An Overview of Secure Name Resolution

DNSSEC, DNSCurve and Namecoin

Matthäus Wander

<matthaeus.wander@uni-due.de>

29C3

Hamburg, December 29, 2012

UNIVERSITÄT DUISBURG ESSEN

Outline

- DNS Spoofing
- DNSSEC
  - Introduction
  - Deployment status
  - Implications
  - Root zone
- DNSCurve
- Namecoin
- Zooko's triangle

Typical DNS Query Path

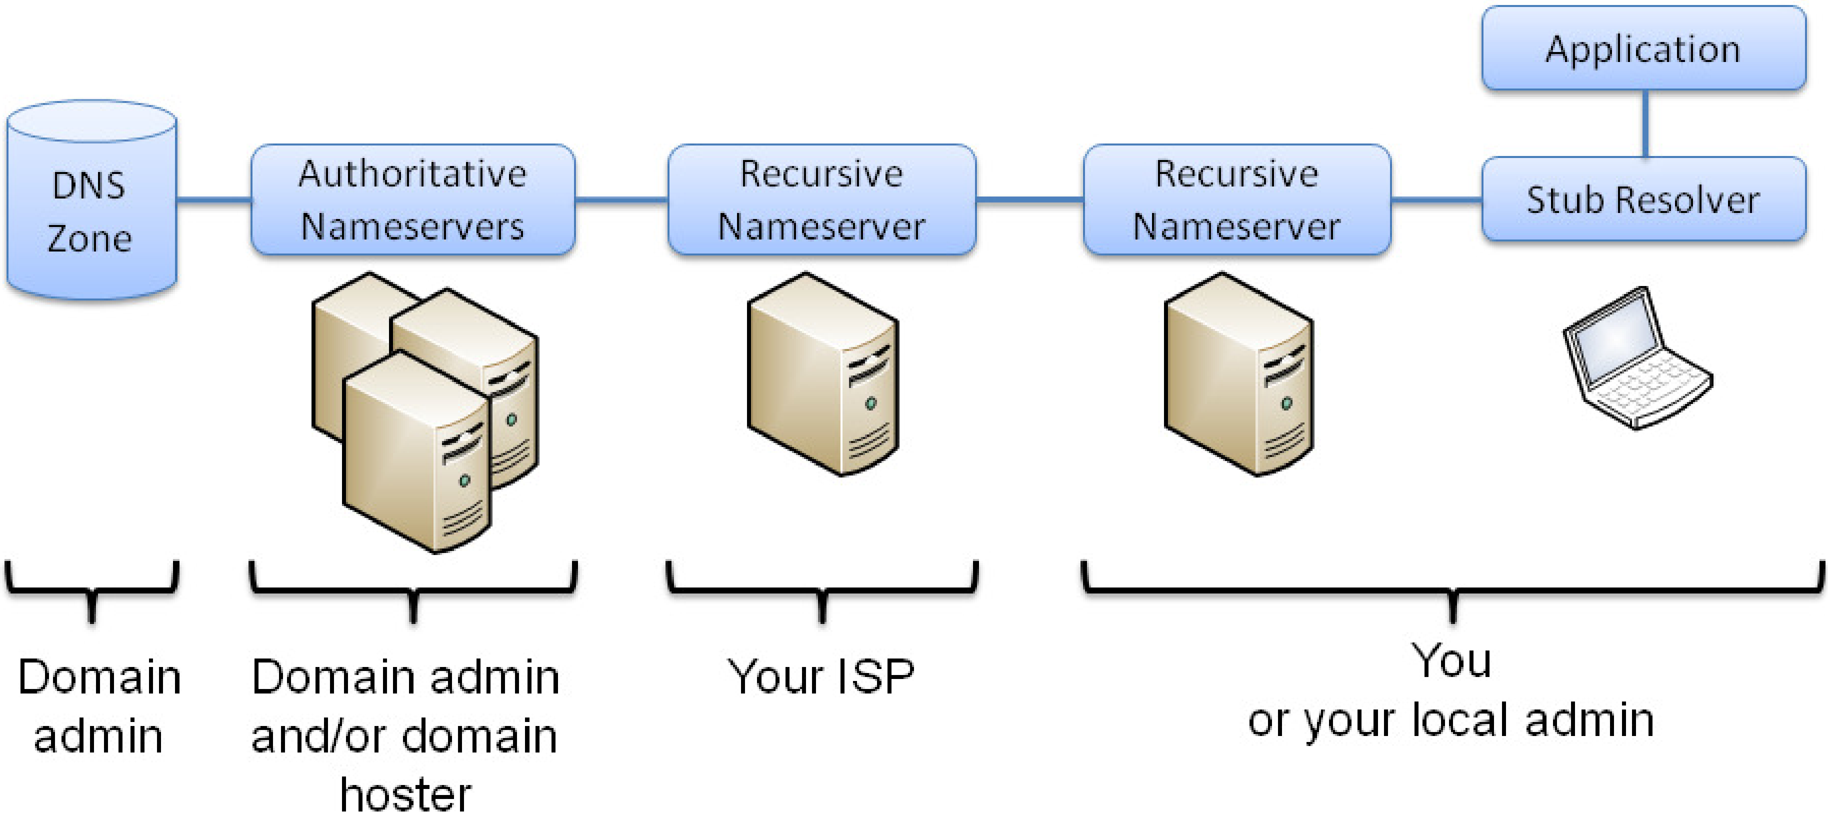

DNS Spoofing

- Attacker wants to spoof DNS response
- Remote UDP spoofing
  - Attacker triggers DNS queries on your machine (e.g. HTML link)
  - Mitigation: put random data into DNS query (transaction ID, source port)
  - Attacker must guess random data to spoof succesful response
  - Vulnerability: expensive attack
- Local UDP spoofing
  - Attacker is in your local network (e.g. Wi-Fi in coffee bar)
  - Mitigation: 
  - Vulnerability: easy attack

DNSSEC

- Domain Name System Security Extensions
- Uses cryptography to achieve data integrity and authenticity
  - Note: not confidentiality, not availability
- Sign resource records with private key
- Publish signatures as RRSIG record

```
example.net.    IN  A      1.2.3.4
example.net.    IN  RRSIG A 5 3 600 20120519... m1TWzfNDMg8NpgTo4i...
```

- Publish public key as DNSKEY record

```
example.net.    IN  DNSKEY   256 3 8 BQEAAAABv5hDo9fIU91cSFaDmnNPg...
```

- Tie DNSKEY with parent zone to create chain of trust

Secure Delegations

- DS record for secure delegation
  - Indicates whether child zone is signed
  - Contains hash of DNSKEY
  - DS record is signed, too

