Interconnections on the Internet:

Exchange Points

Keith Mitchell

CTO, XchangePoint

RIPE NCC Regional Meeting

Dubai, 7th Dec 2003

Outline of Presentation

- Introduction
- Internet Interconnect Principles
- Internet Exchange Technologies
- Internet Exchange Governance Models
- Setting up an Internet Exchange

Speaker's Background

- Founder of UK's first commercial ISP, *PIPEX*, 1992-1996
- Founder and Executive Chairman of London Internet Exchange, *LINX,* 1994-2000
- First chair of RIPE *EIX* Working Group
- Founder and CTO of first pan-European commercial IXP operator, *XchangePoint*, 2000-
- XchangePoint currently operates IXPs in London and Hamburg, shortly deploying in Frankfurt and Amsterdam

What happens at an Internet Exchange Point ?

- Multiple ISPs locate backbone IP router nodes in single building operated by co-location provider
- In-building connections
  - to shared interconnect fabric using Ethernet LAN switching technology
  - over point-to-point private interconnections
- Routing information is exchanged bi-laterally between ISPs using BGP
- Exchange operator may or may not be same organisation as co-location provider
- Co-location provider will generally have other customers:
  - carriers, hosting, ASPs, content distributors

IXP Advantages

- Single large pipe to the IXP more efficient than many smaller pipes to many ISPs

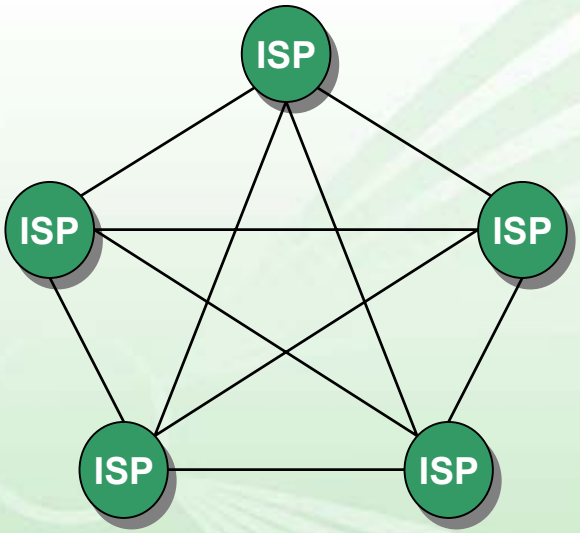

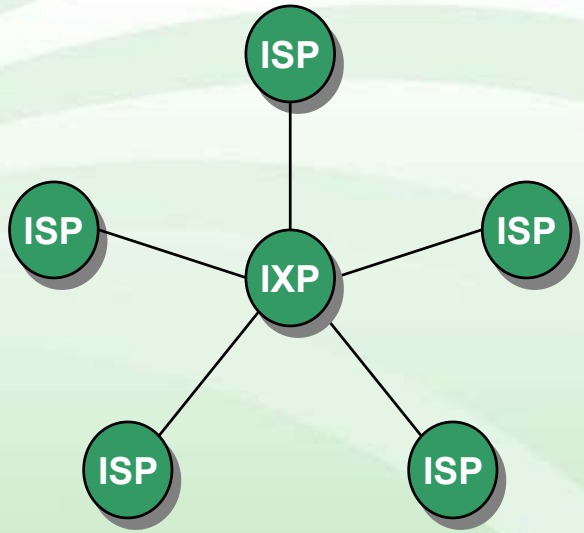

ISP = Internet Service Provider

IXP = Internet eXchange Point

IXP Advantages

- Keeps domestic traffic within a country/region without having to take indirect international route
- Typically 20-35% of traffic can be domestic
- Reduced bandwidth costs
- Improved throughput and latency performance
- Economies of scale
- Commercial basis of traffic exchange between ISPs across IXP usually via cost-saving peering
- Critical mass of ISPs in a single location creates competitive market in provision of capacity, transit and services

Internet Interconnect Principles

Inter-ISP Traffic Exchange

- Peering: Two ISPs agree to provide access to each others' customers
  - commonly no money changes hands: "settlement free"
  - barter of perceived equal value
  - simple commercial agreements
- Transit: One ISP agrees to give another's customers access to the whole Internet
  - they always charge for this !
  - usually volume and/or capacity based
  - wholesale version of what every Internet user buys
- Other models exist

