Date

A Primer on Internet Exchanges Points

Wednesday 1 July 2020, 09:00-11:30 (Bangkok Time)

Aftab Siddiqui
Sr. Internet Technology Manager
Siddiqui@isoc.org

What are we discussing today

Agenda

- The basics of Internet Routing
- Transit, Peering and the Border Gateway Protocol (BGP)
- How an IXP works
- Security considerations, including network routing security
- Best practices in establishing a neutral IX

The basics of Internet Routing

Supporting content outlining
the following section.

How Internet Works?

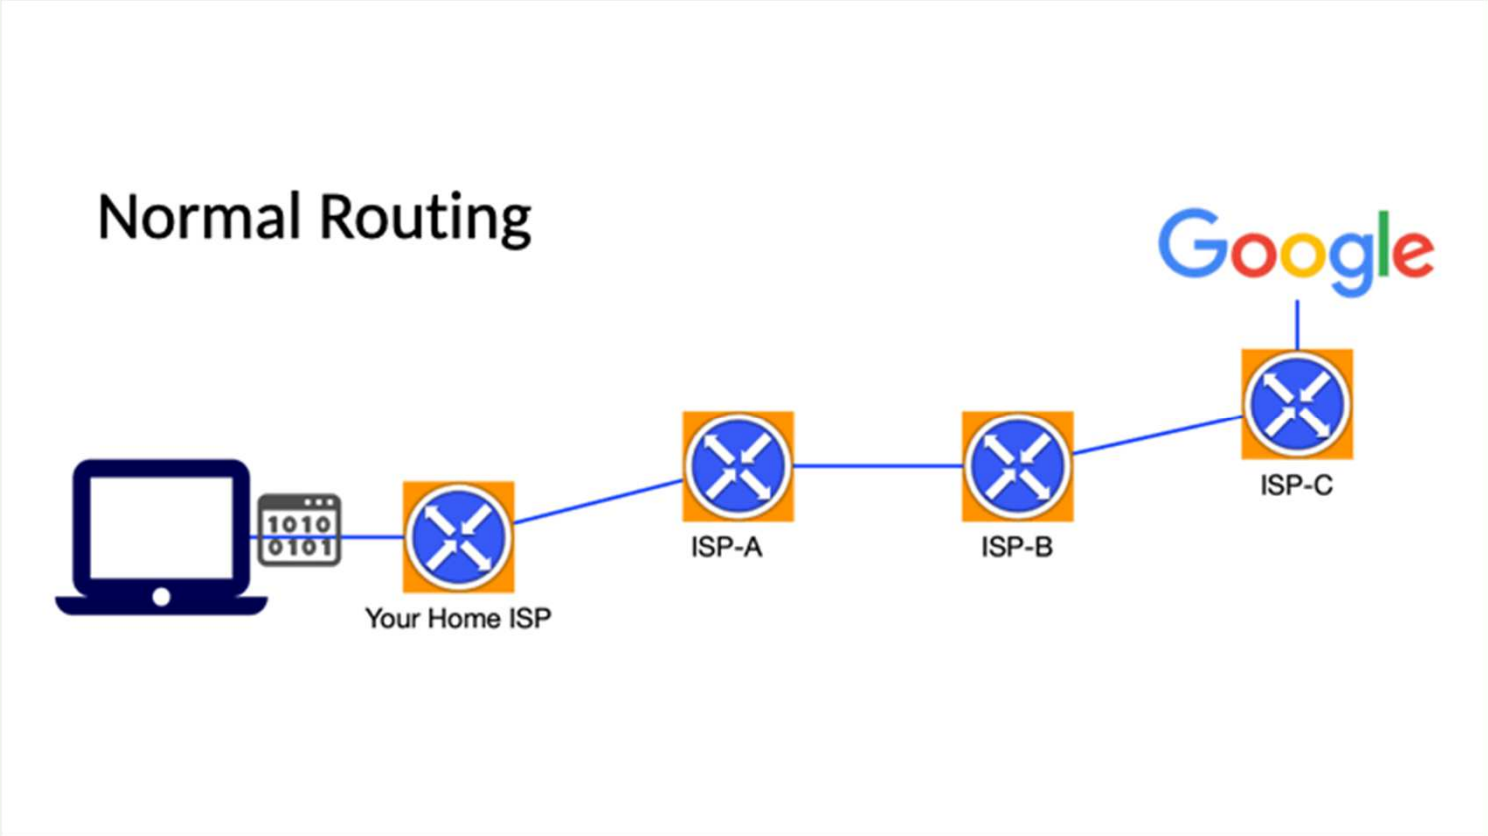

How Internet works?

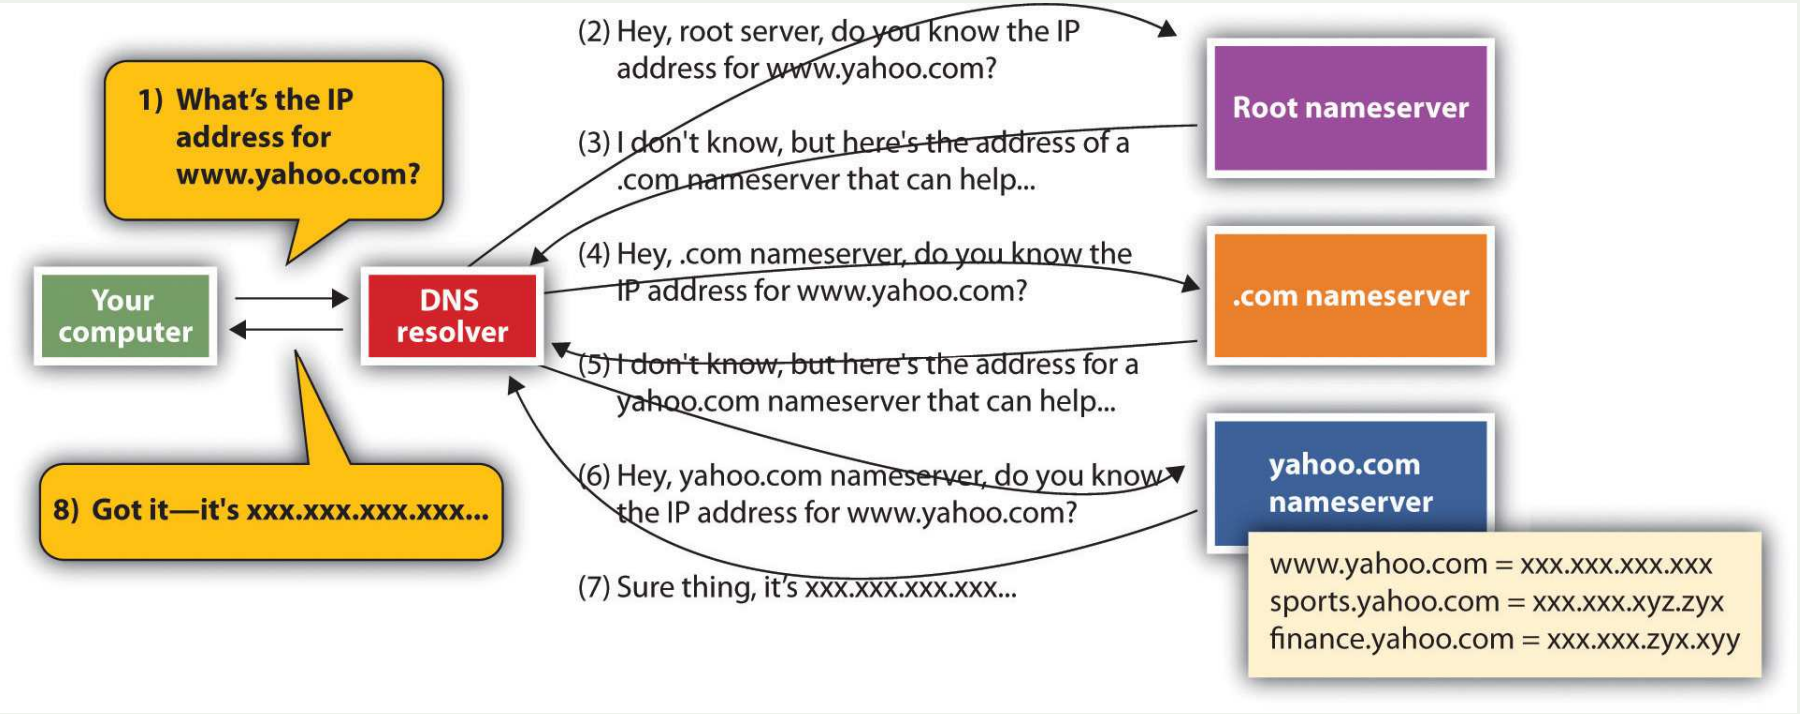

Information Systems by University of Minnesota

How Internet works?

