

Internet Exchange Point

**Gaurab Raj Upadhaya (Limelight Networks) &
Philip Smith (Cisco)**

MENOG 7

Istanbul, 21st October 2010

What is an Internet eXchange Point (IXP) ?

- Internet eXchange Points (IXPs) are the most critical part of the Internet's Infrastructure. It is the meeting point where ISPs interconnect with one another. With out IXPs, there would be no Internet. Interconnecting with other networks is the essence of the Internet. ISPs must interconnect with other networks to provide Internet services.
- Private and Bi-Lateral Peering are considered to be a type of IXP.

Background

- The Internet is a decentralized network of autonomous commercial interests
- Internet Service Providers (ISPs) operate by exchanging traffic at their borders, propagating data from its source to its destination
- This exchange can be settlement-free (“Peering”) or paid (“Transit”)

Why This is Important

- If you have no domestic Internet exchange facility, your domestic ISPs must purchase transit from foreign ISPs
- The large foreign ISPs who sell transit are American, Japanese, and British
- This is an expensive and unnecessary exportation of capital to developed nations at the expense of your domestic Internet industry

Second-Order Benefits of Domestic Exchange

- A strong domestic Internet industry creates high-paying knowledge-worker jobs
- Domestic traffic exchange reduces the importation of Foreign content and cultural values, in favor of domestic content authoring and publishing

A Brief History of Internet Exchanges

First Exchanges

- Metropolitan Area Ethernet

Washington, D.C.

10mb shared FOIRL into assorted switches

No fixed topology

MFS fiber plant

Shared administration

First Exchanges

- Commercial Internet Exchange

Moved from Washington, D.C. to Palo Alto

Layer-3 MMLPA

Commodity DS1 (T1) lines into a Cisco 7010

Not-for-profit industry association

First Exchanges

- MAE-West / Federal Internet Exchange

San Jose / Mountain View

FDDI “dumbbell” ring

Bridged to 10mb Ethernet in many locations

Two locations, two administrations

First Exchanges

- Hong Kong Internet Exchange
 - Chinese University of Hong Kong
 - Single location Ethernet switch
 - Administered by the university
 - First major free exchange

Technological Progression

- Shared 10Base-T / FOIRL Ethernet
- Switched 10mb Ethernet
- Shared FDDI
- Switched FDDI
- 100Base-T / 100Base-FX
- Gigabit Ethernet
- 10Gigabit Ethernet

Other Technologies

- Layer-3 route-servers
- Frame Relay
- ATM
- Wireless Ethernet
- Crossconnect mesh
- DPT

Common Services

- Route-server
- Looking-glass
- Measurement and instrumentation
- Network Time Protocol
- Web cache parent
- News server
- Root server mirror

Common Business Models

- Hosted by a university or government
- Informal
- Industry association
- Neutral for-profit
- Anything else may not be recognized

Size Differentiation

- Municipal
- Large metro-area
- National
- “Regional” (meaning changing)
- Global

Regional Exchange

- Any IXP that has grown to become a prominent IX in a region
 - Easy access to new members
 - Liberal financial and legal framework for overseas companies to connect
- Most ISPs will connect to the domestic IX and then to the Regional IX
- Examples : JPIX, WAIX, VIX, APE, NOTA

Global IXPs

- Over time some cities and their IXPs have become large and taken a role of global IXP.
 - New York / Washington - NYIIX, Equinix-Ashburn
 - California - PAIX, LAIIX
 - London / Amsterdam - LINX, AMSIX
 - Hong Kong - HKIX