Public and Private Interconnect

- Public Interconnect
  - Internet Peering Point (“IPP” or “IXP” or “NAP”)
  - multiple parties connect to shared switched fabric
  - commonly Ethernet based
  - many-to-many connectivity
- Private Interconnect
  - single circuit dedicated between two parties
  - typically used for transit
  - commonly SDH based
  - point-to-point connectivity

The Interconnect Market-Space

	Peering	Transit
Private	High Volume QoS North America	Traditional
Public	Traditional	Evolving

Metro Area IXPs

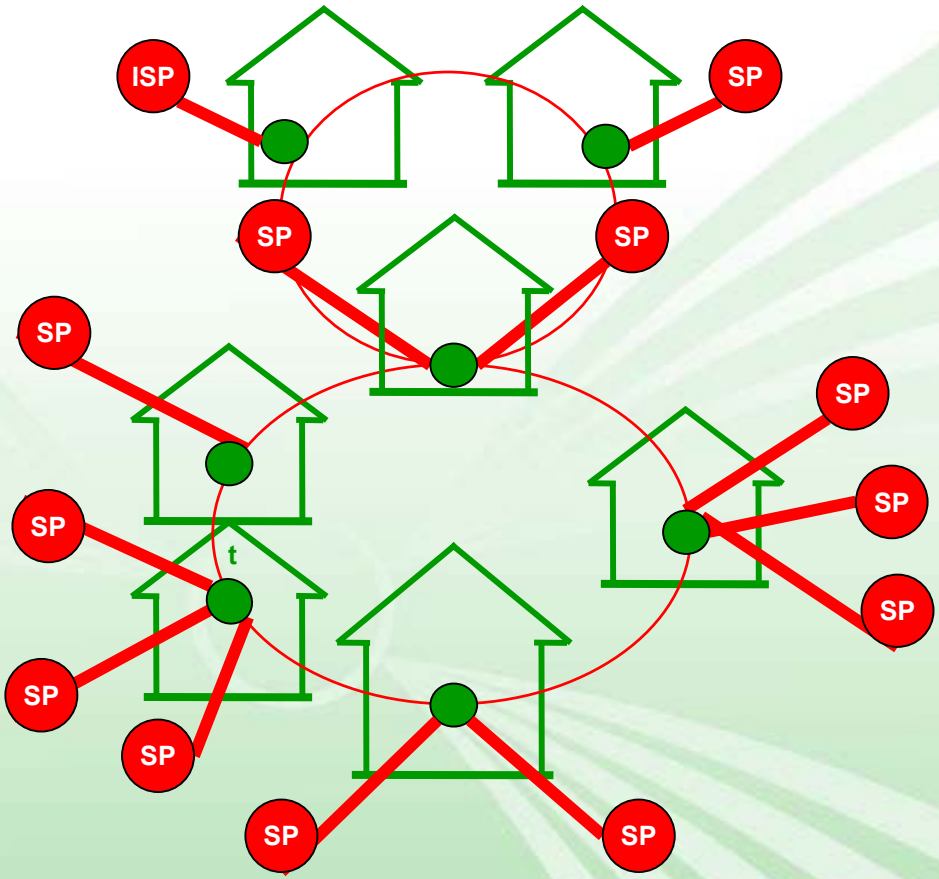

- Many major IXPs interconnect multiple co-location buildings in same metro area
- Inter-building MAN connections increasingly common as co-lo facilities fill up
- Distributed architecture connects ISPs within the metro area
- Ethernet switches for public & private peering and transit interconnect
- Gigabit dark fibre links between co-lo sites

IXP Technologies

IXP Technologies History

- Initially (1992-4):
  - 10Mb/s Ethernet from ISP router to IXP switch
  - FDDI between IXP switches
  - Single switch in single location
- 100Mb/s mostly replaced these 5+ years ago
- Some use of ATM at 155Mb/s and 622Mb/s meantime
- 1Gb/s Ethernet now common access technology
- 1Gb/s Ethernet also used in core of networks, but not quite enough for large IXPs with many 1Gb/s access connections
- 10Gb/s Ethernet increasingly common in IXP cores
- Some limited use of DWDM and MPLS

