | No | Recommended |
| | 2406:6400:0029:0000::/48 | | | | |
| | 2406:6400:002A:0000::/48 | | | | |
| | 2406:6400:002B:0000::/48 | | | | |
| 43 | 2406:6400:002C:0000::/48 | WAN Prefix CS Link R12 Region4 | | No | Recommended |
| | 2406:6400:002D:0000::/48 | | | | |
| | 2406:6400:002E:0000::/48 | | | | |
| | 2406:6400:002F:0000::/48 | | | | |

Training ISP IPV6 Addressing Plan

| Table 7: CS link WAN summarization options | | | |
|--|--------------------------|--|----------------|
| Block# | Prefix | Description | Reverse Domain |
| 24 | 2406:6400:0010:0000::/45 | WAN CS Link Region1 Summary [R2] | |
| 25 | 2406:6400:0010:0000::/46 | WAN CS Link Region1 POP1 Summary [R1] | |
| 26 | 2406:6400:0014:0000::/46 | WAN CS Link Region1 POP2 Summary [R3] | |
| | | | |
| | | | |
| | | | |
| Block# | Prefix | Description | Reverse Domain |
| 29 | 2406:6400:0018:0000::/45 | WAN Prefix CS Link Region2 Summary [R5] | |
| 30 | 2406:6400:0018:0000::/46 | WAN CS Link Region2 POP1 Summary [R4] | |
| 31 | 2406:6400:001C:0000::/46 | WAN CS Link Region2 POP2 Summary [R6] | |
| | | | |
| | | | |
| | | | |
| Block# | Prefix | Description | Reverse Domain |
| 34 | 2406:6400:0020:0000::/45 | WAN Prefix CS Link Region3 Summary [R8] | |
| 35 | 2406:6400:0020:0000::/46 | WAN CS Link Region3 POP1 Summary [R7] | |
| 36 | 2406:6400:0024:0000::/46 | WAN CS Link Region3 POP2 Summary [R9] | |
| | | | |
| | | | |
| | | | |
| Block# | Prefix | Description | Reverse Domain |
| 39 | 2406:6400:0028:0000::/45 | WAN Prefix CS Link Region4 Summary [R11] | |
| 40 | 2406:6400:0028:0000::/46 | WAN CS Link Region4 POP1 Summary [R10] | |
| 41 | 2406:6400:002C:0000::/46 | WAN CS Link Region4 POP2 Summary [R12] | |

Training ISP IPV6 Addressing Plan



Training ISP IPV6 Addressing Plan

| Table 8: Further detail loopback | | | | | |
|----------------------------------|-----------------------------|---------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 20 | 2406:6400:0000:0000::/48 | Loopback | | No | Recommended |
| | | | YES | | |
| 43 | 2406:6400:0000:0000::1/128 | Router1 loopback 0 | YES | No | No |
| 44 | 2406:6400:0000:0000::2/128 | Router2 loopback 0 | YES | No | No |
| 45 | 2406:6400:0000:0000::3/128 | Router3 loopback 0 | YES | No | No |
| 46 | 2406:6400:0000:0000::4/128 | Router4 loopback 0 | YES | No | No |
| 47 | 2406:6400:0000:0000::5/128 | Router5 loopback 0 | YES | No | No |
| 48 | 2406:6400:0000:0000::6/128 | Router6 loopback 0 | YES | No | No |
| 49 | 2406:6400:0000:0000::7/128 | Router7 loopback 0 | YES | No | No |
| 50 | 2406:6400:0000:0000::8/128 | Router8 loopback 0 | YES | No | No |
| 51 | 2406:6400:0000:0000::9/128 | Router9 loopback 0 | YES | No | No |
| 52 | 2406:6400:0000:0000::10/128 | Router10 loopback 0 | YES | No | No |
| 53 | 2406:6400:0000:0000::11/128 | Router11 loopback 0 | YES | No | No |
| 54 | 2406:6400:0000:0000::12/128 | Router12 loopback 0 | YES | No | No |

Training ISP IPV6 Addressing Plan

| Table 9: Further detail transport | | | | | |
|-----------------------------------|---------------------------|------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 21 | 2406:6400:0002:0000::/48 | Purple Transport | | No | Recommended |
| | 2406:6400:0002:0000::1/48 | Router2 fa0/0 | YES | No | No |
| | 2406:6400:0002:0000::2/48 | Router5 fa0/0 | YES | No | No |
| | 2406:6400:0002:0000::3/48 | Router8 fa0/0 | YES | No | No |
| | 2406:6400:0002:0000::4/48 | Router11 fa0/0 | YES | No | No |
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 22 | 2406:6400:0003:0000::/48 | Green Transport | | No | Recommended |
| | 2406:6400:0003:0000::1/48 | Router2 fa0/1 | YES | No | No |
| | 2406:6400:0003:0000::2/48 | Router5 fa0/1 | YES | No | No |
| | 2406:6400:0003:0000::3/48 | Router8 fa0/1 | YES | No | No |
| | 2406:6400:0003:0000::4/48 | Router11 fa0/1 | YES | No | No |

Training ISP IPV6 Addressing Plan

| Table 10: Further detail Infra WAN | | | | | |
|------------------------------------|--------------------------|-----------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 23 | 2406:6400:000E:0000::/48 | WAN Prefix Infra Link | | No | Recommended |
| 55 | 2406:6400:000E:0000::/64 | R2[::1]-R1[::2] | YES | No | No |
| 56 | 2406:6400:000E:0001::/64 | R2[::1]-R3[::2] | YES | No | No |
| 57 | 2406:6400:000E:0002::/64 | R1[::1]-R3[::2] | YES | No | No |
| | 2406:6400:000E:0003::/64 | | | | |
| | 2406:6400:000E:0004::/64 | | | | |
| | 2406:6400:000E:0005::/64 | | | | |
| | 2406:6400:000E:0006::/64 | | | | |
| | 2406:6400:000E:0007::/64 | | | | |
| | 2406:6400:000E:0008::/64 | | | | |
| | 2406:6400:000E:0009::/64 | | | | |
| | 2406:6400:000E:000A::/64 | | | | |
| | 2406:6400:000E:000B::/64 | | | | |
| | 2406:6400:000E:000C::/64 | | | | |
| | 2406:6400:000E:000D::/64 | | | | |
| | 2406:6400:000E:000E::/64 | | | | |
| | 2406:6400:000E:000F::/64 | | | | |
| 58 | 2406:6400:000E:0010::/64 | R5[::1]-R4[::2] | YES | No | No |
| 59 | 2406:6400:000E:0011::/64 | R5[::1]-R6[::2] | YES | No | No |
| 60 | 2406:6400:000E:0012::/64 | R4[::1]-R6[::2] | YES | No | No |
| | 2406:6400:000E:0013::/64 | | | | |
| | 2406:6400:000E:0014::/64 | | | | |
| | 2406:6400:000E:0015::/64 | | | | |
| | 2406:6400:000E:0016::/64 | | | | |
| | 2406:6400:000E:0017::/64 | | | | |
| | 2406:6400:000E:0018::/64 | | | | |
| | 2406:6400:000E:0019::/64 | | | | |
| | 2406:6400:000E:001A::/64 | | | | |
| | 2406:6400:000E:001B::/64 | | | | |
| | 2406:6400:000E:001C::/64 | | | | |
| | 2406:6400:000E:001D::/64 | | | | |
| | 2406:6400:000E:001E::/64 | | | | |
| | 2406:6400:000E:001F::/64 | | | | |
| 61 | 2406:6400:000E:0020::/64 | R8[::1]-R7[::2] | YES | No | No |
| 62 | 2406:6400:000E:0021::/64 | R8[::1]-R9[::2] | YES | No | No |
| 63 | 2406:6400:000E:0022::/64 | R7[::1]-R9[::2] | YES | No | No |
| | 2406:6400:000E:0023::/64 | | | | |
| | 2406:6400:000E:0024::/64 | | | | |
| | 2406:6400:000E:0025::/64 | | | | |
| | 2406:6400:000E:0026::/64 | | | | |
| | 2406:6400:000E:0027::/64 | | | | |
| | 2406:6400:000E:0028::/64 | | | | |
| | 2406:6400:000E:0029::/64 | | | | |
| | 2406:6400:000E:002A::/64 | | | | |
| | 2406:6400:000E:002B::/64 | | | | |
| | 2406:6400:000E:002C::/64 | | | | |
| | 2406:6400:000E:002D::/64 | | | | |
| | 2406:6400:000E:002E::/64 | | | | |
| | 2406:6400:000E:002F::/64 | | | | |
| 64 | 2406:6400:000E:0030::/64 | R11[::1]-R10[::2] | YES | No | No |
| 65 | 2406:6400:000E:0031::/64 | R11[::1]-R12[::2] | YES | No | No |
| 66 | 2406:6400:000E:0032::/64 | R10[::1]-R12[::2] | YES | No | No |
| | 2406:6400:000E:0033::/64 | | | | |
| | 2406:6400:000E:0034::/64 | | | | |
| | 2406:6400:000E:0035::/64 | | | | |
| | 2406:6400:000E:0036::/64 | | | | |
| | 2406:6400:000E:0037::/64 | | | | |
| | 2406:6400:000E:0038::/64 | | | | |
| | 2406:6400:000E:0039::/64 | | | | |
| | 2406:6400:000E:003A::/64 | | | | |

Training ISP IPV6 Addressing Plan

| Table 11: Detail CS link WAN Region 1 | | | | | |
|---------------------------------------|--------------------------|-------------------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 27 | 2406:6400:0010:0000::/48 | WAN Prefix CS Link R1 Region1 | | No | Recommended |
| | 2406:6400:0010:0000::/64 | R1[::1]-CAR1[::2] | Yes | No | No |
| | 2406:6400:0010:0001::/64 | | Yes | No | No |
| | 2406:6400:0010:0002::/64 | | Yes | No | No |
| | 2406:6400:0010:0003::/64 | | Yes | No | No |
| | 2406:6400:0010:0004::/64 | | Yes | No | No |
| | 2406:6400:0010:0005::/64 | | Yes | No | No |
| | 2406:6400:0010:0006::/64 | | Yes | No | No |
| | 2406:6400:0010:0007::/64 | | Yes | No | No |
| | 2406:6400:0010:0008::/64 | | Yes | No | No |
| | 2406:6400:0010:0009::/64 | | Yes | No | No |
| | 2406:6400:0010:000A::/64 | | Yes | No | No |
| | 2406:6400:0010:000B::/64 | | Yes | No | No |
| | 2406:6400:0010:000C::/64 | | Yes | No | No |
| | 2406:6400:0010:000D::/64 | | Yes | No | No |
| | 2406:6400:0010:000E::/64 | | Yes | No | No |
| | 2406:6400:0010:000F::/64 | | Yes | No | No |
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 28 | 2406:6400:0014:0000::/48 | WAN Prefix CS Link R3 Region1 | | No | Recommended |
| | 2406:6400:0014:0000::/64 | R3[::1]-CBR1[::2] | Yes | No | No |
| | 2406:6400:0014:0001::/64 | | Yes | No | No |
| | 2406:6400:0014:0002::/64 | | Yes | No | No |
| | 2406:6400:0014:0003::/64 | | Yes | No | No |
| | 2406:6400:0014:0004::/64 | | Yes | No | No |
| | 2406:6400:0014:0005::/64 | | Yes | No | No |
| | 2406:6400:0014:0006::/64 | | Yes | No | No |
| | 2406:6400:0014:0007::/64 | | Yes | No | No |
| | 2406:6400:0014:0008::/64 | | Yes | No | No |
| | 2406:6400:0014:0009::/64 | | Yes | No | No |
| | 2406:6400:0014:000A::/64 | | Yes | No | No |
| | 2406:6400:0014:000B::/64 | | Yes | No | No |
| | 2406:6400:0014:000C::/64 | | Yes | No | No |
| | 2406:6400:0014:000D::/64 | | Yes | No | No |
| | 2406:6400:0014:000E::/64 | | Yes | No | No |
| | 2406:6400:0014:000F::/64 | | Yes | No | No |

