

Intro to IXP and Routing: IXP Overview


PCTA e-Tech Show 2022

07 April 2022


What is an Internet eXchange Point (IXP)?




- An IXP is a shared physical network infrastructure over which various Autonomous Systems can do easy peering with one another



One physical connection to IXP can be used for interconnections with multiple networks



More cost-effective and scalable, especially with more participants



ASes to be served by IXP include Internet Gateways, Internet Service Providers (ISPs), Research & Education (R&E) Networks, Cloud Service Providers, Content Providers and Content Delivery Network (CDN) Providers

IXPs are Layer-2 Networks



Switched Ethernet

- One physical connection for interconnections with multiple networks
- Only routers are allowed to connect to the switching fabric directly

IXP participants can do direct Bilateral Peering (BLPA) over the layer 2 infrastructure anytime


With Route Server added to the layer 2 infrastructure, IXP participants can also do Multilateral Peering (MLPA) for easier interconnections among everybody

- Traffic exchange is direct and not going through the route server


Those called themselves “IX” but serving layer-3 services are considered as transit providers

- Note that IXPs, transit providers and data centres are not the same things

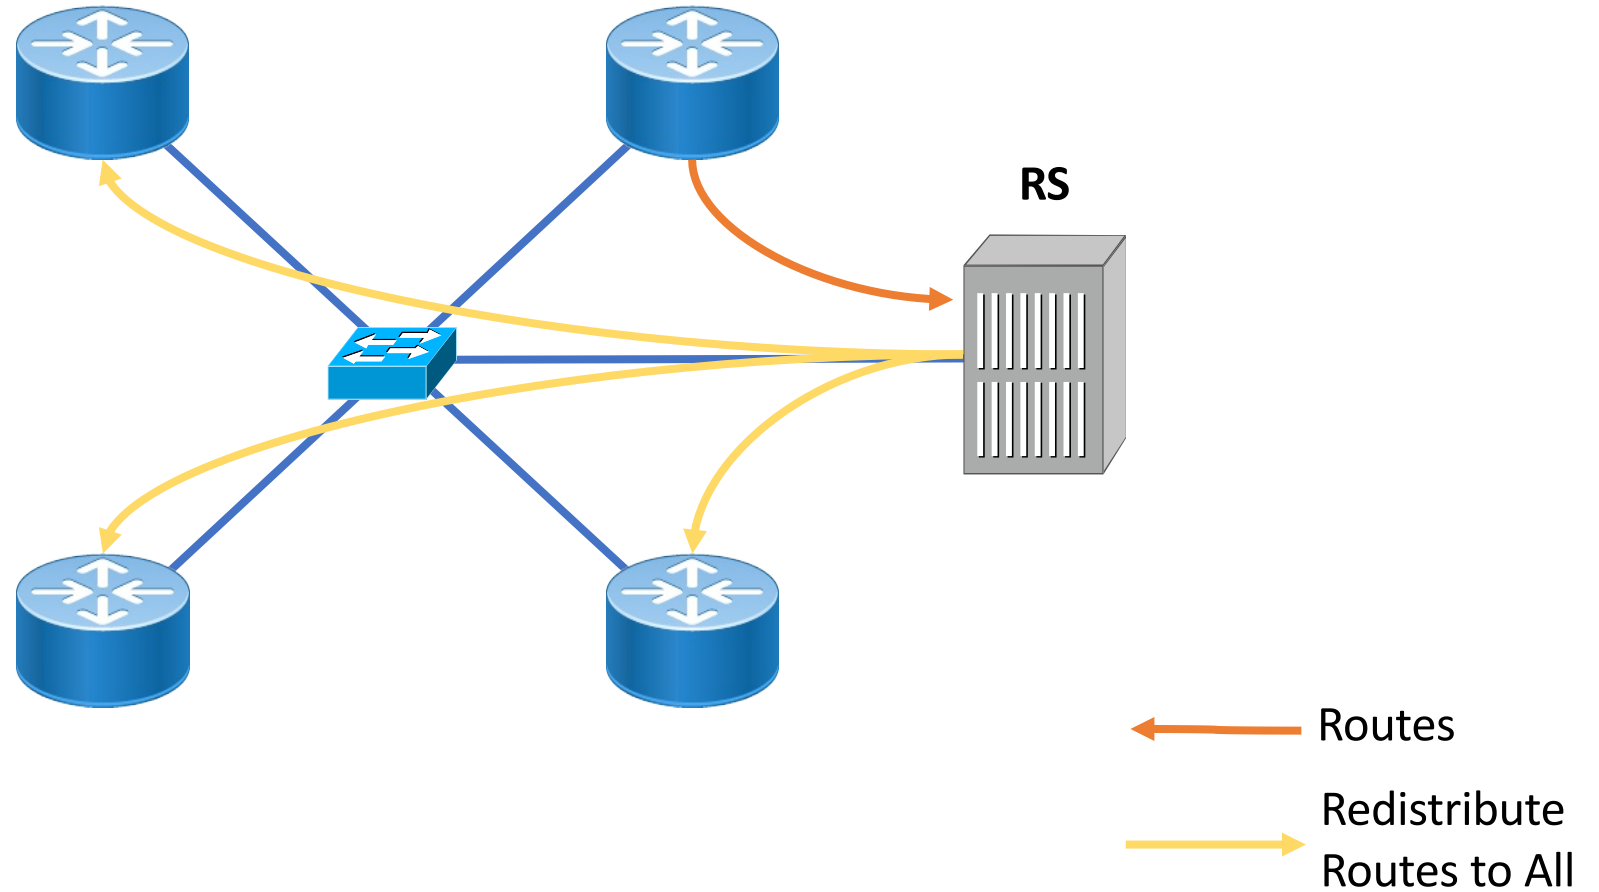
Simplest IXP Topology



IXP Topology with Minimal Switch/Site Resilience



Adding Route Server for Multilateral Peering



Main Benefits of IXP



- One main objective of an IXP is to keep local traffic local
