Internet Exchange Evolution, 1994-2011 & Beyond..

Keith Mitchell, VP Engineering

Internet Systems Consortium

Tor-IX AGM, Toronto, April 2011

Outline of Presentation

- Introduction
- Internet Exchange Background
- Internet Exchanges History
- Internet Exchange Models
- Growing Your Internet Exchange
- Observations, Conclusions

Speaker's Background

- Founder of UK's first commercial ISP, *PIPEX*, 1992-1996
- Founder and Executive Chairman of London Internet Exchange, *LINX,* 1994-2000
- Founder and Director of Nominet UK, 1996-2002
- First chair of RIPE *EIX* Working Group
- Founder and CTO of first pan-European commercial IXP operator, *XchangePoint*, 2000-2005
- Chairman of UK Network Operators' Forum 2004-
- President, DNS-OARC, 2006-2010
- VP of Engineering, Internet Systems Consortium, 2008-

- http://www.linkedin.com/in/keithmitchell

Internet Exchange Background

Typical IXP Properties

- Multiple ISPs locate backbone IP router nodes in single data center 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 (via BGP), and hence traffic, is exchanged bi-laterally between ISPs
- Exchange operator may or may not be same organization as co-location provider:
  - "European" vs "US" models
- Co-location provider will generally have other customers:
  - carriers, hosting, content distributors, NS registries/registrars
- Other models exist

IXP Advantages

- Keeps local traffic within a region without having to take indirect long-haul route
- Typically 20-35% of traffic can be local
- 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

IXP Technologies

- Initially (1992-4):
  - 10 & 100Mb/s Ethernet from ISP router to IXP switch
  - FDDI between IXP switches
  - Single switch in single location
  - Some limited use of ATM
- 1Gb/s Ethernet mostly replaced these 5+ years ago
- 1Gb/s Ethernet now common access technology
- 10Gb/s Ethernet used in core of networks between switches and sites
- 10Gb/s Ethernet increasingly common for access
- Some limited use of WDM, MPLS, optical switching
- 100Gb/s Ethernet is overdue !!
  - Extensive use of trunked Nx10Gb/s at both edge and core
  - See http://www.uknof.org.uk/uknof19/DAmbrosia-100GbE.pdf

Internet Exchanges History

15+ Years of Internet Exchanges

- The London Internet Exchange (LINX) first switched UK to UK Internet traffic on 8th November 1994
- Original LINX switch became permanent exhibit at the Science Museum, London in November 2004 !

Formation of LINX - 1994

- Before 1994, there were 4 ISPs in the UK with their own international connectivity:
  - PIPEX (AS1849)
  - Demon (AS2529)
  - JANET (AS786)
  - EUnet GB (AS1290)
- Various ad-hoc bi-lateral interconnect arrangements were in place
- BTnet (AS2855) entered the market during 1994 as the 5th player
  - as incumbent PTT British Telecom were under strong regulatory pressure to peer but there was much distrust
- Strong economic motivation to keep UK Internet traffic in UK
- At this point bi-lateral peering no longer scaled
- Challenge was to find high-quality *neutral* facility to enable LAN-based peering, following approach pioneered in Amsterdam and Stockholm

Formation of LINX - 1994

- One possibility was to use the national academic network (JANET) PoP at ULCC
  - but issues with 24x7 support
- Fortunately Telehouse operated **neutral** data centre/co-location facility already used by some ISPs
- Typically international links were 2Mb/s so in-building Ethernet interconnect was big deal
- Initially simple 10Mb/s Ethernet hub
- Infrastructure and connectivity established first…
- …finance, governance, legalities came *later*
- First peering traffic shipped on 8th Nov, 1994

Early LINX Growth

- The LINX operated for around a year with no formal arrangements in place
- Then members 6 & 7, INSnet and Xara, wanted to join –
  - had own international circuits
  - but how ?
- Simple MoU and website quickly drafted up
- Followed by establishment of membership-based legal entity
  - *"Public Company Limited by Guarantee"*
  - *Six founder Directors appointed*
- LINX was operated on volunteer basis until 1996, when I had crazy idea to become world's first IXP employee...

Evolution of LINX 1994-2010

- Incorporated as not-for-profit membership organisation 1995
- Hired first full-time employee 1996
- Over 50 members in 1997
- Added 6 staff and office in 1997
- Multiple data centres in London metro area 1998
- Over 1Gb/s traffic 1999
- Over 100 members 2000
- XchangePoint established as commercial company by LINX founders late 2000
- Over 100Gb/s of IXP traffic in London 2006

LINX in 2011

- 390 members
- 15 new applications in 2011
- 818 connected member ports
- 333 member-facing 10GigE ports
- over 873 Gb/sec of peak traffic
- 50 member countries

