Internet Exchanges: Enabling Local Online Communities

Keith Mitchell

Chair, United Kingdom Network Operators' Forum

NOTACON3

Cleveland, April 2006

Outline of Presentation

UKNOF

- Introduction
- Internet Interconnect Principles
- Internet Exchange History
- Internet Exchange Models
- Internet Exchange Security
- Setting up an Internet Exchange
- Regional Internet Exchanges

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*, 2000-2005
- Chairman of UK Network Operators' Forum 2004-
- New resident of Cleveland, OH !

Internet Interconnect Principles

What happens at an Internet Exchange Point ?

UKNOF

- 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
- Co-location provider will generally have other customers:
  - carriers, hosting, content distributors, NS registries/registrars

IXP Advantages

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

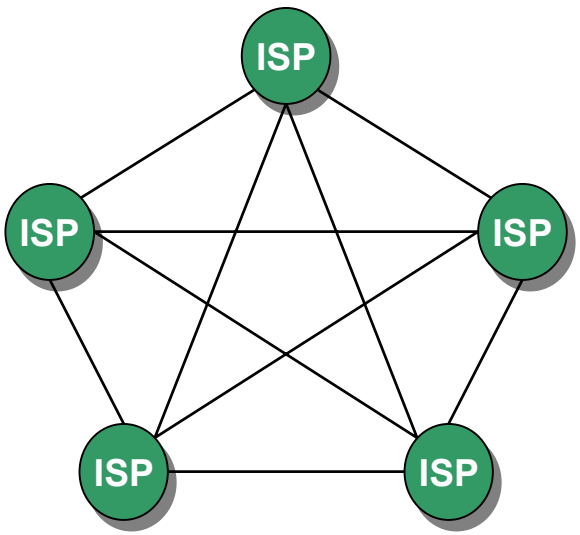

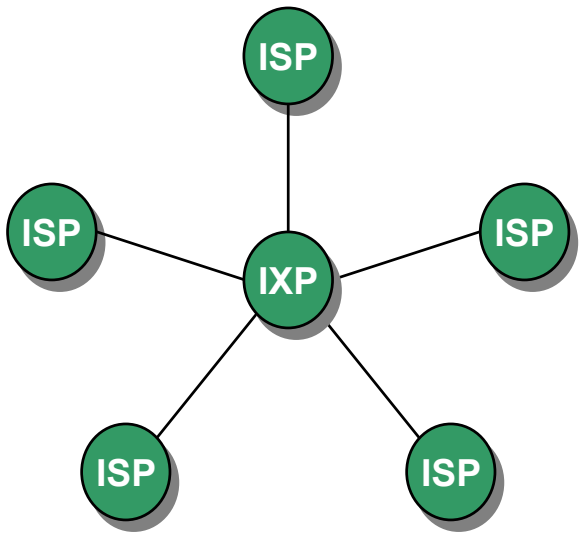

ISP = Internet Service Provider

IXP = Internet eXchange Point

IXP Advantages

UKNOF

- 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

Inter-ISP Interconnect

UKNOF

- 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
- Public Interconnect:
  - Internet Exchange Point ("IXP" or "NAP")
  - multiple parties connect to shared switched fabric
  - commonly Ethernet based
  - open, many-to-many connectivity
  - traffic exchange between consenting pairs of participants
- Other models exist

IXP Technologies

UKNOF

- Initially (1992-4):
  - 10Mb/s Ethernet from ISP router to IXP switch
  - FDDI between IXP switches
  - Single switch in single location
- 100Mb/s Ethernet mostly replaced these 6+ years ago
- Some use of ATM meantime
- 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 and MPLS

Gigabit Ethernet

UKNOF

- Cost-effective and simple high bandwidth
- Most common technology for many ISPs accessing major IXPs
- Works well for local and metropolitan distances
- Proven and deployed at most major IXPs
- Almost universally used for IXP inter-switch links
- Technology is mature and price dropping
  - e.g. 1Gb/s over copper
- Cost-effective high-performance switches available from various vendors:
  - Cisco, Extreme, Force10, Foundry

