

#### **K-root and DNSSEC**

Anand Buddhdev RIPE NCC



# **RIPE and RIPE NCC**

- Réseaux IP Européens
  - Community of all Internet users in Europe
  - Now includes the Middle East
  - Unregistered, informal community
- RIPE Network Coordination Centre
  - Registered as a Dutch non-profit organisation
  - Provides services to the RIPE Community, as well as the Internet
  - 116 employees



#### **RIPE NCC**

- One of the five Regional Internet Registries
- Provides IP address and AS number resources to Europe and Middle-East regions
- Organises RIPE meetings to facilitate development of policies and exchange of ideas
- Provides additional services related to IP address and AS number resources
- Membership model
  - Local Internet Registries



# **Additional Services**

- RIPE Database
  - Public database for storing Internet resource information
- Training
  - Educate LIRs on resource management and how to use RIPE NCC's services
- Information Services
  - Hostcount
  - Test Traffic Measurement (TTM) and DNSMON
  - Routing Information Service (RIS)



#### The RIPE NCC DNS Team



Anand Buddhdev







Anand Buddhdev

Snow BV, 4 March 2010

http://www.ripe.net



## **Additional DNS Services**

- Manage RIPE NCC's zones (ripe.net, etc)
- Reverse DNS for IPv4 and IPv6 address space
- DNSSEC
- Secondary DNS for ccTLDs and other registries
  - Support developing countries
- ENUM
  - Uses DNS to map telephone numbers to resources
- AS 112 instance
  - Mops up reverse DNS queries for RFC 1918 address space



#### **Features of DNS**

- Translates names to IP addresses
- Invented around 25 years ago
- Uses UDP as its primary transport low overhead
- Essential for a working Internet
- "512 bytes should be enough"



# **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
  - Kaminsky-style attacks
  - Misdirect clients
  - Steal personal data (passwords, account numbers)



# **The Solution: DNSSEC**

- Security extensions to DNS
- 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



# **DNSSEC** at the **RIPE NCC**

- One of the first DNSSEC deployments in 2005
- All forward and reverse zones signed
- Uses open-source tools developed at the RIPE NCC
- Helped gain a lot of experience



#### **Consequences of DNSSEC**

- DNS responses carry signatures and are bigger
  - 512 bytes no longer enough
- IETF created DNS extensions to allow for larger packets (EDNS0)
  - Clients can solicit bigger responses of up to 4096 bytes
  - In theory everything should just work
  - The reality is very different !



#### **Large DNS Packets**

- Some devices and software still enforce the 512byte 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 Unvalidatable 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
  - Single trust anchor to begin validation chain



## **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
  - Sends minimal responses to DNSKEY queries
- 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



#### **Reply-size Tester**

0 0		
For more information see	: http://k.root-servers.org/replysizetest	RIPE
Test results		
for resolver: 212.54.40.2	25	
Announced buffer size:	: 1400 bytes	
Measured buffer size:	1388 bytes	
EDNS enabled:	yes	
DNSSEC enabled:	yes	

#### Your resolver announced a buffer size bigger than the largest packet that it can receive.

Note: There will always be a difference between the announced and measured buffer size because of the algorithm used. However this difference should not exceed 300 bytes.

For detailed explanations about these messages see: http://k.root-servers.org/replysizetest



# **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





Anand Buddhdev

Snow BV, 4 March 2010