Improving the Internet Infrastructure

Philip Smith

3 December 2007

Topics

- Internet Operations Groups
- Registry System
- IXPs
- Service Provider Security
- Root Nameserver Operations
- Training

Internet Operations Groups

menog
MIDDLE EAST
NETWORK OPERATORS GROUP

CNNOG

JANOG

SANOG

NZNOG

Internet Operations Groups

- Where network engineers and operators meet their colleagues
  - Peering & Business relationships
  - Industry relationship
  - Technology discussions
  - Operational best practices
  - Compare experiences (supplier, operational,…)
  - Purchasing decisions influenced
  - Routing software feature requests worked out
  - Jobs fair
  - **Keeping the Internet Working**

Regional Internet Operations Groups

- NANOG – North America
- APRICOT – Asia & Pacific Region
- SANOG – South Asia
- MENOG – Middle East
- PacNOG – Pacific Islands
- RIPE – Europe
- AfNOG – Africa
- LACNOG – Latin America

Country Network Operations Groups

- NZNOG – New Zealand
- JANOG – Japan
- CNNOG – China
- PhNOG – Philippines
- AusNOG – Australia
- SWINOG – Switzerland
- UKNOF – United Kingdom
- …

New NOGs

- NOG creation is a recent phenomenon
  - Local Language
  - Local Culture – Internet is not just American culture
  - Local Needs
- SANOG and NZNOG are common models
  - Too much temptation to introduce bureaucracy in newer NOGs
- Potential newcomers:
  - Central Asia
  - Caribbean
  - Latin America

The Registry System

Regional Internet Registries

- Responsible for distribution of:
  - IPv4 and IPv6 address space
  - AS numbers
- 5 RIRs
  - AfriNIC, APNIC, ARIN, LACNIC, RIPE NCC
- Membership driven
  - LIRs: most are ISPs or other service providers
- Politics higher up, e.g. ICANN, etc
  - Very little relevance to day to day Internet operations

Regional Internet Registries

- Policies
  - Membership driven
- Minimum IPv4 allocation is ~/21
  - (if you can justify a /22 you can get a /21)
  - (AfriNIC and LACNIC minimum is /22)
- Minimum IPv6 allocation is a /32
- ASN assigned if connecting to two different autonomous networks

Regional Internet Registries

- Work together to try and ensure that allocation policies are approximately aligned globally
  - There will be local variations
- Examples:
  - Initial IPv6 allocation policy
  - 4-byte ASN policy
- Success of these two seems to encourage some dubious policy proposals aiming for global consensus

Regional Internet Registries

- All hold two meetings per year
- For AfriNIC, APNIC and ARIN, one meeting held with regional NOG
- LACNIC hold meetings along with other Internet related organisations, e.g. IPv6 Task Force, NAPLA (LA IXP forum), etc
- RIPE NCC also hold Regional Meetings

  Moscow, Dubai, Bahrain, Doha,…
- APNIC Policy Showcases

  SANOG, NZNOG, etc

Internet Exchange Points

Internet Exchange Points

- Technical:
  - An Ethernet switch in a co-lo facility
  - ISPs bring routers, and peer with each other
- Business:
  - **The creator of the local Internet economy**
  - Avoiding paying upstream transit provider to carry local traffic
  - Avoiding RTTs which impede “doing business”
- Political:
  - Monopoly & state telcos don’t like them
  - IXPs without regulator support are doomed to failure

Internet Exchange Points

- Activities:
  - Well established for many years in Europe, North America and many parts of SE and NE Asia
  - African IX activity increasing
  - South Asia activity increasing
  - Latin America still sees most peering in Miami, USA
  - Middle East and Pacific Islands has discussion
- Issues
  - IXP still confused with monopoly transit provider or ISP transit service
  - Regional IXP is still the dream of those who don't understand what an IXP is

Internet Exchange Points

- Operations:
  - Biggest IXPs (LINX, AMS-IX, etc) are using high end 10GigE Switches, handling several 100Gbps of traffic
  - Smallest IXPs are still using typical 24 port 10/100 managed desktop switches
- Significance:
  - Maybe not “critical infrastructure” but vital for Internet economy
  - More than “just a switch”
- Getting started:
  - 90% political, 10% technical
  - Latter is simple Ethernet switch and BGP peering between participants

Euro-IX

- Euro-IX

  Not a European Region IXP!!

  Consortium of mostly European IXPs (+ some others)

  Meetings typical see 40+ IXes represented

  Technical & operational forum for advice, sharing & exchange of ideas, best practices, etc

- Cisco is patron of Euro-IX

  Along with Foundry, Force10 and Glimmerglass

Service Provider Security

Service Provider Security

- 1990s saw rapid growth of Internet
  - Getting established and financial profit came before quality and professional service
- Early 2000s saw bigger threats to Internet infrastructure
  - DOS against routers and high profile servers/services
  - Packet amplification attacks
- Responses
  - Formation of the ISP Security Community (NSP-SEC)
  - Development of more techniques and robust network design to thwart abuse of Internet infrastructure

Service Provider Security

- NSP-SEC
  - Global reach
  - Web of trust – membership by invitation/recommendation only
  - Open to key members of ISP security operations team only
  - Key security personnel of vendors participate (e.g. Cisco PSIRT)
- Regional NSP-SECs forming too
  - e.g. Japan, China,…
  - Every major region needs one – no ISP is an island

Anycast Root Nameservers & DNS

Anycast DNS

- Anycast:
  - Multiple instances of the identical service visible in multiple parts of the Internet
  - Individual devices share the same global IP address
  - Routing system chooses service closest to the end-user
- DNS Anycast Advantages
  - Insulates DNS against DOS attacks
  - Improves DNS lookup performance
  - Located at IXPs meaning low latency to end users

Anycast DNS

- DNS Root Nameservers

  Many of the operators now anycast the DNS service (e.g. F-root, I-root are visible in many parts of the world)

- GTLD and CCTLD Nameservers

  Many cctld and gtld operators now anycast their DNS services (e.g. Verisign, PCH)

Training

Training

- NOGs

  Many NOGs have workshops (e.g. ISP Routing, BGP Multihoming, Scalable Services, Network Management, DNS & DNSSEC, etc)

  Many NOGs have tutorials (e.g. Routing, IPv6, BGP Techniques, Multihoming, BGP Troubleshooting, MPLS, etc)

- Many other organisations organise their own events:

  The RIRs

  NSRC – Network Startup Resource Center (www.nsrc.org)

  AIT – Asian Institute of Technology

  Cisco (ISP and Security Workshops)

  Team Cymru (Security Workshops)

Training

- So much training available
  - So many venues – http://ws.edu.isoc.org/calendar
  - Most is cost recovery ($100 per day) or free; compare with professional courses ($1000 per day)
  - Most is very high quality and relevant; compare with professional courses which simply teach technology skills
- Yet ISP management deny these training opportunities to their technical staff
  - **Doing so denies their business the opportunity of success**

Summary

- Internet Infrastructure
  - Is taken for granted by too many
  - Is cared for by too few
- End-users only see services and when those services are working/failing
- Every ISP is responsible for their piece of the Infrastructure