K-root and DNSSEC

Anand Buddhdev
RIPE NCC

RIPE NCC

- One of the five Regional Internet Registries
- Provides IP address and AS number resources to Europe and Middle-East regions
- Operates the parent reverse DNS zones for allocations from IANA
- Operates K-root, one of the 13 root DNS servers
  - Anycast cluster of 18 instances
  - http://k.root-servers.org/

Features of DNS

- Translates names to IP addresses
- Invented around 25 years ago
- Uses UDP as its primary transport – low overhead
- Original spec limited payload to 512 bytes
- Essential for a working Internet

Security concerns of DNS

- UDP based – address spoofing
- Neither transport nor content is secure
- Protocol design limitations
  - 16-bit query ID
  - 512 bytes of payload
- Fast hardware and networks make attacks easy
  - Misdirect clients
  - Steal personal data (passwords, account numbers)

One solution: DNSSEC

- DNSSEC = DNS SECurity
- Introduces cryptographic security for content
- Been in development within IETF for about 10 years
- Uses Public Key Cryptography
  - Content is signed by private key
  - Clients on the Internet have the public key for validation
- Learn more in a DNSSEC tutorial

Consequences of DNSSEC

- Security comes at a price
  - DNS responses carry signatures and are bigger
  - Many responses are bigger than 512 bytes
  - Clients would have had to fall back to TCP
- IETF created DNS extensions to allow for larger packets (EDNS0)
  - In theory, it allows DNS speakers to use 4 kB buffers
  - The reality is quite different!

Large DNS Packets

- Some devices and software still enforce the 512-byte limit on DNS and/or UDP packets
- Path MTU limits cause packet fragmentation
  - Some firewalls block fragments
  - Originating servers don't always get back "fragmentation needed" messages due to ICMP filtering
- TCP fallback not practical because of a large number of queries
  - TCP not suitable in anycast setups

DNSSEC in the DNS Root Zone

- The IETF considers DNSSEC to be mature enough to be deployed in the root zone
- In 2009, NTIA asked Verisign and ICANN to sign the root zone
- Much work going on, with progress updates at http://www.root-dnssec.org/
- Verisign and ICANN co-ordinating deployment with root-server operators

Staged Roll-out

- Prevents a “big bang” situation
- Clients which have problems can switch to another root server
- Allows people time to upgrade software and networks while still receiving DNS service
- Allows Verisign, ICANN, root-server operators and researchers to gauge the effects and make informed decisions

DURZ

- Deliberately Unverifiable Root Zone
- Signed zone with dummy keys
- Ensures that no-one depends upon it
- Can be withdrawn quickly without breaking service
- Real keys will be published after all root servers are serving a signed root zone

K-root Preparation

- Upgrade to NSD 3.2.4
  - Has options for tuning TCP connection limits and buffer sizes
  - Clears the DF (don't fragment) bit on response packets – allows routers to fragment large packets
- Network upgrades
  - Upgrade to Gigabit Ethernet ports at global instances
- Co-operation with NLNet Labs on load testing of our K-root setup

Monitoring and Data Collection

- Upgraded DSC to report TCP connection rates
- Enhanced pcap filter to capture TCP queries and responses
- Special pcap filter to capture just **priming queries**
- Mini-DITL runs to upload pcap data to OARC before and after each root-server publishes signed zone
- Reply-size tester deployed at global instances

Reply-size Testing

- Code by Duane Wessels of OARC
- dig +short txt test.rs.ripe.net [@resolver]
- Hidden HTML element on www.ripe.net triggers the same query
- Java application on labs.ripe.net to perform the same test
- Helps users to figure out a reasonable buffer size for their resolvers

Tuning EDNS buffer size

- BIND and Unbound default is 4096 bytes
- For BIND 9, use “edns-udp-size n;” in options clause in named.conf
- For Unbound 1.4.0+, use “edns-buffer-size: n” in unbound.conf
- Allow TCP/53 connections through your firewall

Non-DNSSEC-aware Resolvers

x.x.x.x lacks EDNS, defaults to 512

x.x.x.x summary bs=512,rs=486,edns=0,do=0

- These resolvers are unaware of DNSSEC
- Will continue to receive DNS responses without signatures
- PowerDNS recursor, djbdns
- BIND with “dnssec-enable no;” in options clause

Public Awareness

- Articles on RIPE Labs and in Member Update
- Presentations at technical meetings and conferences
- Outreach to ISPs and network community

Questions?