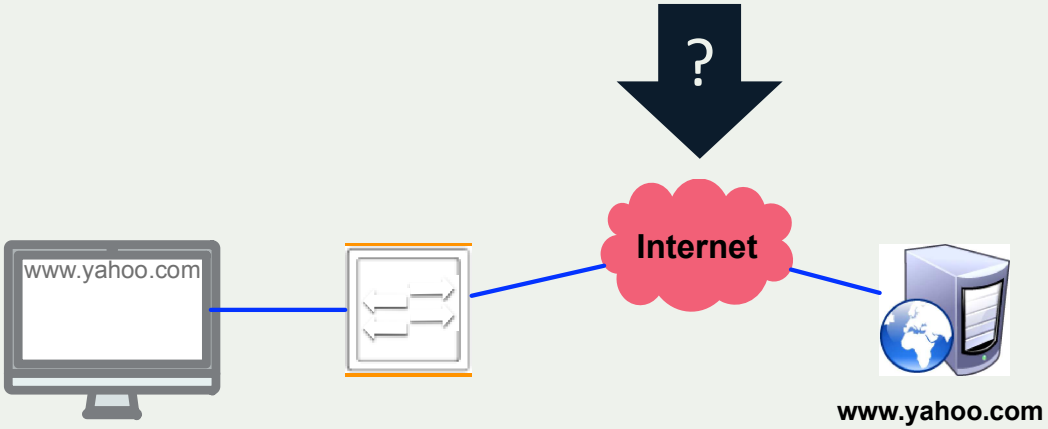

The Basics: How Routing Works

There are ~69,000 networks (Autonomous Systems) across the Internet, each using a unique Autonomous System Number (ASN) to identify itself to other networks.

Routers use Border Gateway Protocol (BGP) to exchange “reachability information” - networks they know how to reach.

Routers build a “routing table” and pick the best route when sending a packet, typically based on the shortest path.

Internet Growth – Routing Table

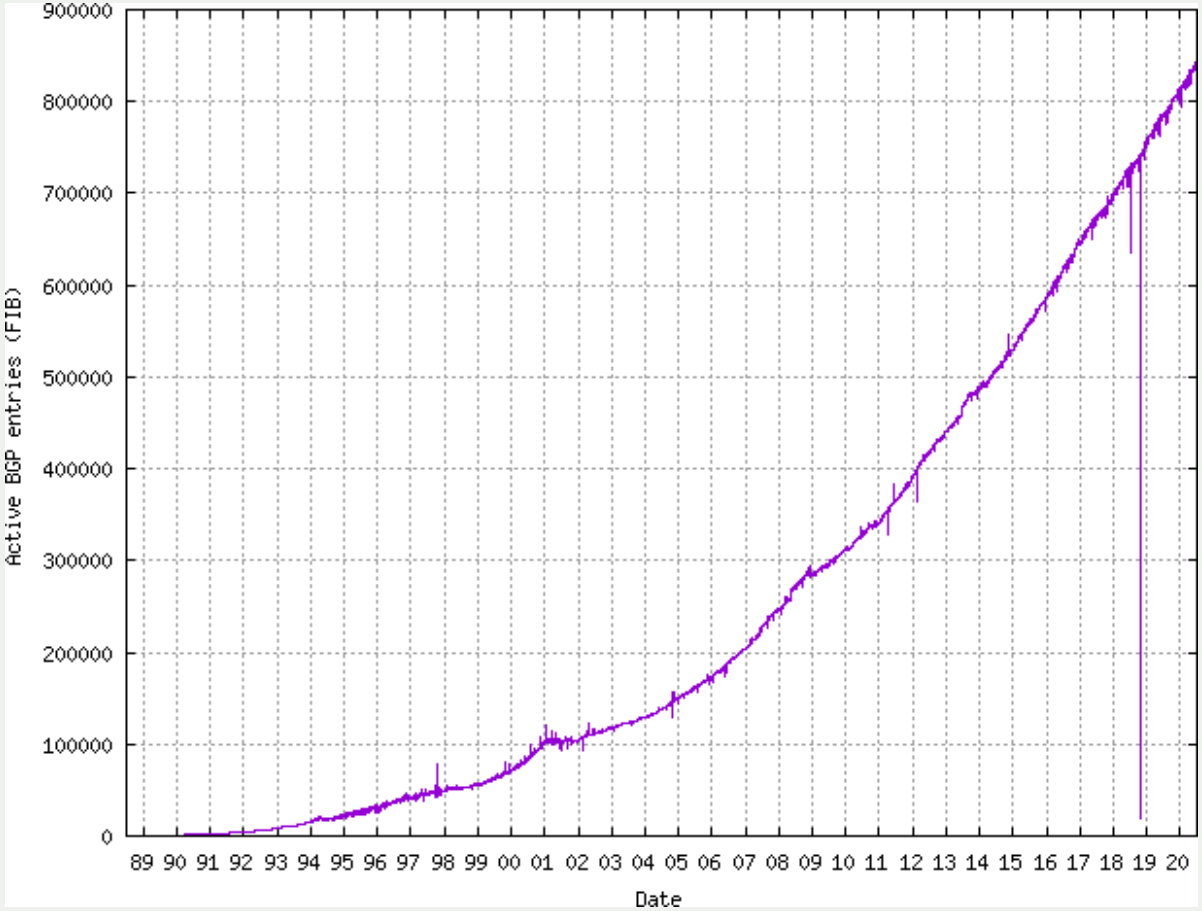

www.cidr-report.org

Internet Growth – ASN (Networks)

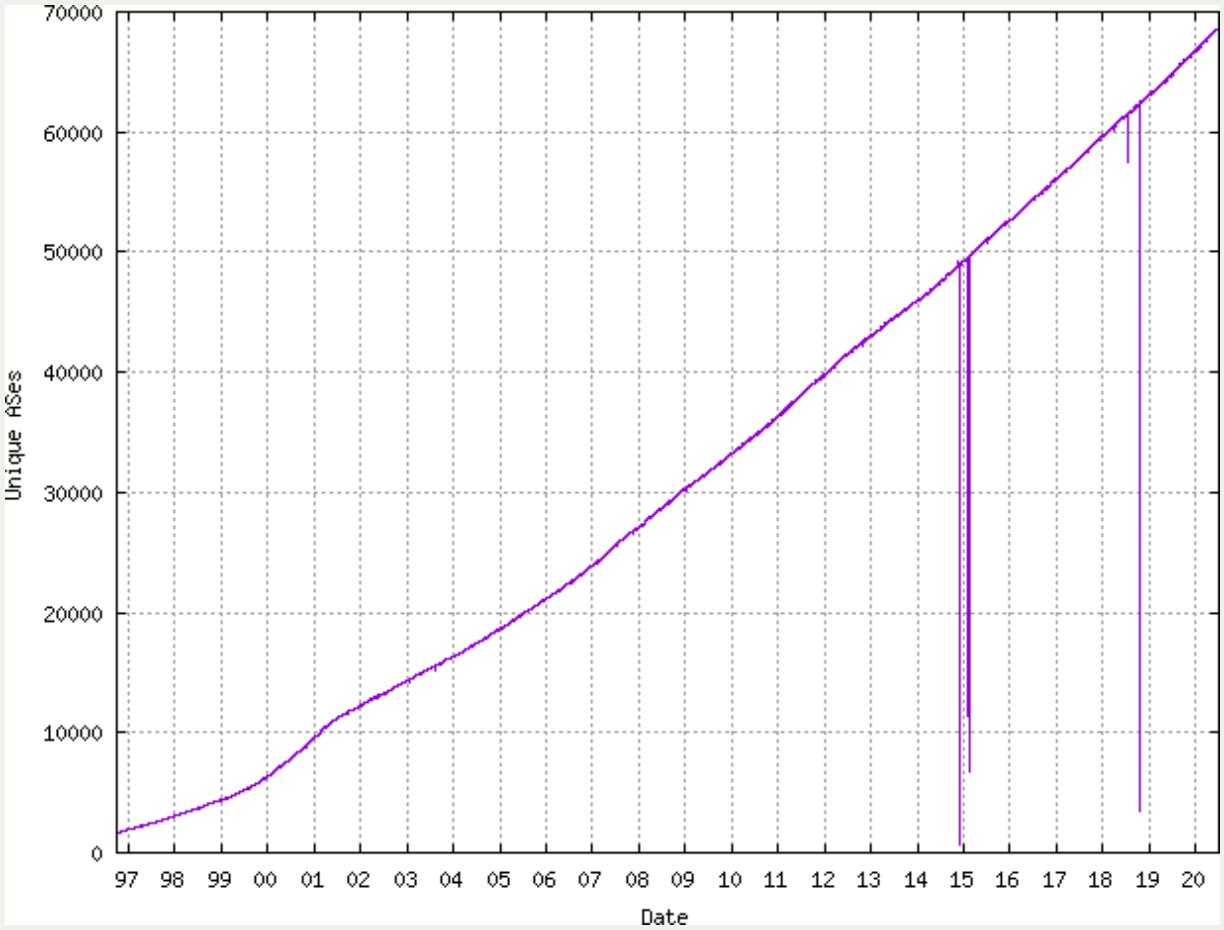

www.cidr-report.org

Some Definitions

Router

find path

forward packet, forward packet, forward packet, forward packet…. Something wrong…