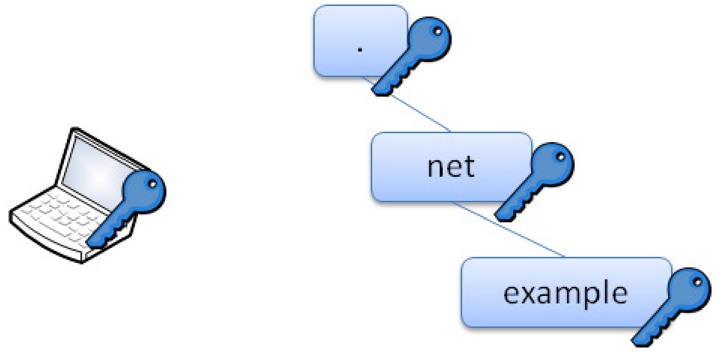

- Resolver must know a trust anchor (root key) beforehand

```
verteiltesysteme.net.      IN   NS    ns1.verteiltesysteme.net.
verteiltesysteme.net.      IN   NS    ns2.verteiltesysteme.net.
verteiltesysteme.net.      IN   DS    61908 5 1 3497D121F4C91369E95DC73D8...
verteiltesysteme.net.      IN   DS    61908 5 2 2F87866A60C3603F447658AC3...
verteiltesysteme.net.      IN   RRSIG DS 8 2 86400 20130103051550 2012122...

ns1.verteiltesysteme.net.  IN   A     134.91.78.139
ns2.verteiltesysteme.net.  IN   A     134.91.78.141
```

```
verteiltesysteme.net.      IN   DNSKEY 257 3 5 BQEAAAABy5oBPRz/mSEcFYXlcL...
```

Secure Denial of Existence

- DNSSEC signs resource records, not responses
- Negative responses (NXDOMAIN) have no records
- Sort names in canonical order
- Sign proof of non-existence

```
ftp    IN  NSEC   mail
```

- How to avoid zone disclosure?
- Sign salted hashes of domain names

```
3a45  IN  NSEC3  78a1
```

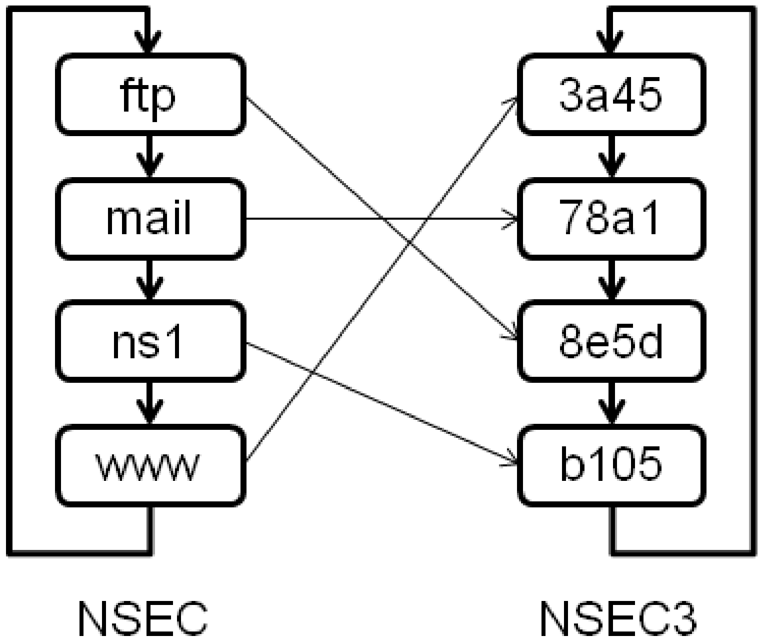

- Note: hash values can be reversed by offline dictionary attack [1]