Training ISP IPV6 Addressing Plan

| Table 12: Detail CS link WAN Region 2 | | | | | |
|---------------------------------------|--------------------------|-------------------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 32 | 2406:6400:0018:0000::/48 | WAN Prefix CS Link R4 Region2 | | No | Recommended |
| | 2406:6400:0018:0000::/64 | R4[::1]-CAR2[::2] | Yes | No | No |
| | 2406:6400:0018:0001::/64 | | Yes | No | No |
| | 2406:6400:0018:0002::/64 | | Yes | No | No |
| | 2406:6400:0018:0003::/64 | | Yes | No | No |
| | 2406:6400:0018:0004::/64 | | Yes | No | No |
| | 2406:6400:0018:0005::/64 | | Yes | No | No |
| | 2406:6400:0018:0006::/64 | | Yes | No | No |
| | 2406:6400:0018:0007::/64 | | Yes | No | No |
| | 2406:6400:0018:0008::/64 | | Yes | No | No |
| | 2406:6400:0018:0009::/64 | | Yes | No | No |
| | 2406:6400:0018:000A::/64 | | Yes | No | No |
| | 2406:6400:0018:000B::/64 | | Yes | No | No |
| | 2406:6400:0018:000C::/64 | | Yes | No | No |
| | 2406:6400:0018:000D::/64 | | Yes | No | No |
| | 2406:6400:0018:000E::/64 | | Yes | No | No |
| | 2406:6400:0018:000F::/64 | | Yes | No | No |
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 33 | 2406:6400:001C:0000::/48 | WAN Prefix CS Link R6 Region2 | | No | Recommended |
| | 2406:6400:001C:0000::/64 | R6[::1]-CBR2[::2] | Yes | No | No |
| | 2406:6400:001C:0001::/64 | | Yes | No | No |
| | 2406:6400:001C:0002::/64 | | Yes | No | No |
| | 2406:6400:001C:0003::/64 | | Yes | No | No |
| | 2406:6400:001C:0004::/64 | | Yes | No | No |
| | 2406:6400:001C:0005::/64 | | Yes | No | No |
| | 2406:6400:001C:0006::/64 | | Yes | No | No |
| | 2406:6400:001C:0007::/64 | | Yes | No | No |
| | 2406:6400:001C:0008::/64 | | Yes | No | No |
| | 2406:6400:001C:0009::/64 | | Yes | No | No |
| | 2406:6400:001C:000A::/64 | | Yes | No | No |
| | 2406:6400:001C:000B::/64 | | Yes | No | No |
| | 2406:6400:001C:000C::/64 | | Yes | No | No |
| | 2406:6400:001C:000D::/64 | | Yes | No | No |
| | 2406:6400:001C:000E::/64 | | Yes | No | No |
| | 2406:6400:001C:000F::/64 | | Yes | No | No |

Training ISP IPV6 Addressing Plan

| Table 13: Detail CS link WAN Region3 | | | | | |
|--------------------------------------|--------------------------|-------------------------------|------------|-----|--------------|
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 37 | 2406:6400:0020:0000::/48 | WAN Prefix CS Link R7 Region3 | | No | Recommended |
| | 2406:6400:0020:0000::/64 | R7[::1]-CAR3[::2] | Yes | No | No |
| | 2406:6400:0020:0001::/64 | | Yes | No | No |
| | 2406:6400:0020:0002::/64 | | Yes | No | No |
| | 2406:6400:0020:0003::/64 | | Yes | No | No |
| | 2406:6400:0020:0004::/64 | | Yes | No | No |
| | 2406:6400:0020:0005::/64 | | Yes | No | No |
| | 2406:6400:0020:0006::/64 | | Yes | No | No |
| | 2406:6400:0020:0007::/64 | | Yes | No | No |
| | 2406:6400:0020:0008::/64 | | Yes | No | No |
| | 2406:6400:0020:0009::/64 | | Yes | No | No |
| | 2406:6400:0020:000A::/64 | | Yes | No | No |
| | 2406:6400:0020:000B::/64 | | Yes | No | No |
| | 2406:6400:0020:000C::/64 | | Yes | No | No |
| | 2406:6400:0020:000D::/64 | | Yes | No | No |
| | 2406:6400:0020:000E::/64 | | Yes | No | No |
| | 2406:6400:0020:000F::/64 | | Yes | No | No |
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 38 | 2406:6400:0024:0000::/48 | WAN Prefix CS Link R9 Region3 | | No | Recommended |
| | 2406:6400:0024:0000::/64 | R9[::1]-CBR3[::2] | Yes | No | No |
| | 2406:6400:0024:0001::/64 | | Yes | No | No |
| | 2406:6400:0024:0002::/64 | | Yes | No | No |
| | 2406:6400:0024:0003::/64 | | Yes | No | No |
| | 2406:6400:0024:0004::/64 | | Yes | No | No |
| | 2406:6400:0024:0005::/64 | | Yes | No | No |
| | 2406:6400:0024:0006::/64 | | Yes | No | No |
| | 2406:6400:0024:0007::/64 | | Yes | No | No |
| | 2406:6400:0024:0008::/64 | | Yes | No | No |
| | 2406:6400:0024:0009::/64 | | Yes | No | No |
| | 2406:6400:0024:000A::/64 | | Yes | No | No |
| | 2406:6400:0024:000B::/64 | | Yes | No | No |
| | 2406:6400:0024:000C::/64 | | Yes | No | No |
| | 2406:6400:0024:000D::/64 | | Yes | No | No |
| | 2406:6400:0024:000E::/64 | | Yes | No | No |
| | 2406:6400:0024:000F::/64 | | Yes | No | No |

Training ISP IPV6 Addressing Plan

Table 14: Detail CS link WAN Region 4

| Block# | Prefix | Description | PTR Record | SOR | Registration |
|--------|--------------------------|--------------------------------|------------|-----|--------------|
| 42 | 2406:6400:0028:0000::/48 | WAN Prefix CS Link R10 Region4 | | No | Recommended |
| | 2406:6400:0028:0000::/64 | R10[::1]-CAR4[::2] | Yes | No | No |
| | 2406:6400:0028:0001::/64 | | Yes | No | No |
| | 2406:6400:0028:0002::/64 | | Yes | No | No |
| | 2406:6400:0028:0003::/64 | | Yes | No | No |
| | 2406:6400:0028:0004::/64 | | Yes | No | No |
| | 2406:6400:0028:0005::/64 | | Yes | No | No |
| | 2406:6400:0028:0006::/64 | | Yes | No | No |
| | 2406:6400:0028:0007::/64 | | Yes | No | No |
| | 2406:6400:0028:0008::/64 | | Yes | No | No |
| | 2406:6400:0028:0009::/64 | | Yes | No | No |
| | 2406:6400:0028:000A::/64 | | Yes | No | No |
| | 2406:6400:0028:000B::/64 | | Yes | No | No |
| | 2406:6400:0028:000C::/64 | | Yes | No | No |
| | 2406:6400:0028:000D::/64 | | Yes | No | No |
| | 2406:6400:0028:000E::/64 | | Yes | No | No |
| | 2406:6400:0028:000F::/64 | | Yes | No | No |
| Block# | Prefix | Description | PTR Record | SOR | Registration |
| 43 | 2406:6400:002C:0000::/48 | WAN Prefix CS Link R12 Region4 | | No | Recommended |
| | 2406:6400:002C:0000::/64 | R12[::1]-CBR4[::2] | Yes | No | No |
| | 2406:6400:002C:0001::/64 | | Yes | No | No |
| | 2406:6400:002C:0002::/64 | | Yes | No | No |
| | 2406:6400:002C:0003::/64 | | Yes | No | No |
| | 2406:6400:002C:0004::/64 | | Yes | No | No |
| | 2406:6400:002C:0005::/64 | | Yes | No | No |
| | 2406:6400:002C:0006::/64 | | Yes | No | No |
| | 2406:6400:002C:0007::/64 | | Yes | No | No |
| | 2406:6400:002C:0008::/64 | | Yes | No | No |
| | 2406:6400:002C:0009::/64 | | Yes | No | No |
| | 2406:6400:002C:000A::/64 | | Yes | No | No |
| | 2406:6400:002C:000B::/64 | | Yes | No | No |
| | 2406:6400:002C:000C::/64 | | Yes | No | No |
| | 2406:6400:002C:000D::/64 | | Yes | No | No |
| | 2406:6400:002C:000E::/64 | | Yes | No | No |
| | 2406:6400:002C:000F::/64 | | Yes | No | No |

Training ISP IPV6 Addressing Plan

| Table 15: Customer block Region 1 | | | | | |
|-----------------------------------|--------------------------|-------------------------------|-------------|------------|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| 7 | 2406:6400:8000:0000::/35 | Customer block Region 1 | | | |
| | 2406:6400:8000:0000::/40 | Customer block POP1 [R1] | | >= /48 Yes | Yes |
| | 2406:6400:8100:0000::/40 | | | | |
| | 2406:6400:8200:0000::/40 | | | | |
| | 2406:6400:8300:0000::/40 | | | | |
| | 2406:6400:8400:0000::/40 | | | | |
| | 2406:6400:8500:0000::/40 | | | | |
| | 2406:6400:8600:0000::/40 | | | | |
| | 2406:6400:8700:0000::/40 | | | | |
| | 2406:6400:8800:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:8900:0000::/40 | | | | |
| | 2406:6400:8A00:0000::/40 | | | | |
| | 2406:6400:8B00:0000::/40 | | | | |
| | 2406:6400:8C00:0000::/40 | | | | |
| | 2406:6400:8D00:0000::/40 | | | | |
| | 2406:6400:8E00:0000::/40 | | | | |
| | 2406:6400:8F00:0000::/40 | | | | |
| | 2406:6400:9000:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:9100:0000::/40 | | | | |
| | 2406:6400:9200:0000::/40 | | | | |
| | 2406:6400:9300:0000::/40 | | | | |
| | 2406:6400:9400:0000::/40 | | | | |
| | 2406:6400:9500:0000::/40 | | | | |
| | 2406:6400:9600:0000::/40 | | | | |
| | 2406:6400:9700:0000::/40 | | | | |
| | 2406:6400:9800:0000::/40 | Customer block POP2 [R3] | | >= /48 Yes | Yes |
| | 2406:6400:9900:0000::/40 | | | | |
| | 2406:6400:9A00:0000::/40 | | | | |
| | 2406:6400:9B00:0000::/40 | | | | |
| | 2406:6400:9C00:0000::/40 | | | | |
| | 2406:6400:9D00:0000::/40 | | | | |
| | 2406:6400:9E00:0000::/40 | | | | |
| | 2406:6400:9F00:0000::/40 | | | | |