 - Important to local Internet development
- Helps bypass 3rd-party network infrastructure for easy interconnection and **direct** traffic exchange among participating networks
 - Reduced cost – cheaper connectivity
 - Enhanced network performance – faster speed
 - Reduced latency – lower delay
- Helps encourage development of more local content and local applications
 - Helps local data centre business and other businesses
- Everybody is benefited
 - The gain for each may be different but all will gain
 - At the end, it is the most important that end users or consumers are benefited
- Often considered as Critical Internet Infrastructure locally, regionally or globally

Benefits of IXP - Example



A Fijian ISP in Suva accessing content at the University of the South Pacific in Suva

- Packet travels $> 25,000\text{km}$
- Physical distance $< 10\text{km}$
- Adding long latency
- Possibly high jitter too
- Using expensive submarine capacity

Return path had similar issue



Internet eXchange Point (IXP)

- Need a location or facility that ISPs can access and can connect to each other over a common shared media
 - Eg: Ethernet switch
- Should be a NEUTRAL venue
- Needs to have multiple telco circuit providers and/or allow any licenced provider to install services
- Needs controlled environment & access



Internet eXchange Points


- Variety of shapes and sizes
 - Commercial
 - Community
 - Tbps to Mbps
 - Single location or Metropolitan Area scoped
 - Purely a traffic exchange
 - Value added services

- Layer 2 exchange point
 - Ethernet Switches (100Gbps/10Gbps/1Gbps/100Mbps)