Potential Secure Path of DNSSEC

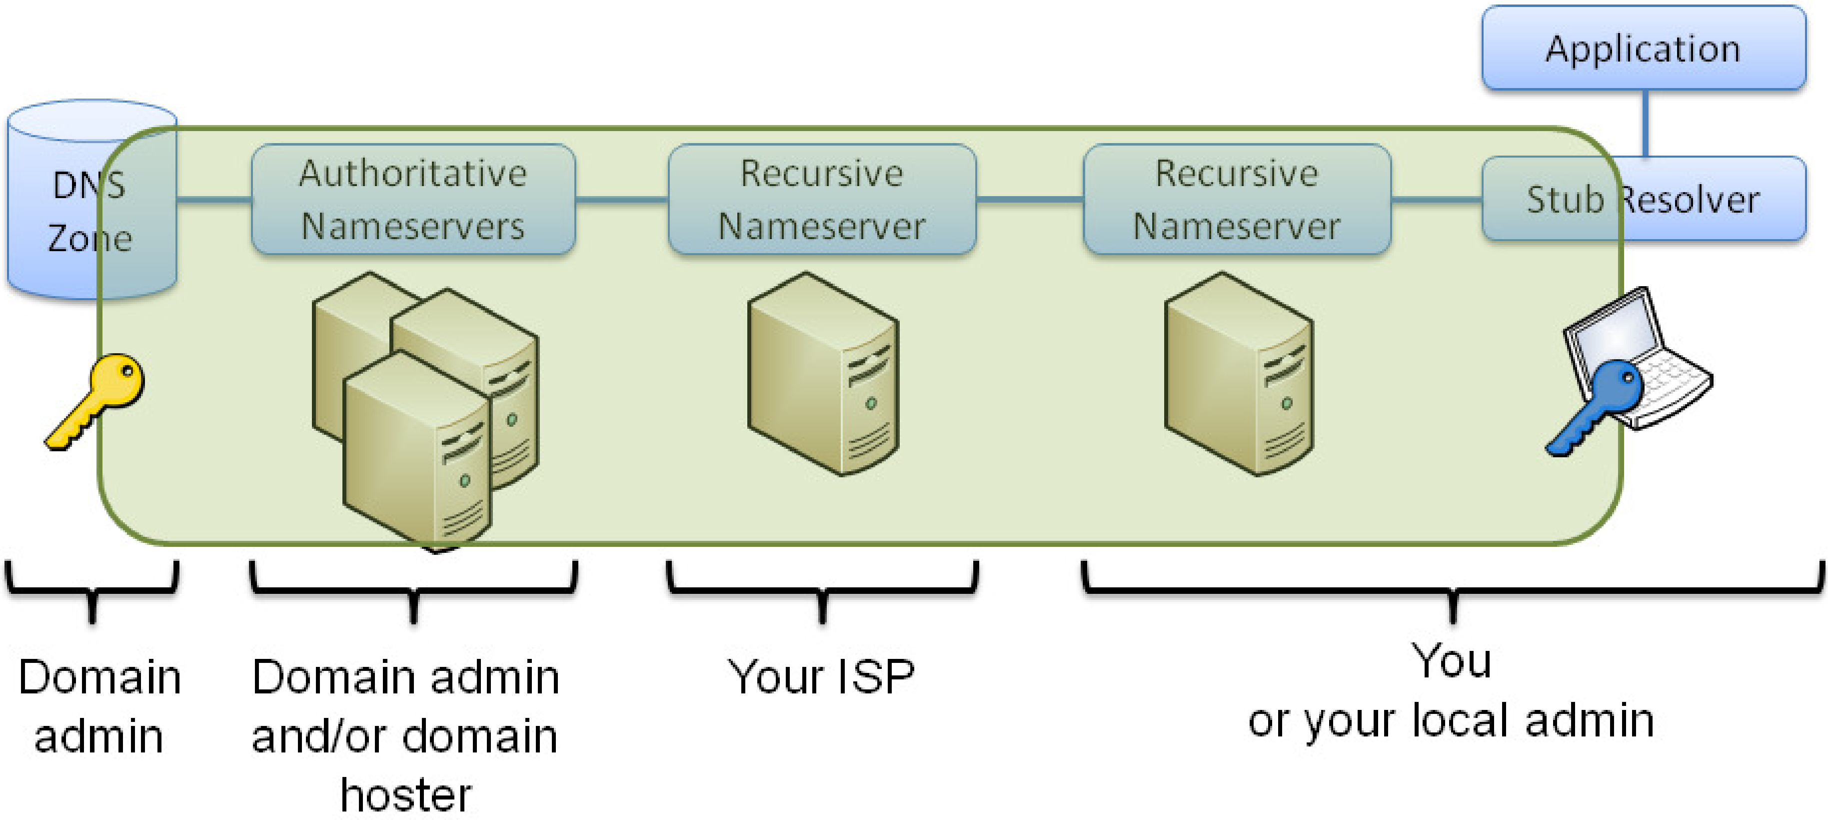

DNSSEC Deployment: Signed Zones

- Root zone is signed since July 2010
- 98/316 top-level domains are signed (31%) [2]
    - 10 more are signed without secure delegation in root

TLD	Signed	Total	Percentage	Reference
br	352k	3M	11%	[3]
com	139k	100M	0.1%	[4]
cz	380k	1M	38%	[5]
net	29k	15M	0.2%	[4]
nl	1.3M	5.1M	26%	[6] [7]
se	148k	1.3M	12%	[8]

Table 1: Number of signed second-level domains for selected TLDs

DNSSEC Deployment: Stub Resolvers

Stub Resolver	Built-in Valid.
Android 4.2	no
FreeBSD 9	no
GNU libc 2.16	no
iOS 6.0	no
Mac OS X 10.8	no
OpenBSD 5.2	no
Windows Phone 7	no
Windows XP SP3	no
Win Vista SP2	no
Windows 7 SP1	no, reads AD
Windows 8	no, reads AD

Alternatives:

- Run local nameserver
  - BIND, Unbound, dnssec-trigger
- Validating resolver libs are available
  - to link your application against it
- BIND9 on Debian 7 has validation enabled
  - expect name resolution problems

- AD flag ≙ *"server authenticated data successfully"*
  - like an inverted *evil bit* ☺ [9]
  - basically meaningless in insecure local networks

DNSSEC-capable Resolvers

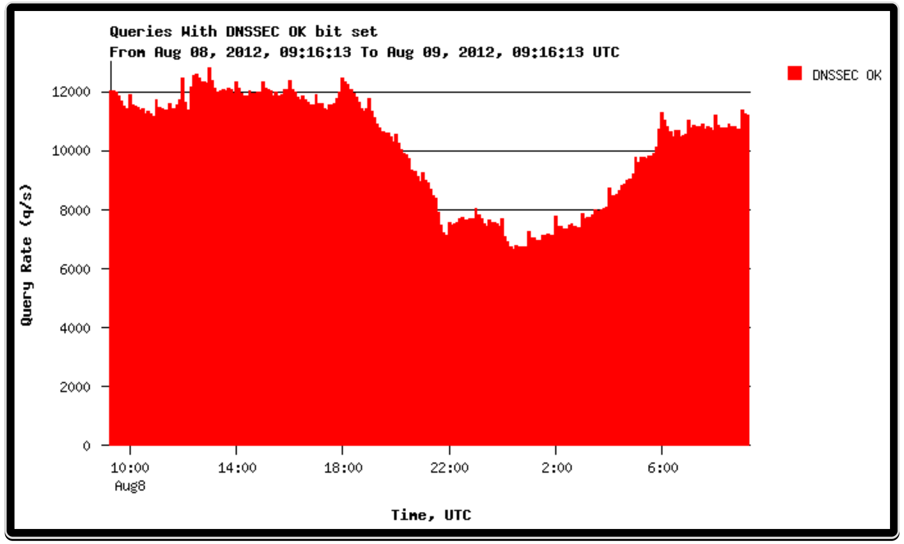

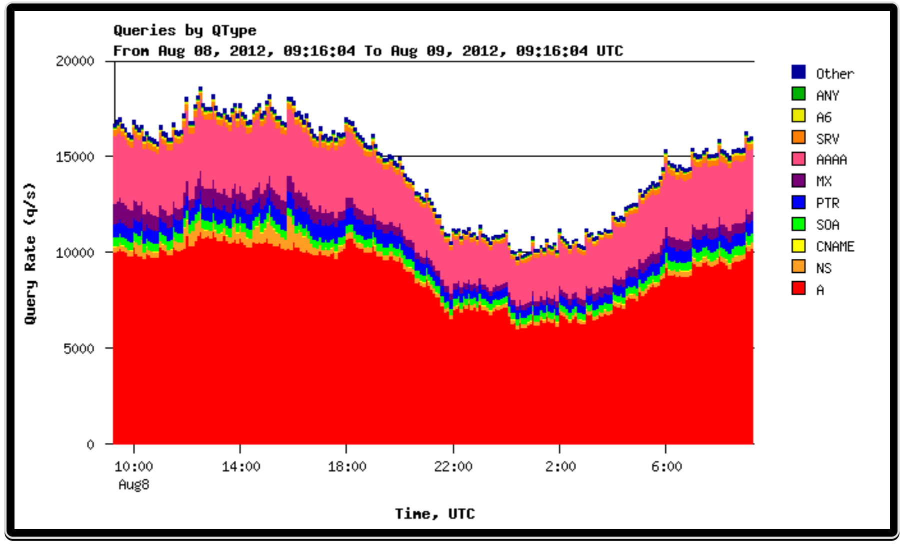

Figure 1: K-root nameserver statistics [10]

- ~70% of queries at K-root have DNSSEC OK (DO) flag set
- DO flag ≙ resolver claims to be DNSSEC-capable
- Note: says nothing about validation

