# An Analysis of the Development of IXPs

btNOG 5

4th June 2018

# IXP is the Keystone to E-Commerce



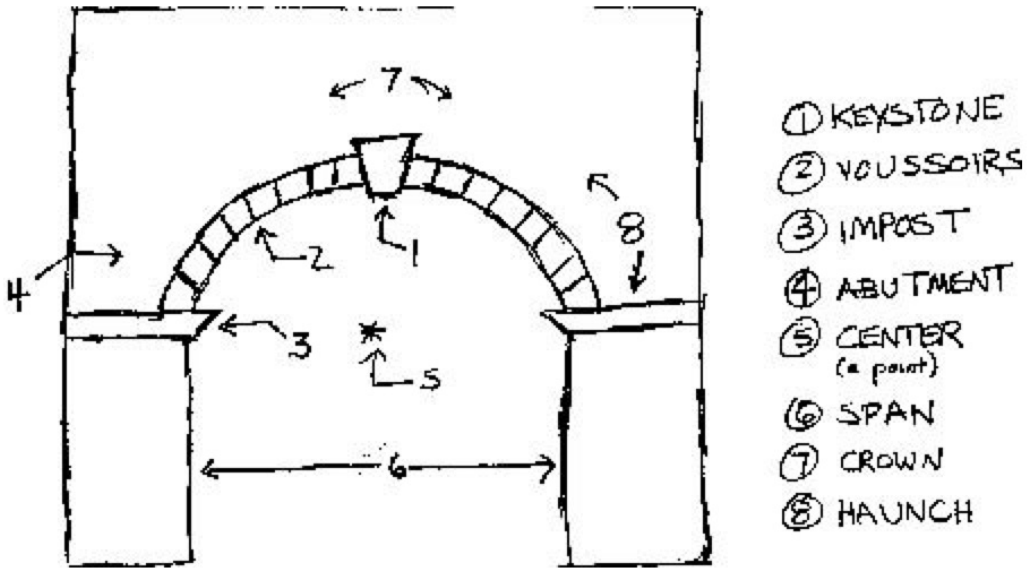


2002 Cisco ISP/IXP Workshop Slide

# Metcalfe's Law

- Metcalfe's law states:
  - the effect of a **telecommunications network** is **proportional to the square** of the number of connected users of the system ($n^2$)
  - Originally to describe ethernet, but now commonly applied to the global Internet
  - Source: https://en.wikipedia.org/wiki/Metcalfe%27s_law

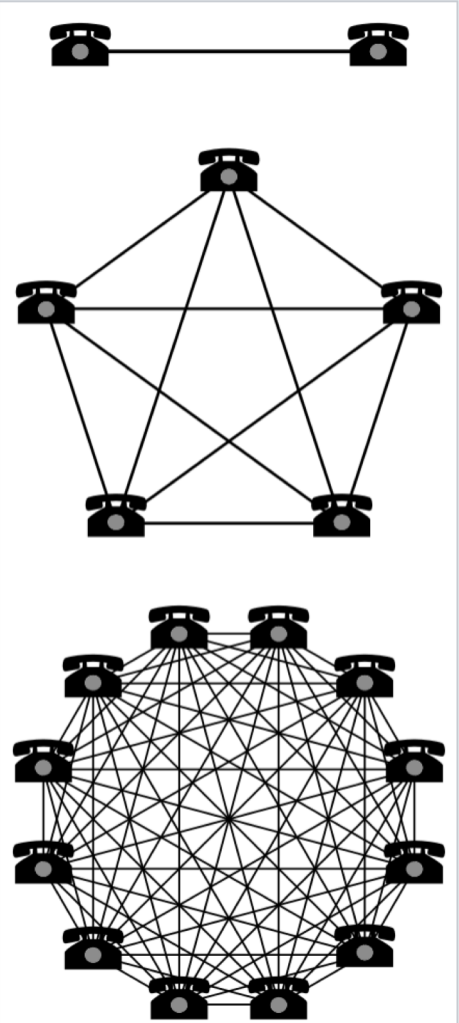

Two telephones can make only one connection, five can make 10 connections, and twelve can make 66 connections.