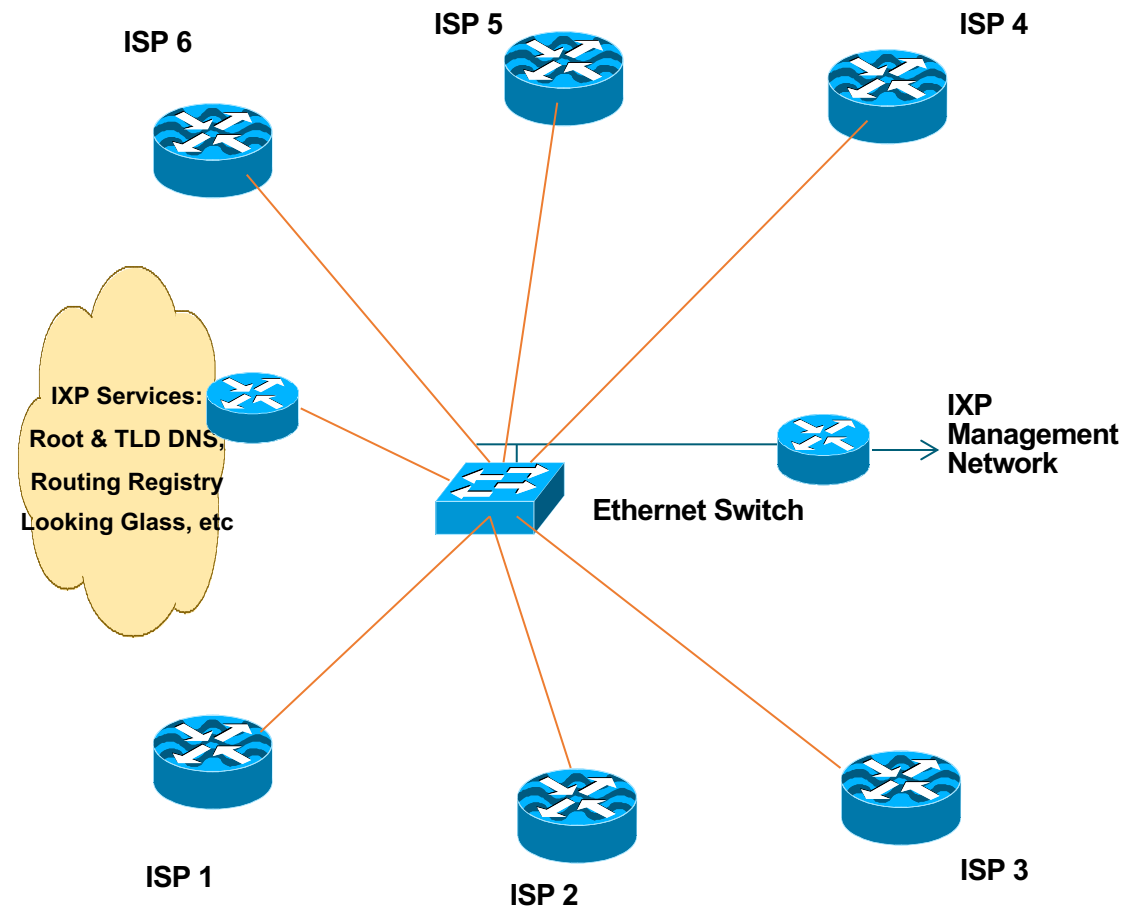
Internet eXchange Point




- Border routers in different Autonomous Systems



Layer 2 Exchange



Layer 2 Exchange





Layer 2 Exchange

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



Layer 2 Exchange

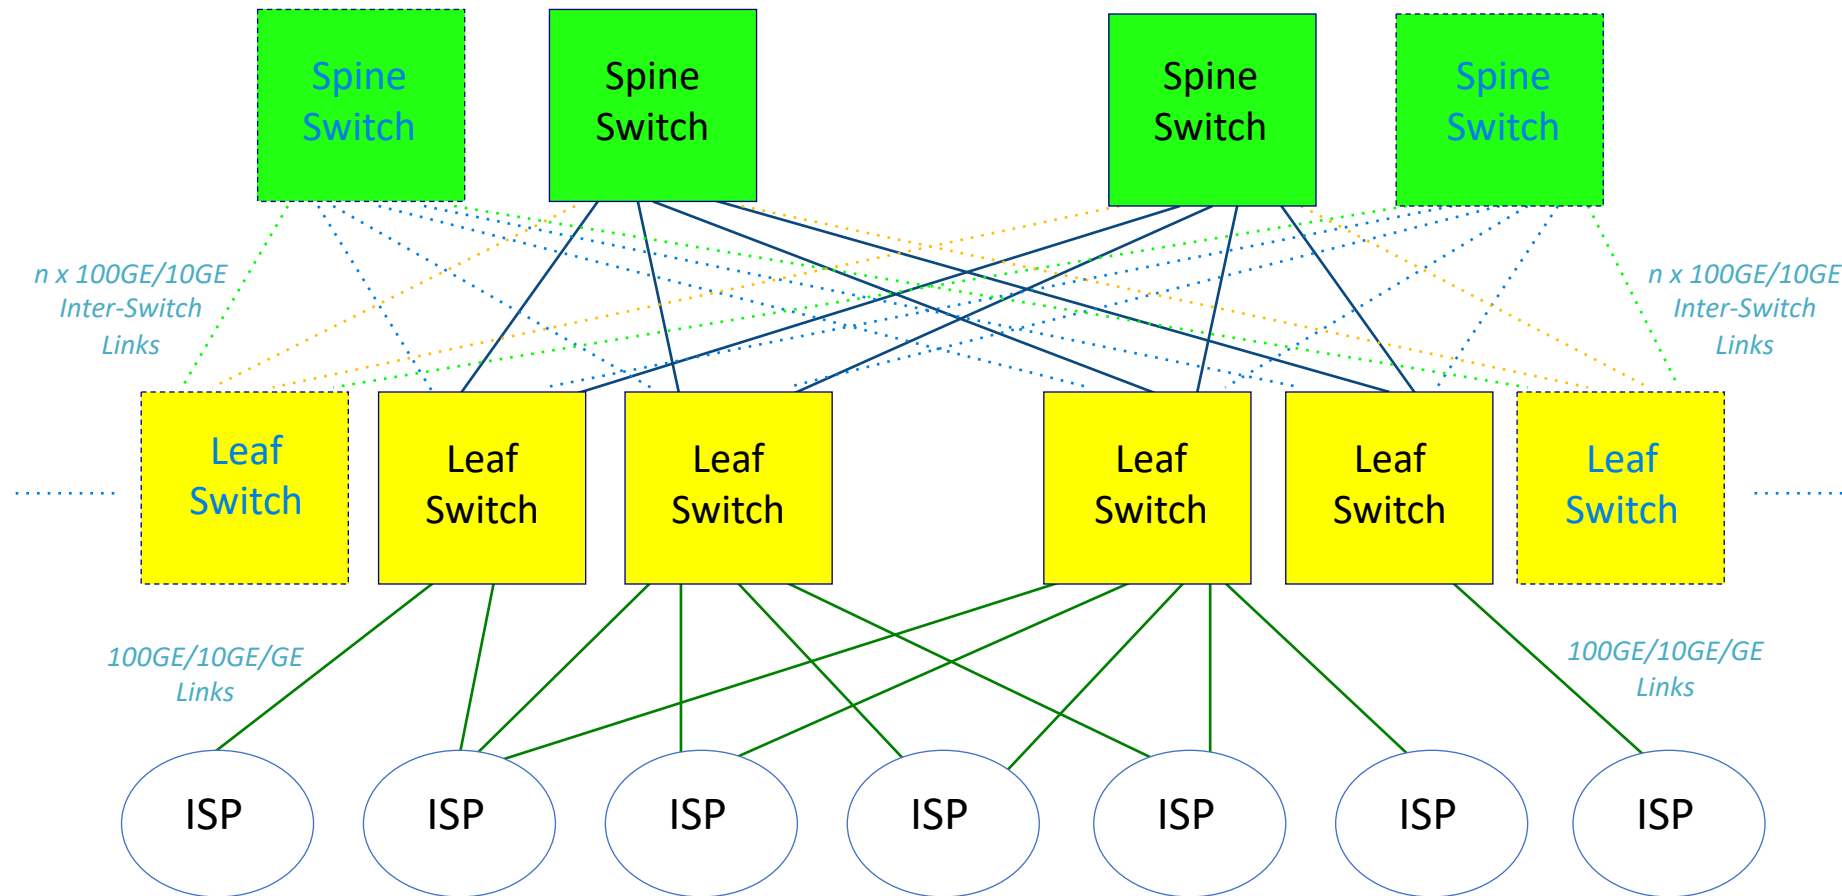
- Requires neutral IXP management
 - Usually funded equally by IXP participants
 - 24x7 cover, support, value add services
- Secure and neutral location
- Configuration
 - Private address space if non-transit and no value add services
 - Otherwise public IPv4 (/24) and IPv6 (/48, /56, /64)
 - ISPs require ASN, basic IXP does not
 - Route Servers need ASN

