MAGPI: Advanced Services IPv6, Multicast, DNSSEC

Shumon Huque
MAGPI GigaPoP & Univ. of Pennsylvania

MAGPI Technical Meeting
April 19th 2006, Philadelphia, PA

Outline

- A description of advanced services that we offer or are in the process of offering
  - IPv4 Multicast
  - IPv6
  - IPv6 Multicast
  - DNS Security
  - Jumbo Frames (time permitting)

Multicast in MAGPI

- Native Multicast routing (PIM-SM)
- Local Rendezvous Point
- External peerings with:
  - Abilene (Internet2 backbone network)
  - ~ dozen MAGPI customers
- Looking for opportunities to peer with commercial ISPs

IPv4 Multicast

- Protocols:
  - M-BGP (Multi-protocol BGP)
  - MSDP (Multicast Source Discovery Protocol)
  - PIM-SM (Protocol Independent Multicast - Sparse Mode)
    - Consider Bootstrap or anycast for RP redundancy
  - IGMP v2
  - IGMP v3 (for source specific multicast)
- Recommendations:
  - Deploy within your network first and test
  - Arrange for inter-domain peering with us
  - Establish debugging methodology (you'll need it!)
  - Attend an Internet2 IP Multicast workshop

IPv4 Multicast debugging

- Establish a debugging methodology!
  - Lots of new control protocols
  - Lack of practice
  - Inverted paradigm
  - Receiver driven communication model
  - Symptom may be far from problem
  - Same symptom can have may different causes, at different places in the path
  - Not many good debugging tools

IPv4 Multicast debugging

- Establish direction
- Establish group address(es)
- Have on path:
  - Constantly active source
  - Constantly active receiver
- Know how to examine multicast routing state on your equipment
- Have contacts at peer networks (for debugging interdomain operation)

Multicast Addressing

- RFC 3171
- http://www.iana.org/assignments/multicast-addresses
- Categories
  - 224.0.0.0/24 - Local network control (not forwarded)
  - 224.0.1.0/24 - Internetwork control block
  - 232.0.0.0/8 - SSM
  - 239.0.0.0/8 - Administratively scoped
  - 233.x.y.0/24 - GLOP (x.y encodes the AS#)
  - 224.2.0.0/16 - SDR/SAP Block

Multicast: useful tools

- NLANR Multicast beacon
  - http://dast.nlanr.net/projects/Beacon/
- Asmping/ssmping
  - http://www.venaas.no/multicast/ssmping/
- Abilene router node proxy
  - http://ratt.uits.iu.edu/routerproxy/abilene/
- Iperf
  - http://dast.nlanr.net/projects/iperf/
- Mtrace
  - ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/mtrace-5.2.tar.gz

Multicast: useful tools

- Internet2 Multicast Working Group:
  - http://multicast.internet2.edu/
- Multicast debugging handbook
  - http://imj.ucsb.edu/mdh/index.php
- Debugging Multicast using Abilene looking glass
  - http://www.accessgrid.org/agdp/guide/looking_glass.html

IPv6

- Next generation Internet Protocol
- Over a decade old!
- Not much deployment in the US (yet)
  - Research & Education networks
    - Abilene and various GigaPoPs, some universities
    - vBNS, ESNet, DREN
  - Very few commercial ISPs
    - NTT-Verio, Level-3

IPv6 benefits

- 128 bit addresses
- Scalable routing
- Fix inequities in address allocation
- Built-in security?
  - IPSEC implementation mandatory
- No NAT
  - Restore e2e architectural model of the Internet
- Neighbor discovery / Router solicitation
- Autoconfiguration

IPv6 Issues

- Multi-homing
  - Multi6, shim6, pi addresses
- Mobility
- Deployment incentives
  - Perhaps OMB or interop with Asia-Pacific
  - Or a killer app
- Some folks feel:
  - No significant architectural progress
  - Needed locator/identifier split (8+8/GSE etc)

IPv6 in MAGPI

- Native IPv6 routing deployed
- External v6 peering with Abilene
- External peerings with 2 connectors:
  - University of Penn & Princeton University
- Routing protocols:
  - Intra-domain: IS-IS
  - Inter-domain: M-BGP
- A few test services in MAGPI:
  - Web, traceroute, ping server, NTP