DNSSEC Deployment: Clients

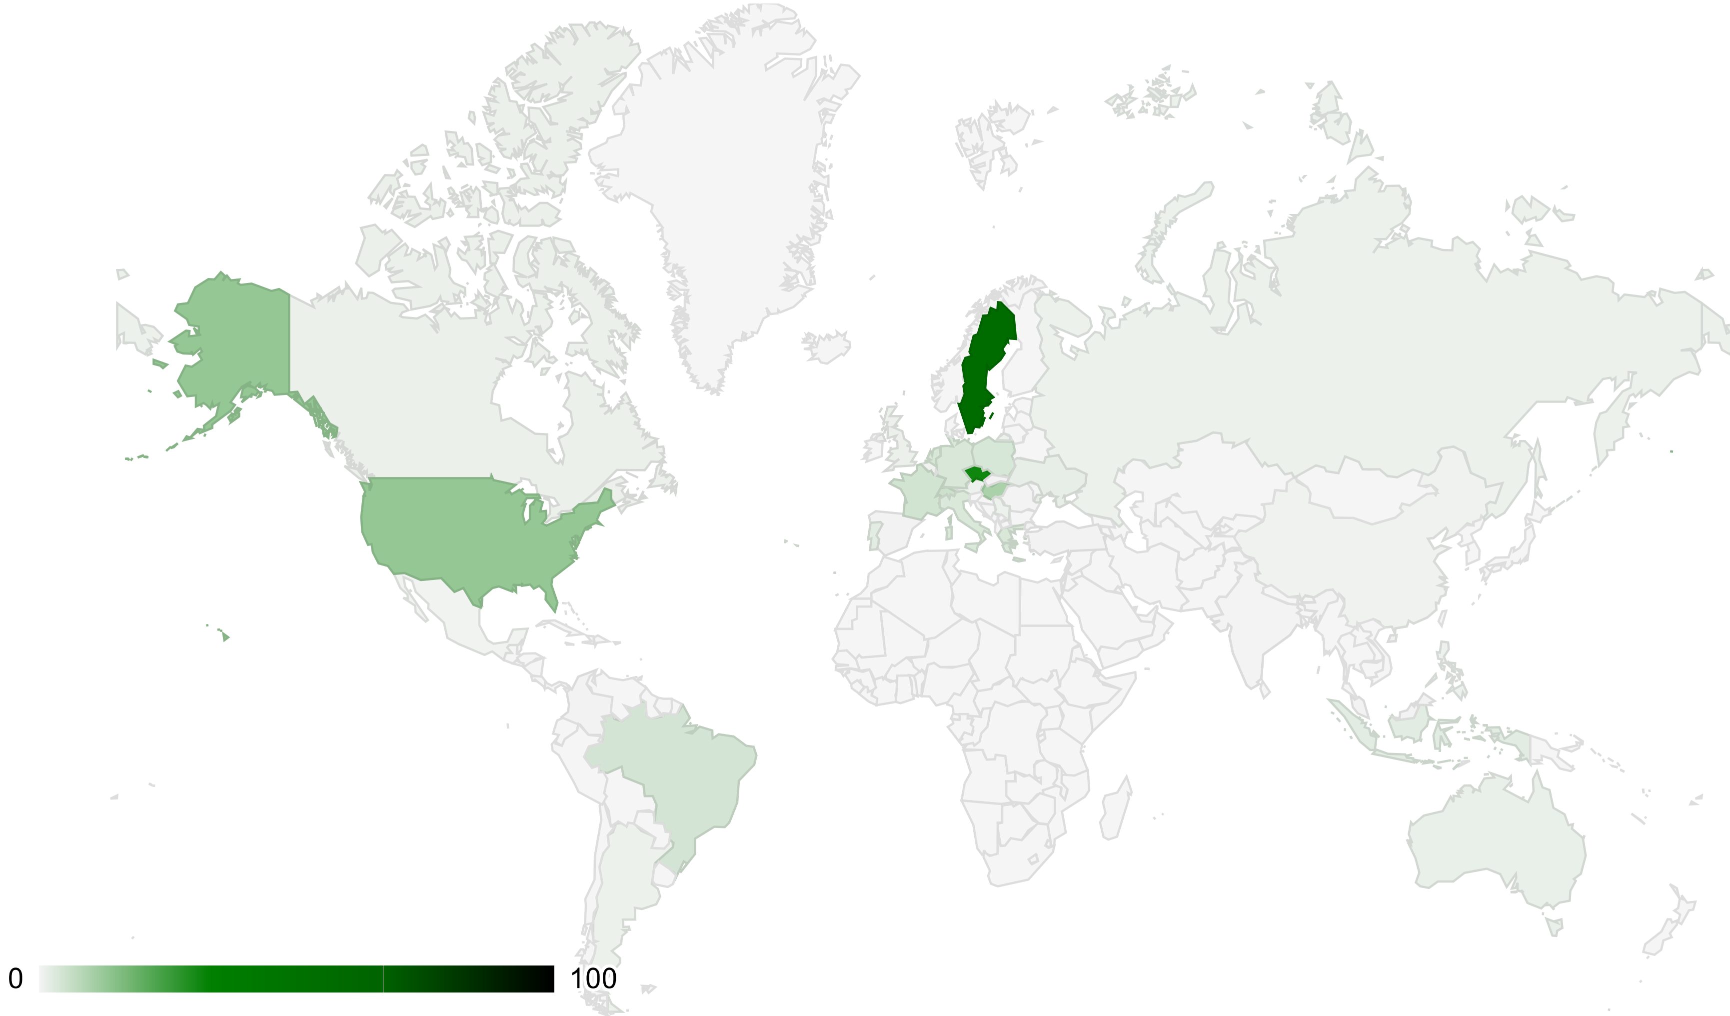

http://dnssec.vs.uni-due.de

Other tests: http://test.dnssec-or-not.net, http://dnssectest.sidn.nl

Implications of DNSSEC Deployment

- DNSSEC adds security but also complexity
  - Bert Hubert (PowerDNS): *“we keep finding DNSSEC corner cases that make the authors of the very RFCs swoon.”*
  - Roy Arends (Nominet UK): *“I have yet to be swooned by any of the DNSSEC corner cases you've found.”* [dns-operations]
- Validation failures look like general DNS failures
  - Unlike HTTPS no security warning and no way to override error
  - Stub resolver interface lacks validation information