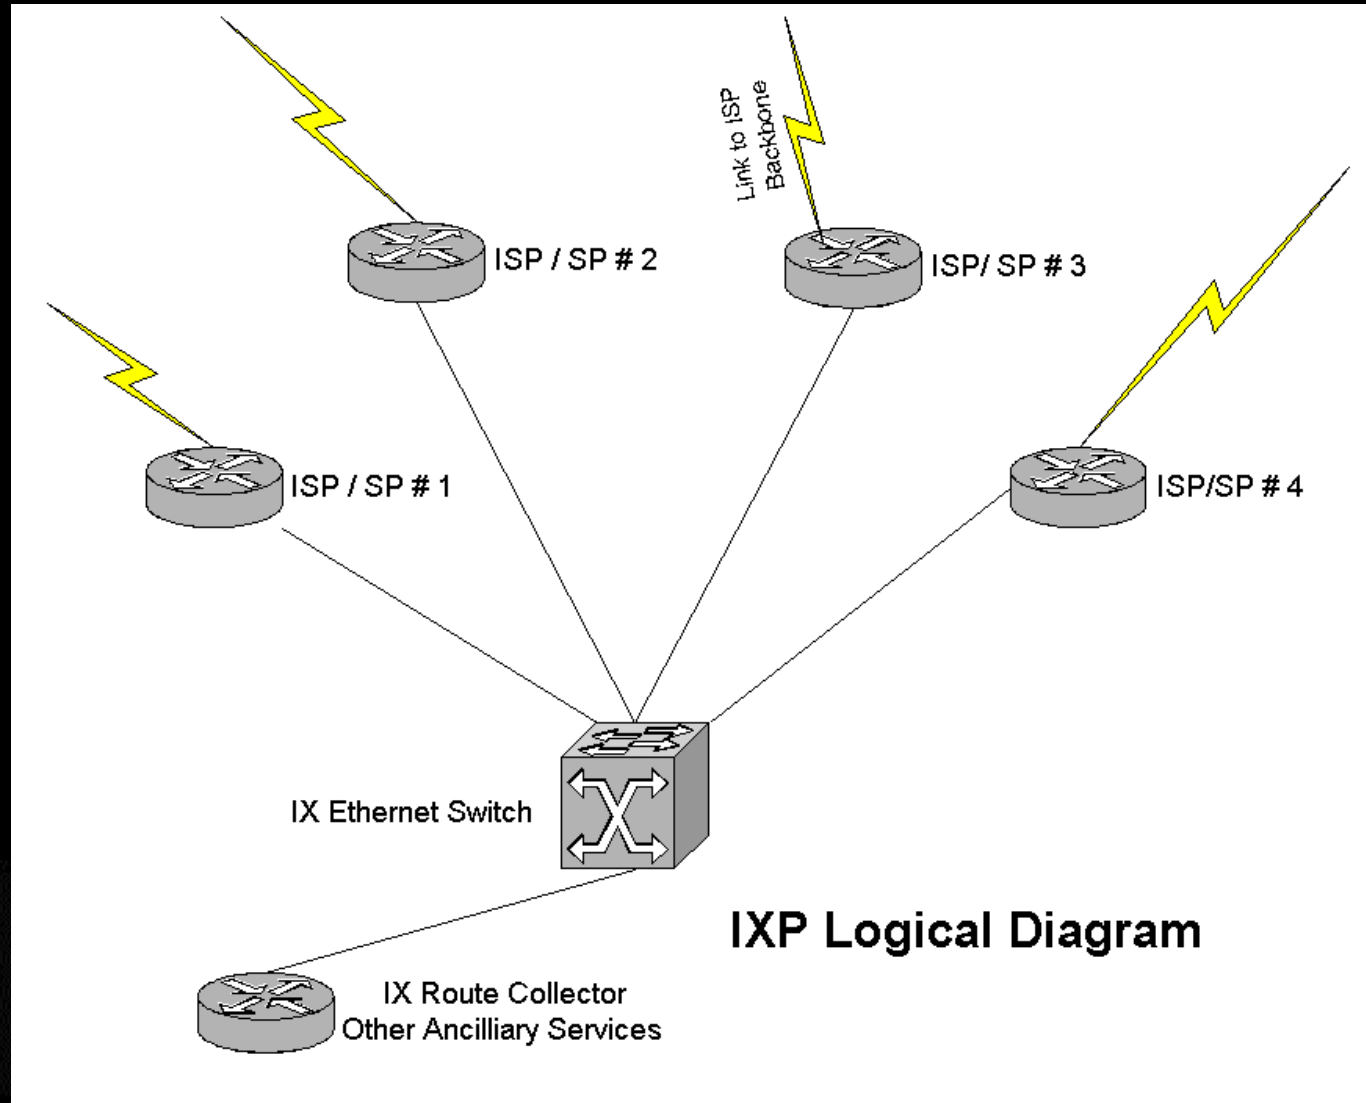
Peering / Transit Differentiation

- New concept
 - You pay when buying "transit", even if you have a BGP 'peering' session.
 - You don't pay when 'peering', even if it doesn't involve a BGP 'peering' session.
- Very different pricing
 - Free Vs. \$\$\$
- Very different competitiveness

Types of IXPs

*How can ISPs connect to each other
?*

Typical IXP.....