find alternate path

forward packet, forward packet, forward packet, forward packet

repeat until powered off

Some Definitions

Routing vs Forwarding

Routing = building maps and giving directions

Forwarding = moving packets between interfaces according to the "directions"

Routing Protocol

A protocol implemented by network routers to exchange the information necessary to determine the best route for data transfer. [Oxford Dictionary of Computing].

Dynamic routing as its name suggest dynamically discover network destinations and how to get to them. It allows routing tables in routers to change if a route on the router goes down or if a new network is added. Dynamic Routing protocols can be classified into different groups according to their characteristics.

- **Function**:
  - Interior Gateway Protocol (IGP) [RIP, EIGRP, OSPF, IS-IS]
  - Exterior Gateway Protocol (EGP) [BGP]
- **Process**:
  - Distance vector protocol [RIP, EIGRP]
  - Link-state protocol [OSPF, IS-IS]
  - Path-vector protocol [BGP]

Routing Protocol

IGP (Interior Routing Protocol): It is used for exchanging routing information between routers within an autonomous system (for example, a system of corporate local area networks). [OSPF, RIP, ISIS]

EGP (Exterior Routing Protocol): Exterior Gateway Protocols handle routing outside an Autonomous System. It is used for exchanging routing information between two neighbor routers (in different Autonomous System). EGP is commonly used between router on the Internet to exchange routing table information. [BGP]

Autonomous System (AS): It is a group of networks under a single administrative control which could be an Internet Service Provider (ISP) or a large Enterprise Organization.

How Routing Works?

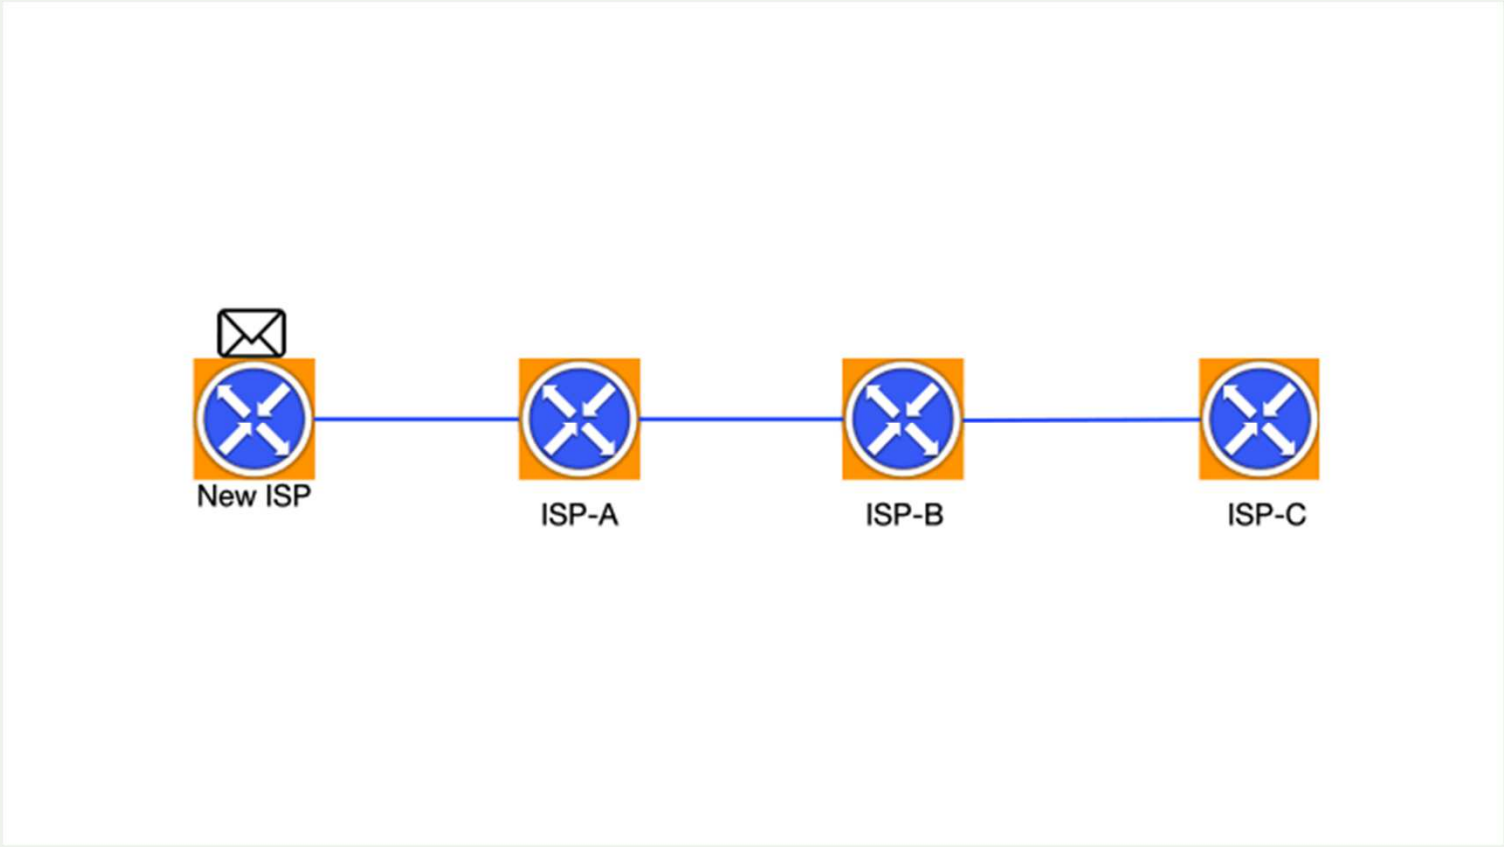

Transit, Peering and the Border Gateway Protocol (BGP)

Some Definitions

Peering – exchanging routing information and traffic directly with peer, usually free

Transit – carrying traffic across a network, usually for a fee

Default – where to send traffic when there is no explicit match in the routing table

DFZ – Default Free Zone

In the context of Internet routing, the **default-free zone** (DFZ) refers to the collection of all Internet autonomous systems (AS) that do not require a **default** route to route a packet to any destination.

