Keeping Local Traffic Local by Doing Local Peering

Che-Hoo Cheng
Infrastructure & Development Director, APNIC
PacNOG25@2019-12-02

How Does Internet Operates?

- Internet is a network of networks, composed of networks of ISPs and users
- User networks connect to ISPs
- Small ISPs connect to large ISPs
- Various networks (large or small) are **inter**connected with one another to form Internet

Autonomous Systems

- A network on Internet is called Autonomous System (AS) which is represented by AS Number (ASN)
  - ASN is unique around the world
    - APNIC is in charge of ASN assignment for AP region
  - Used together with BGP (Border Gateway Protocol) for interconnections with multiple networks (or multi-homing)
  - Networks having ASNs can be more independent, or portable
    - Together with portable IP addresses
    - Like what APNIC members are enjoying…

Ordinary Transit Model

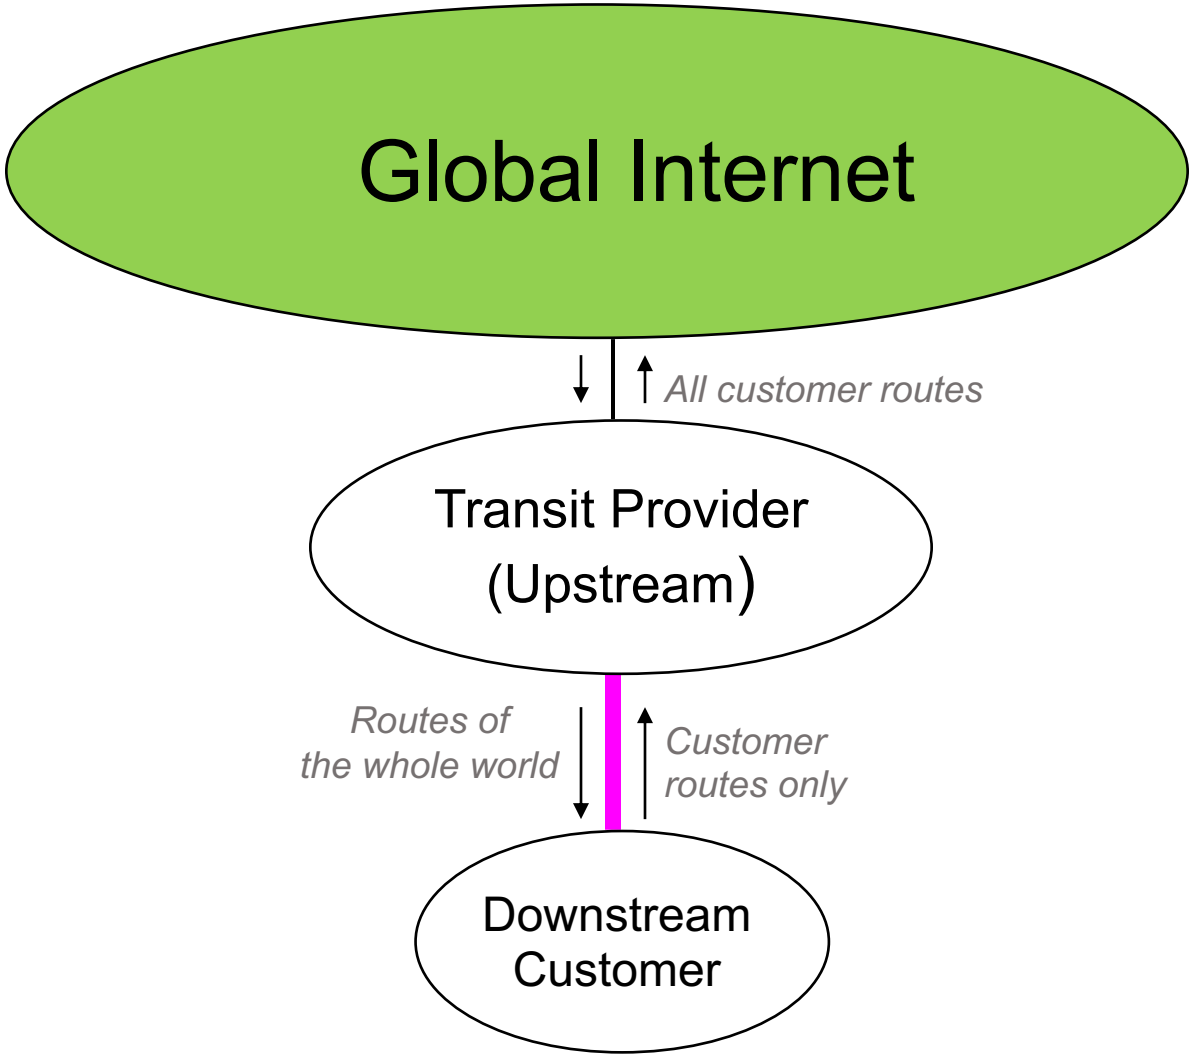

Transit in General

- Networks pay transit providers to get to the whole Internet
  - Can connect to multiple transit providers for resilience and portability
- A few very large ISPs act as transit providers for the whole world (the so-called tier-1 networks) which do not need to pay others to get full Internet connectivity
  - Other ISPs must be transit customers of those tier-1 networks directly or indirectly in order to gain full connectivity
- Networks on Internet are trying to bypass transit providers as much as possible
  - for lower cost and higher performance