Training ISP IPV6 Addressing Plan

| Table 16: Summarization oprions customer block Region 1 | | | |
|---|--------------------------|-------------------------------|----------------|
| Block# | Prefix | Description | Reverse Domain |
| | 2406:6400:8000:0000::/35 | Customer block Region 1 [R2] | |
| | 2406:6400:8000:0000::/37 | Customer block POP1 [R1] | |
| | 2406:6400:8800:0000::/37 | Customer block future use/POP | |
| | 2406:6400:9000:0000::/37 | Customer block future use/POP | |
| | 2406:6400:9800:0000::/37 | Customer block POP2 [R3] | |

Training ISP IPV6 Addressing Plan

| Table 17: Detail customer block Region 1 | | | | | |
|--|--------------------------|-------------------------------|-------------|-----|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| | 2406:6400:8000:0000::/40 | 1st Customer block POP1 [R1] | | | |
| | 2406:6400:8000:0000::/48 | 1st Customer prefix POP1 [R1] | | Yes | Yes |
| | 2406:6400:8001:0000::/48 | | | | |
| | 2406:6400:8002:0000::/48 | | | | |
| | 2406:6400:8003:0000::/48 | | | | |
| | 2406:6400:8004:0000::/48 | | | | |
| | 2406:6400:8005:0000::/48 | | | | |
| | 2406:6400:8006:0000::/48 | | | | |
| | 2406:6400:8007:0000::/48 | | | | |
| | 2406:6400:9800:0000::/40 | 1st Customer block POP2 [R3] | | | |
| | 2406:6400:9800:0000::/48 | 1st Customer prefix POP2 [R3] | | Yes | Yes |
| | 2406:6400:9801:0000::/48 | | | | |
| | 2406:6400:9802:0000::/48 | | | | |
| | 2406:6400:9803:0000::/48 | | | | |
| | 2406:6400:9804:0000::/48 | | | | |
| | 2406:6400:9805:0000::/48 | | | | |
| | 2406:6400:9806:0000::/48 | | | | |
| | 2406:6400:9807:0000::/48 | | | | |

APNIC



Training ISP IPV6 Addressing Plan

| Table 18: Customer block Region 2 | | | | | |
|-----------------------------------|--------------------------|-------------------------------|-------------|------------|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| 8 | 2406:6400:a000:0000::/35 | Customer block Region 2 | | | |
| | 2406:6400:A000:0000::/40 | Customer block POP1 [R4] | | >= /48 Yes | Yes |
| | 2406:6400:A100:0000::/40 | | | | |
| | 2406:6400:A200:0000::/40 | | | | |
| | 2406:6400:A300:0000::/40 | | | | |
| | 2406:6400:A400:0000::/40 | | | | |
| | 2406:6400:A500:0000::/40 | | | | |
| | 2406:6400:A600:0000::/40 | | | | |
| | 2406:6400:A700:0000::/40 | | | | |
| | 2406:6400:A800:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:A900:0000::/40 | | | | |
| | 2406:6400:AA00:0000::/40 | | | | |
| | 2406:6400:AB00:0000::/40 | | | | |
| | 2406:6400:AC00:0000::/40 | | | | |
| | 2406:6400:AD00:0000::/40 | | | | |
| | 2406:6400:AE00:0000::/40 | | | | |
| | 2406:6400:AF00:0000::/40 | | | | |
| | 2406:6400:B000:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:B100:0000::/40 | | | | |
| | 2406:6400:B200:0000::/40 | | | | |
| | 2406:6400:B300:0000::/40 | | | | |
| | 2406:6400:B400:0000::/40 | | | | |
| | 2406:6400:B500:0000::/40 | | | | |
| | 2406:6400:B600:0000::/40 | | | | |
| | 2406:6400:B700:0000::/40 | | | | |
| | 2406:6400:B800:0000::/40 | Customer block POP2 [R6] | | >= /48 Yes | Yes |
| | 2406:6400:B900:0000::/40 | | | | |
| | 2406:6400:BA00:0000::/40 | | | | |
| | 2406:6400:BB00:0000::/40 | | | | |
| | 2406:6400:BC00:0000::/40 | | | | |
| | 2406:6400:BD00:0000::/40 | | | | |
| | 2406:6400:BE00:0000::/40 | | | | |
| | 2406:6400:BF00:0000::/40 | | | | |


Training ISP IPV6 Addressing Plan

| Table 19: Summarization oprions customer block Region 2 | | | |
|---|--------------------------|-------------------------------|----------------|
| Block# | Prefix | Description | Reverse Domain |
| | 2406:6400:A000:0000::/35 | Customer block Region 2 [R5] | |
| | 2406:6400:A000:0000::/37 | Customer block POP1 [R4] | |
| | 2406:6400:A800:0000::/37 | Customer block future use/POP | |
| | 2406:6400:B000:0000::/37 | Customer block future use/POP | |
| | 2406:6400:B800:0000::/37 | Customer block POP2 [R6] | |

Training ISP IPV6 Addressing Plan

| Table 20: Detail customer block Region 2 | | | | | |
|--|--------------------------|-------------------------------|-------------|-----|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| | 2406:6400:A000:0000::/40 | 1st Customer block POP1 [R4] | | | |
| | 2406:6400:A000:0000::/48 | 1st Customer prefix POP1 [R4] | | Yes | Yes |
| | 2406:6400:A001:0000::/48 | | | | |
| | 2406:6400:A002:0000::/48 | | | | |
| | 2406:6400:A003:0000::/48 | | | | |
| | 2406:6400:A004:0000::/48 | | | | |
| | 2406:6400:A005:0000::/48 | | | | |
| | 2406:6400:A006:0000::/48 | | | | |
| | 2406:6400:A007:0000::/48 | | | | |
| | 2406:6400:B800:0000::/40 | 1st Customer block POP2 [R6] | | | |
| | 2406:6400:B800:0000::/48 | 1st Customer prefix POP2 [R6] | | Yes | Yes |
| | 2406:6400:B801:0000::/48 | | | | |
| | 2406:6400:B802:0000::/48 | | | | |
| | 2406:6400:B803:0000::/48 | | | | |
| | 2406:6400:B804:0000::/48 | | | | |
| | 2406:6400:B805:0000::/48 | | | | |
| | 2406:6400:B806:0000::/48 | | | | |
| | 2406:6400:B807:0000::/48 | | | | |

Training ISP IPV6 Addressing Plan



Training ISP IPV6 Addressing Plan

| Table 21: Customer block Region 3 | | | | | |
|-----------------------------------|--------------------------|-------------------------------|-------------|------------|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| 9 | 2406:6400:c000:0000::/35 | Customer block Region 3 | | | |
| | 2406:6400:C000:0000::/40 | Customer block POP1 [R7] | | >= /48 Yes | Yes |
| | 2406:6400:C100:0000::/40 | | | | |
| | 2406:6400:C200:0000::/40 | | | | |
| | 2406:6400:C300:0000::/40 | | | | |
| | 2406:6400:C400:0000::/40 | | | | |
| | 2406:6400:C500:0000::/40 | | | | |
| | 2406:6400:C600:0000::/40 | | | | |
| | 2406:6400:C700:0000::/40 | | | | |
| | 2406:6400:C800:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:C900:0000::/40 | | | | |
| | 2406:6400:CA00:0000::/40 | | | | |
| | 2406:6400:CB00:0000::/40 | | | | |
| | 2406:6400:CC00:0000::/40 | | | | |
| | 2406:6400:CD00:0000::/40 | | | | |
| | 2406:6400:CE00:0000::/40 | | | | |
| | 2406:6400:CF00:0000::/40 | | | | |
| | 2406:6400:D000:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:D100:0000::/40 | | | | |
| | 2406:6400:D200:0000::/40 | | | | |
| | 2406:6400:D300:0000::/40 | | | | |
| | 2406:6400:D400:0000::/40 | | | | |
| | 2406:6400:D500:0000::/40 | | | | |
| | 2406:6400:D600:0000::/40 | | | | |
| | 2406:6400:D700:0000::/40 | | | | |
| | 2406:6400:D800:0000::/40 | Customer block POP2 [R9] | | >= /48 Yes | Yes |
| | 2406:6400:D900:0000::/40 | | | | |
| | 2406:6400:DA00:0000::/40 | | | | |
| | 2406:6400:DB00:0000::/40 | | | | |
| | 2406:6400:DC00:0000::/40 | | | | |
| | 2406:6400:DD00:0000::/40 | | | | |
| | 2406:6400:DE00:0000::/40 | | | | |
| | 2406:6400:DF00:0000::/40 | | | | |

Training ISP IPV6 Addressing Plan

| Table 22: Summarization oprions customer block Region 3 | | | |
|---|--------------------------|-------------------------------|----------------|
| Block# | Prefix | Description | Reverse Domain |
| | 2406:6400:c000:0000::/35 | Customer block Region 3 [R8] | |
| | 2406:6400:C000:0000::/37 | Customer block POP1 [R7] | |
| | 2406:6400:C800:0000::/37 | Customer block future use/POP | |
| | 2406:6400:D000:0000::/37 | Customer block future use/POP | |
| | 2406:6400:D800:0000::/37 | Customer block POP2 [R9] | |

Training ISP IPV6 Addressing Plan

| Table 23: Detail customer block Region 3 | | | | | |
|--|--------------------------|-------------------------------|-------------|-----|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| | 2406:6400:C000:0000::/40 | 1st Customer block POP1 [R7] | | | |
| | 2406:6400:C000:0000::/48 | 1st Customer prefix POP1 [R7] | | Yes | Yes |
| | 2406:6400:C001:0000::/48 | | | | |
| | 2406:6400:C002:0000::/48 | | | | |
| | 2406:6400:C003:0000::/48 | | | | |
| | 2406:6400:C004:0000::/48 | | | | |
| | 2406:6400:C005:0000::/48 | | | | |
| | 2406:6400:C006:0000::/48 | | | | |
| | 2406:6400:C007:0000::/48 | | | | |
| | 2406:6400:D800:0000::/40 | 1st Customer block POP2 [R9] | | | |
| | 2406:6400:D800:0000::/48 | 1st Customer prefix POP2 [R9] | | Yes | Yes |
| | 2406:6400:D801:0000::/48 | | | | |
| | 2406:6400:D802:0000::/48 | | | | |
| | 2406:6400:D803:0000::/48 | | | | |
| | 2406:6400:D804:0000::/48 | | | | |
| | 2406:6400:D805:0000::/48 | | | | |
| | 2406:6400:D806:0000::/48 | | | | |
| | 2406:6400:D807:0000::/48 | | | | |