Internet Connectivity: Transit

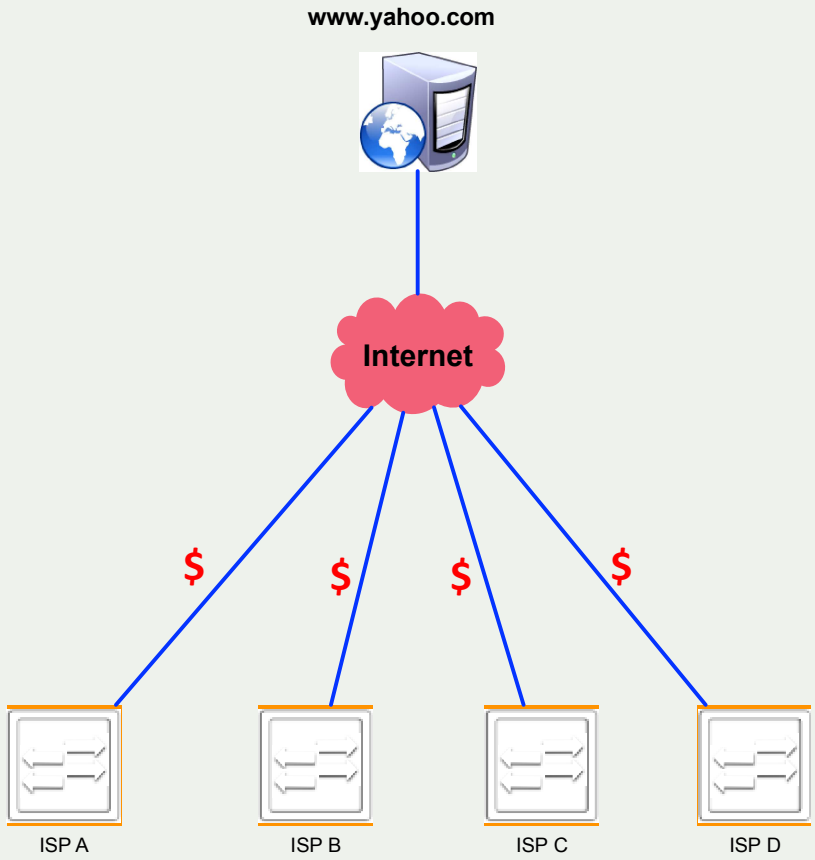

Internet Connectivity: Transit

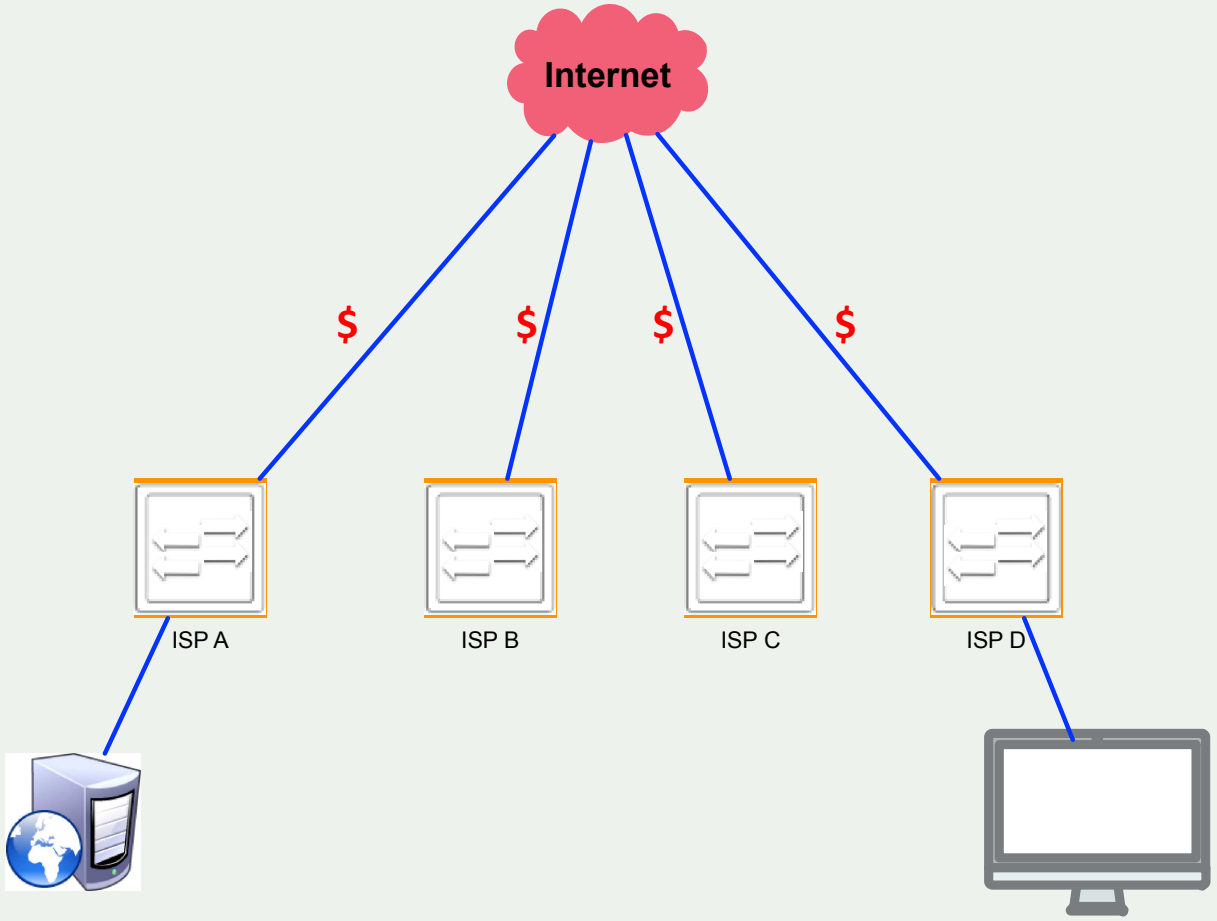

Internet Connectivity: Peering

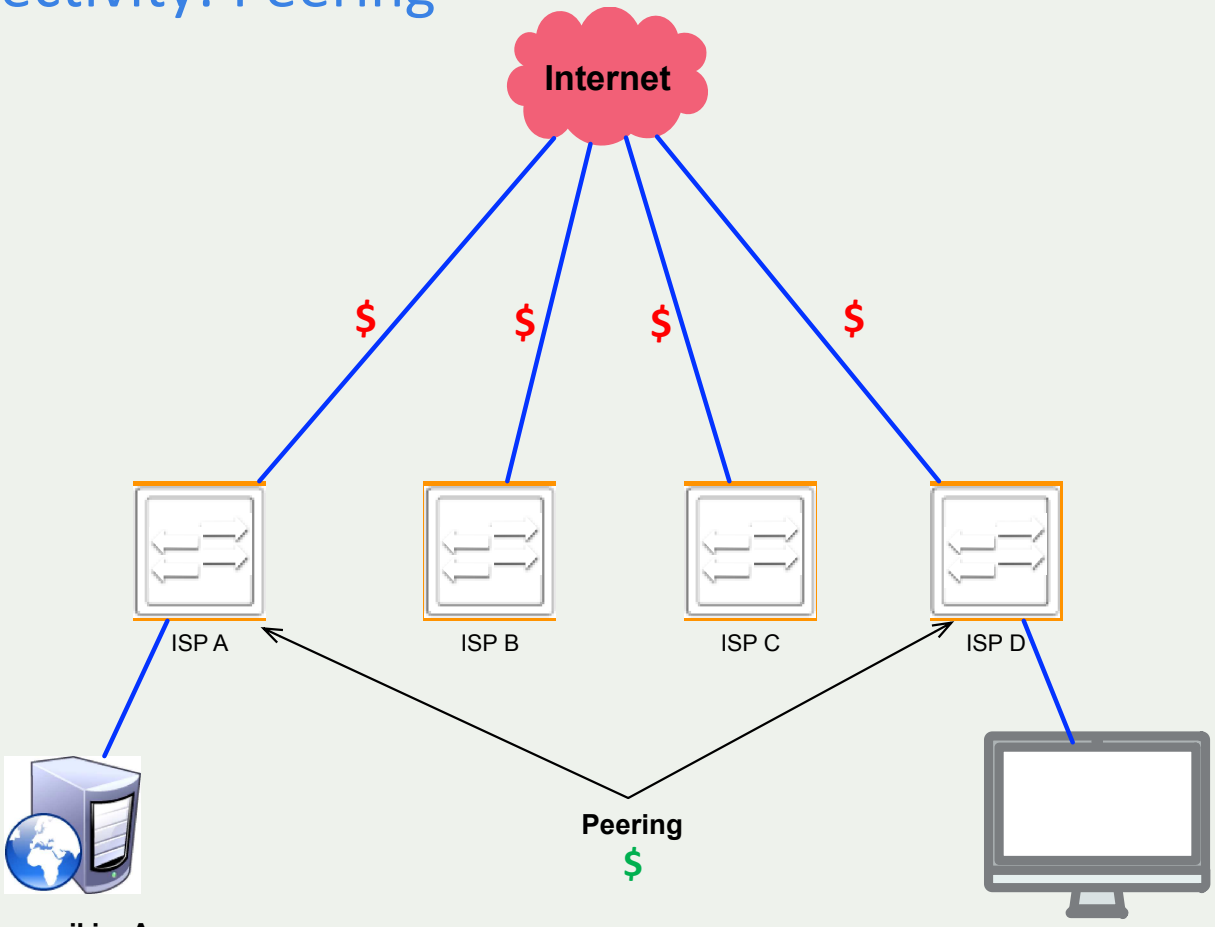

Transit vs Peering

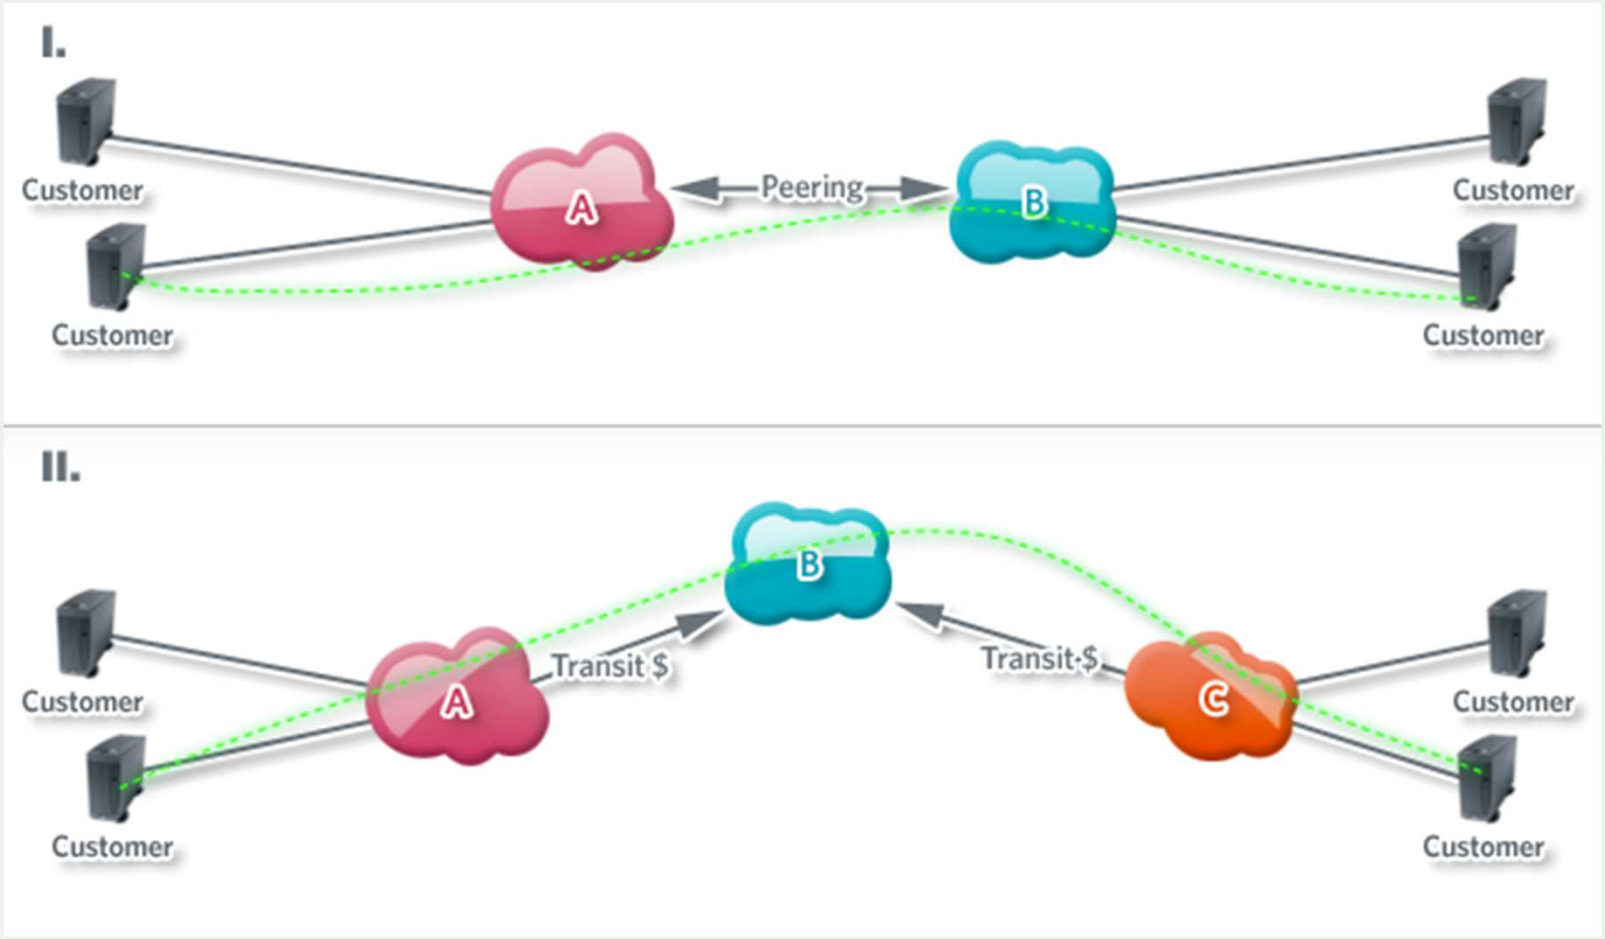

arstechnica.com