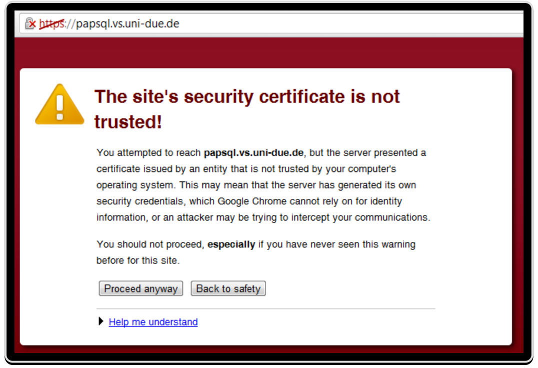

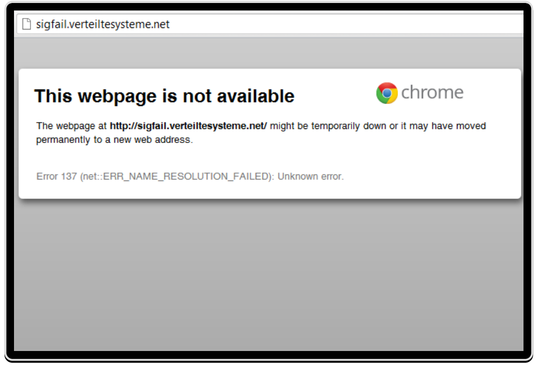

DNSSEC-related Outages

Date	Domain	Reason	Reference
2012-12-27	mil	signatures expired	[dnssec-deployment]
2012-12-07	arpa	APNIC reverse lookups failed after hardware fault	[dnssec-deployment]
2012-01-18	nasa.gov	KSK rollover failed	[11]
2011-07-25	nist.gov	no valid DNSKEY record	[dnssec-deployment]
2011-06-15	co.th	rollover from NSEC to NSEC3 failed	[dnssec-deployment]
2011-01-03	gi	signatures expired	[dnssec-deployment]
2010-10-07	be	signatures expired	[dnssec-deployment]
2010-09-15	mozilla.org	DS published before signed zone was online	[dnssec-deployment]
2010-09-11	uk	inconsistent ZSK after hardware fault	[12]

more: http://dns.comcast.net

- NASA.gov outage perceived by users: “Comcast Blocks Customer Access to NASA.gov” [13]
    - Comcast uses negative trust anchors (manual validation exemptions)

System Time vs. DNSSEC

- Keys do not expires
- Signatures have absolute validity periods
  - in addition to relative TTL from legacy DNS
  - typically on the order of days or weeks
- Desync system time → DNSSEC DoS
- Bootstrap system time via (S)NTP — how to resolve pool.ntp.org?
- Unsigned NTP domain name doesn't help
  - Root and top-level domain are signed
- Set up Anycast cloud as NTP fallback when DNS pool fails?

Amplification Attacks

- CPU load increases on validators but not that much on servers
  - offline + incremental signing
- Network load increases significantly
- Problem: DDoS'ers abuse public DNS for amplification attacks
  - becomes even more effective with DNSSEC (1:10 → 1:60)
- Cause: IP spoofing from botnets
- Solution: filter spoofed traffic near source (e.g. BCP 38)
  - Still too many networks with IP spoofing
- DNS-specific countermeasure: DNS rate limiting
- Trade-off: effective filtering vs. collateral damage

DNS Rate Limiting

- Naive approach: iptables rate limiting (usually **bad**)
  - either specific to one attack or easy to abuse (lock-out victim)
- Better approach: DNS Response Rate Limiting [14]
  - assumption: resolvers have a cache and retry in case of lost packet
  - track state for identical responses per IP address block
  - filter more than n identical responses per sec (n=5)
  - slip truncated response every m filtered packets to force TCP (m=2)
- Note: rate limiting protects amplification **targets** (not amplifiers)
  - Use overprovisioning + Anycast to protect your authoritative servers
- Note: not applicable for recursive servers
  - Use IP-based access control

ISP Wildcard Redirect

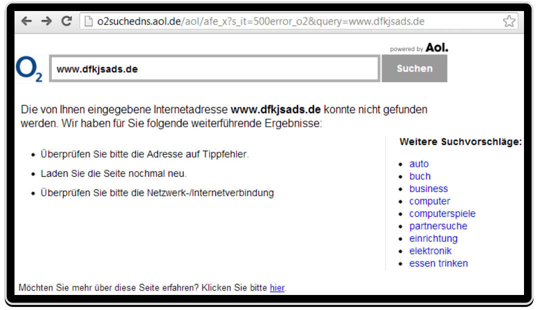

- NXDOMAIN redirection: point non-existent domain name to advertisement web page
- Redirection by ISP (aka: lie to your customer)
  - validating ISP: can add redirect after validation
  - valdiating client: will get SERVFAIL instead of NXDOMAIN
  - looks identical to user

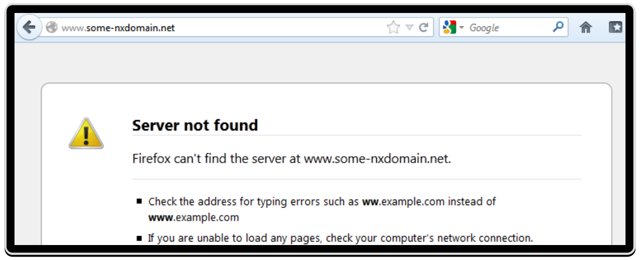

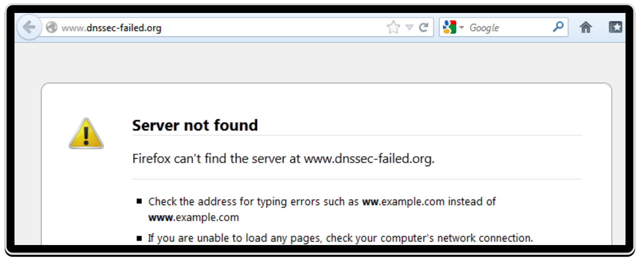

- Redirection by TLD operator (aka: VeriSign Site Finder)
  - would work: wildcards still possible with DNSSEC (but ugly [15] [16])

ISP Censorship Redirect

- Government-mandated ISP redirection
  - e.g. Zensurursula attack discussed in Germany
  - validating ISP: can add redirect after validation
  - validating client: will get SERVFAIL instead of A record
  - blocking still works but without notice
- If you are affected by this, do not use your ISP forwarders
- In general more reliable to run resolver without forwarders
  - allows to scatter retries among all authoritative servers
  - non-validating forwarders may cache bogus delegations

DNS Injection

- DNS injection: deep packet inspection to spoof DNS response
- Widely used within mainland China [17]
  - coarse-grained filter may match e.g. twitter.com.example.net
  - any source and destination IP addresses
- Affects also other countries which transit Chinese ASes
- With Anycast in root and TLD your packets take strange routes
- Study suggests open resolvers from 109 countries are affected
  - original packets do not seem to be suppressed
- DNSSEC validation protects from unsuppressed DNS injection
- With suppression validating resolver will retry another nameserver
  - will succeed if you have uncensored route to another nameserver

X.509 vs. DNSSEC

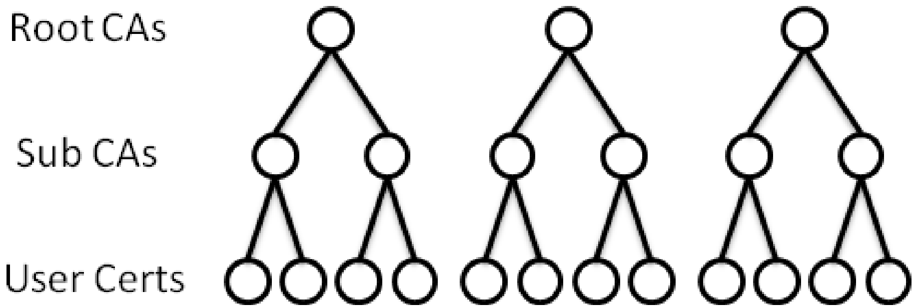

- 650 CA organizations [18]
- 1500 CA certificates
- Trusted by Microsoft or Mozilla