APNIC



Training ISP IPV6 Addressing Plan

| Table 24: Customer block Region 4 | | | | | |
|-----------------------------------|--------------------------|-------------------------------|-------------|------------|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| 10 | 2406:6400:e000:0000::/35 | Customer block Region 4 | | | |
| | 2406:6400:E000:0000::/40 | Customer block POP1 [R10] | | >= /48 Yes | Yes |
| | 2406:6400:E100:0000::/40 | | | | |
| | 2406:6400:E200:0000::/40 | | | | |
| | 2406:6400:E300:0000::/40 | | | | |
| | 2406:6400:E400:0000::/40 | | | | |
| | 2406:6400:E500:0000::/40 | | | | |
| | 2406:6400:E600:0000::/40 | | | | |
| | 2406:6400:E700:0000::/40 | | | | |
| | 2406:6400:E800:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:E900:0000::/40 | | | | |
| | 2406:6400:EA00:0000::/40 | | | | |
| | 2406:6400:EB00:0000::/40 | | | | |
| | 2406:6400:EC00:0000::/40 | | | | |
| | 2406:6400:ED00:0000::/40 | | | | |
| | 2406:6400:EE00:0000::/40 | | | | |
| | 2406:6400:EF00:0000::/40 | | | | |
| | 2406:6400:F000:0000::/40 | Customer block future use/POP | | >= /48 Yes | Yes |
| | 2406:6400:F100:0000::/40 | | | | |
| | 2406:6400:F200:0000::/40 | | | | |
| | 2406:6400:F300:0000::/40 | | | | |
| | 2406:6400:F400:0000::/40 | | | | |
| | 2406:6400:F500:0000::/40 | | | | |
| | 2406:6400:F600:0000::/40 | | | | |
| | 2406:6400:F700:0000::/40 | | | | |
| | 2406:6400:F800:0000::/40 | Customer block POP2 [R12] | | >= /48 Yes | Yes |
| | 2406:6400:F900:0000::/40 | | | | |
| | 2406:6400:FA00:0000::/40 | | | | |
| | 2406:6400:FB00:0000::/40 | | | | |
| | 2406:6400:FC00:0000::/40 | | | | |
| | 2406:6400:FD00:0000::/40 | | | | |
| | 2406:6400:FE00:0000::/40 | | | | |
| | 2406:6400:FF00:0000::/40 | | | | |

Training ISP IPV6 Addressing Plan

| Table 25: Summarization oprions customer block Region 4 | | | |
|---|--------------------------|-------------------------------|----------------|
| Block# | Prefix | Description | Reverse Domain |
| | 2406:6400:e000:0000::/35 | Customer block Region 4 [R11] | |
| | 2406:6400:E000:0000::/37 | Customer block POP1 [R10] | |
| | 2406:6400:E800:0000::/37 | Customer block future use/POP | |
| | 2406:6400:F000:0000::/37 | Customer block future use/POP | |
| | 2406:6400:F800:0000::/37 | Customer block POP2 [R12] | |

Training ISP IPV6 Addressing Plan

| Table 26: Detail customer block Region 4 | | | | | |
|--|--------------------------|--------------------------------|-------------|-----|--------------|
| Block# | Prefix | Description | Reverse DNS | SOR | Registration |
| | 2406:6400:E000:0000::/40 | 1st Customer block POP1 [R10] | | | |
| | 2406:6400:E000:0000::/48 | 1st Customer prefix POP1 [R10] | | Yes | Yes |
| | 2406:6400:E001:0000::/48 | | | | |
| | 2406:6400:E002:0000::/48 | | | | |
| | 2406:6400:E003:0000::/48 | | | | |
| | 2406:6400:E004:0000::/48 | | | | |
| | 2406:6400:E005:0000::/48 | | | | |
| | 2406:6400:E006:0000::/48 | | | | |
| | 2406:6400:E007:0000::/48 | | | | |
| | 2406:6400:F800:0000::/40 | 1st Customer block POP2 [R10] | | | |
| | 2406:6400:F800:0000::/48 | 1st Customer prefix POP2 [R10] | | Yes | Yes |
| | 2406:6400:F801:0000::/48 | | | | |
| | 2406:6400:F802:0000::/48 | | | | |
| | 2406:6400:F803:0000::/48 | | | | |
| | 2406:6400:F804:0000::/48 | | | | |
| | 2406:6400:F805:0000::/48 | | | | |
| | 2406:6400:F806:0000::/48 | | | | |
| | 2406:6400:F807:0000::/48 | | | | |

APNIC



Training ISP IPV4 Addressing Plan

Summary parent block IPV4

| Block# | Prefix | Size | Description |
|--------|-------------|------|------------------|
| 1 | 172.16.0.0 | /19 | Parent block |
| 2 | 172.16.0.0 | /20 | Infrastructure |
| 3 | 172.16.16.0 | /20 | Customer network |

Training ISP IPV4 Addressing Plan

Detail DC infrastructure block IPV4

| Block# | Prefix | Size | Description | SOR | Register |
|--------|------------|------|-------------------------|-----|-------------|
| 2 | 172.16.0.0 | /20 | Infrastructure | | |
| | | | | | |
| 4 | 172.16.0.0 | /23 | Router2 DC summary net | | |
| 5 | 172.16.0.0 | /24 | Router2 DC | No | Recommended |
| | | | | | |
| 6 | 172.16.2.0 | /23 | Router5 DC summary net | | |
| 7 | 172.16.2.0 | /24 | Router5 DC | No | Recommended |
| | | | | | |
| 8 | 172.16.4.0 | /23 | Router8 DC summary net | | |
| 9 | 172.16.4.0 | /24 | Router8 DC | No | Recommended |
| | | | | | |
| 10 | 172.16.6.0 | /23 | Router11 DC summary net | | |
| 11 | 172.16.6.0 | /24 | Router11 DC | No | Recommended |

Training ISP IPV4 Addressing Plan

Detail infrastructure WAN block IPV4

| | | | | | |
|----|--------------|-----|-----------------|----|----------|
| 12 | 172.16.10.0 | /24 | WAN prefix | | Optional |
| 13 | 172.16.10.0 | /30 | Router2-1 WAN | No | |
| 14 | 172.16.10.4 | /30 | Router2-3 WAN | No | |
| 15 | 172.16.10.8 | /30 | Router1-3 WAN | No | |
| | | | | | |
| 16 | 172.16.10.24 | /30 | Router5-4 WAN | No | |
| 17 | 172.16.10.28 | /30 | Router5-6 WAN | No | |
| 18 | 172.16.10.32 | /30 | Router4-6 WAN | No | |
| | | | | | |
| 19 | 172.16.10.48 | /30 | Router8-7 WAN | No | |
| 20 | 172.16.10.52 | /30 | Router8-9 WAN | No | |
| 21 | 172.16.10.56 | /30 | Router7-9 WAN | No | |
| | | | | | |
| 22 | 172.16.10.72 | /30 | Router11-10 WAN | No | |
| 23 | 172.16.10.76 | /30 | Router11-12 WAN | No | |
| 24 | 172.16.10.80 | /30 | Router10-12 WAN | No | |

Training ISP IPV4 Addressing Plan

Detail customer link WAN block

| Block# | Prefix | Size | Description | SOR | Register |
|--------|---------------|------|------------------------|-----|----------|
| | 172.16.11.0 | /26 | WAN CS Link Region1 | | |
| | 172.16.11.0 | /27 | WAN CS Link POP1 [R1] | | |
| | 172.16.11.0 | /30 | R1[::1]-CAR1[::2] | No | No |
| | 172.16.11.4 | /30 | | | |
| | 172.16.11.32 | /27 | WAN CS Link POP2 [R3] | | |
| | 172.16.11.32 | /30 | R3[::33]-CBR1[::34] | No | No |
| | 172.16.11.36 | /30 | | | |
| | | | | | |
| | 172.16.11.64 | /26 | WAN CS Link Region2 | | |
| | 172.16.11.64 | /27 | WAN CS Link POP1 [R4] | | |
| | 172.16.11.64 | /30 | R4[::65]-CAR2[::66] | No | No |
| | 172.16.11.68 | /30 | | | |
| | 172.16.11.96 | /27 | WAN CS Link POP2 [R6] | | |
| | 172.16.11.96 | /30 | R6[::97]-CBR2[::98] | No | No |
| | 172.16.11.100 | /30 | | | |
| | | | | | |
| | 172.16.11.128 | /26 | WAN CS Link Region3 | | |
| | 172.16.11.128 | /27 | WAN CS Link POP1 [R7] | | |
| | 172.16.11.128 | /30 | R7[::129]-CAR3[::130] | No | No |
| | 172.16.11.132 | /30 | | | |
| | 172.16.11.160 | /27 | WAN CS Link POP2 [R9] | | |
| | 172.16.11.160 | /30 | R9[::161]-CBR3[::162] | No | No |
| | 172.16.11.164 | /30 | | | |
| | | | | | |
| | 172.16.11.192 | /26 | WAN CS Link Region4 | | |
| | 172.16.11.192 | /27 | WAN CS Link POP1 [R10] | | |
| | 172.16.11.192 | /30 | R10[::193]-CAR4[::194] | No | No |
| | 172.16.11.196 | /30 | | | |
| | 172.16.11.224 | /27 | WAN CS Link POP2 [R12] | | |
| | 172.16.11.224 | /30 | R12[::225]-CBR4[::226] | No | No |
| | 172.16.11.228 | /30 | | | |

Training ISP IPV4 Addressing Plan

Detail infrastructure block Transport & Loopback IPV4

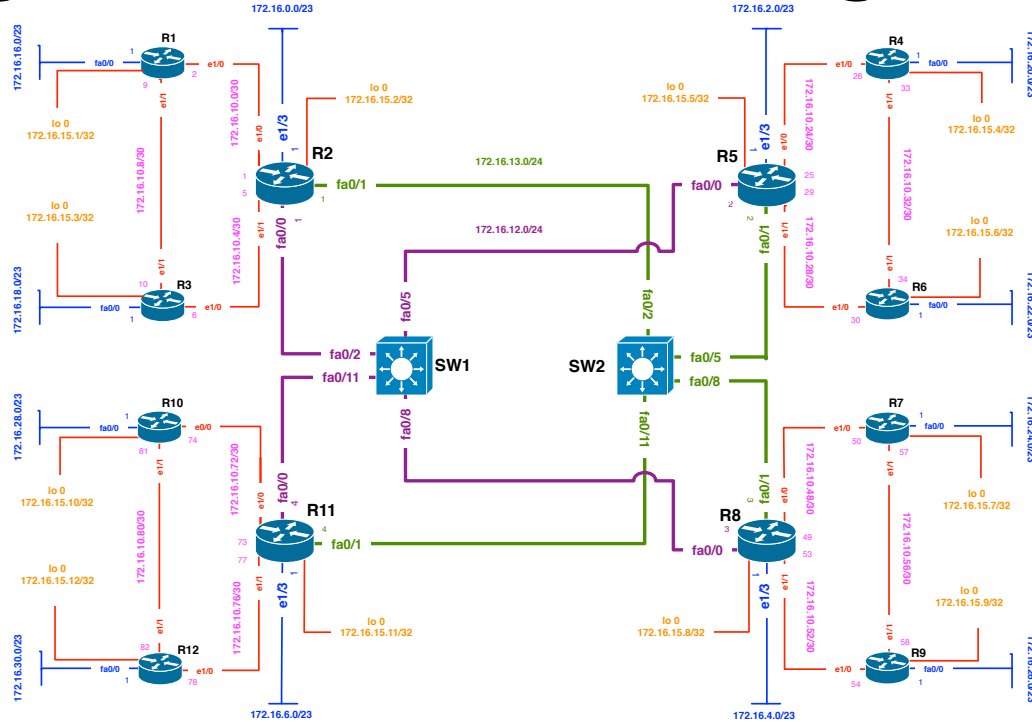
| | | | | | |
|----|-------------|-----|-----------------------|----|--|
| 25 | 172.16.12.0 | /24 | Transport link PURPLE | No | |
| 26 | 172.16.13.0 | /24 | Transport link GREEN | No | |
| | | | | | |
| 27 | 172.16.15.0 | /24 | Loopback | No | |