Transit vs Peering

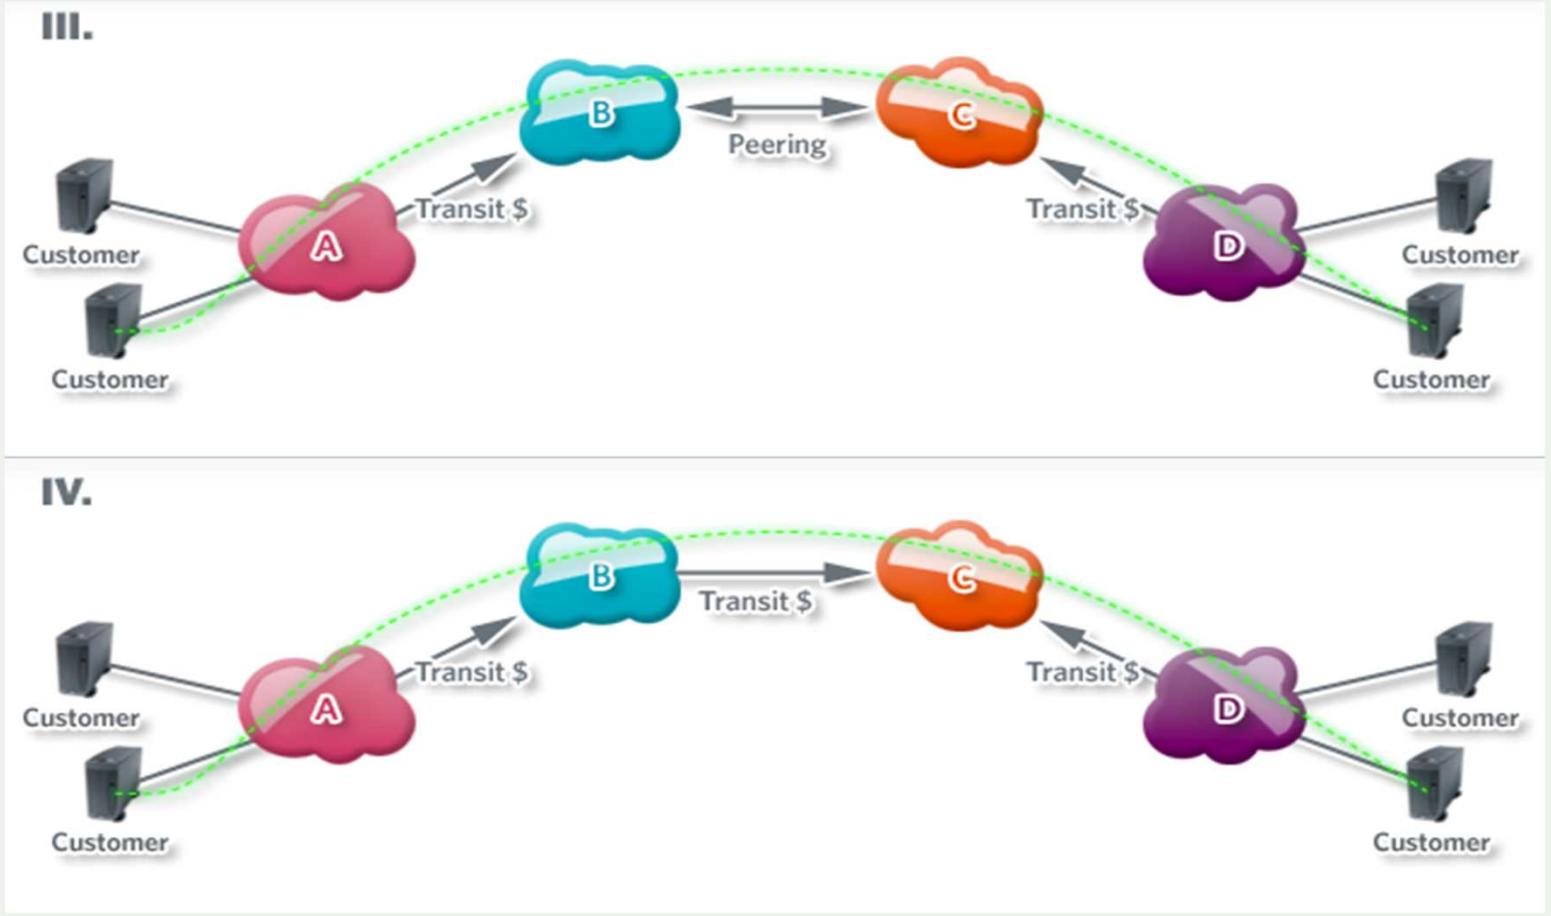

arstechnica.com

Border Gateway Protocol (BGP)

Border Gateway Protocol (**BGP**) is a standardized exterior gateway protocol designed to exchange routing and reachability information among autonomous systems (AS) on the Internet.