# N-squared Interconnect

- For large numbers of network operators, direct links to each other does not scale

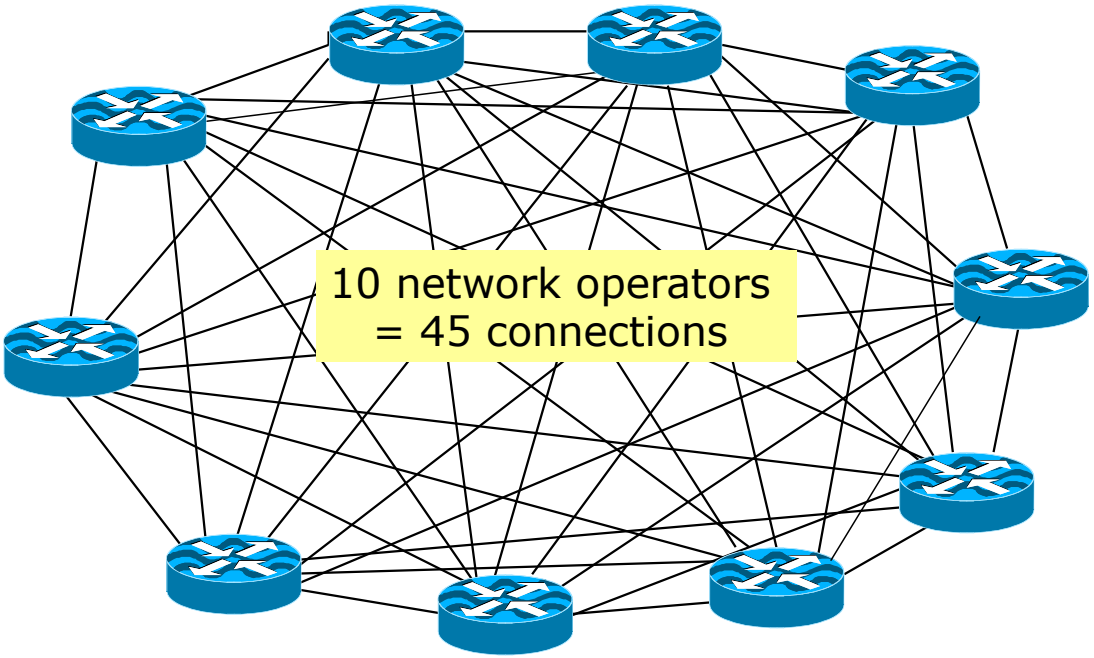

# IXP

- With an IXP the network operator routers connect to each other via the IXP fabric