Training ISP IPV4 Addressing Plan

Detail customer block

| Block# | Prefix | Size | Description | SOR | Register |
|--------|-------------|------|----------------------|-----|----------|
| 28 | 172.16.6.0 | /20 | Customer network | | |
| | | | | | |
| 29 | 172.16.16.0 | /22 | Router2 summary net | | |
| 30 | 172.16.16.0 | /23 | Router1 CS network | Yes | Must |
| 31 | 172.16.18.0 | /23 | Router3 CS network | Yes | Must |
| | | | | | |
| 32 | 172.16.20.0 | /22 | Router5 summary net | | |
| 33 | 172.16.20.0 | /23 | Router4 CS network | Yes | Must |
| 34 | 172.16.22.0 | /23 | Router6 CS network | Yes | Must |
| | | | | | |
| 35 | 172.16.24.0 | /22 | Router8 summary net | | |
| 36 | 172.16.24.0 | /23 | Router7 CS network | Yes | Must |
| 37 | 172.16.26.0 | /23 | Router9 CS network | Yes | Must |
| | | | | | |
| 38 | 172.16.28.0 | /22 | Router11 summary net | | |
| 39 | 172.16.28.0 | /23 | Router10 CS network | Yes | Must |
| 40 | 172.16.30.0 | /23 | Router12 CS network | Yes | Must |

Training ISP IPV4 Addressing Plan



Training ISP IPv4 Address Plan

Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise: Basic Routing, Interface & OSPF**
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise: BGP Peering**
- IXP Design Considerations
- **Hands On Lab Exercise: IXP Configuration**
- Route Collectors & Servers
- IXP BCP and What can go wrong?

OSPF

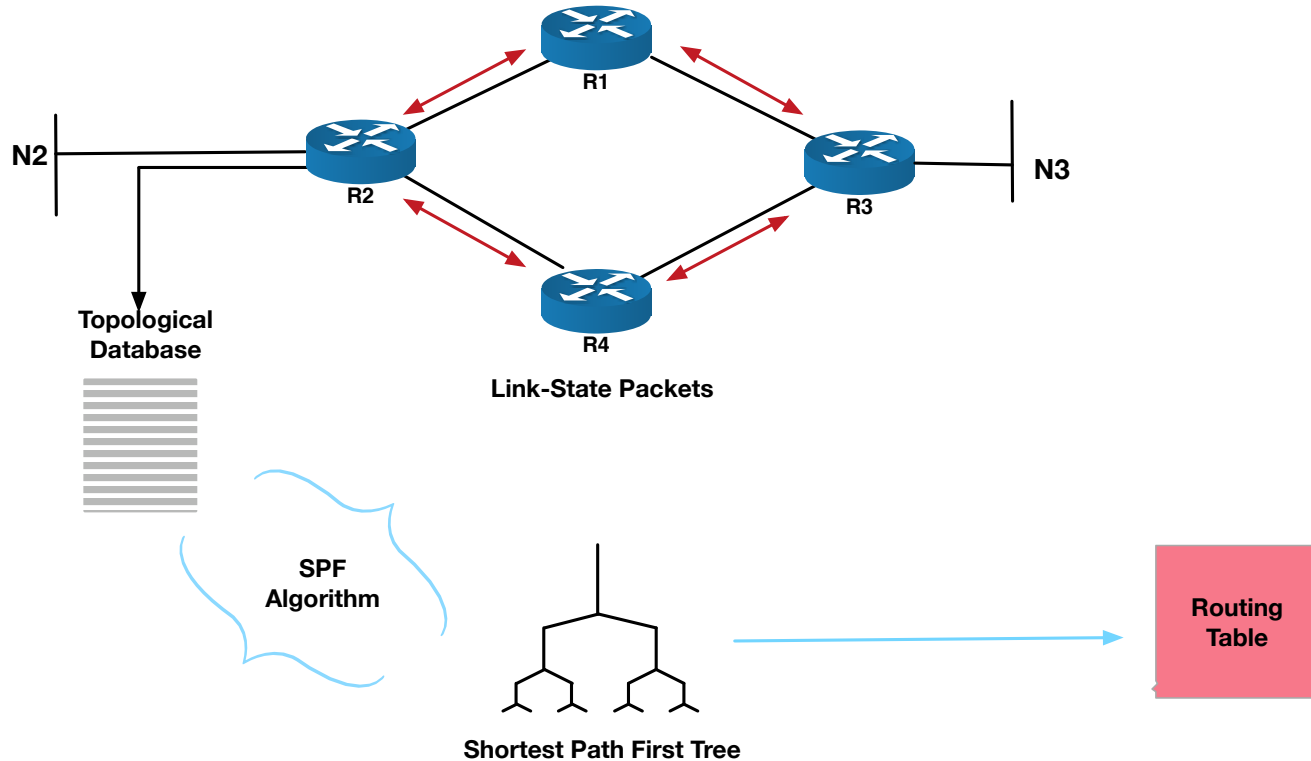
- **O**pen **S**hortest **P**ath **F**irst
- Link State Protocol or SPF technology
- Developed by OSPF working group of IETF (RFC 1247)
- Comes with two version
 - OSPFv2 (IPv4) standard described in (RFC 2328)
 - OSPFv3 (IPv6) standard described in (RFC 2740)

OSPF

- Designed for
 - TCP/IP environment
 - Fast convergence
 - Route redistribution
 - Variable length subnet masks (VLSM)
 - Dis-contiguous subnets
 - Incremental updates
 - Route authentication
- **OSPF runs on IP, Protocol 89**

<http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml>

Link State Routing Protocol



Link State Routing Protocol Functions



- All routers have the same view
- Do not send full routing table on periodic interval
- Use Shortest Path First (SPF) algorithm to select best path from topology table
- Send very small periodic (Hello) message to maintain link condition
- Send triggered update instantly when network change occur

Link State Data Structure

- **Neighbor Table**
 - List of all recognized neighboring router to whom routing information will be interchanged
- **Topology Table**
 - Also called LSDB which maintain list of routers and their link information i.e network destination, prefix length, link cost etc
- **Routing table**
 - Also called forwarding table contain only the best path to forward data traffic

Link State Routing Protocol Advantages

- **Low bandwidth utilization**
 - Only changes propagated
 - Uses multicast on multi-access broadcast network
- **Fast Convergence**
 - Detection Plus LSA/SPF (Known as the Dijkstra Algorithm)
- **Finding a new route**
 - LSA flooded throughout area
 - Acknowledgement based
 - Topology database synchronized
 - Each router derives routing table to destination network



Overview

IXP Workshop


- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- **BGP Routing Protocol Operation- Make the IXP Works**
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

What is Border Gateway Protocol?

- BGP:
 - A path vector routing protocol to exchange routing information between different Autonomous System (AS)
 - ASes are the building block of BGP operational unites
 - AS is a collection of routers with a common routing policy
 - Specification is defined in RFC4271

What is an Autonomous System (AS)

- An AS is a collection of networks with same routing policy
- Usually under a single administrative control unit
- A public AS is identified by a unique number called AS number
- Around 32000 ASes are visible on the Internet now



BGP features

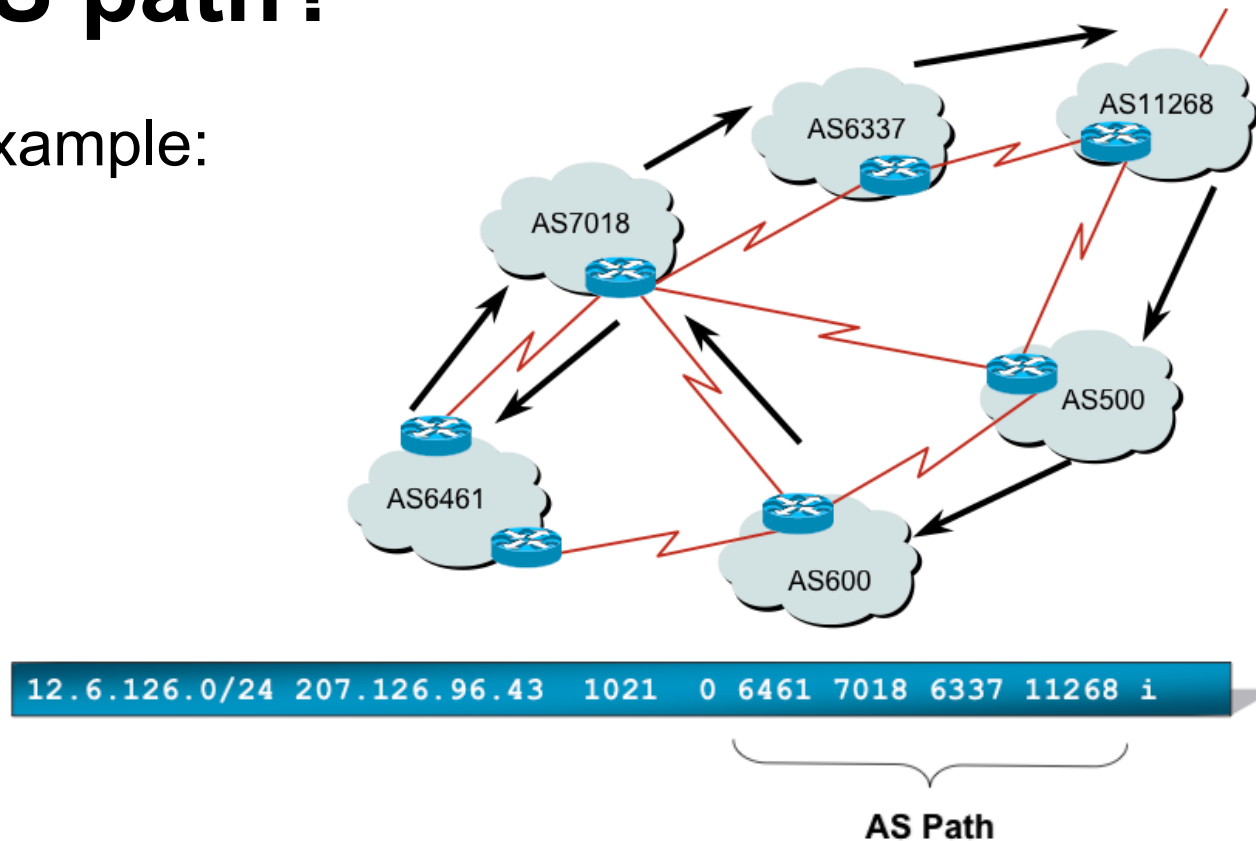
- Path Vector Routing Protocol
- Send incremental updates to peers
- Runs over TCP –Port 179
- Select path based on routing policy/ organization's business requirement
- Support Classless Inter Domain Routing (CIDR) concept
- Widely used in today's Internet Backbone
- Current BGP version is MP-BGP

What is Path Vector Routing Protocol

- A path vector routing protocol is used to span different autonomous systems
- It defines a route as a collection of a number of AS that it passes through from source AS to destination AS
- This list of ASes are called AS path and used to avoid routing loop
- AS path is also used to select path to destination

What is AS path?