IPv6 in MAGPI

- Our (provider allocated) block:
  - 2001:0468:1800::/40
- Customer delegations are /48 sized
  - Enough to number 65,536 /64 subnets
- Current allocations:
  - 2001:468:1800::/48 MAGPI Infrastructure
  - 2001:468:1802::/48 University of Pennsylvania
  - 2001:468:1804::/48 Princeton University
- Native IPv6 peerings only
- DNS answers over IPv4 today
  - IPv6 transport for DNS planned for near future

IPv6 Resources

- Internet2 IPv6 Working Group
  - http://ipv6.internet2.edu/

IPv6 Multicast

- PIM-SM (with v6 support)
- MLD (Multicast Listener Discovery) v2
- MBGP
- For Interdomain ASM:
  - Static RP-group mapping
  - Embedded RP (RFC 3956)
  - No MSDP!
  - BGMP?

IPv6 Multicast

- Multicast addresses:
  - FF00::/8
- Further details:
  - http://www.iana.org/assignments/ipv6-multicast-addresses
  - RFC 4291: IPv6 Addressing Architecture
  - RFC 3306: Unicast prefix based multicast addresses
  - RFC 3956: Embedded RP

IPv6 Multicast in MAGPI

- Coming soon :-)
- We deployed an IPv6 Multicast network for the Fall 2005 Internet2 member meeting
  - Using static RP at Renatar (France)
  - Streaming video demos between Philly, New York and Norway
- Upcoming research project with a major cable provider in the area

DNSSEC

- The problem:
  - DNS data published by the registry is being replaced on it's path between the server and the client
  - Bogus data is being inserted into caching resolvers (cache poisoning)
  - This can happen in multiple places in the DNS architecture
    - Some places more vulnerable to attack than others
    - Vulnerable software often makes it easier

DNSSEC

- Goals
  - Verify authenticity of DNS “data”
- Operation:
  - Registry signs data and publishes it securely on authoritative name servers
  - DNS clients (remote caching resolvers and possibly stub resolvers) validate any queried data

DNSSEC

- Additional potential benefits:
  - Secure public key exchange:
    - SSHFP, IPSECKEY, CERT resource records

DNSSEC

- Object security, not channel security
  - Authenticate the DNS data itself
    - Registry cryptographically signs the data
    - Security-aware resolvers verify the signature
- Transaction/channel security:
  - TSIG, SIG(0), IPSEC etc
  - This may or may not be important depending on how the endpoint obtains the DNS responses
- DNSSEC doesn’t provide:
  - Confidentiality or Authorization

DNSSEC setup tasks

- Setup zones
- Create keys:
  - ZSK: zone signing key pairs
  - KSK: key signing key pairs
- Sign entire zones with ZSK
- Sign ZSK with KSK
- Safeguard private keys
- Secure zone transfers between authority servers
- Arrange secure delegation with parent and children zones
- Establish key maintenance/rollover procedures

DNS Deployment issues

- DNS software support
- Additional processing requirements
- New technology: chicken and egg
- Automated key rollover and distribution
- Zone enumeration possibility
- Universal vs Islands Of Trust
  - Trust anchor maintenance costs
- How to get root and TLDs signed?
  - Should we use DLV registries to start?

A MAGPI DNSSEC record

- [live demo here]
- New resource records
  - DNSKEY, RRSIG, NSEC, DS
  - NSEC3 (coming)

MAGPI DNSSEC plans

- Sign all MAGPI DNS data
  - (a few zones have already been signed)
- Trust anchor distribution:
  - Publish on secure web page
  - Publish in DLV registry (which one?)
  - Exchange with other I2 institutions directly?
- Work with Internet2 pilot on
  - Getting .edu TLD signed
  - What about ARIN (in-addr.arpa) and Verisign (.net)

Internet2 DNSSEC Pilot

- Coming soon ..

DNSSEC Resources

- A good general website:
  - http://www.dnssec.net/
- Internet2 workshop:
  - http://dnssec-nm.secret-wg.org
- Protocol specifications:
  - RFC's 4033, 4034, 4035
- Threat analysis of DNS:
  - RFC 3833

Questions or comments?

- Shumon Huque
  - shuque -@- isc.upenn.edu