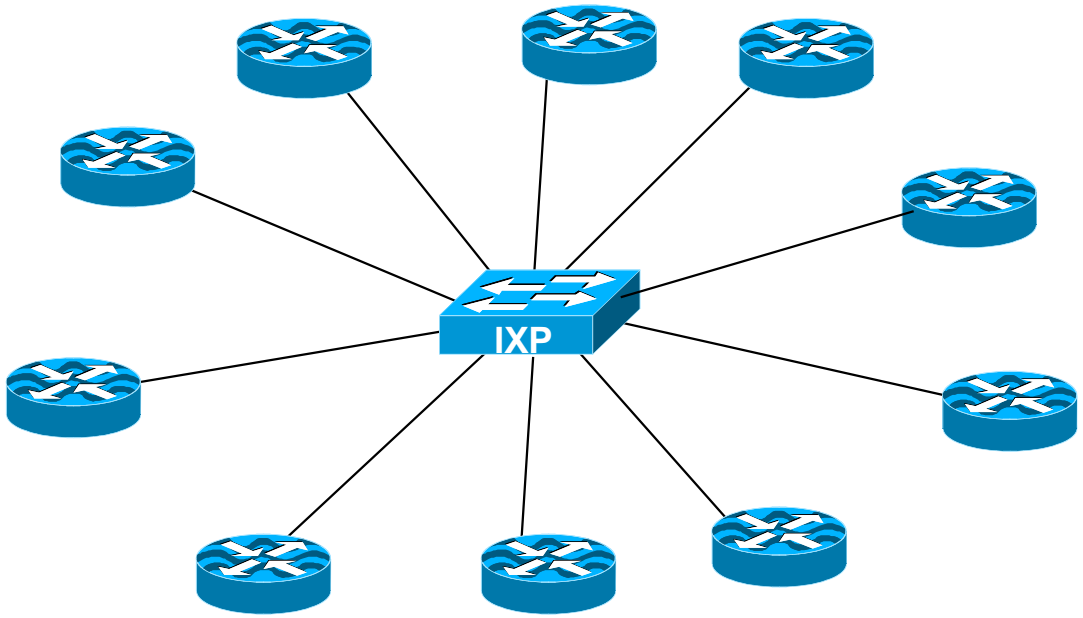

# A Bit of History…

- NSFnet – one major backbone
  - US “National Science Foundation” funded
  - Connected academic & research institutions
  - Also connected “private company” networks, under acceptable use policy (AUP), at network access points
  - AUP: No commercial activity

- Four Network Access Points (NAPs)
  - Chicago – run by Ameritech
  - San Francisco – run by PacBell
  - New York – run by Sprint
  - Vienna (Virginia) – run by MFS

# More History…

- Private companies needed to interconnect their networks too
  - Requirement to send “commercial traffic”
  - Could not cross NSFnet due to AUP

- Resulted in the first “commercial Internet Exchanges” in the early 1990s:
  - FIX-E (Virginia) was the first true IXP, FIX-W (Bay Area) also followed

- Leading to:
  - MAE-East – Virginia + CIX-West – Bay Area

# More History still…

- End of the NSFnet in 1995:
  - Meant move towards commercial Internet
  - Private companies selling their bandwidth
    - Transit / Peering model we know today
- These NAPs were among the original “exchange points”
  - NAP operators were providing commercial Internet access as well
  - All NAPs were replaced by neutral/commercial IXPs
- A global Distributed GIX proposed in mid 1990s
  - But never happened (planned to be CIX-West, MAE-East, SE-GIX and a Paris IX)

# Latency

2002 slide!

- In 1990s, Europe was using US for interconnects – 100ms to 200ms away!
- Up to early 2000s, Asia was using US and Europe for interconnects – 200ms to 900ms (satellite) away