- An AS path example:



BGP Traffic Arrangement Definition

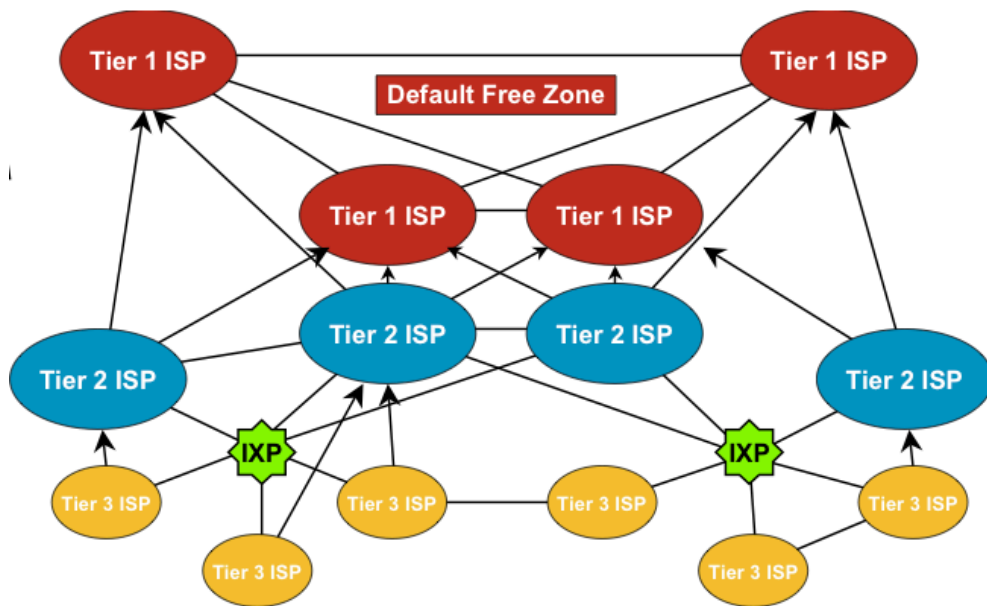
- Transit
 - Forwarding traffic through the network usually for a fee
 - I.e Internet service from upstream ISP
- Peering
 - Exchanging traffic without any fee
 - I.e Connection in an IXP
- Default
 - Where to send traffic if there no explicit route match in the routing table

What is Default Free Zone?

- Default free zone is made up of Tier One ISP routers which have explicit routing information about every part of the Global Internet
- So there is no need of default route
- If there is no destination network match, then that prefix is still not announced/ used by any ISP yet

ISP Hierarchical Connection

- Connectivity Diagram:



BGP General Operation

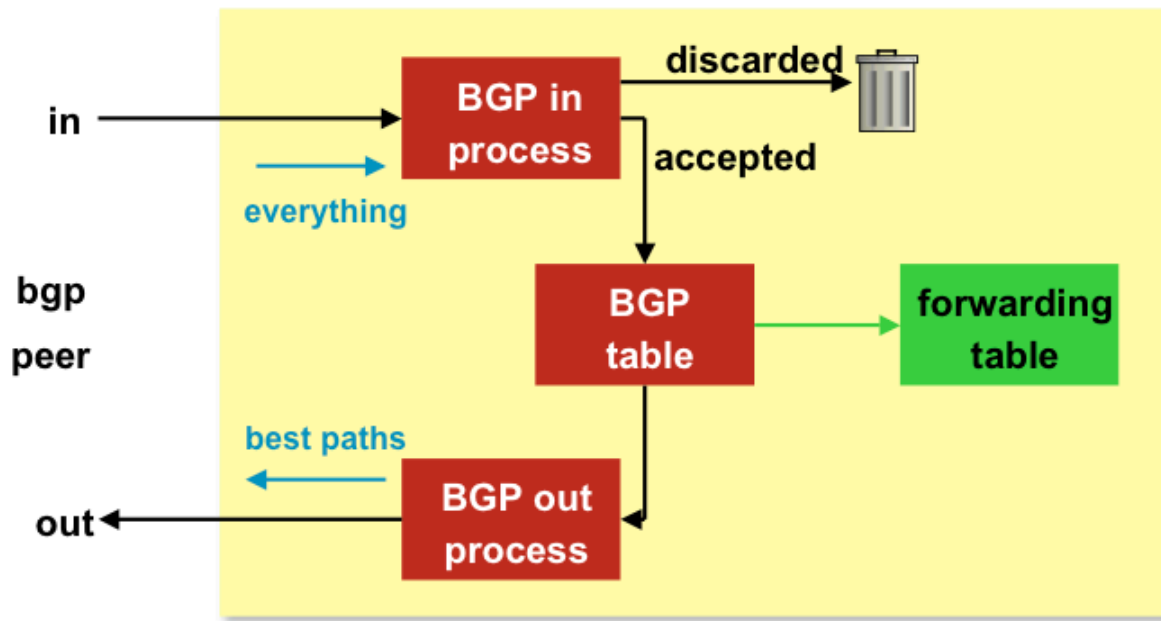
- BGP maintain 3 database i.e Neighbor Table, BGP Table and Forwarding Table
- Learns multiple paths via internal and external BGP speakers
- Picks the best path and installs them on the forwarding tables
- Best path is sent to external BGP neighbors
- Policies are applied by influencing the best path selection

Constructing the Forwarding Table

- BGP “In” process
 - Receives path information from peers
 - Results of BGP path selection placed in the BGP table “best path” flagged
- BGP “Out” process
 - Announce “best path” information to peers
- Best path installed in forwarding table if:
 - Prefix and prefix length are equal
 - Lowest protocol distance

Constructing the Forwarding Table

Flowchart:



BGP Terminology

- Neighbor
 - Any two routers that have formed a TCP connection to exchange BGP routing information are called peers or neighbors
- iBGP
 - iBGP refers to the BGP neighbor relationship within the same AS.
 - ☐ The neighbors do not have to be directly connected.
- eBGP
 - When BGP neighbor relationship are formed between two peers belongs to different AS are called eBGP.
 - ☐ EBGP neighbors by default need to be directly connected.

Building Neighbor Relationship

- After adding BGP neighbor:
 - Both router establish a TCP connection and send open message
 - If open message is accepted then both send keepalive message to each other to confirm open message
 - After both confirm open message by sending keepalive message they establish BGP neighbor relationship and exchange routing information

BGP message type

- Open Message
 - To establish BGP neighbor relationship
- Keepalive message
 - Only contain message header to maintain neighbor relationship. Sent every periodic interval
- Update message
 - Contain path information. One update message contain one path information. Multiple path need multiple update message to be sent
- Notification message
 - Sent when an error condition occur and BGP connection closed immediately

BGP Open message

- Open message contain:
 - BGP Version number
 - AS number of the local router
 - BGP holdtime in second to elapse between the successive keepalive message
 - BGP router ID which is a 32 bit number. Usually an IPv4 address is used as router ID
 - Optional parameters i.e types, length and value encoded. An example optional parameter is session authentication info

BGP Keepalive Message

- Send between BGP peers after every periodic interval (60 Sec)
- It refresh hold timer from expiration (180sec)
- A keepalive message contain only the message header

BGP Update Message

- An update message contain:
 - Withdrawn routes: a list contain address prefix that are withdrawn from service
 - Path attributes: includes AS path, origin code, local pref etc
 - Network-layer reachability information: includes a list of address prefix reachable by this path

BGP Notification message

- Only sent when an error condition occur and detected in a network and BGP connection is closed immediately
- Notification message contain an error code, an error subcode, and data that are related to that error

BGP Neighbor Relationship States

- BGP neighbor goes through following steps:
 - Idle: Router is searching its routing table to reach the neighbor
 - Connect: Router found route and completed TCP three-way handshake
 - Open Sent: Open message sent with the parameter for BGP session
 - Open Confirm: Router receive agreement on the parameter to establish BGP session
 - Established: Peering is established and routing information exchange began


Troubleshoot BGP Neighbor Relation

- Idle:
 - The router can not find address of the neighbor in its routing table
- Active:
 - Router found address of the neighbor in its routing table sent open message and waiting for the response from the neighbor
- Cycle between Active/Idle
 - Neighbor might peer with wrong address
 - Does not have neighbor statement on the other side
 - BGP open message source IP address does not match with remote side neighbor statement or no route to source IP address

iBGP Peering

- BGP peer within the same AS
- Not required to be directly connected
- iBGP peering require full mesh peering
 - Within an AS all iBGP speaker must peer with other iBGP speaker
 - They originate connected network
 - Pass on prefixes learned from outside AS
 - They do not forward prefixes learned from other iBGP peer

iBGP Peering with Loopback Interface




- If iBGP speakers has multiple connection then it is advisable to peer with loopback
- Connected network can go down which might loose iBGP peering
- Loopback interface will never go down

iBGP Neighbor Update Source

- This command allows the BGP process to use the IP address of a specified interface as the source IP address of all BGP updates to that neighbor
- A loopback interface is usually used as it will never goes down as long as the router is operational
- All BGP message will use the referenced interface as source of the messages


eBGP Peering




- Peering with BGP speaker in different AS
- Peers should be directly connected and share same WAN link
- eBGP neighbors are usually routed through connected network

BGP Next Hop Behavior

- BGP is an AS-by-AS routing protocol not a router-by router routing protocol.
- In BGP, the next hop does not mean the next router it means the IP address to reach the next AS
 - I.e Router A advertise 150.10.0.0/16 and 160.10.0.0/16 to router B in eBGP with next hop 150.10.1.1
 - Router B will update Router C in iBGP keeping the next hop unchanged



iBGP Next Hop



- Next hop is iBGP router loopback address
- Recursive route look-up
- Loopback address need to announce through IGP (OSPF)