IXP Governance and Commercial Models

Internet Exchanges in Europe

- IXP operators are typically:
  - neutral
  - nonprofit membership organisations
  - do not run hosting/co-location facilities
  - not same organization as co-location provider
- Major cities, e.g. London, Amsterdam, Frankfurt, Paris
  - switch pan-European traffic
  - have multiple exchange operators
  - have multiple co-location facilities
  - each have several 100s of Gb/s of traffic
- Usually one smaller national exchange per country for domestic traffic

Internet Exchanges in US

- Major IXP operators typically:
  - data center providers
    - e.g. Equinix, Tel/X, Telehouse
  - run co-location facilities
  - are not ISPs themselves (neutral)
  - IXP is run as one more service within data center
- Main IXPs in major metro areas e.g.
  - SF Bay area
  - Washington DC
  - New York
  - Chicago
  - Los Angeles

Internet Exchanges in US

- Many small regional IXPs
  - typically volunteer membership organizations
  - informal governance
  - mostly local ISPs
  - lower traffic volumes
- Examples:
  - CMH-IX (Columbus, OH)
  - PITX (Pittsburgh, PA))

Regional Internet Exchanges

- "Second Wave" of IXPs in late 90s following successful growth of (supra-) national exchanges
- Examples include:
  - MaNAP (Manchester, England)
  - Scot-IX (Edinburgh, Scotland)
  - HH-CIX (Hamburg, Germany)
- Usually set up to balance over-centralization caused by incumbent IXP
- Lower joining threshold, i.e. not just Tier-1 ISPs
- Often had support from local government development agencies
  - Seen as way for local economy to enjoy benefits of dot-com boom
- Not all have survived…..

Importance of IXP Neutrality

- In most markets, IXPs are a natural monopoly
  - problem of trust between competitors
  - risks of capture, abuse and conflicts of interest
  - risks of regulatory intervention
- Successful IXPs are not usually:
  - owned, operated or housed by a single ISP or carrier
  - ISPs or wholesale IP ("*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

Successful 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 metro area co-lo sites
- May be present at multiple co-lo sites and providers

Governance
Commercial Models

- Operated by public sector national academic network
  - e.g. GIGAPIX, CATNIX
- Not-for-profit membership associations of participating ISPs
  - e.g. LINX, AMS-IX Amsterdam, TorIX, SIX Seattle
  - Over 90% of the 400+ IXPs globally work this way !
- Service within commercial co-location operator
  - e.g. Equinix, Tel/X, Telehouse
- Companies whose shareholders are participating ISPs
  - e.g. MIX, JPIX, JPNAP

Governance Pros & Cons

- IMHO, the Internet works best when there is a balance between competition and co-ordination
- Commercial IXPs can be more flexible, less sensitive to short-term problems, but will always be tempted to be compromise neutrality in return for revenue
- Nonprofit IXPs can work very well, but need to build critical mass to be viable and have governance overhead
- Volunteer IXPs are very resource efficient, but not well positioned to meet SLA requirements, and are vulnerable to capture by vested interests or to apathy
- Public sector/subsidised IXPs can serve local interests very well, but can create monopoly and may be open to political influence

Growing Your 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 Growth Challenges

- Large player infrastructure and organization centralization outside IXP's country/region
- Finding sites of suitable quality and neutrality
- Costs of intra-region local circuits to common interconnect site
- Ensuring all potential participants have sufficient routing etc technical clue
- Cost of entry-level technical resources
  - less of a problem than it used to be
- Political interference
- Dropping cost of transit impacts viability….
- Lack of scalable, affordable high-capacity technologies

Peering vs Transit

- The cost of wholesale Internet connectivity ("transit") has plunged since the dot-com bust
- This means the purely cost-based savings of peering are much less
- Leaves less money to pay for kit and connection to IXP
- Large established IXPs have sufficient critical mass to survive
- But this makes life harder for smaller IXPs
- e.g. nonprofit Cape Town (ZA), Manchester (UK) IXPs had to lay off all their staff in 2005
- What non cost-based benefits are there from IXPs ?

Keeping it Stable

- Fundamental issue at IXPs is that many parties, each managing their own backbones, are sharing a common Ethernet medium/subnet
- There are different approaches possible to protecting this in a membership organisation environment:
  - MoU based *(what is allowed/prohibited)*
  - SLA based *(what is supported)*
  - Technical prevention measures
  - Monitoring techniques

Stability Risks at IXPs

- Broadcast storms
- Unauthorised connection of layer-2 switching devices
- Flaky layer-2 link 'resilience' protocols
- Failure of switches to contain traffic to correct destination ports
- Non-scalable non-unicast traffic
- ARP spoofing
- Unauthorised static routing/next-hop
- Hijacking of routing resources