Layer 2 Exchange

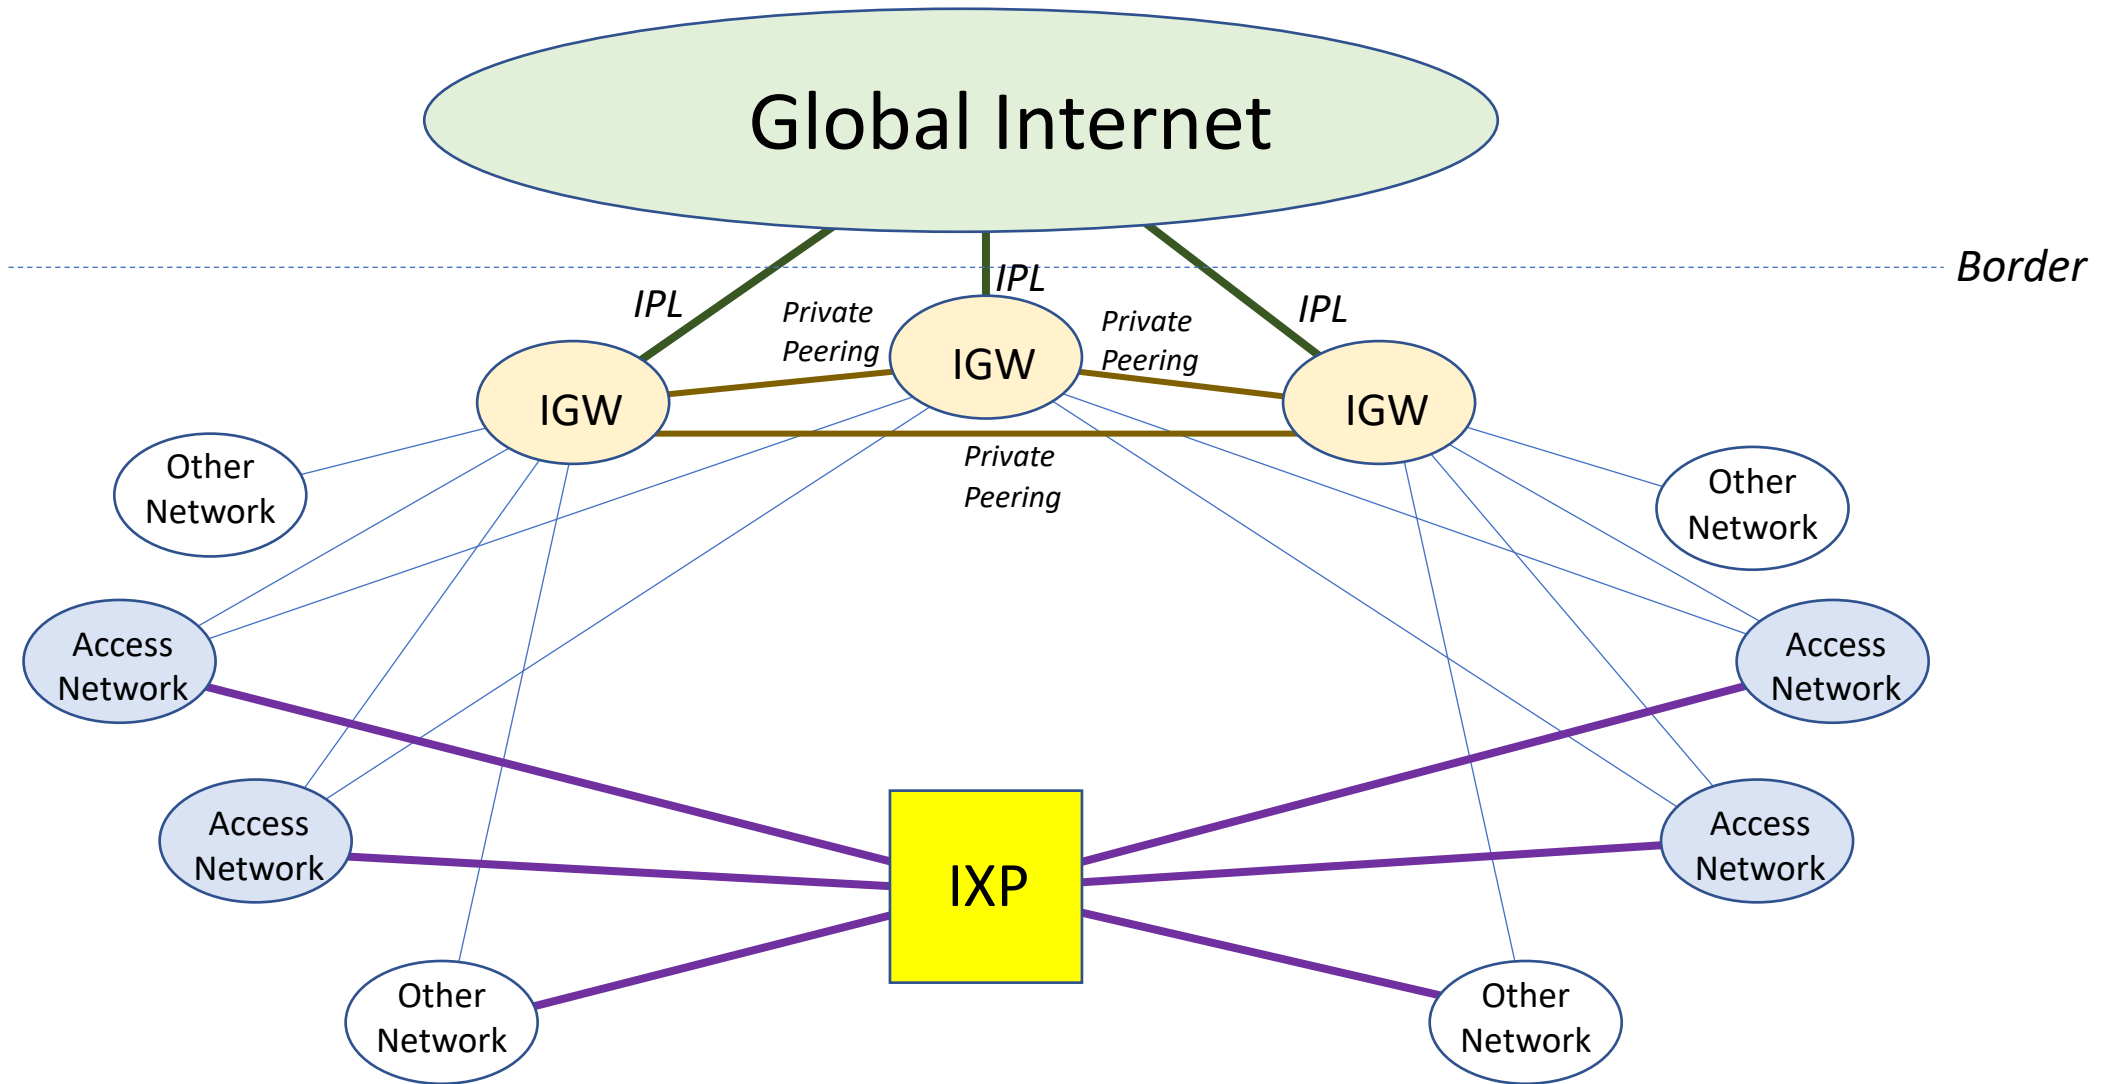


- Network Security Considerations
 - LAN switch needs to be securely configured
 - Management routers require AAA authentication, vty security
 - IXP services must be behind router(s) with strong filters

Spine-and-Leaf Architecture for Beyond 2 Switches



Possible Scenario for Improving Local Peering



BGP and Peering at IXPs



Peering partners establish connection using BGP

```
router bgp 4608
  neighbor equinix.sg1.v4 peer-group
  neighbor equinix.sg1.v4 update-source GigabitEthernet0/0/0.101
  neighbor 27.111.228.226 remote-as 4775
  neighbor 27.111.228.226 peer-group equinix.sg1.v4
  neighbor 27.111.228.226 description equinix.sg1:globetelecom
```

PeeringDB Entry

Set up a record of your ASN on www.peeringdb.com and tell everyone where you are (at which IXPs and/or data centres) and that you are willing to do BLPAs.