BGP Synchronous Rule

- BGP do not use or advertise any route to an external neighbor learned by iBGP until a matching route has been learned from an IGP i.e OSPF or static
- It ensure consistency of information throughout the AS
- Avoid black hole route within an AS
- It is safe to turn off if all routers with in the AS run full-mesh iBGP
- Advisable to disable this feature (BCP)



Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- **BGP Attributes and Path Selection Process- Send Traffic Through IXP**
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

BGP Attributes

BGP metrics are called path attributes. Here is the classifications BGP attributes:

Well-known Mandatory

- AS-Path
- Next-hop
- Origin

Well-known Discretionary

- Local preference
- Atomic aggregate

Optional Transitive

- Community
- Aggregator

Optional Non-Transitive

- Multi-exit-discriminator (MED)

Well-Known Attributes

- Must be recognized by all compliant BGP implementations
- Are propagated to other neighbors

Well-Known Mandatory Attributes

- Must be present in all update messages
 - ***AS Path***
 - ***Next-hop***
 - ***Origin***

Well-Known Discretionary Attributes

- May be present in update messages
 - ***Local preference***
 - ***Atomic aggregate***

Optional Attributes

- Recognized by some implementations (could be private) expected not to be recognized by everyone
- Recognized optional attributes are propagated to other neighbors based on their meaning

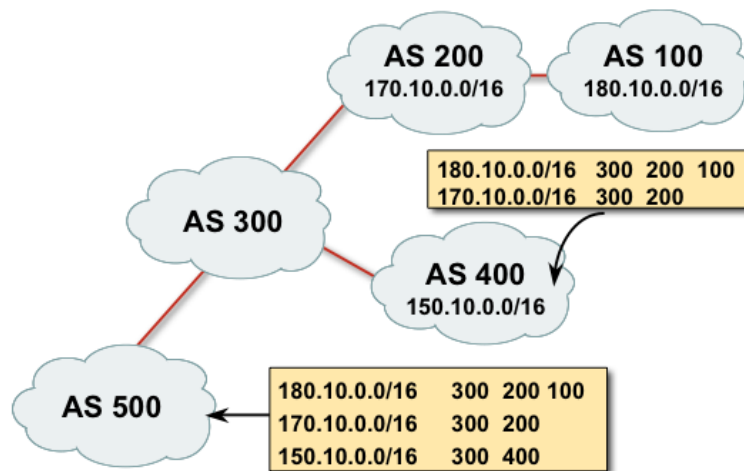
Optional Transitive Attributes

- If not recognized, are marked as partial and propagated to other neighbors
 - **Community**
 - **Aggregator**

Optional Non Transitive attributes

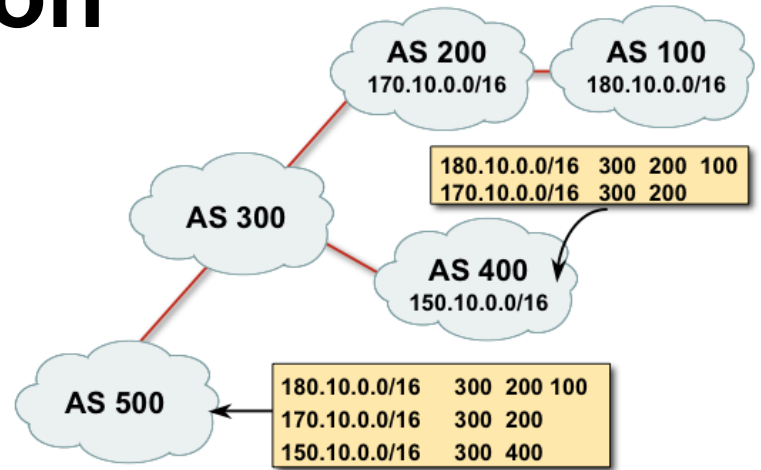
- Discarded if not recognized
 - **Multi Exit Discriminator (MED)**

AS Path Attribute




- Sequence of ASes a route has traversed
- Used for
 - Loop detection
 - Path metrics where the length of the AS Path is used as in path selection

AS Path Loop Detection



- 180.10.0.0/16 is not accepted by AS100 as the prefix has AS100 in its AS-PATH
- This is loop detection in action

AS Path Attribute (2 byte and 4 byte)



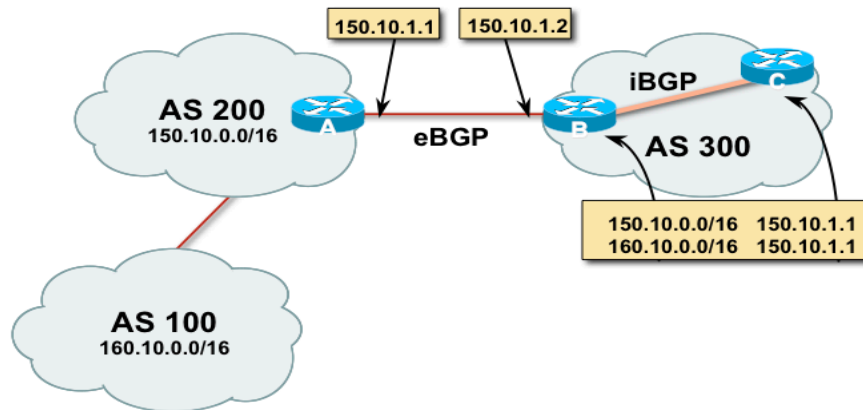
- Internet with 16-bit and 32-bit ASNs
 - 32-bit ASNs are 65536 and above
 - AS-PATH length maintained

AS Path and AS4 Path Example

Router5:


| Network | Next Hop | Metric | LocPrf | Weight | Path |
|------------------|-------------------|-----------------------------------|------------------------|-------------------------|------|
| *> 2001::/32 | 2406:6400:F:41::1 | | | | |
| | | | | 0 23456 38610 6939 | I |
| * i | 2406:6400:D::5 | 0 100 | 0 45192 4608 4826 6939 | i | |
| *> 2001:200::/32 | 2406:6400:F:41::1 | | | | |
| | | | | 0 23456 38610 6939 2500 | i |
| * i | 2406:6400:D::5 | 0 100 0 45192 4608 4826 6939 2500 | i | | |

eBGP Next Hop



- The IP address to reach the next AS
 - Router A advertises 150.10.0.0/16 and 160.10.0.0/16 to router B in eBGP with next hop 150.10.1.1 (Change it to own IP)
 - Router B will update Router C in iBGP keeping the next hop unchanged
- Well known mandatory attribute

iBGP Next Hop



- Next hop is iBGP router loopback address
- Recursive route look-up
- Loopback address need to announce through IGP (OSPF)
- iBGP send update next-hop unchanged

Next Hop Best Practice

- IOS default is for external next-hop to be propagated unchanged to iBGP peers
 - This means that IGP has to carry external next-hops
 - Forgetting means external network is invisible
 - With many eBGP peers, it is unnecessary extra load on IGP
- ISP Best Practice is to change external next-hop to be that of the local router
 - neighbor x.x.x.x next-hop-self

Next Hop Self Configuration

- Next hop default behavior can be changed by using next-hop-self command
- Forces all updates for this neighbor to be advertised with this router as the next hop
- The IP address used for next-hop-self will be the same as the source IP address of the BGP packet

BGP Origin Attribute

- The origin attribute informs all autonomous systems how the prefix introduced into BGP
- Well known mandatory attribute
- Three values: IGP, EGP, incomplete
 - IGP generated by BGP network statement
 - EGP generated by EGP
 - Incomplete redistributed from another routing protocol

BGP Origin Attribute Example

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

r RIB-failure, S Stale


Origin codes: i - IGP, e - EGP, ? - incomplete

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|--------------|-------------------|--------|--------|--------|-----------------|
| *> 2001::/32 | 2406:6400:F:41::1 | 0 | 23456 | 38610 | 6939 i |
| * i | 2406:6400:D::5 | 0 | 100 | 0 | 45192 4608 4826 |
| 6939 i | | | | | |

BGP Local Preference Attribute

- Local preference is used to advertise to IBGP neighbors only about how to leave their AS (Outbound Traffic).
- Paths with highest preference value are most desirable
- Local preference attribute is well-known and discretionary and is passed only within the AS
- Cisco Default Local Pref is 100

BGP Local Preference Attribute



- For destination 160.10.0.0/16 Router A advertise local pref 500 and Router B advertise local pref 800 in iBGP
- 800 will win best path (Router B)


BGP Local Pref Attribute Example

```
Network      Next Hop    Metric LocPrf Weight  Path
*> 2001::/32  2406:6400:F:41::1
                                     0 23456 38610 6939
i
* i          2406:6400:D::5  0    100   0 45192 4608 4826
6939 i
*> 2001:200::/32  2406:6400:F:41::1
                                     0 23456 38610 6939 2500 i
* i          2406:6400:D::5    0 100  0 45192 4608 4826
6939 2500 i
```

BGP MED Attribute

- MED is used to advertise to EBGP neighbors about how to exit their AS to reach networks owned by this AS (Incoming traffic).
- MED is sent to EBGP neighbors only.
- The paths with the lowest MED value are the most desirable
- The MED attribute is optional and non transitive

BGP MED Attribute



- For prefix 120.68.1.0/24 Router B send MED 1000 and router A send MED 2000 to eBGP neighbor
- Incoming traffic from AS200 will choose Router B since lowest MED will win

BGP MED Example

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---------|----------|--------|--------|--------|------|
|---------|----------|--------|--------|--------|------|

```
*> 2001::/32 2406:6400:F:41::1
```

```
0 23456 38610 6939 i
```

```
* i 2406:6400:D::5 0 100 0 45192 4608 4826 6939  
i
```

```
*> 2001:200::/32 2406:6400:F:41::1
```

```
0 23456 38610 6939 2500 i
```

```
* i 2406:6400:D::5 0 100 0 45192 4608 4826 6939  
2500 i
```

BGP Community Attribute

- Community is a tagging technique to mark a set of routes
- Upstream service provider routers can then use these flags to apply specific routing policies (i.e local preference etc) within their network
- Represented as two 16 bit integers (RFC1998)
- Common format is <local-ASN>:xx
- I.e 0:0 to 0:65535 and 65535:0 to 65535:65535 are reserved
- Very useful in applying policies within and between ASes
- Optional & transitive attribute

BGP Route Selection Process

- Step 1: Prefer highest weight (local to router)
- Step 2: Prefer highest local preference (global within AS)
- Step 3: Prefer route originated by the local router
- Step 4: Prefer shortest AS path
- Step 5: Prefer lowest origin code (IGP < EGP < incomplete)
- Step 6: Prefer lowest MED (from other AS)
- Step 7: Prefer EBGP path over IBGP path
- Step 8: Prefer the path through the closest IGP neighbor
- Step 9: Prefer oldest route for EBGP paths
- Step 10: Prefer the path with the lowest neighbor BGP router ID



Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise: BGP Peering**
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- **IXP Design Considerations**
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- IXP BCP and What can go wrong?