BGP exchanges routing and reachability information between autonomous systems (ASN) on the Internet. Each **BGP** speaker (Router), which is called a “peer” in BGP terms, exchanges routing information (routing table) with its neighbors (peers) in the form of prefix announcements.

BGP – Border Gateway Protocol

Path-vector Routing Protocols: A path vector protocol defines a route as a pairing between a destination and the attributes of the path to that destination.

1.0.166.0/24 203.202.143.34 0 7474 7473 38040 23969 } Routing Information

Some Definitions

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

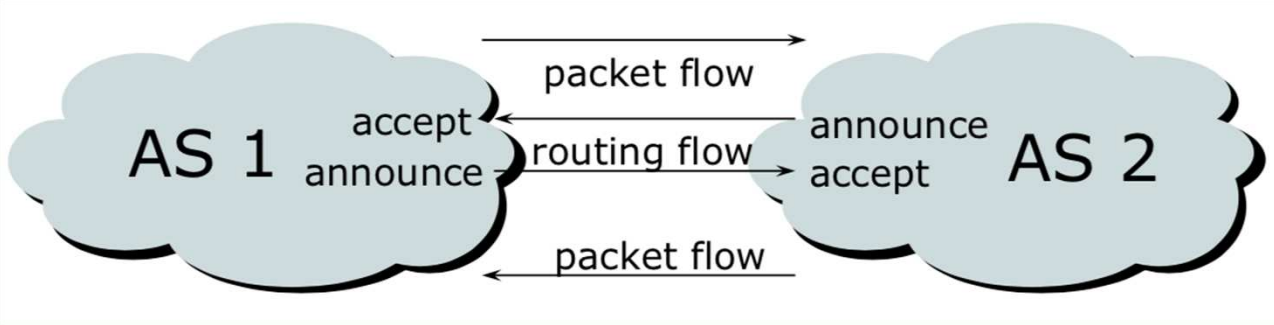

BGP – Border Gateway Protocol

General Operation

- Learns multiple paths via internal and external BGP speakers
- Picks the best path and installs it in the routing table (RIB)
- Best path is sent to external BGP neighbours
- Policies are applied by influencing the best path selection

BGP – Border Gateway Protocol

Constructing the Forwarding Table

- BGP “in” process
  - receives path information from peers
  - results of BGP path selection placed in the BGP table n “best path” flagged
- BGP “out” process
  - announces “best path” information to peers
  - Best path stored in Routing Table (RIB)
- Best paths in the RIB are installed in forwarding table (FIB) if:
  - prefix and prefix length are unique
  - lowest “protocol distance”

How an IXP Works

What is it?

Simply defined, it is a place where ISPs meet to exchange IP traffic via BGP (peering)

IXPs (Internet Exchange Points) are an access solution that wouldn't be possible without people working together. From the very start, an IXP is the result of partnership, collaboration, and trust. They represent the very best of what can happen when people work in the same way the Internet does – interconnecting networks and people.

How it works?

For an IXP to work it needs: a switch, routers, servers, a neutral locations, appropriate power sources, cooling, security, and technical experts to run and manage the IXP.

Bilateral Peering (without IXP)

All ISPs can peer with each bilaterally without an IXP, but the only problem is that it doesn't scale. If you add 100 more ISPs to peer then it will be hard to manage and certainly not cost effective

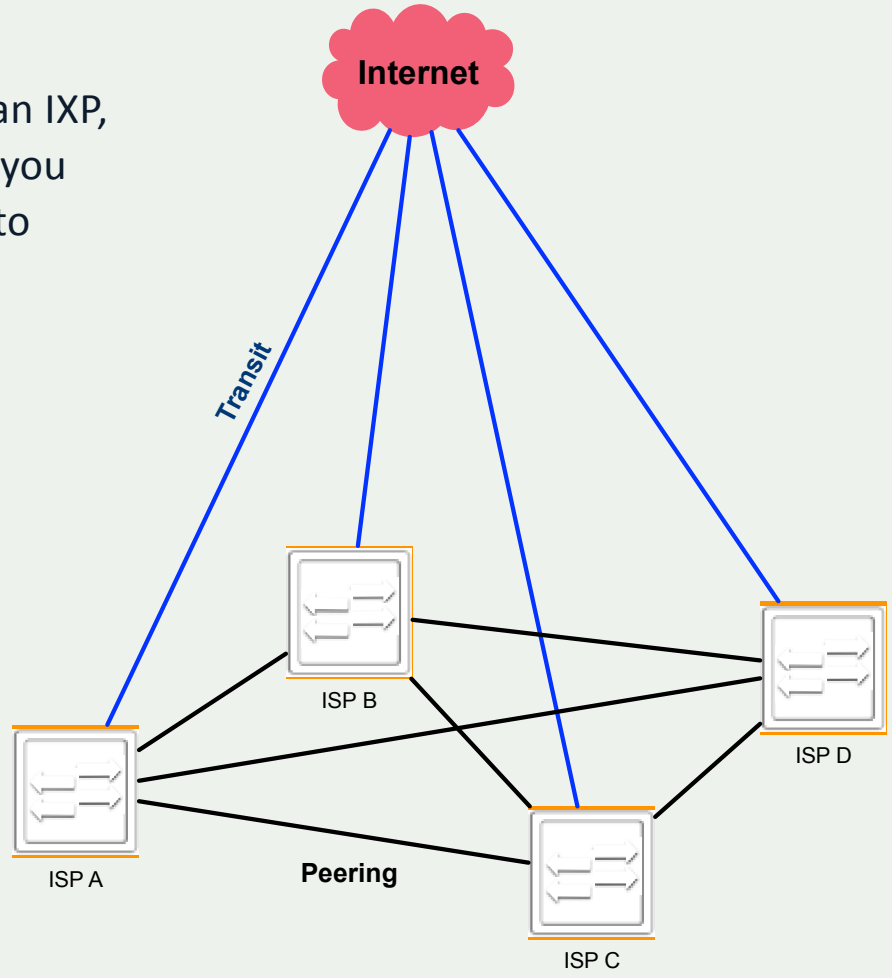

Peering with IXP

If you are connected to an IXP, then its easier to just peer with the route-server. Also you can pick the peers for bilateral peering using the same IXP fabric and physical connectivity.