Also use it to find your potential BLPAs peers

APNIC Silver Sponsor

Organization	APNIC
Also Known As	Asia Pacific Network Information Centre
Long Name	
Company Website	https://www.apnic.net
ASN	4608
IRR as-set/route-set	AS-4608
Route Server URL	
Looking Glass URL	
Network Type	Non-Profit
IPv4 Prefixes	30
IPv6 Prefixes	30
Traffic Levels	100-1000Mbps
Traffic Ratios	Mostly Outbound
Geographic Scope	Asia Pacific
Protocols Supported	<input checked="" type="checkbox"/> Unicast IPv4 <input type="checkbox"/> Multicast <input type="checkbox"/> IPv6 <input type="checkbox"/> Never via route servers
Last Updated	2021-08-10T03:33:06Z
Public Peering Info Updated	2021-08-22T00:13:30
Peering Facility Info Updated	2020-07-15T04:41:32Z
Contact Info Updated	2020-06-24T23:16:56Z
Notes	

Peering Policy Information

Peering Policy	
General Policy	Open
Multiple Locations	Not Required
Ratio Requirement	No
Contract Requirement	Required

Contact Information

Role	Name	Phone	E-Mail
NOC	Network Operations	+61738583100	noc@apnic.net
Technical	Peering Team	+61738583100	peering@apnic.net

Public Peering Exchange Points

Exchange ASN	IPv4 IPv6	Speed RS Peer
EdgeIX - Brisbane 4608	103.136.103.23 2001:df0:680:4::17	10G ☺
Equinix Singapore 4608	27.111.230.31 2001:de8:4::4608:1	1G ☺
Equinix Sydney 4608	45.127.172.3 2001:de8:6::4608:1	1G ☺
IX Australia (Brisbane QLD) 4608	218.100.76.36 2001:7fa:11:2:0:1200:0:1	10G ☺
MegalX Brisbane 4608	103.26.70.9 2001:dea:0:20::9	1G ☺
MegalX Sydney 4608	103.26.69.18 2001:dea:0:10::112	1G ☺

Private Peering Facilities

Facility ASN	Country City
Interactive Brisbane 4608	Australia Eight Mile Plains
NEXTDC B1 4608	Australia Brisbane



Some IXPs in the Philippines



MANILA	
ComclarkIX MIX GETAFIX GIX/NIGX	PCTA-IX PHIX PhOpenIX VIX

CEBU
GETAFIX GIX PhOpenIX VIX

DAVAO
GIX PhOpenIX




*From the latest "IXP & DC landscape"



Unfortunately, a lot of IXP participants do not make the best use of the IXP(s) they have connected

IXP Participants without enough knowledge and skills may disrupt the operations of IXP from time to time

IXPs need to do a lot of education or push to their participants



So, IXP engineers would be busy and dedicated resources would be needed

Depending on the size, volunteering type of operations mode may not be able to sustain for too long

If volunteers are to be used

- There should be multiple engineers from multiple companies / organisations for resilience
- Help from universities should be considered

IXP Development Work of APNIC



APNIC strongly believes IXPs help Internet development

- After all, IXPs serve and benefit APNIC members
- In fact, IXPs need IP addresses and ASNs and so are APNIC members themselves

Do more on helping those developing economies

- Especially those which do not have any IXP yet
- Or those which their only IXP is not functioning well

Training and Technical Assistance work primarily

- Not just for IXP operators but also for IXP participants
- Also help talk to major stakeholders to convince them of the benefits of having a local IXP while maintaining neutrality
- May need help of Community Trainers and Consultants from time to time

Having been supporting IXP development in Fiji, PNG, Vanuatu, Mongolia, Bhutan, Myanmar, some cities in India and others

IXP Development Package of APNIC



Training & Technical Assistance as the minimum


- Technical & non-technical

Other possible support items (on case-by-case basis according to individual needs):


- Ethernet switch
- Root Server anycast instance
- Route Server
- ROV & IPv6 deployment support
- IXP Manager
- RIPE Atlas Anchor
- CSIRT Establishment
- Honeypot of HoneyNet Project for Analysis
- BGP Route Collection for Analysis

APIX & MANRS Memberships are recommended to all IXPs

To Add Value



ROUTE SERVER / ROA AND RPKI




DOMAIN NAME INFRASTRUCTURE
(ROOT/TLD SERVERS)



SHARED CDN CACHES

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
 - provide a public view of the Routing Information available at the IXP



Route Servers



- 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
- Optionally uses Policy registered in IRR