- X.509: all root CAs and sub CAs are fully trusted
- DANE/TLSA: put TLS certificate into DNS [19]
- DNSSEC: Trust is limited to domain
  - .com can't mess with .org
- DNS root can mess with anyone
  - Pro: trust in root limited to one organization
  - Con: power concentrated in one organization

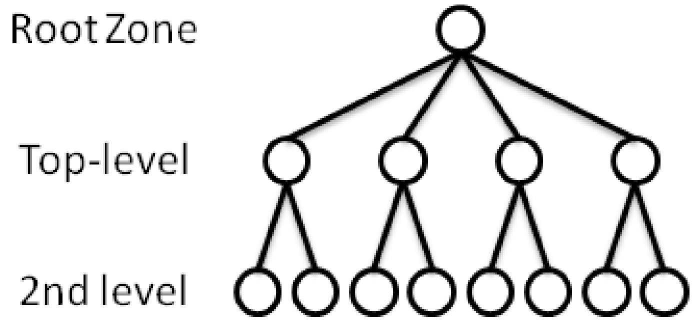

Trust Anchor

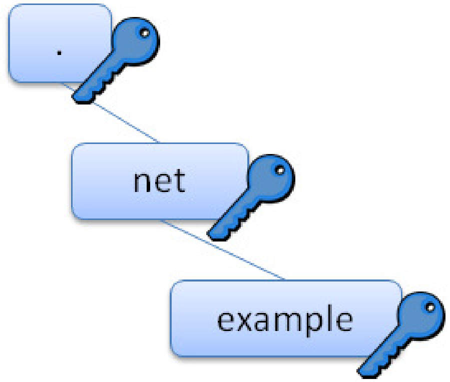

- Who can forge your 2nd-level domain?
  - Root zone operator
  - Registry/TLD operator
  - Registrar
- Configure other trust anchors in your resolver
  - for specific domains (if you don't trust the operators mentioned above)
  - for alternative DNS roots
- Automatic rollover of trust anchors [20]
  - add second DNSKEY to zone, wait some weeks, remove first DNSKEY
  - works if resolvers are online regularly and private key is not lost
  - does not initially retrieve trust anchor

Root Zone

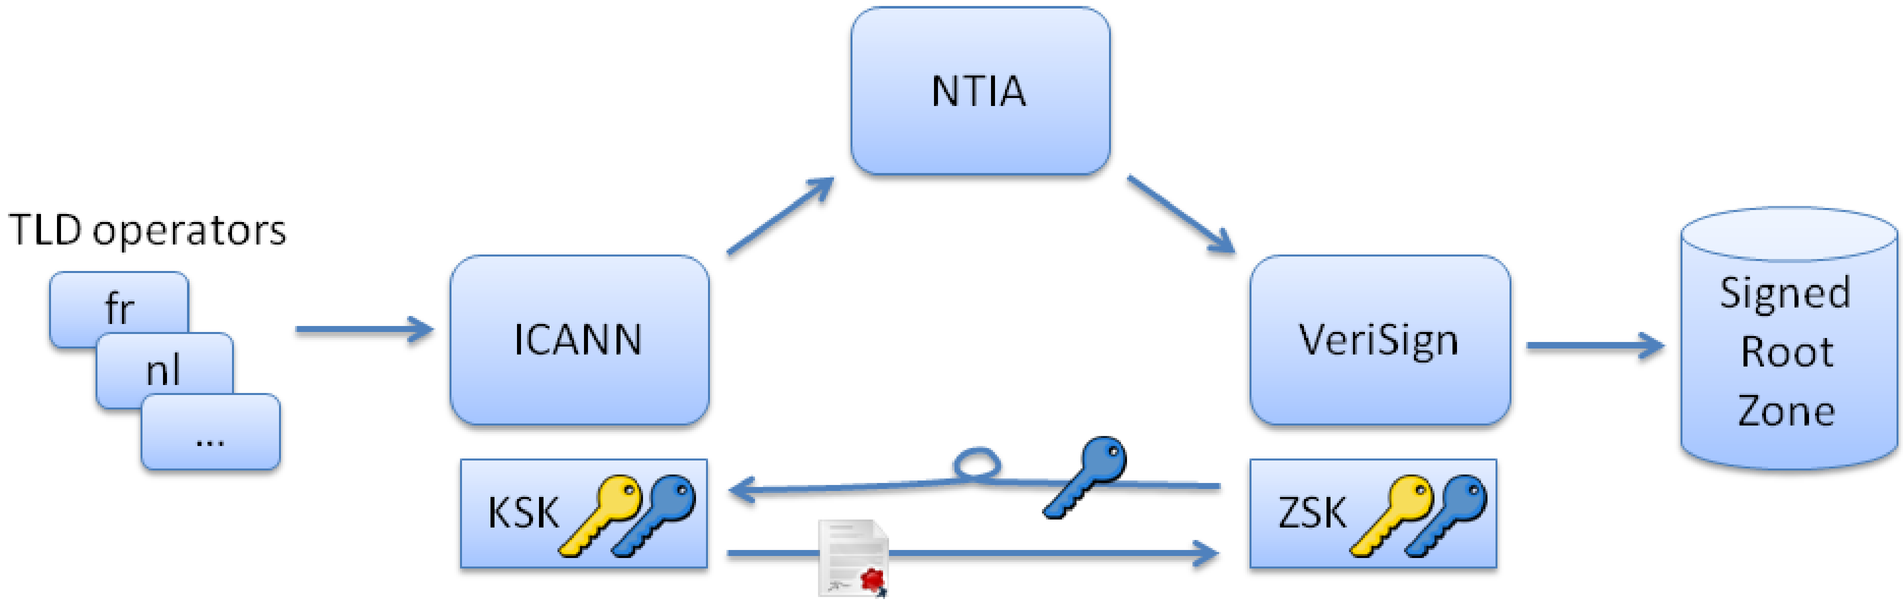

- IANA Functions Operator: ICANN
- Root Zone Administrator: NTIA (US government)
- Root Zone Maintainer: VeriSign
  - Also operates A-root and J-root

ICANN KSK Facilities

Figure 2: [21] [22]

- Two facilities in commercial data centers
  - West: 1920 E Maple Ave, El Segundo, CA 90245
  - East: 18155 Technology Dr, Culpeper, VA 22701
- Create and store KSK, sign ZSK
- ICANN, VeriSign & trusted community representatives