Internet Exchanges History

A decade+ 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

UKNOF

- LINX was set up through voluntary co-operation between 5 founder ISPs:
  - PIPEX, Demon, JANET, BT, EUnet GB
  - In 1994, these were only UK ISPs with their own international connectivity !
- Located in **neutral** data centre/co-location facility, Telehouse
- Initially simple 10Mb/s Ethernet hub
- Infrastructure and connectivity established first…
- …finance, governance, legalities came *later*

Evolution of LINX 1994-2000

UKNOF

- Incorporated as not-for-profit membership organisation 1995
- Hired first full-time employee 1996
- Over 50 members 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

Internet Exchanges in Europe

UKNOF

- IXP operators are typically:
  - neutral
  - nonprofit membership organizations
  - 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 to 10s of Gb/s of traffic
- Usually one smaller national exchange per country for domestic traffic

Internet Exchanges in US

UKNOF

- Major IXP operators typically:
  - data center providers
    - e.g. Equinix, Switch & Data, Terremark
  - run co-location facilities
  - are not ISPs themselves (neutral)
  - IXP is run one 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

UKNOF

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

- These IXPs outside the largest cities are the ones I mean when referring to “regional” IXPs later on

Regional Internet Exchanges

UKNOF

- “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…..

IXP Governance and Commercial Models

Importance of IXP Neutrality

UKNOF

- 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 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

UKNOF

- 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. BNIX, GIGAPIX, CATNIX
- Not-for-profit membership associations of participating ISPs
  - e.g. LINX, AMS-IX Amsterdam, SIX Seattle
  - Over 90% of the 400+ IXPs globally work this way !
- Service within commercial co-location operator
  - e.g. Equinix, Switch & Data, Terremark, IX Europe
- Companies whose shareholders are participating ISPs
  - e.g. MIX, JPIX

Governance Pros & Cons

UKNOF

- 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
- 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

Setting up an Internet Exchange

Getting Started

UKNOF

- 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 Customer Requirements

UKNOF

- 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 “registrar” of NRO member registry e.g. **ARIN,** LACNIC, RIPE NCC
- Router(s) which can do BGP
  - most Cisco/Juniper routers
  - also open-source based *nix PC platforms (bgpd, quagga)
- Space in one of the co-lo facilities at which it is present
- This does not have to be expensive
  - I do this myself (AS24865) !

IXP Resources

UKNOF

- This is no longer rocket science !
  - lots of help available if you want it
- Global IXP Directory
  - http://www.ep.net
- Packet Clearing House
  - http://www.pch.net
- Euro-IX Association of IXP Operators
  - http://www.euro-ix.net
- RIPE EIX (European Internet eXchange) Working Group
  - http://www.ripe.net/ripe/wg/eix/

UKNOF

Internet Exchange Security

Security at IXPs

UKNOF

- I am no security expert, but…☺
  …topic likely to be of interest to this audience
- Biggest issue at IXPs is that many parties, each managing their own backbones, are sharing a common Ethernet medium/subnet
- Once upon a time dumb hubs or even thick Ethernet cables were used for the IXP fabric, with all the attendant wire-tapping risks
- Today, cost-effective Ethernet switching avoids this, but does not eliminate all risks…

Risks at IXPs

UKNOF

- Broadcast storms
- Unauthorised connection of layer-2 switching devices
- 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

Broadcast Storms

UKNOF

- IXP operators biggest nightmare ☹
- Layer 2 bridging loops caused by failure of spanning-tree (STP) usually implicated
- Many IXPs prefer manual restoration of paths precisely to avoid STP software problems like this
- **Essential** to contain and control connection of 3rd-party switches
- Generally shuts down entire IXP and makes root cause determination very difficult
- Excessive broadcast traffic can also burn CPU resource on connected routers with knock-on effects