Risk Prevention & Detection

- It is very important to have a clear policy for what is and is not acceptable traffic, e.g.
  - Appendix 1 at http://www.linx.net/govern/mou.html
- ..and even more important to pro-actively monitor and enforce it
  - tools such as IXPwatch, RMON exist to do this
  - NetFlow, sFlow can detect abnormal traffic patterns
- Dedicated routers are generally easier to secure than general-purpose server boxes running routing software
- Much is preventable with appropriate filtering in switches

IXP Other Benefits

- IXPs are the logical place to locate, and hence attract, other Internet infrastructure resources
  - "peering magnets"
  - e.g. top-level name servers, time servers, performance measurement tools, research projects, public benefit content
  - ISC is involved with all of these and is open to working with IXP operators everywhere !
- Can enable new high-bandwidth, low latency applications
- Improved technical co-ordination and knowledge sharing
- Center of expertise for Internet technology
- Co-ordination of security, infrastructure protection, abuse response activities

IXP Other Benefits

- Increase diversity and resilience for participants
  - e.g. mutual backup arrangements
- Reduce latency for users and applications
  - e.g. gaming, multimedia
- Efficient multicast possibilities
- Multi-site IXPs can provide point-to-point and point-to-multipoint metro Ethernet services
- Build stakeholder community which can engage in other activities promoting local interests
  - Trade association, lobbying

Other Roles for IXPs

- Can create market for out-of-region transit providers to sell services to entire community of local ISPs at single cost-effective location
- Convenient point for regional academic/ research/nonprofit operator(s) to manage interconnect arrangements
- Potential for hand-off/resale of dial-up and unbundled DSL services
  - via L2TP over Ethernet VLANs
- Local-loop for wide-area Ethernet over MPLS circuit providers

Causes of IXP Failure

- Inability to provide reliable service or cope with traffic/member growth
- Exclusive arrangements with co-lo providers which subsequently go out of business
- Failure to build critical member/traffic mass before seed funding/goodwill runs out
- Incomplete set of resources
- Nonprofits can't easily borrow or raise funds so are vulnerable to cash-flow crunches
- Acquisition or capture by non-neutral operator
- Market consolidation to outside of region
- Lack of well-defined **need** – there is no point in creating an IXP for the sake of it

Observations, Conclusions

Optimal Distance & Scope

- What is the ideal number of IXPs in the world ?
- How big should they be ?
  - # of participants ?
  - Geographic area ?
  - Traffic share
  - Revenue, staff, etc….
- How far apart should they be ?
- What is the correct balance between technical quality and economic viability ?
- Does it make sense to have multiple operators competing in the same metro area ?
- Can multiple exchanges per country help overcome the bandwidth barrier ?

Optimal Distance & Scope

- This is a lot to do with local conditions
- Gigabit Ethernet can go >80km, the regional scope of most IXPs will be smaller than this
- Minimum magic number is 5 participants
- Multiple transit providers (at least 3) serving the region from outside it
- Multi-site IXPs need one or more of:
  - Several times more participants than sites
  - Low-cost dark fibre between sites
- There are no magic formulas for revenue, staff, traffic, SLA, competition - these all need to be tailored to the local community and its needs

Some Things We Got Right

- Some of these were by happy accident rather than design !
- Splitting co-lo provider and IXP provider
  - made it easier to add further co-lo facilities when the first one filled up
  - avoided being captured by any one co-lo provider
- But, it was a battle to:
  - convince members inter-building interconnect would not compete with them
  - get necessary investment in new co-lo facilities (at first !)
  - fend off every co-lo wannabe (later !)

Some Things We Got Right

- Staff employed by the IXP legal entity itself:
  - Avoided capture risk and neutrality compromise of outsourcing this
  - Made staff accountable for delivering good service
  - Ensured there were motivated advocates for growing the IXP and keeping barrier to entry low

Some Things We Got Right

- Communication
  - Regular F2F member meetings
  - Visible reachable contacts
  - Timely & accurate operational information
- Dedicated sales & marketing staff
- Carefully constrained trade association/public policy activities

IXP Resources

- This is no longer rocket science !
  - lots of help available if you want it
- Internet Systems Consortium
  - http://www.isc.org
  - several major IXP founders in our senior staff ☺
- Euro-IX Association of IXP Operators
  - http://www.euro-ix.net
  - allow non-EU members
- RIPE EIX (European Internet eXchange) Working Group
  - http://www.ripe.net/ripe/wg/eix/
- Packet Clearing House
  - http://www.pch.net

Contact Details

E-mail: keith@isc.org

Phone: +1 650 423 1348

Web: http://www.isc.org

http://www.linkedin.com/in/keithmitchell

Questions?

- While you're thinking of questions:
  - If you want to peer with F-Root, send mail to peering@isc.org
  - We host public-benefit organizations through our Hosted@ and SNS-PB programs. Contact {hosted,sns}@isc.org
  - Remember ISC is a public-benefit and survives through donations, forum memberships, SNS-COM and support contracts.
  - We appreciate the help and support of Tor-IX and its members, and need this to keep doing good work!