Ordinary Peering Model

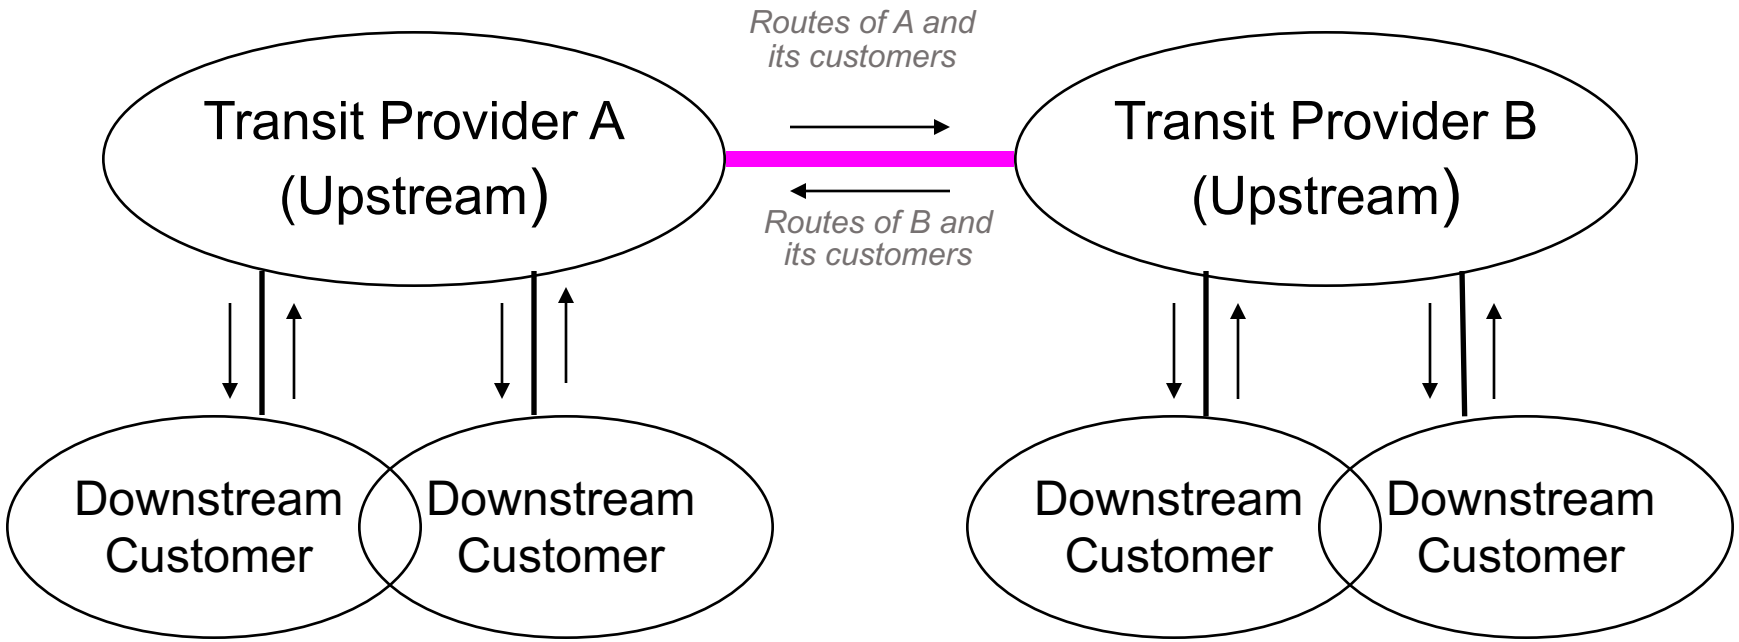

Peering in General

- ASes are interconnected/peered at Internet exchanges points (IXPs) or privately
- Interconnection/peering is among ISPs / data centres / content providers / cloud services providers which have different ASNs using BGP protocol

For mutual benefits	• **For higher performance, lower latency and lower cost** • **Usually no settlement between peers and cost is shared**
Local Peering	• **Local-to-local traffic do NOT need to route through overseas** • **Important to local Internet development**
Between 2 ASes	• BLPA (Bi-Lateral Peering Agreement)
Among > 2 ASes	• MLPA (Multi-Lateral Peering Agreement)