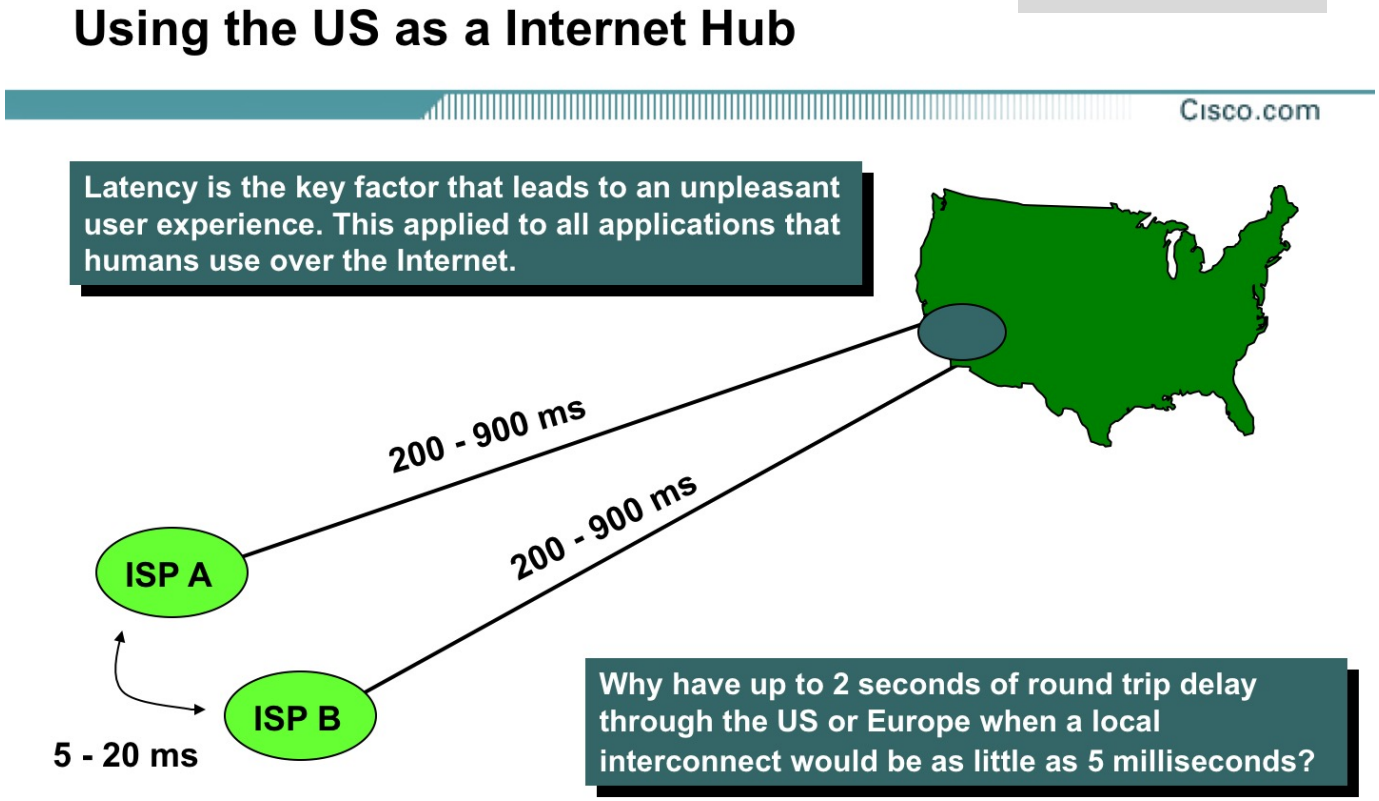

# The IXes

- SE-GIX formed in Stockholm in 1993
  - Three network operators interconnected
  - Latency reduction, performance gains, local traffic stays local
  - (Proposed to be part of the D-GIX)

- LINX formed in London in 1994
  - Five UK operators interconnected
  - Latency reduction, performance gains, local traffic stays local
  - (Proposed to be part the D-GIX when Paris fell through)

# The IXes

- HKIX was formed in Hong Kong in 1995
  - For intra-Hong Kong traffic
  - Within a decade, more than 60 ISPs were participating
  - Latency reduction, performance gains, local traffic stays local
- AMSIX (Amsterdam) and DE-CIX (Frankfurt) in 1996
  - Followed LINX model
  - Latency reduction, performance gains, local traffic stays local
- JINX (Johannesburg) in 1996
  - Initially for ISPA members only, later open to all
  - Latency reduction, performance gains, local traffic stays local

# Transit

- Paying another network operator for access to the Internet
- Significant operational cost
  - Data/traffic charges
  - Physical connectivity charges
- Transit provider determines onward connectivity, including
  - Diversity of service, and
  - Quality of service

# Peering

- Peering takes place between two network operators
  - To exchange traffic between each other's customers
- Minimises operational cost
  - **Peering is for free**
- Provides:
  - Improved customer experience (reduced latency, increased bandwidth)
  - Access to each other's hosted content

# Internet Exchange Point

- What it is:
  - A neutral location with unrestricted access where network operators freely interconnect their networks to exchange traffic
- What is the physical IX:
  - An ethernet switch
- How does it work:
  - IX Host provides the switch (IX fabric) and rack space
  - Network Operators interconnect via the IX fabric

# Why an IXP: Costs

- Internet transit costs money
  - For physical media, and
  - For Data/Traffic
  - (Even though transit costs are reducing all the time)

- IXP is “almost free”
  - Local access fibre
  - Optical interface costs
  - Data/Traffic is free
  - Contribution to IX operation costs (varies from IXP to IXP)

# Why an IXP: Capacity

- Transit often involves limited capacity
  - Operators run paid-for transit links “almost full”
  - Reduces user experience (more waiting, slower responses)

