Securing BGP using DNSSEC

Lutz Donnerhacke

db089309: 1c1c 6311 ef09 d819 e029 65be bfb6 c9cb

IKS

The Problem

- BGP announcements can be „fat fingered“
- Prominent examples: AS7007, YouTube
- Most common usage: Spam injection
- Proposed solutions like s-BGP, so-BGP, ps-BGP, SIDR are too detailed
- DNS(SEC) mapping by Tony Li and Randy Bush was a theoretical proposal

IKS

Design principles

- Don´t solve everything, only „config errors“
- KISS principle: Minimal impact to routers
- Deployment: Use existing software, do not modify protocols on the wire
- Goal: Verify a BGP prefix and path **after** the injection point (missing filters)
- Allow cold start, allow private exceptions

IKS

Verification process

- BGP Updates contain prefixes with paths
  - 2003::/19 ”1273 3320 i” ”5400 3356 3320 i”
- Check origin (prefix to AS check)
  - Is 3320 allowed to announce by 2003::/19 ?
- Check path (peering/upsteam/...)
  - Does 3320 export ”3320” to peer 3356 ?
  - Does 3356 import ”3320” from peer 3320 ?
  - Does 3356 export ”3356,3320” to peer 3356 ?
  - Does 5400 import ”3356,3320” from 3356 ?

IKS

Testbed

- Verify spec by real world data
  - Full IANA allocates
  - Almost all RIPE allocates
  - Almost all IRDB (RIPE) assignments
  - Spec changes are result of real problems
- Simulation of RIPE region
  - Running 260 DNS servers (IPv6, single host)
  - ~ 15000 AS, 1700 IPv6, 70000 IPv4 zones, 1GB compressed zone data

IKS

Using the Testbed

- Stub/slave zone
  - zone "bgp.arpa" {
    type stub;
    masters { 2001:4bd8::3:0; };
    };
  - DS 48622 5 1 5EC22EB16EB4E6E94889BF249EE82920608D3558
- Signed root
  - http[s]://www.iks-jena.de/leistungen/keys.txt
- Real tests should use a remote resolver w/ signed root and the AD bit from DNSSEC

IKS

Mapping into bgp.arpa.

- AS# in asdot+ format
  - 15725 > 0.15725 > 5.2.7.5.1.0.as.bgp.arpa.
  - 3.10 > 3.00010 > 0.1.0.0.0.3.as.bgp.arpa.
- Prefix mapping like in-addr/ip6.arpa
  - 217.17.192.0/20 > 192/20.17.217.ipv4.bgp.arpa.
  - 2003::/20 > 0/20.3.0.0.2.ipv6.bgp.arpa.
- Different namespace to follow allocations and assignments as closely as possible

IKS

Mapping into bgp.arpa.

- Route origin
  - 192/20.17.217.ipv4.bgp.arpa. ASSET 15725
  - $ORIGIN 0/19.3.0.0.2.ipv6.bgp.arpa.
    0/19 ASSET 3320
    0/20 ASSET 3320
- Peering information
  - $ORIGIN unicast.ipv4.3.0.0.0.0.3.as.bgp.arpa.
    5539.import ASSET ANY
    5539.export ASSET 3.3
    6695.import ASSET as-decix.5.9.6.6.0.0.as.bgp.arpa.
    6695.export ASSET 3.3

IKS

ASSET Ressouce Record

- Deaggregated AS-Sets are huge
- References allow „auto-update“
- Aggregated AS-Sets are still large
- Fallback to TXT (exploiting NSEC)
- Transition mode: Allow, warn, and reject
- Everybody can create the origin and the peering information DNS right now

IKS

DNSSEC as PKI

- No crypto processing in routing devices
  - External validating resolver => AD bit
- Delegate maintainance to special Ops
- Private address space: Local zones
- Private peerings: Local secondaries
- No special software necessary

IKS

Bootstrap and Self-DDoS

- Self DDoS
  - „All“ routers ask when route flaps
  - Utilize peers cache by asking routers
- Boot1: Waiting for DNS before verifying
  - Accept routes and postphone verification
  - Redistribute only verified routes
- Boot2: Chaining peers
  - UUCP like DNS query routing

IKS

Securing BGP using DNSSEC

Questions?

IKS