Key Ceremony

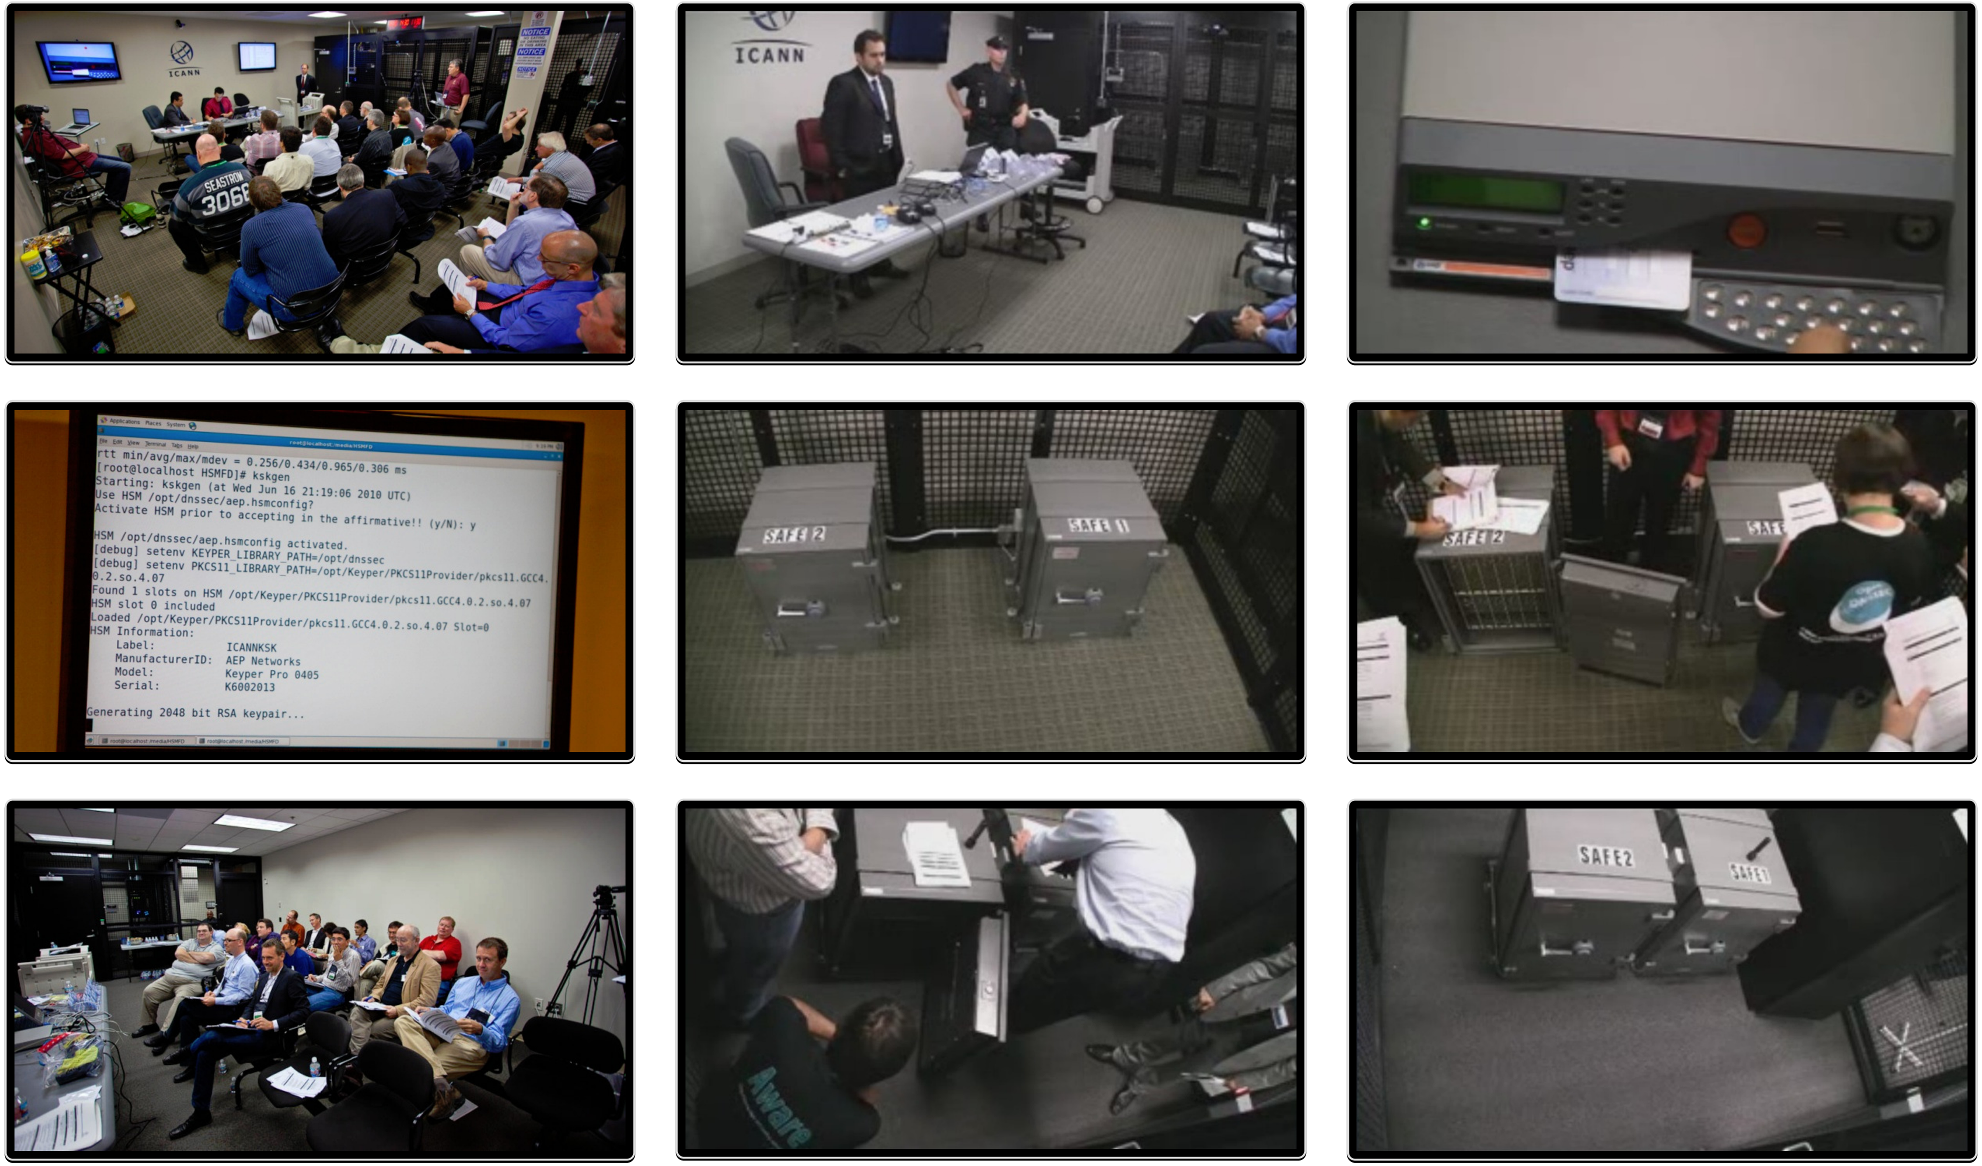

Figure 3: KSK Ceremonies 1 & 2, June & July 2010 [23] [24]