Example of Not Having Local Peering

- A Fijian ISP in Suva accessing content at the University of the South Pacific in Suva
- Packet travels > 25,000km
- Physical distance < 10km
- Adding long latency
- Possibly high jitter too
- Using expensive submarine capacity
- Return path had similar issue

Private Peering

- A form of BLPA having dedicated point-to-point connection between 2 ASes
- Using cross-connect or local loop or IPL to interconnect
  - Cost is usually shared between 2 peers
- May have multiple connections between 2 ASes for resiliency
- Not quite cost-effective
  - Spare bandwidth cannot be used for other traffic
- Not very scalable
  - ${}_nC_2$ physical connections for n ASes to peer fully with one another

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

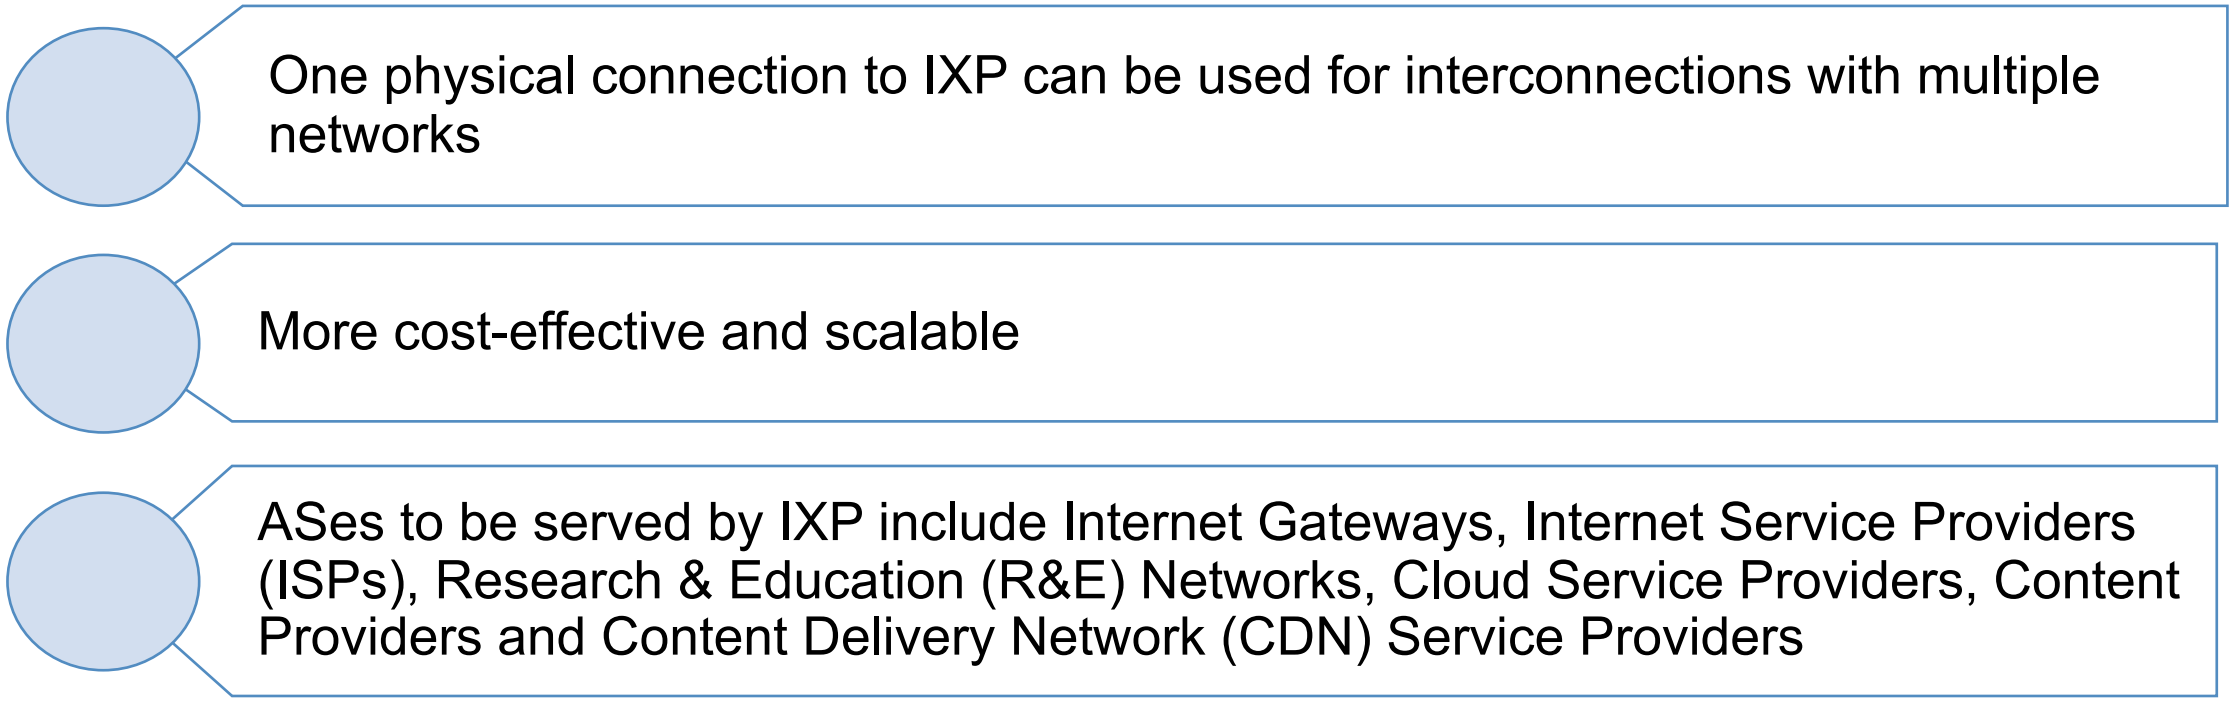

Simplest IXP Topology

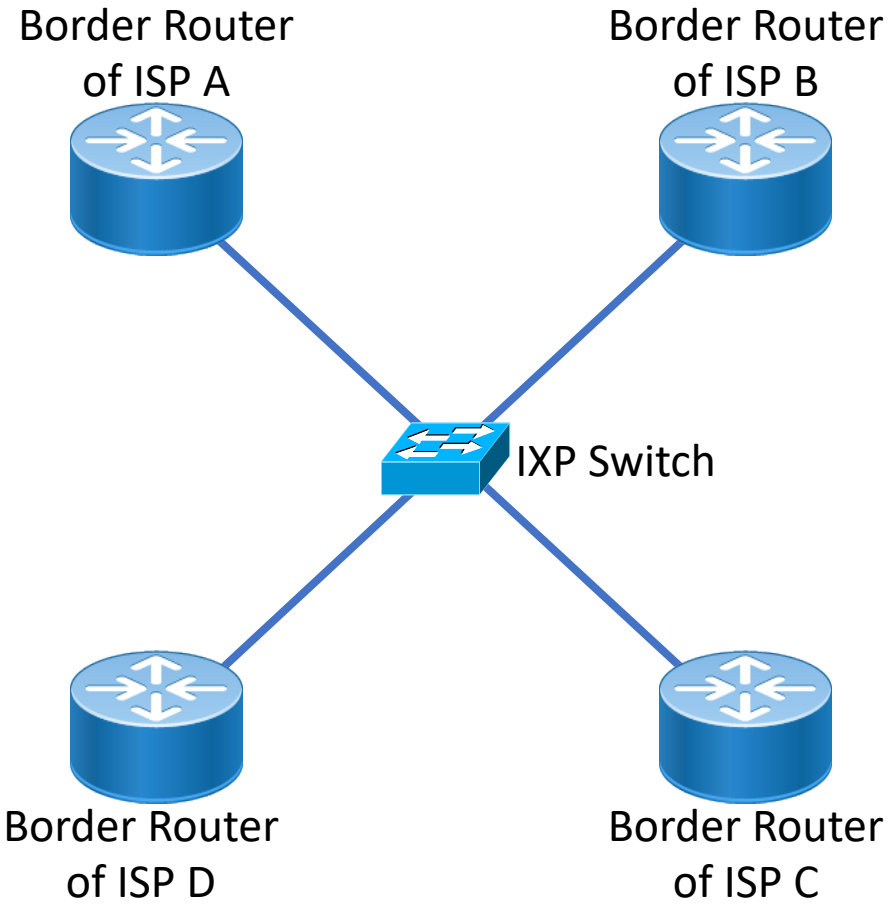

Benefits of IXP

- 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

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 not going through the route server but direct
- Those called themselves “IXes” but serving layer-3 services are mostly transit providers

Value and Attractiveness of an IXP

- Proportional to the number of different networks (ASNs) connected and also the traffic volume
- Snowball effect after reaching critical mass
  - The initial period usually is the hardest
    - Most will take wait-and-see approach
  - Gradually will have good mix of networks of different types
    - E.g. Eyeballs vs Content

Evolution

- IXP development is an evolutionary process
  - Can be done step by step
- It can be improved over time, but picking the right initial neutral organisation / governance model and a neutral site at the start is important for future success

Possible Steps for IXP Development

- Can be gradual, step by step

- Layer-2 network is the bare minimal
  - Can use private IP addresses if small amount of participants
- Public IP addresses next
  - Legal entity issue
- Site resilience is IMPORTANT while equipment resilience is already included
  - Has to have site resilience sooner or later
- Route server(s) with ASN follows
  - RPKI consideration
- Other value added services
  - DNS: Root / TLDs / Recursive
  - Shared Content Caches?