IX Peering Model

- BLPA (Bi-Lateral Peering Agreement)
 - IX will only provide layer two connection/switch port to ISPs
 - Every ISPs will arrange necessary peering arrangement with others by their mutual business understanding.
- MLPA (Multi-Lateral Peering Agreement)
 - IX will provide layer two connection/switch port to ISPs
 - Each ISP will peer with a route server on the IX.
 - Route server will collect and distribute directly connected routes to every peers.


IXP Operating Cost

- Access link
- Link maintenance
- Utility
- Administration

IXP Cost Model

- Not for profit
- Cost sharing
- Membership based
- Commercial IX

IXP Network Diagram



Services to Offer

- Route Collector
 - Route collector shows the reachability information available at the exchange
 - Technical detail covered later on
- Looking Glass
 - One way of making the Route Collector routes available for global view (e.g. www.traceroute.org)
 - Public or members only access

Services to Offer

- Content Redistribution/Caching
 - For example, Akamised update distribution service
- Network Time Protocol
 - Locate a stratum 1 time source (GPS receiver, atomic clock, etc) at IXP
- Routing Registry
 - Used to register the routing policy of the IXP membership (more later)



Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- **Route Collectors & Servers**
- IXP BCP and What can go wrong?

Introduction to Route Collectors

What routes are available at the IXP?


What is a Route Collector?

- Usually a router or Unix system running BGP
- Gathers routing information from service provider routers at an IXP
 - Peers with each ISP using BGP
- Does **not** forward packets
- Does **not** announce any prefixes to ISPs

Purpose of a Route Collector

- To provide a public view of the Routing Information available at the IXP
 - Useful for existing members to check functionality of BGP filters
 - Useful for prospective members to check value of joining the IXP
 - Useful for the Internet Operations community for troubleshooting purposes
 - E.g. www.traceroute.org

Route Collector at an IXP



Route Collector Requirements

- Router or Unix system running BGP
 - Minimal memory requirements – only holds IXP routes
 - Minimal packet forwarding requirements – doesn't forward any packets
- Peers eBGP with every IXP member
 - Accepts everything; Gives nothing
 - Uses a private ASN
 - Connects to IXP Transit LAN
- “Back end” connection
 - Second Ethernet globally routed
 - Connection to IXP Website for public access

Route Collector Implementation

- Most IXPs now implement some form of Route Collector
- Benefits already mentioned
- Great public relations tool
- Unsophisticated requirements
 - Just runs BGP

Introduction to Route Servers

How to scale very large IXPs


What is a Route Server?

- Has all the features of a Route Collector
- But also:
 - Announces routes to participating IXP members according to their routing policy definitions
- Implemented using the same specification as for a Route Collector

Features of a Route Server


- Helps scale routing for large IXPs
- Simplifies Routing Processes on ISP Routers
- Optional participation
 - Provided as service, is **NOT** mandatory
- Does result in insertion of RS Autonomous System Number in the Routing Path
- Optionally uses Policy registered in IRR

Diagram of N-squared Peering Mesh




- For large IXPs (dozens for participants) maintaining a larger peering mesh becomes cumbersome and often too hard

Peering Mesh with Route Servers



- ISP routers peer with the Route Servers
 - Only need to have two eBGP sessions rather than N

RS based Exchange Point Routing Flow



Advantages of Using a Route Server

- Helps scale Routing for very large IXPs
- Separation of Routing and Forwarding
- Simplify Routing Configuration Management on ISPs routers

Disadvantages of using a Route Server

- ISPs can lose direct policy control
 - If RS is only peer, ISPs have no control over who their prefixes are distributed to
- Completely dependent on 3rd party
 - Configuration, troubleshooting, etc...
- Insertion of RS ASN into routing path
 - Traffic engineering/multihoming needs more care
- These are major disadvantages
 - Usually out-weigh the advantages

Typical usage of a Route Server

- Route Servers may be provided as an **OPTIONAL** service
 - Most common at large IXPs (>50 participants)
 - Examples: LINX, TorIX, AMS-IX, etc
- ISPs peer:
 - Directly with significant peers
 - With Route Server for the rest

Things to think about...

- Would using a route server benefit you?
 - Helpful when BGP knowledge is limited (but is NOT an excuse not to learn BGP)
 - Avoids having to maintain a large number of eBGP peers
 - But can you afford to lose policy control? (An ISP not in control of their routing policy is what?)



Overview

IXP Workshop

- What is an Internet Exchange Point (IXP)?
- What is the value of Peering?
- How to build an IXP?
- How Internet works & Routing Protocol Basic
- **Hands On Lab Exercise:** Basic Routing, Interface & OSPF
- BGP Routing Protocol Operation- Make the IXP Works
- BGP Attributes and Path Selection Process- Send Traffic Through IXP
- **Hands On Lab Exercise:** BGP Peering
- IXP Design Considerations
- **Hands On Lab Exercise:** IXP Configuration
- Route Collectors & Servers
- **IXP BCP and What can go wrong?**

Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

What can go wrong?

Concept

- Some Service Providers attempt to cash in on the reputation of IXPs
- Market Internet transit services as “Internet Exchange Point”
 - “We are exchanging packets with other ISPs, so we are an Internet Exchange Point!”
 - So-called Layer-3 Exchanges — really Internet Transit Providers
 - Router used rather than a Switch
 - Most famous example: SingTelIX

What can go wrong?

Competition

- Too many exchange points in one locale
 - Competing exchanges defeats the purpose
- Becomes expensive for ISPs to connect to all of them
- An IXP:
 - is **NOT** a competition
 - is **NOT** a profit making business

What can go wrong?

Rules and Restrictions

- IXPs try to compete with their membership
 - Offering services that ISPs would/do offer their customers
- IXPs run as a closed privileged club e.g.:
 - Restrictive membership criteria (closed shop)
- IXPs providing access to end users rather than just Service Providers
- IXPs interfering with ISP business decisions e.g. Mandatory Multi-Lateral Peering

What can go wrong?

Technical Design Errors

- Interconnected IXPs
 - IXP in one location believes it should connect directly to the IXP in another location
 - Who pays for the interconnect?
 - How is traffic metered?
 - Competes with the ISPs who already provide transit between the two locations (who then refuse to join IX, harming the viability of the IX)
 - Metro interconnections work ok (e.g. LINX)

What can go wrong?

Technical Design Errors

- ISPs bridge the IXP LAN back to their offices
 - “We are poor, we can’t afford a router”
 - Financial benefits of connecting to an IXP far outweigh the cost of a router
 - In reality it allows the ISP to connect any devices to the IXP LAN — with disastrous consequences for the security, integrity and reliability of the IXP

What can go wrong?

Routing Design Errors

- Route Server implemented from Day One
 - ISPs have no incentive to learn BGP
 - Therefore have no incentive to understand peering relationships, peering policies, &c
 - Entirely dependent on operator of RS for troubleshooting, configuration, reliability
 - RS can't be run by committee!
- Route Server is to help scale peering at LARGE IXPs

What can go wrong?

Routing Design Errors

- iBGP Route Reflector used to distribute prefixes between IXP participants
- Claimed Advantage (1):
 - Participants don't need to know about or run BGP
- Actually a Disadvantage
 - IXP Operator has to know BGP
 - ISP not knowing BGP is big commercial disadvantage
 - ISPs who would like to have a growing successful business need to be able to multi-home, peer with other ISPs, etc — these activities require BGP

What can go wrong?

Routing Design Errors (cont)

- Route Reflector Claimed Advantage (2):
 - Allows an IXP to be started very quickly
- Fact:
 - IXP is only an Ethernet switch — setting up an iBGP mesh with participants is no quicker than setting up an eBGP mesh

What can go wrong?

Routing Design Errors (cont)

- Route Reflector Claimed Advantage (3):
 - IXP operator has full control over IXP activities
- Actually a Disadvantage
 - ISP participants surrender control of:
 - Their border router; it is located in IXP's AS
 - Their routing and peering policy
 - IXP operator is single point of failure
 - If they aren't available 24x7, then neither is the IXP
 - BGP configuration errors by IXP operator have real impacts on ISP operations

What can go wrong?

Routing Design Errors (cont)

- Route Reflector Disadvantage (4):
 - Migration from Route Reflector to “correct” routing configuration is highly non-trivial
 - ISP router is in IXP’ s ASN
 - Need to move ISP router from IXP’ s ASN to the ISP’ s ASN
 - Need to reconfigure BGP on ISP router, add to ISP’ s IGP and iBGP mesh, and set up eBGP with IXP participants and/or the IXP Route Server

More Information

Exchange Point Policies & Politics

- AUPs
 - Acceptable Use Policy
 - Minimal rules for connection
- Fees?
 - Some IXPs charge no fee
 - Other IXPs charge cost recovery
 - A few IXPs are commercial
- Nobody is obliged to peer
 - Agreements left to ISPs, not mandated by IXP

Exchange Point etiquette

- Don't point default route at another IXP participant
- Be aware of third-party next-hop
- Only announce your aggregate routes
 - Read RIPE-399 first
www.ripe.net/docs/ripe-399.html
- Filter! Filter! Filter!
 - And do reverse path check

Exchange Point Examples

- LINX in London, UK
- TorIX in Toronto, Canada
- AMS-IX in Amsterdam, Netherlands
- SIX in Seattle, Washington, US
- PA-IX in Palo Alto, California, US
- JPNAP in Tokyo, Japan
- DE-CIX in Frankfurt, Germany
- HK-IX in Hong Kong
- ...
- All use Ethernet Switches

Features of IXPs (1)

- Redundancy & Reliability
 - Multiple switches, UPS
- Support
 - NOC to provide 24x7 support for problems at the exchange
- DNS, Route Collector, Content & NTP servers
 - ccTLD & root servers
 - Content redistribution systems such as Akamai
 - Route Collector – Routing Table view

Features of IXPs (2)

- Location
 - neutral co-location facilities
- Address space
 - Peering LAN
- AS Number
 - If using Route Collector/Server
- Route servers (optional, for larger IXPs)
- Statistics
 - Traffic data – for membership

More info about IXPs

- <http://www.pch.net/documents>
 - Another excellent resource of IXP locations, papers, IXP statistics, etc
- <http://www.telegeography.com/ee/ix/index.php>
 - A collection of IXPs and interconnect points for ISPs

Summary

- L2 IXP – most commonly deployed
 - The core is an ethernet switch
 - ATM and other old technologies are obsolete
- L3 IXP – nowadays is a marketing concept used by wholesale ISPs
 - Does not offer the same flexibility as L2
 - Not recommended unless there are overriding regulatory or political reasons to do so
 - **Avoid!**

Structure of the Course

Day 3 : Stage 4

- Building a Demo IXP
 - Some presentation on Route Server
 - Will connect network on the IX

Day 2 : Stage 3

- Building BGP Concept
 - Introduction to BGP
 - BGP Path control
 - Hands-On Exercise

Day 1 : Stage 2

- Building the concept of Routing
 - Routing Introduction
 - How Internet Works?
 - Glue it together with Internet context
 - Some Hand-on Exercise

Day 1 : Stage 1

- Demystifying IXP Concept
 - What is IXP?
 - Value of Peering
 - How to Build an IXP?

Thank You!