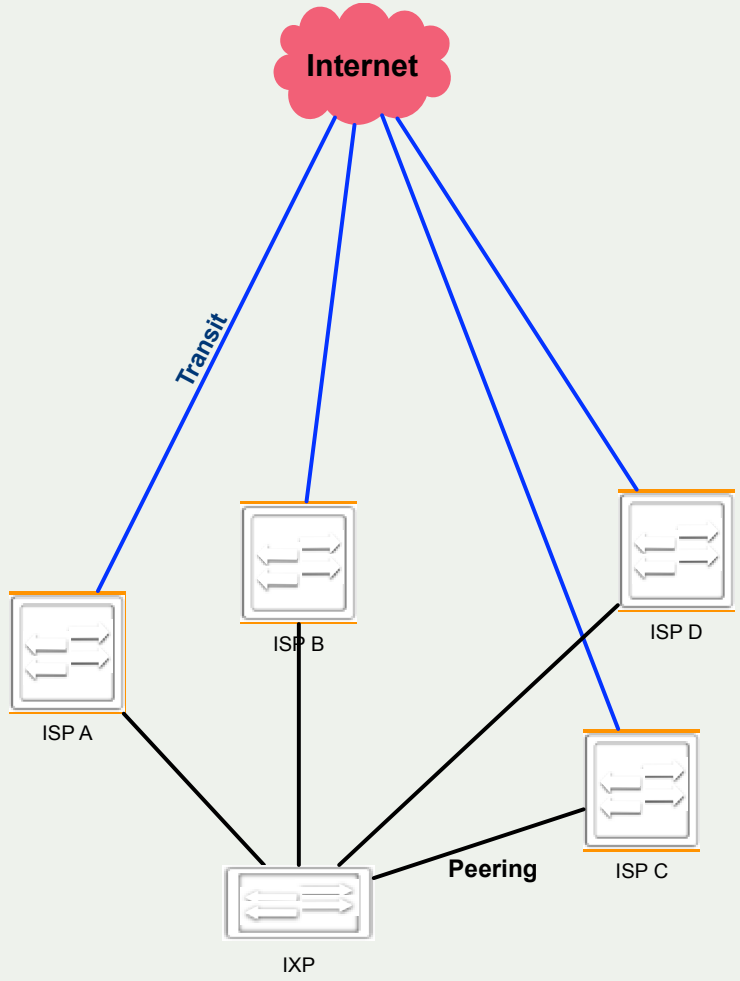

How it looks like?

An Internet Exchange Point (IXP) is a physical and usually neutral location where different IP networks meet to exchange local traffic via a switch.

Different Kinds of IXPs

IXPs fall roughly into five categories:

- non-profit organizations,
- associations of ISPs,
- operator-neutral for-profit companies,
- university or government agencies, and
- informal associations of networks.

Where to establish an IXP

- The location of the IXP is very important.
- The IXP location should be neutral and low cost.
- In considering the IXP location the following factors should be considered;
  - Space
  - Environmental Control
  - Security
  - Power
  - Access to terrestrial Infrastructure
  - Cabling
  - Support

Who should own it?

- Most IXPs are owned and managed neutrally with respect to all operators (members and non-members).
- Many ISPs have expressed strong feelings about the importance of neutrality of IXPs.
- IXPs generally abstain from carrying out any activity that may compete with member business activities or opportunities.
- If an IXP competes with members/customers it could lose their support.
- The key point is that the **ownership** and **management** of the IXP should always remain neutral.

Who should own it?

- Many IXPs begin with donations of equipment, rack space, labour and other assistance - that is part of the co-operative nature of most start up IXPs.
- In case of donations, written agreements are necessary to define the transaction and ownership thereafter.
- Neutrality can be at stake if one member owns parts of the IXP
- Therefore the IXP should always maintain ownership and responsibility of it’s assets.

Business Case

Internet exchange points (IXPs) can improve Internet quality and affordability in local communities. IXPs help strengthen local Internet connectivity, develop local Internet industry, improve competitiveness, and serve as a hub for technical activity.

IXP Peering and Interconnection Policies

- **Open peering Policy**
  - Develop an open peering policy to encourage non-traditional members such as CDNs, Government, Academia, Banks, etc to peer
  - Initiate Strategies to grow membership – marketing, public seminars, “tell the story why is this a good business
- **Regional Interconnection Policy**
  - Encourage members to exploit the cross-border interconnection opportunities by negotiating fair contracts with Infrastructure Operators and International bandwidth providers.
  - Assist operators and members take advantage of regional Interconnection opportunities and become Regional carriers
- **Transit Policy**
  - The ability to attract carriers and transit providers at an IXP is important to grow the value and traffic at an IXP.
  - This policy is subject to national regulations on Internet transit.

IXP Peering Agreements

A Peering agreement is applicable to all members who choose to peer an IXP. There are three main peering agreements implemented by IXPs

- **Bilateral Peering (BLPA)**
  - This agreement requires every member to have a bilateral agreement with each member at the IXP.
  - It is commonly found in most developed IXPs due to the right of refusal, to peer granted to each member, to peer with any member for commercial purposes
- **Multi-lateral Peering (MLPA)**
  - This agreement requires every member to peer multi-laterally with every member at an IXP. The agreement is not often enforced.
  - The MLPA is useful in trying to grow peering value at an IXP
  - A Mandatory version of this agreement (MMLPA) is commonly found in startup IXP in developing worlds due to its ability to force incumbents to peer with others at the IXP. MMLPA is enforced using a Route-Server IXP Topology
- **Hybrid Peering Agreement (HPA)**
  - This agreement supports both the BLPA and MLPA.
  - Eliminates the need for a competing IXP in the same location to serve unsupported policy.
  - A Route-Server is often used for MLPA members.

Innovation Catalyst

- The IX serves as a catalyst for innovation of new products and services for the local community.
- Therefore stakeholders should promote initiatives that encourage local innovation of content and application development
- Implement monitoring tools to provide traffic stats for the benefit of members and measurement of growth over period of time for researcher

Capacity Building Activities

- Members to understand economics of Peering and Transit, the IXP value curve and value of peering.
- Neutralizing the asymmetry of information between members and the rest of the world.
- Facilitate the development of a pool of skilled peering coordinators and network routing engineers.
- Ensuring that the IXP and the members operate stable and robust network infrastructure services based on best practices.