IP/ASN Resources for an IXP

- Considered as Critical Infrastructure under APNIC Policy
  - Using public IP addresses and ASN is recommended
  - IPv4: /24
  - IPv6: /48
  - ASN: 1 (for route server)
- But IXP may need another network to provide transit
  - Own servers such as network management & monitoring
  - DNS anycast servers: Authoritative or Cache/Resolving/Recursive
  - Shared Content Caches for Participants
  - Usually small

Neutral Location is Good Starting Point

- May choose one of the followings as starting point:
  - University
  - Technology Park
  - Carrier Neutral Data Center
  - Government Data Center
- Having multiple carrier options with easy access is important
- Should maintain neutrality continuously
- Expansion to multiple sites within the same metro area can be done gradually, coupled with growth

Governance

- Multi-stakeholder bottom-up approach is the preferred approach for maximum acceptance of the community
- Government support is also important
- Be as inclusive as possible in order to provide maximum benefits to the whole community which it serves
- Should be fair and consistent to every participant
- Should be open and transparent as much as possible

IXP Models

Developed economies vs Developing economies

Non-profit vs Commercial

Subsidized vs Self-financed

Government-led vs Industry-led

- No one single model which can suit all situations
- Relative Neutrality is important

Commercial vs Non-Profit

- Commercial set-up is free to do anything
  - No need to care about neutrality too much
  - IXP is mostly a service to help other business
- Non-profit set-up tends to be more cautious
  - Neutrality is more important, at least to the target participants
  - Tend to be more independent
  - Tend to offer fewer services

Advanced / Developed Economies

- IXPs are business
  - Even for not-for-profit set-up
  - Less government involvement
- Multiple IXPs
  - Keen competition
- But if they cannot keep intra-economy traffic local, someone needs to step up
  - Government? Industry group? Customer pressure?

Developing Economies

- Some do not have any IXPs yet
- Local traffic does not stay local
  - A lose-lose situation for everybody
- IXPs can help Internet development a lot
  - Better to be non-for-profit set-up
  - May need to start with subsidized model
  - May not be a business at all
  - Help from government is mostly needed
  - Active participation of the biggest players is also very important

Examples of Pacific Islands

- Far from any other places
- External connectivity is very expensive
  - More submarine cables are being built for them
- Small markets because of small population
- Usually just a few ISPs but they may not be interconnected locally
- Local traffic across ISPs usually routed through US or Australia
- Local IXP is very much needed for helping Internet development
- Observed immediate benefits on Day 1 of set-up of Fiji-IXP
  - Much improved latency and high volume of traffic

Before Fiji-IXP was set up

- A Fijian ISP in Suva accessing content at the University of the South Pacific in Suva
- Packet travels > 25,000km
- Physical distance < 10km
- Adding long latency
- Possibly high jitter too
- Using expensive submarine capacity
- Return path had similar issue

Politics Involved in Early IXP Development

- Usually larger ISPs like IXP less than smaller ISPs because smaller ISPs are mostly target customers of larger ISPs
- Larger ISPs refuse to connect to IXP making the value of IXP lower
- There are multiple possible mitigation options for that but in any case, larger ISPs need to collaborate
  - E.g. separating access networks from Internet gateway or transit network
- If hurting the goal of **“Keeping Local Traffic Local”**, then it is lose-lose to everybody
- Government involvement may help or may hurt the case
  - It depends on the relationship between the industry and the government
  - Forcing large ISPs to do peering may not achieve the expected outcomes
- But having an IXP is NOT a magic wand to solve all the issues
  - But collaborative spirit is

Government Funding for IXPs?

- Is it good or bad?
- More needed during infancy stage of IXP development
- But for long-term, it is probably better to have bottom-up industry-led governance for IXP
  - Align with bottom-up multi-stakeholder approach
  - Need to have a long-term sustainable financial model

Which Models Can Sustain?

Pure Business Model	• IXP alone cannot make big money • Or IXP may just be a value added service
Subsidized Model	• Government funding may or may not be more reliable
Model relying on sponsorship and/or volunteers	• Most risky as sponsorship or support of volunteers is not guaranteed
Membership-based Model	• Open Membership vs Closed Membership • Proper governance is important • Most neutral but still need to have good financial model for long-term sustainability

Geography

- IXP usually is not expanded beyond a metro area so as to avoid competing with IXP participants and to maintain neutrality
- Should start with the biggest city first and gradually set up separate infrastructure in other bigger cities one by one

To Leverage the Position of & to Add Value to an IXP

- **Domain Name Infrastructure:** DNS infrastructure is very important to Internet operations so Root/TLD DNS server instance(s) should be connected directly to IXP for direct peering in order to benefit all participants for better DNS performance and resilience
- **Shared Cache:** Connecting cache servers of popular content to the IXP will help everyone save bandwidth, but the cost of the bandwidth for cache-fill has to be properly shared by the ISPs benefited
  - Different cache service providers have different supported models
  - Need to think about long-term sustainability
- **NOTE:** Transit for the above should NOT be used for providing usual transit service to IXP participants so as to maintain neutrality