- Local interconnect at IXP
  - Bandwidth as large as that of the chosen IXP port
  - Quality of experience dependent on operator infrastructure only

# Why an IXP: Latency

- Transit often involves higher latency
  - Reduces relative throughput
  - Reduces user experience (more waiting, slower responses)

- Local interconnect at IXP
  - Negligible latency
  - Quality of experience dependent on operator infrastructure, not interconnect

# Why an IXP: Local Economy

- Content is located where most end-users have best (biggest bandwidth & lowest latency) access to it

- With no Internet Exchange Point:
  - Content is located out of country or out of region
- With an Internet Exchange Point:
  - Content clusters around the IXP

- Content hosting and related Internet businesses locate closest to where the maximum number of users can access the content

# Internet Exchange Point Features

- Unrestricted access:
  - Neutral location, neutral operator
  - Many connectivity access options (fibre etc)
- Maximum bandwidth:
  - Ethernet switch
    - 1G, 10G, N*10G, 100G ports
    - 400G and 1T now “in the works”
- Scaling:
  - Redundant switches
  - Redundant sites

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center

# IXP Detail

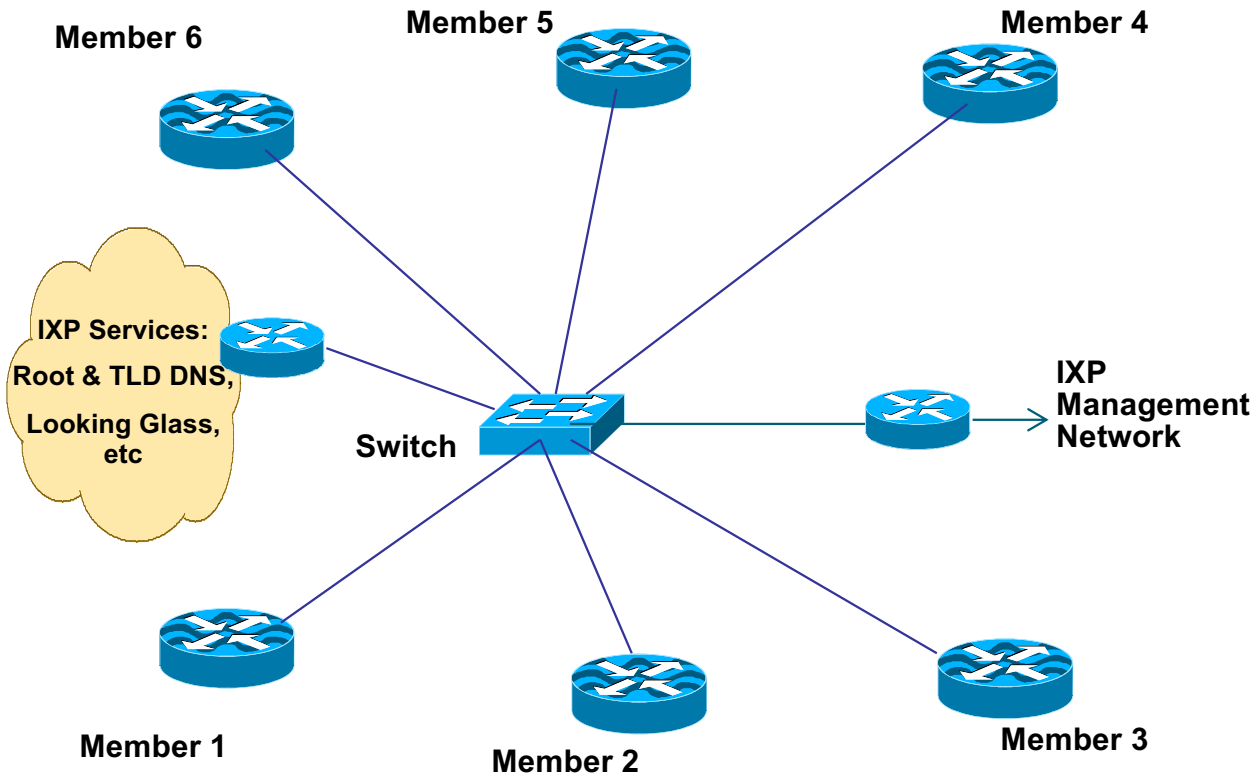
Member 6
Member 5
Member 4
IXP Services:
Root & TLD DNS,
Looking Glass,
etc
Switch
IXP
Management
Network
Member 1
Member 2
Member 3