Security Considerations

Routing Security

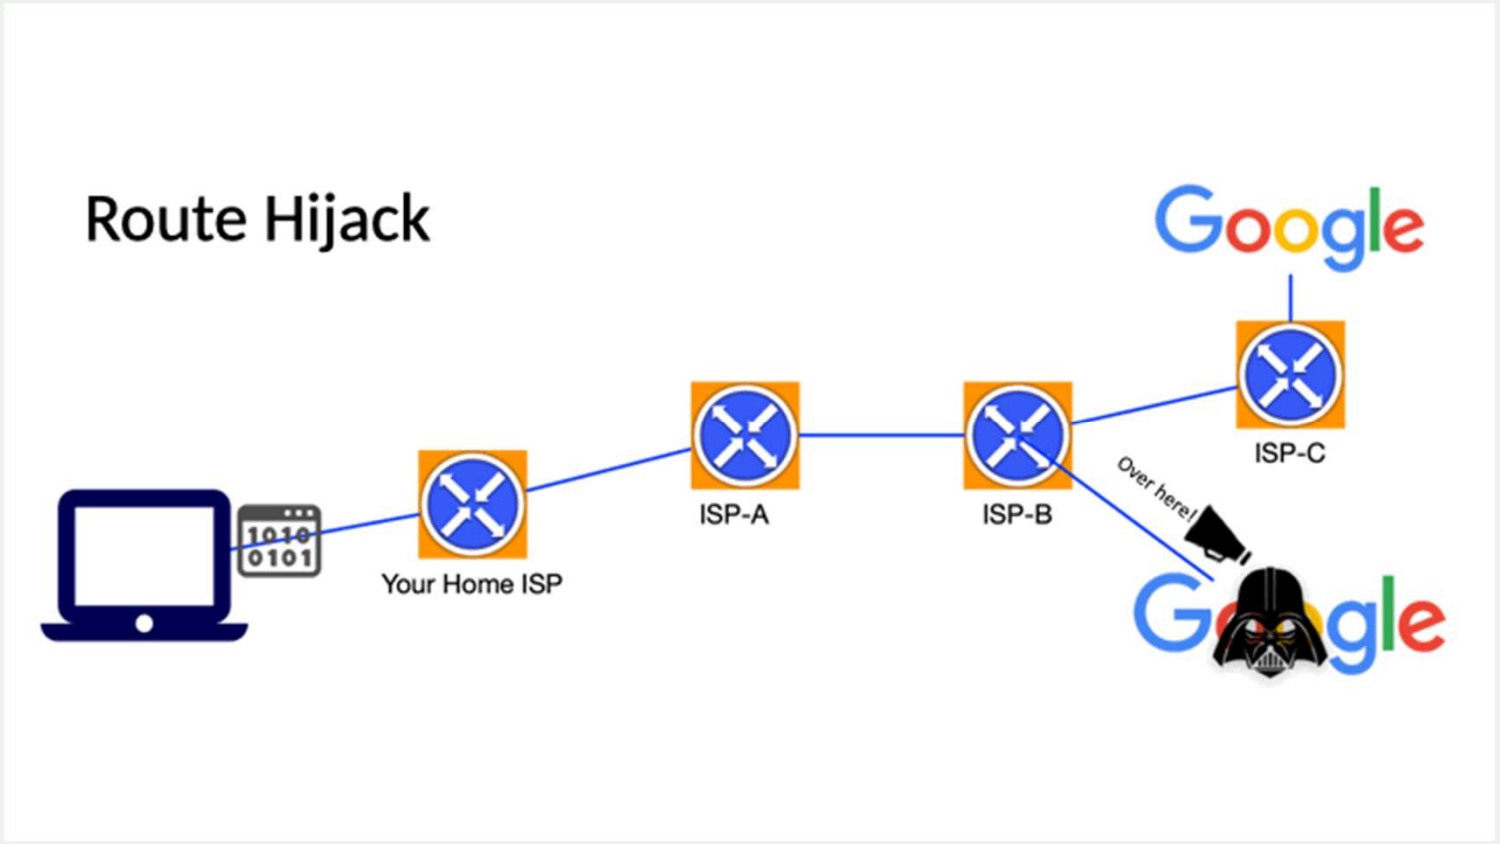

Routing Security

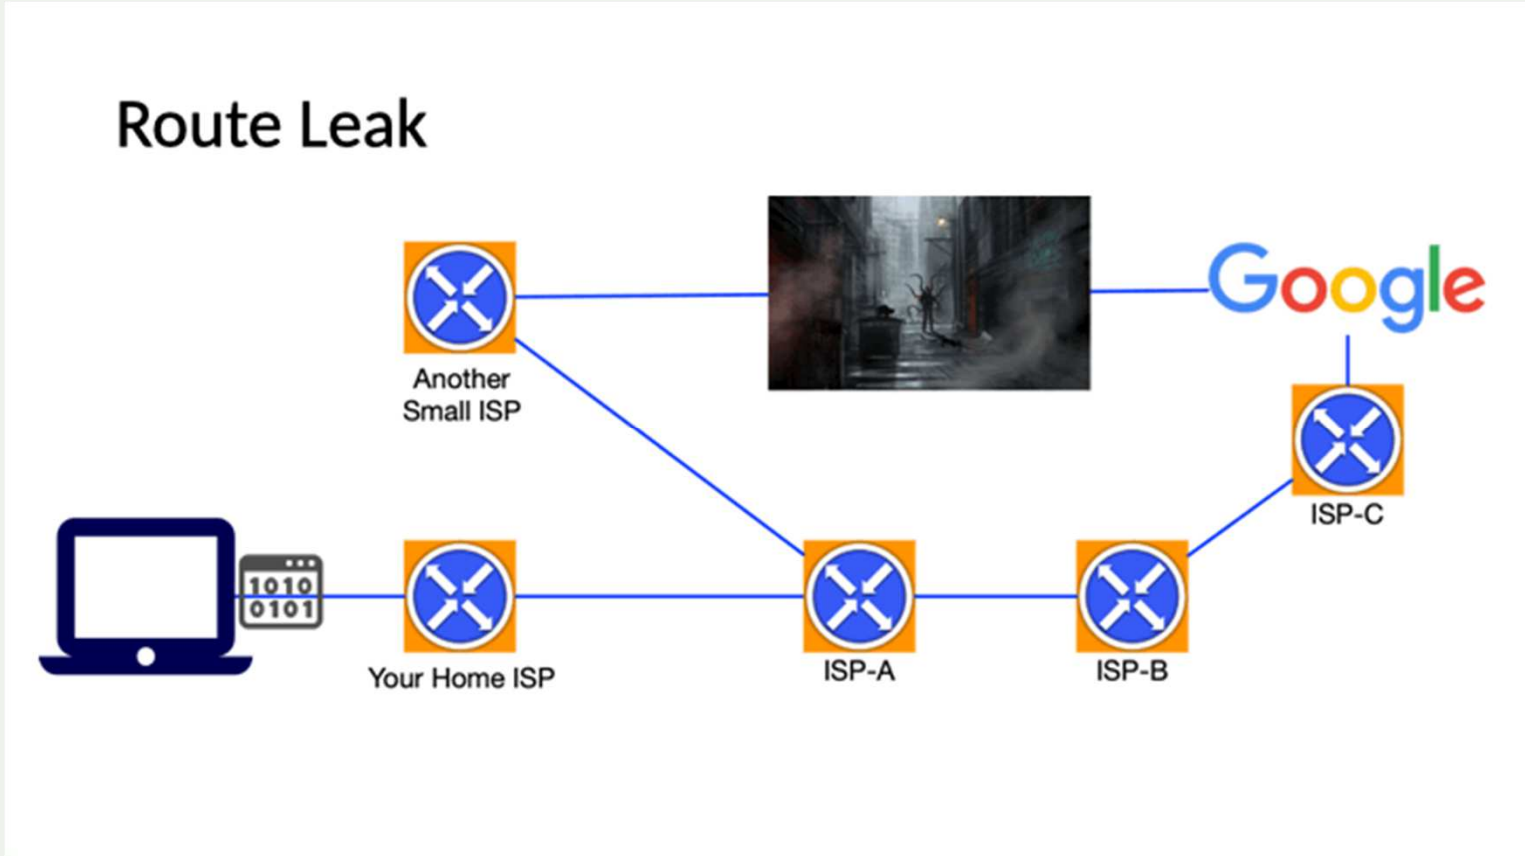

Security at IXP

The IXP implements filtering of route announcements at the Route Server based on routing information data (IRR and/or RPKI). Based on the outcome of the validation process, the invalid announcements are filtered in accordance with the IXP published policy.

IXPs using a Route Server to facilitate multilateral peering should use it to validate received route announcements from a peer and subsequently filter them to other peers. Special purpose cases, such as research projects, are out of scope for this requirement.

Security at IXP

The IXP should have a published policy of traffic not allowed on the peering fabric and performs filtering of such traffic.

Commonly, filtering applies to:

- Not allowed Ethernet frame formats
- Not allowed Ethertypes
- Link-local protocols, such as IRDP, ICMP redirects, Discovery protocols (CDP, EDP), VLAN/trunking protocols (VTP, DTP), BOOTP/DHCP, etc.
- Restricted by the MAC port security configuration

While not strictly routing, applying hygiene on Layer 2 can ensure the smooth operation of the platform and contribute to the stability of the IXP infrastructure and routing.

Best Practices - Discussion

MANRS

Mutually Agreed Norms for Routing Security (MANRS) is a global initiative, supported by the Internet Society, that provides crucial fixes to reduce the most common routing threats.

MANRS outlines four simple but concrete actions that network operators should take:

Filtering – Ensure the correctness of your own announcements and of announcements from your customers to adjacent networks with prefix and AS-path granularity

Anti-spoofing – Enable source address validation for at least single-homed stub customer networks, your own end-users, and infrastructure

Coordination – Maintain globally accessible up-to-date contact information

Global Validation – Publish your data, so others can validate routing information on a global scale

http://www.manrs.org

Discussion Points

- Neutrality - Location
- Ownership
- Business Model
- Governance
- More…

Discussion Points

- Support IXPs and speed their development with information and communications technology (ICT) policy objectives that promote an enabling environment for interconnection via policy and regulatory frameworks.
- Provide as much policy and regulatory transparency as possible to encourage regional and international entities to participate in the local interconnection and peering environment.

Discussion Points

- Encourage competitive access to wired and wireless connections, which will help lower the costs associated with connecting to an IXP.
- Promote local investment opportunities via tax holidays, and reduced duties on the equipment needed to build IXPs and operator networks (e.g., switches, routers, and servers). Provide clear guidance about local business rules and practices.

Discussion Points

- Foster relationships with IXPs to learn more about local interconnection environments and the sustainability and technical management of an IXP.
- Work with existing IXPs and expert organizations to avoid the mistakes that other IXPs have made, obtain start-up assistance and equipment donations, and learn more about training and human-capacity development opportunities.

Acknowledgement and Attribution

Some content and information in this presentation originally developed and maintained by the following

Dr Philip Smith (NSRC)

Mark Tinka

APNIC Training Team

CIDR Report

Thank you.

Aftab Siddiqui

Sr. Internet Technology Manager
siddiqui@isoc.org

Quai de l'île 13
CH-1204 Geneva
Switzerland

11710 Plaza America Drive
Suite 400
Reston, VA 20190, USA

Rambla Republica de Mexico 6125
11000 Montevideo,
Uruguay

66 Centrepoint Drive
Nepean, Ontario, K2G 6J5
Canada

Sin El Fil, Dekwaneh Highway
Aramex Building, 2nd Floor
Beirut, Lebanon

Science Park 400
1098 XH Amsterdam
Netherlands

internetsociety.org
@internetsociety

9 Temasek Boulevard
#09-01 Suntec Tower Two
Singapore 038989