IXP Participants

- 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
- IXP operators need to do a lot of education or push to their participants
- So, IXP engineers would be busy and dedicated resources would be needed
  - Volunteering type of operations mode cannot sustain for too long

Scalability Issue

- IXPs were not supposed to have any packet loss in its infrastructure
  - And with very low latency too
- Become an issue when IXP grow beyond one switch
  - Due to not enough ports or expanding to multiple sites
- Inter-switch links are the risk
  - Over-subscription or not?
- Spine-and-leaf architecture helps a bit but not for all cases
  - Need to determine how much bandwidth from leaf to spine anyway
  - Still not ideal if there are adjacent leaf switches at one site
    - All traffic among 2 adjacent leaf switches has to go to the spine first?

Spine-and-Leaf Architecture

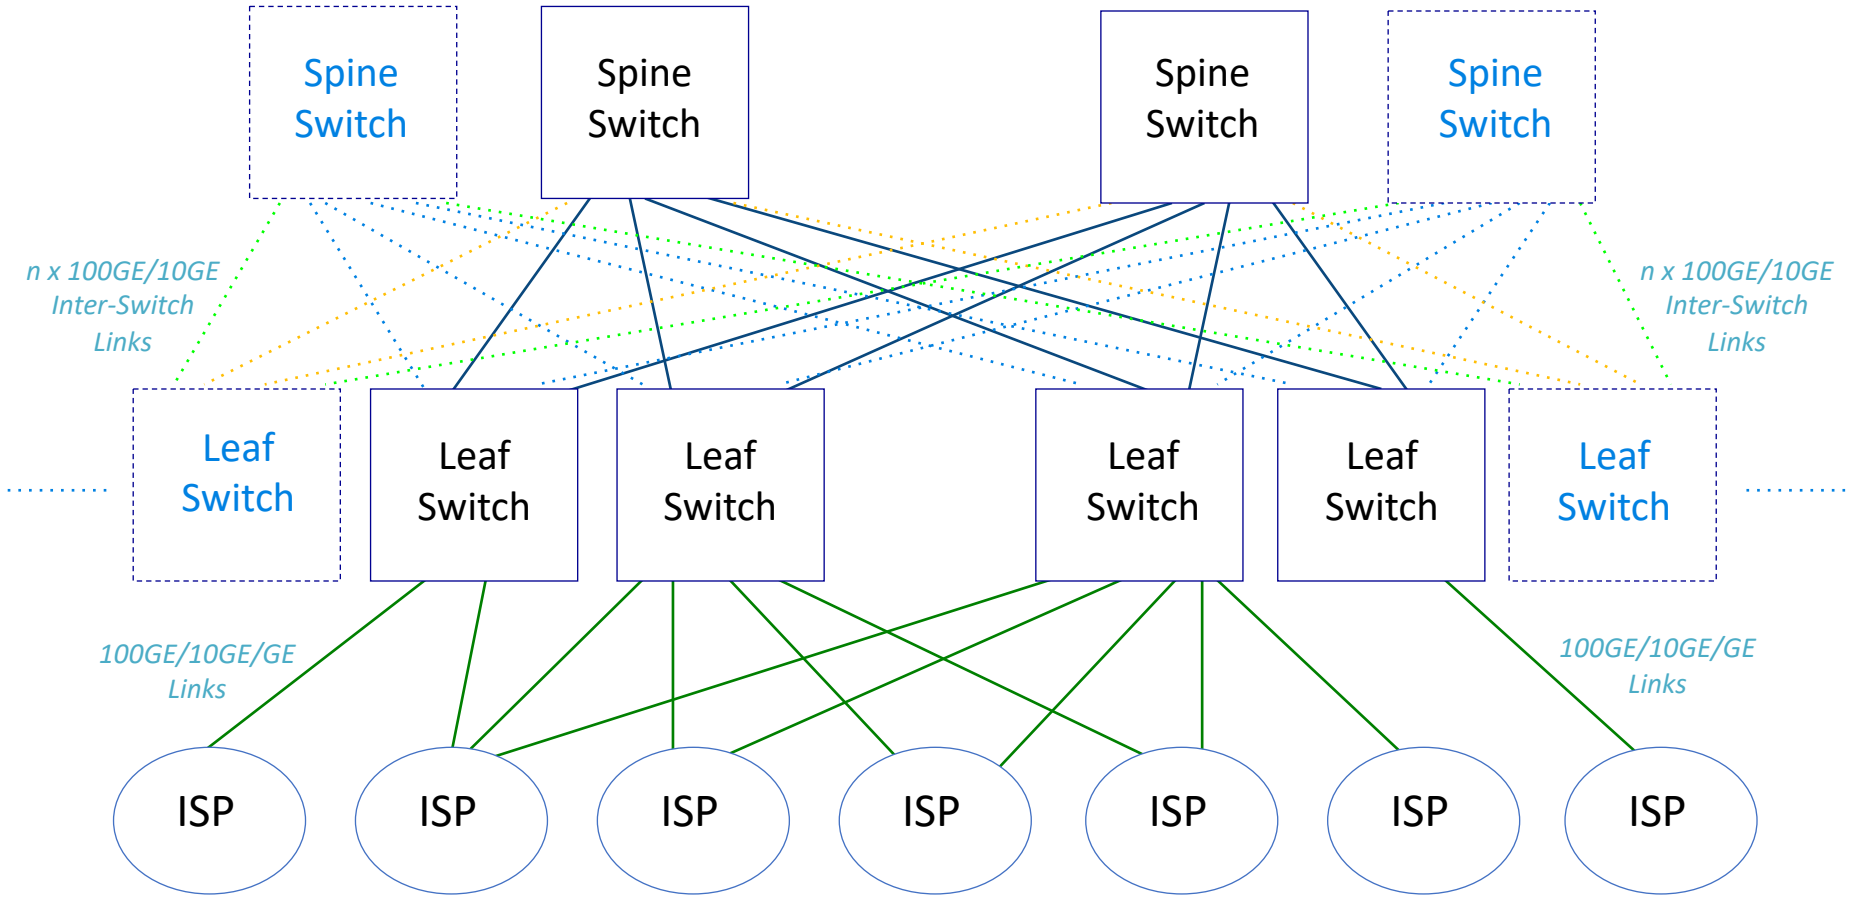

IXP Development Work of APNIC

- APNIC strongly believes IXPs help Internet development
  - That is why we support APIX and related activities
  - 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
- 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

- Providing Training & Technical Assistance is the minimum
- Will tailor-make support according to individual needs
- Other possible support items (on a case-by-case basis):
  - Ethernet switch
  - Root Server anycast instance
  - Route Server
  - ROV & IPv6 deployment support
  - IXP Manager
  - RIPE Atlas Anchors
  - CSIRT Establishment
  - Honeypot of Honeynet Project for Analysis
  - BGP Route Collection for Analysis
- APIX Membership is recommended to all IXPs