Types of IXPs


- Two Types of IXes commonly exists
 - Layer 2
 - Layer 3
- Other ways also practiced
 - Ring
 - Private Internet Connect
 - GigaPoP
 - Hybrid

Layer 2 IXPs

- Layer 2
 - Uses a common network medium like Ethernet, 10/100/1000Base technologies)
 - Members bring their own routers and circuits from their backbone.
 - No Transit or customer connections
 - Members of the IXP determine who they peer with. You do not have to peer with everyone.
 - MAE-E, MAE-W, LINX, AMS-IX, NSPIXP, HKIX, PAIX

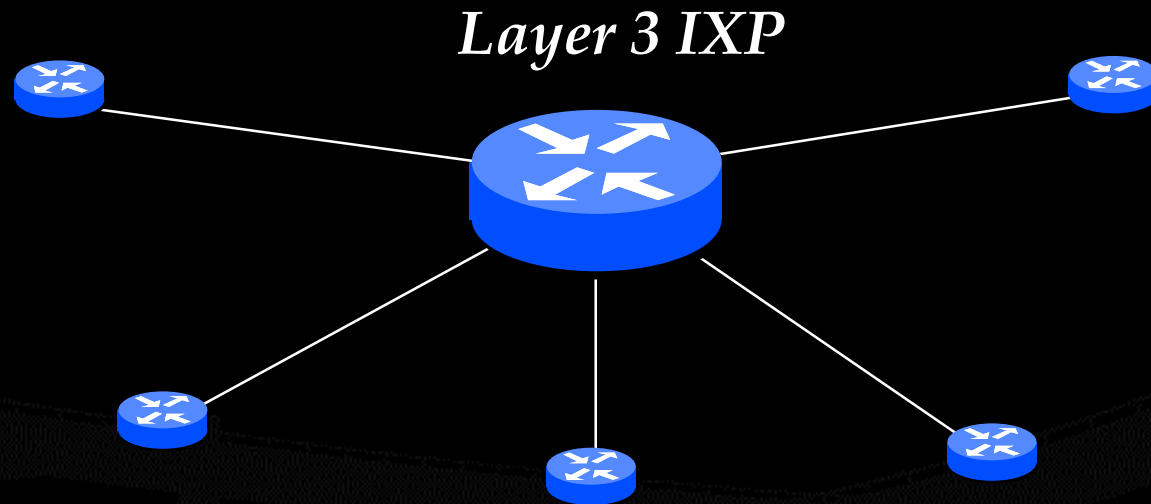
Layer 2 IXPs

- **Layer 2 Exchange** - All traffic is exchanged *outside* routers that are connected to a shared media (i.e. Ethernet 10/100/1000BaseTX)



Layer 3 IXPs

- **Layer 3 Internet eXchange Point** - All traffic is exchanged *inside* a router.



Layer 3 IXPs

- Layer 3 IXPs limit the autonomy of the members.
- Someone has to manage the router in the middle.
- Create business issues, as ISPs don't have control with whom they can peer with

Layer 2 versus Layer 3

- Layer 3
 - IXP team requires good BGP knowledge
 - Rely on 3rd party for BGP configuration
 - Less freedom on who peers with whom
 - Could potentially compete with IXP membership
 - Easier to distribute over wide area
- 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

Design and Construction of Local and Regional Exchange Facilities

Determining Need

- Sufficient end-user base?
- No existing facility to build upon?
- Sufficient degree of locally-destined traffic?

Geographic Location

- User population
- Fiber facilities or rights-of-way
- Founding participants

Density

- Centralized in one room
- Campus of adjacent buildings
- MAN
- Frame or ATM cloud

Building Management

- Telco hotel
- University computing or telecommunications facility
- City emergency services facility

In-Building Facilities

- Pathways
- Power
- Cooling
- Access and security

Services

- Switch fabric
- Crossconnects
- Route-server
- Remote hands
- NTP
- Web caching

Business Structure

- Incorporated or unincorporated?
- Staffed or volunteer?
- Non-profit or for-profit?
- Cooperative or external ownership?
- Cost-recovery (predictive or actuals), ad-hoc, or market pricing?

Policies

- BLP, MLPA or MMPLA?
- Mandatory looking-glass?
- Routing and switch-port information public or members-only?
- Secrecy in the event of security problems, failures, or mistakes
- Extensible switch fabric?

Common scenario for Peering

- ASN is required for BGP
 - APNIC/RIPE NCC members can get one for free
 - Non members can make a request through their upstream provider and get one for free
 - Yet, others can pay a one time US\$500 to get one (APNIC region).
- Members provide their own connectivity equipment at both ends and can have redundant links in place

IXP specific routes

- You only announce the following to your peers
 - Your 'own' addresses/routes
 - Your customer routes
- You do not announce
 - Your upstream routes, or full routing tables
 - Default Routes
 - Your other peers route
 - Things that do not belong to you

Any Question?

- It's all about experience !!

Where to go for more information?

- NOGs
 - SANOG, APRICOT, NANOG, RIPE etc.
- Peering Forums
 - Organized by IX Operators

Thanks to the following

- Bill Woodcock, Research Director, PCH
- Philip Smith, Sr. Consulting Engineer, Cisco
- Thanks to all of you as well