# Who can join an IXP?

- Requirements are very simple:
  - Any organisation which operates their own autonomous network
- Member needs to have:
  - Their own address space
  - Their own AS number
  - Their own transit arrangements

# Who can join an IXP?

- Members include:
  - Commercial Network Operators
  - Academic & Research networks
  - Internet infrastructure operators (eg Root/ccTLDs)
  - Content Providers
  - Content Distribution Services
  - Broadcasters and media
  - Government Information networks

# Benefits: Europe

- Internet success story is because of IXPs
  - Early paid transit model was expensive, restrictive, low bandwidth, and high latency
- London:
  - LINX resulted in major companies locating their content with ISPs connected to LINX, before connecting to LINX themselves
  - Now one of the world's largest interconnects
- Amsterdam & Frankfurt
  - AMS-IX and DE-CIX established soon after LINX, following similar model

# Benefits: Nepal

- In 2002, NPTelecom, Mercantile & Worldlink connected to the Internet via satellite
  - Nepalese content was hosted in Europe and the US
  - No domestic content
  - No incentive for local content

- NPIX established
  - Most ISPs connected at launch, big performance improvements, and new business opportunities for Nepalese content development
  - Incumbent reluctantly joined later

# Benefits: Singapore

- Until 2007, Singapore was largely bypassed
  - Interconnect available (Equinix DC) but little used
  - SOX operated from 2001 but few members
- Taiwanese Earthquake in December 2006
  - 8 out of 9 fibre cables from Hong Kong to Japan and beyond were cut
  - Operators realised that they needed to look west
- Today Singapore is the interconnect for South and South East Asia
  - Domestic and international interconnects
  - Popular location for major international content providers
  - Popular location for global Tier-1 operators

# Benefits: Vanuatu

- VIX launched in 2012, first IXP in the Pacific
- Prior to submarine fibre, local traffic was over satellite between the 5 Network Operators
- Now the IX brings low latency, high bandwidth, and on-island content caches

# The IXP Success Story

- Neutral location
  - Anyone can install fibre or other connectivity media to access the IXP
    - Without cost or regulations imposed by location
- Secure location
  - Thorough security, like any other network data centre
- Accessible location
  - Easy/convenient for all participants to access
- Expandable location
  - IXPs result in Internet growth, and increasing space requirements within the facility

# The IXP Success Story

- Operation:
  - Requires neutral IXP management
- Funding:
  - All costs agreed and covered equally by IXP participants
  - Hosting location often contributes – the IXP brings them more business
- Availability:
  - 24x7 cover provided by hosting location

# The IXP Success Story

- Industry Standards documented by Euro-IX, the European IXP Association
  - Contributed to by the Euro-IX members
  - https://www.euro-ix.net/en/
- IXP BCP
  - General overview of the infrastructure, operations, policies and management of the IXP
- IXP Website BCP
  - Description of what an IXP website should contain

# What can go wrong?

- Network Operators using the “IX” to market their transit services
  - Cashing in on the good name of the IXP ☹
- Internet Gateways being called IXPs
  - Cashing in on the good name of the IXP ☹
  - IGs are commercial for-profit transit services
  - Whether they provide Local or International connectivity

# What can go wrong?

- IXPs pricing the membership fee out of reach of operators
  - IXP is meant to benefit all members
- Multiple IXPs serving the same locality
  - An IXP is not a competition
  - Too expensive for network operators to connect to all of them

# What can go wrong?

- IXP trying to compete with membership
  - Offering services the member would normally sell to customers
  - IXP services need to be agreed by all members

- IXP run as a closed privileged club
  - The only membership requirements are address space, ASN and independent transit
  - Competition regulators usually take a dim view of restrictive practices

# What can go wrong?

- IXPs charging for traffic
  - This competes with members
  - Charging for the switch port is common
  - Annual membership fee / cost contribution is common

- Mandatory Peering
  - Forcing all members to peer, against their own business requirements
  - Drives potential members away

# What can go wrong?

- Interconnecting IXPs
  - IXP in one locality connecting its LAN to an IXP in another locality
  - Competes with members
  - Who pays for the link and the traffic?