Routing and Switching at IXPs

- ISPs perform Layer-3 IP routing over wide-area using routers connected by SDH and/or ATM circuits
- IXPs perform layer-2 switching over local/metro area, usually using Ethernet
- ISPs interconnecting at IXPs exchange IP routing information using BGP (Border Gateway Protocol)
- Most IXPs operate public peering – single Ethernet subnet allowing open, many-to-many, traffic exchange between consenting pairs of participants
- Ethernet VLANs or ATM VCs can be used at some IXPs to provide circuit-switched private interconnects or closed user groups

Gigabit Ethernet

- Cost-effective and simple high bandwidth
- Most common technology for many ISPs accessing major IPPs
- Works well for local and metropolitan distances
- Proven and deployed at most major IPPs
- Almost universally used for IPP inter-switch links
- Technology is mature and price dropping
- High-performance switches available from various vendors:
  - Cisco, Extreme, Foundry

IXP Customer Requirements

- Your own Autonomous System (AS) number
  - you need this if you take service from >1 ISP anyway
- Your own IP address space
  - need to become member registry of RIPE NCC
- Router(s) which can do BGP
  - most medium/large Cisco/Juniper routers
- Space in one of the co-lo facilities at which it is present

IXP Commercial and Governance Models

Importance of IXP Neutrality

- In most markets, IXPs are a natural monopoly
  - Problem of trust between competitors
  - Risks of abuse and conflicts of interest
- Successful IXPs are not usually:
  - Owned, operated or housed by a single ISP or carrier
  - ISPs or IP routing transit providers
  - National or international backbones
- Co-location facility neutrality:
  - normally (mainly in Europe) these are buildings operated by independent commercial companies
  - though sometimes (mainly in US) co-los operate IXPs
  - IXPs tend not to be in carrier co-lo facilities

Some IXP Neutrality Principles

- Does not compete with its ISP members/customers
- Does not discriminate between its ISP members/customers
- Does not move traffic between cities or countries
- Does not make exclusive arrangements with:
  - ISPs
  - Carriers
  - Co-lo Providers
- Does not provide IP transit routing
- Does not take share of ISPs' transit revenues
- Only interconnects between co-lo sites
- May be present at multiple co-lo sites and providers

Governance/Commercial Models

- Operated by public sector national academic network
  - BNIX, GIGAPIX
- Not-for-profit membership associations of participating ISPs
  - LINX, AMS-IX
- Service within commercial co-location operator
  - Equinix, PAIX, NY-IIX
- Companies whose shareholders are participating ISPs
  - MIX, JPIX
- Independent neutral commercial companies
  - XchangePoint, N-IX

Internet Exchanges in Europe

- IXP operators are typically:
  - neutral
  - not-for-profit membership organisations
  - do not run hosting/co-location facilties
  - not same organisation as co-location provider
- Major cities, e.g. London, Amsterdam, Frankfurt, Paris
  - switch pan-european traffic
  - have multiple exchange operators
  - have multiple co-location facilties
  - each have several to 10s of Gb/s of traffic
- Usually one smaller national exchange per country for domestic traffic

Setting up an Internet Exchange

Getting Started

- Key to IXP viability and growth is *critical mass*
- Usually need at least 5 ISPs to get started
- Getting competitors to co-operate is not always easy !
- But demonstrable common benefits should win out in the end
- For associations, simple MoU good starting point
- Commercial operators will often use discounting strategies to attract initial group of ISPs
- Generally best to concentrate on getting traffic moving as first priority, and concentrate on the paperwork/ politics/PR later

IXP Resources

- RIPE EIX (European Internet eXchange) Working Group
  - http://www.ripe.net/ripe/wg/eix/
- Euro-IX Association of IXP Operators
  - http://www.euro-ix.net
- Global IXP Directory
  - http://www.ep.net
- Packet Clearing House
  - http://www.pch.net

Contact Details

Keith Mitchell

Presentation:

- **http://www.xchangepoint.net/info/RIPENCC-Dubai.ppt**

E-mail: **keith@xchangepoint.net**

Phone: **+44 20 7395 6020**