Route Origin Validation (ROV) at IXP – via Route Server and/or Shared Validator

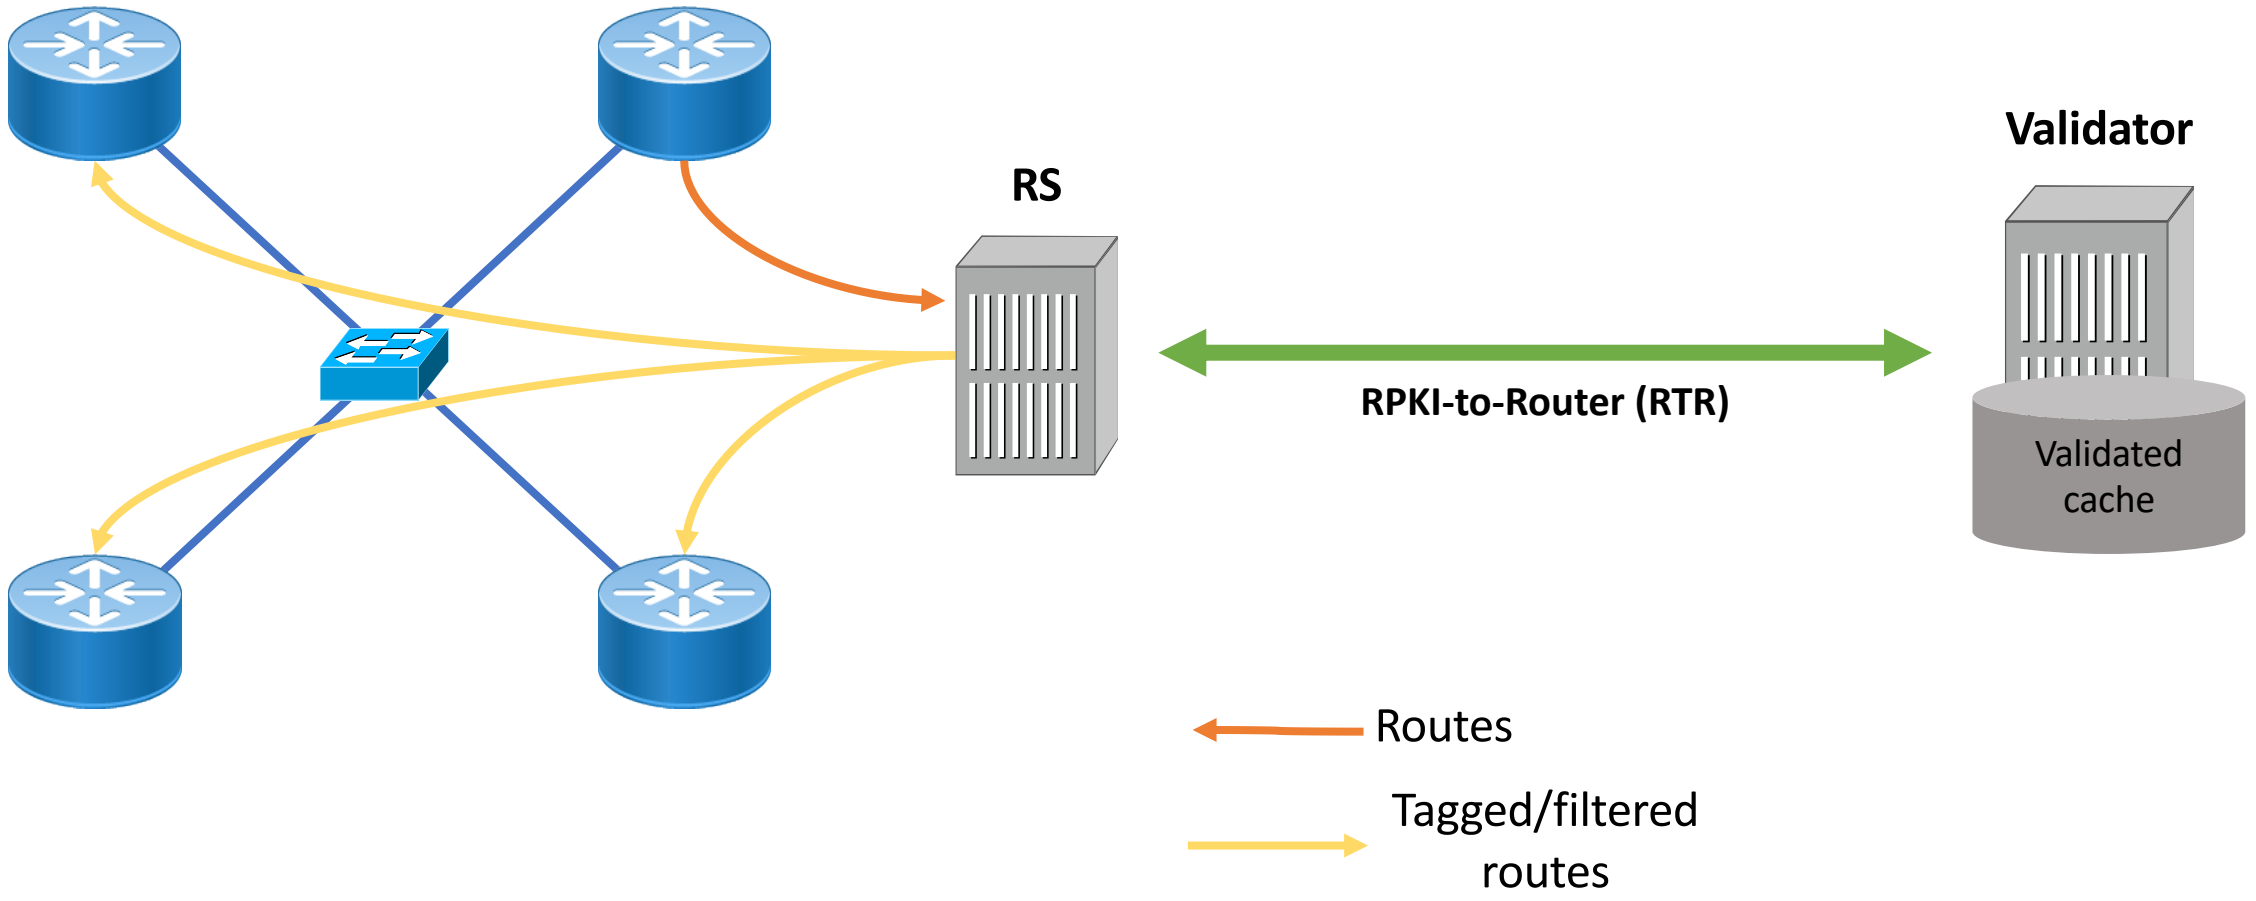

Other Help & Support by APNIC

- APNIC also provides help & support to:
  - APIX
  - Peering Asia
  - Peering Forums hosted by not-for-profit IXPs
  - NOGs (which IXPs usually support)
- APNIC also sponsors:
  - PeeringDB
  - IXP-DB
  - IXP Manager

Final Remarks

- IXPs will continue to play a key role for easier interconnections among networks for **local peering** in order to achieve **“Keeping Local Traffic Local”**
  - Especially for developing economies
  - But IXP is NOT a magic wand to solve all the issues
    - Collaborative spirit is
- Need to find a suitable model for long-term sustainability
- Relative neutrality is important
  - So better to maintain it as much as possible
- After all, **“Keeping Local Traffic Local”** is the most important thing

Thank You!