2017 KSK Rollover Monitoring Plan

Version: 2016-09-15

Contents

Introduction 3

Expected Anomalies 3

Measurement Infrastructure 4

- Software 4
- Traffic Capture Location 4
- Hardware 4

Introduction

This document describes the plan to monitor the traffic towards root name servers to evaluate the effects of changing the root zone KSK. The results of this monitoring effort will help ICANN understand the scale of possible failures due to rolling the KSK.

During the rollover of the current KSK to the new KSK, ICANN plans to capture and analyze traffic to l.root-servers.net in near real time to understand the scale of resolvers that have problems with the KSK rollover. To make sure that the sample is statistically representative, a comparison will be made with traffic to b.root-servers.net.

A subset of resolvers that send traffic to l.root-servers.net currently perform validation using DNSSEC. Some of these resolvers are able to follow the KSK rollover by implementing the protocol described in RFC 5011. In other cases, the resolvers' operators will configure a new trust anchor manually. But some resolvers will not have the root trust anchor updated and validation will fail. It is important to first establish a baseline to understand the behavior before rolling the key. There might be anomalies during this time, but measuring before the roll makes sure that these anomalies are part of the baseline. The underlying premise is that validating resolvers that do not have the new KSK will behave differently than those that do.

Expected Anomalies

During the rollover process, the response size for a DNS query for the DNSKEY resource records for the root zone will change. At some point, there will be more than two keys in the root DNSKEY resource record set. This larger packet size might lead to incomplete fragments received by the resolver, fallback to TCP, or both. These effects are referred to as "response size effects".

Resolvers that have a "stale" trust anchor (meaning a trust anchor that has not been updated for some reason during the rollover process) will not be able to validate the root zone once the new KSK begins signing the key set. This situation will also have an effect on root server traffic because these resolvers might increase their query rate. These are referred to as "stale key effects".

Stale key effects are measured based on the volume of requests and the number of unique AS numbers that are associated with the source IP addresses of those requests. In this way, the scale of impact can be understood.

Measurement Infrastructure

To cope with the sheer volume of requests and to process these requests in a timely fashion, a new measuring infrastructure is deployed. This measurement infrastructure is a cluster of 8 servers, managing 1 Petabyte of data, spread over two datacenters (LAX and IAD). Currently, this cluster holds 20% of all L-root-server traffic. The DNS-Engineering team is working to bring the remaining 80% of L-root data to this new infrastructure. Additionally, this cluster contains historic B-root traffic and traffic for NS.ICANN.ORG which is responsible for 3 top level domains (UG, INT and MUSEUM).

Software

The software that can measure this impact is Nominet's Turing version 1.2. Version 1.1 has been evaluated.

Traffic Capture Location

Traffic to L-root is captured locally at L-root nodes. Compressed traffic captures are then moved to the storage location for further processing. This storage location is either ICANN's LAX or DC data center, depending on availability and bandwidth. Traffic for B-root is stored at the ICANN LAX datacenter due to the proximity of USC-ISI (the operators of B-root).

Hardware

The measuring infrastructure consists of two sets of four servers, one set for the Washington DC location (IAD) and one set for the Los Angeles location (LAX). A single machine is a Dell PowerEdge R730xd, with 128 TByte storage plus 1.6 TByte solid state drives. Each server contains an Intel Xeon Processor, running at 2.4 GHz with an

onboard cache of 25MB. This processor contains 10 hyperthreaded cores (combined: 20 threads). Each server contains 64 GB internal memory.

The total capacity of this cluster is 160 parallel threads running at 2.4 GHz, internal memory of 512 Gigabyte with a combined storage of over 1 Petabyte.

One World, One Internet

ICANN.ORG