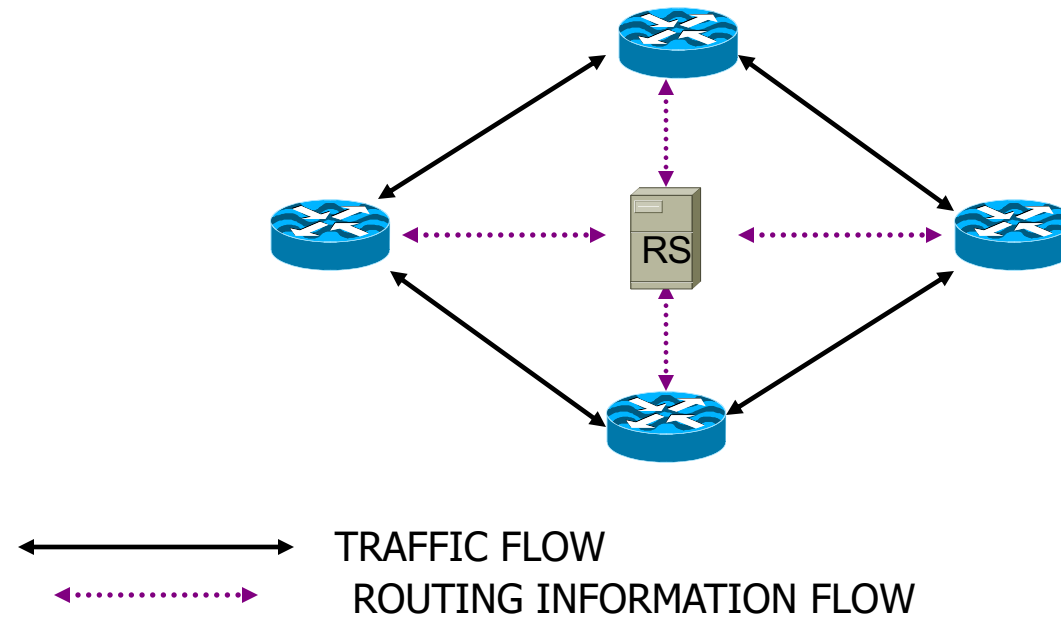
N-squared Peering Mesh vs Peering with Route Server

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

- ISP routers peer with the Route Servers → Only need to have two eBGP sessions



RS based Exchange Point Routing Flow





Advantages & Disadvantages of Using a Route Server

[Advantages]

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

[Disadvantages]

- ISPs can lose direct policy control
 - If RS is only peer, ISPs have no control over who their prefixes are distributed to
 - Some IXPs provide community based filtering option
- Completely dependent on 3rd party
 - Configuration, troubleshooting, etc...



Typical usage of a Route Server

- Route Servers may be provided as an **OPTIONAL** service
 - Most common at large IXPs (>50 participants)
 - Examples: LINX, HKIX, 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?
 - 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?)

Connecting to an IXP

Cost tied to circuit size (not byte count)



- Peering is typically settlement free
 - No charge for the traffic exchanged

- Cost to peer
 - Router interface
 - Circuit to the peering fabric
 - Charges imposed by the IXP
 - All fixed, either capital expenditure or monthly recurring fee

Choosing an IXP



- Some markets have more than one
- Even if there is only one IXP it might appear in multiple locations
 - E.g. LINX is built on two rings through multiple data centres across London
- Best location might be dictated by availability of IPLC, transit, or other factors



Which IXP?

- How many routes are available?
 - How many other operators/providers are at the IX?
 - What is the traffic to and from these destinations, and how much will it reduce the transit cost?
- What is the cost of co-lo space?
 - Availability of power, type of cabinet, ...
- What is the cost of a circuit to the location?
 - If similar to transit costs are you getting a benefit?
- What is the cost of remote-hands?
 - For maintenance purposes to avoid serious outages

Remote locations



- If building to a remote location
- Make sure remote hands work at times when it's important to you
 - Their 9-5 is not normally your office hours
- Check the skill set of the remote hands
 - Maybe engage a local consultant to help



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



QUESTIONS?

Connect with Me!
Email: shane@apnic.net
Socials: [@irrashai](#)

Thank You!