- Technical errors
  - IXP is an ethernet switch, not a router
  - Members must only connect using routers (no switches)
  - Routing design errors (e.g. not using BGP for route exchange)

# IXP Creation

- No economy or circumstance is unique or different
  - The first excuse for not creating an IXP is "we don't need one"
  - The second excuse for not creating an IXP is "oh, it is different here"
- Every locality has its differences
  - But surely every locality wants to:
    - Keep local traffic local
    - Improve network performance and QoS
    - Improve local Internet economy
  - The available technology is the same for every network operator everywhere
  - There is no excuse for not improving the local Internet

# Eco System Development

- IXP association usually created
  - Formed by members who have a port on the IXP
- IXP association members meet regularly
- IXP Technical community could also meet
  - Network operators meeting, involving network and systems operations technicians & engineers
  - Aligned with IXP Association/member meetings
- IXP could facilitate the creation of a Network Operators Group
  - The same technicians & engineers are involved in both!

# Local Internet Exchange Point

- This is a public peering point serving the local Internet industry
- Local:
  - where it becomes cheaper to interconnect with other operators at a common location than it is to pay transit to another operator to reach the same consumer base
  - can mean different things in different regions!

# Regional Internet Exchange Point

- A “local” Internet Exchange Point which has grown to become very popular outside the local area
  - Easy access
  - Inexpensive access
  - Favourable regulatory environment
  - Neutral co-location facilities

- This helps attract network operators from outside the local area
  - Regional Providers peer with each other and sell transit to smaller operators
  - Many show up at several of these Regional IXPs

# Regional Internet Exchange Point

- Also where local operators peer with operators from outside the locality
  - They don't compete in each other's markets
  - Local ISPs don't have to pay transit costs
  - ISPs from outside the locality don't have to pay transit costs
  - Quite often operators of disparate sizes and influences will happily peer – all to defray transit costs

- Singapore and Hong Kong are considered the regional interconnects for SE Asia

# Industry Associations

- IX-F
  - The Internet Exchange Federation
  - http://www.ix-f.net/
  - The federation of Internet Exchange Associations
- Euro-IX
  - The European Internet Exchange Association
  - https://www.euro-ix.net/en/
  - Members from Europe, associate members from around the world
  - Detailed information documented by member IXPs:
    - On how to start an IXP
    - What the IXP Best Practices are

# Industry Associations

- APIX
  - Asia Pacific Internet Exchange association
  - Meets twice a year, during APRICOT and APNIC conferences
  - http://apix.asia
- Af-IX
  - The African IXP Association
  - Meets along with the African Peering Forum
  - http://www.af-ix.net/
- LAC-IX
  - The Latin American & Caribbean IX Association
  - http://www.lac-ix.org/

# Activities

- Almost every country needs an IXP
  - (Even, each major city needs an IXP)
  - Will grow the domestic Internet economy

- Over the years, many activities to help improve interconnection
  - From the start there have many organisation based efforts
  - Peering Simulation Game at various NOGs in early 2000s
  - Today, the various Peering Fora

# Activities

- Many Peering Fora now
  - From the Global Peering Forum
  - To regional events (AfPIF, EPF, Asia Peering Forum, etc)
  - To Country Peering events

- Peering Fora are there to encourage and help operators to interconnect
  - Privately (direct cross connect)
  - Publicly (at IXPs)

# Summary

- Peering is vital for:
  - The growth of the Internet economy
  - Improvement in user experience by reduction of latency and increase of throughput
- IXPs are a fundamental part of the Peering EcoSystem
- Without peering and without IXPs, the Internet would be a very different place today