Switch Containment Failure

UKNOF

- Most Ethernet switches are designed to only forward traffic with a particular destination MAC address to the specified port
- There are few circumstances where this breaks down
  - occasional software/hardware failure
  - cheap switches have limited CAM table size – traffic which spoofs many source MAC addresses can overflow this and cause it to behave like a hub, flooding all traffic to all ports
- Best way to avoid this is to stick to switches from vendors who's core business is making switches/routers
- Good switches can filter/limit MAC addresses per port

Problem Broadcast Traffic

UKNOF

- In normal IXP operation, the only MAC-layer broadcast traffic should be ARP, and there should only be a few of these per second
- ….anything more is an abnormal condition
- Only routers are connected to an IXP, there should be no need for layer 3 IP broadcast (DHCP, SMB, IGP, device discovery etc) traffic
- ARP itself is open to spoofing abuse, some IXPs use static IP/ARP mappings to avoid this
- Switch vendors are gradually improving filtering and monitoring to prevent and detect “bad” broadcasts

Routing Exploits

UKNOF

- Bad ISPs in the past have tried various abuses at IXPs:
  - “Default-dumping” – static routing all outbound traffic across the IXP to an unsuspecting participant
  - Static routing – as above but for some routes only
  - “Next-hop” spoofing - causing traffic to go to a router other than the intended destination
- In one particular case, an ISP was getting free carriage US⇔UK by static tunnelling over other ISPs between two common IXPs !
- BGP sessions protected by MD5 authentication

Prevention & Detection

UKNOF

- It is very important to have a clear policy for what is and is not acceptable traffic, e.g.
  - http://www.xchangepoint.net/custinfo/AUP.php
  - MoU Appendix 1 at http://www.linx.net
- ..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

“Regional” Internet Exchanges

Regional IXP Challenges

UKNOF

- Large player infrastructure and organization centralization outside region
  - especially RBOC and cable operators
- Finding site of suitable quality and neutrality
- Costs of intra-region local loop 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….

Peering vs Transit

UKNOF

- The cost of wholesale Internet connectivity (“transit”) has plunged since the dot-com bust
  - $100s per Mb/s per month to <$10
  - Consolidation, commoditization
- This means the purely cost-based savings of peering are much less
- Leaves less money to pay for kit and connection to IXP
- Large IXPs have sufficient critical mass to survive
- But this makes life harder for regional IXPs
- e.g. nonprofit Cape Town (ZA), Manchester (UK) IXPs had to lay off all their staff last year
- What non cost-based benefits are there from regional IXPs ?

Regional IXP Other Benefits

UKNOF

- Are logical place to locate, and hence attract, other Internet infrastructure resources
  - e.g. top-level name servers, time servers, performance measurement tools, research projects
- 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

Regional IXP Other Benefits

UKNOF

- 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 regional 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
  - e.g. XchangePoint/PacketExchange

Causes of IXP Failure

UKNOF

- 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 mass before seed funding/goodwill runs out
- Incomplete set of resources
- 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

Optimal Distance & Scope

UKNOF

- 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 ?

Optimal Distance & Scope

UKNOF

- This is a lot to do with local conditions
- Gigabit Ethernet can go >50 miles, most “regions” 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 (<$5k/year) dark fiber 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

Summary

UKNOF

- Building an IXP is not hard, resources to do this widely available
- There are some cost benefits, but many less tangible community benefits
- The hard parts are:
  - building critical mass
  - keeping it viable
  - coping with growth and abuse
- Sharing with others can be educational and fun !
- Maybe someday, every city will have one…..

Contact Details

Presentation:

http://www.smoti.org/pres/NOTACON3-IXP.pdf

E-mail: **keith@uknof.org**

Phone: **+1 216 255 6587**

Web: **http://www.keithmitchell.co.uk**