Access to Root KSK

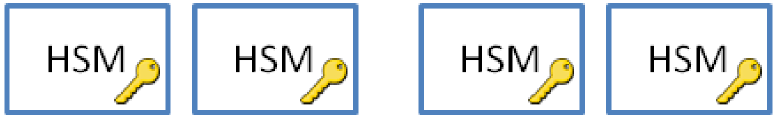

Access to Root KSK (2)

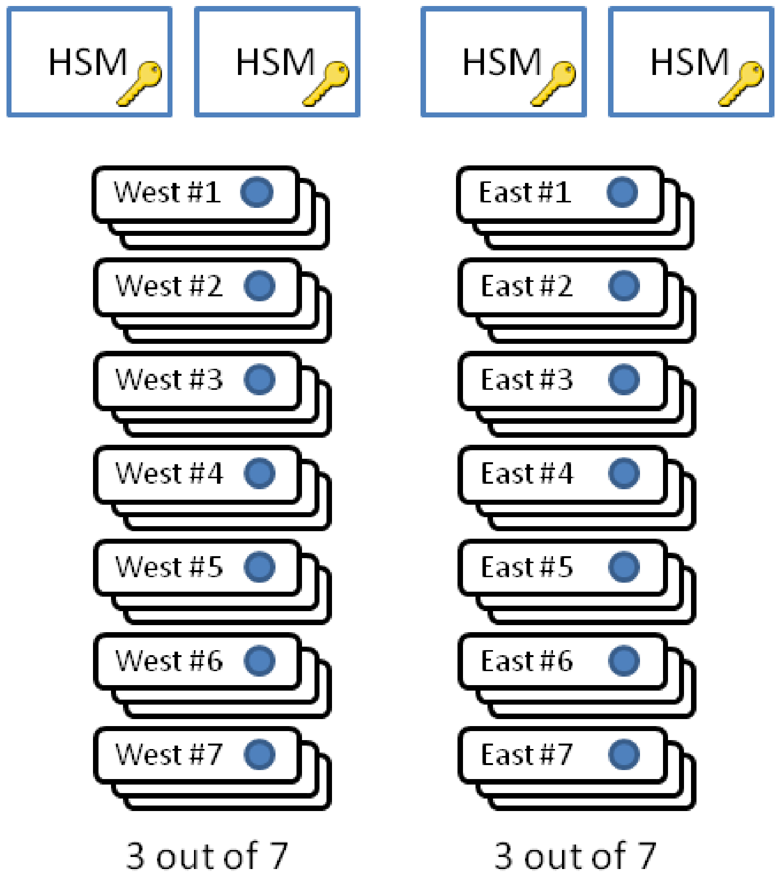

Access to Root KSK (3)

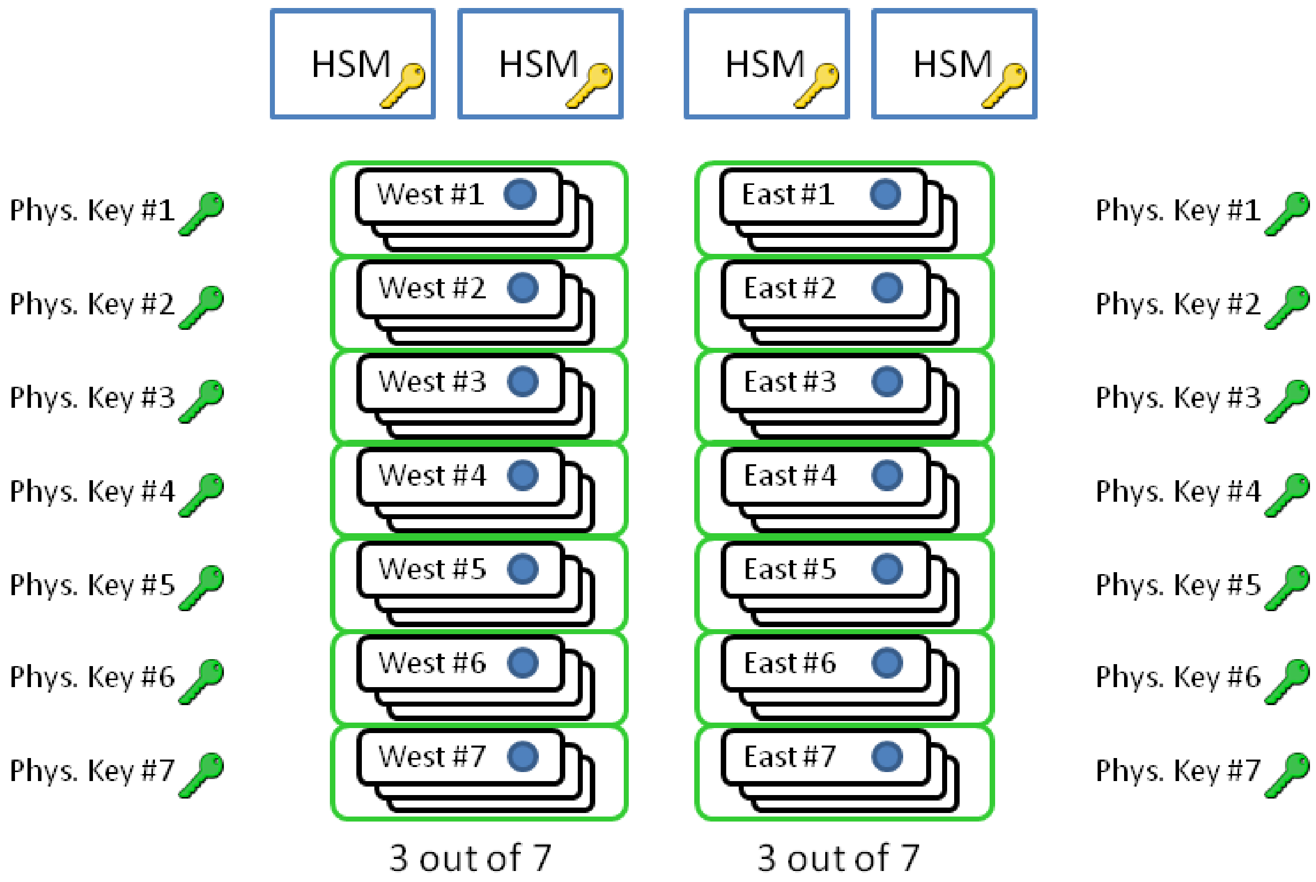

Access to Root KSK (4)

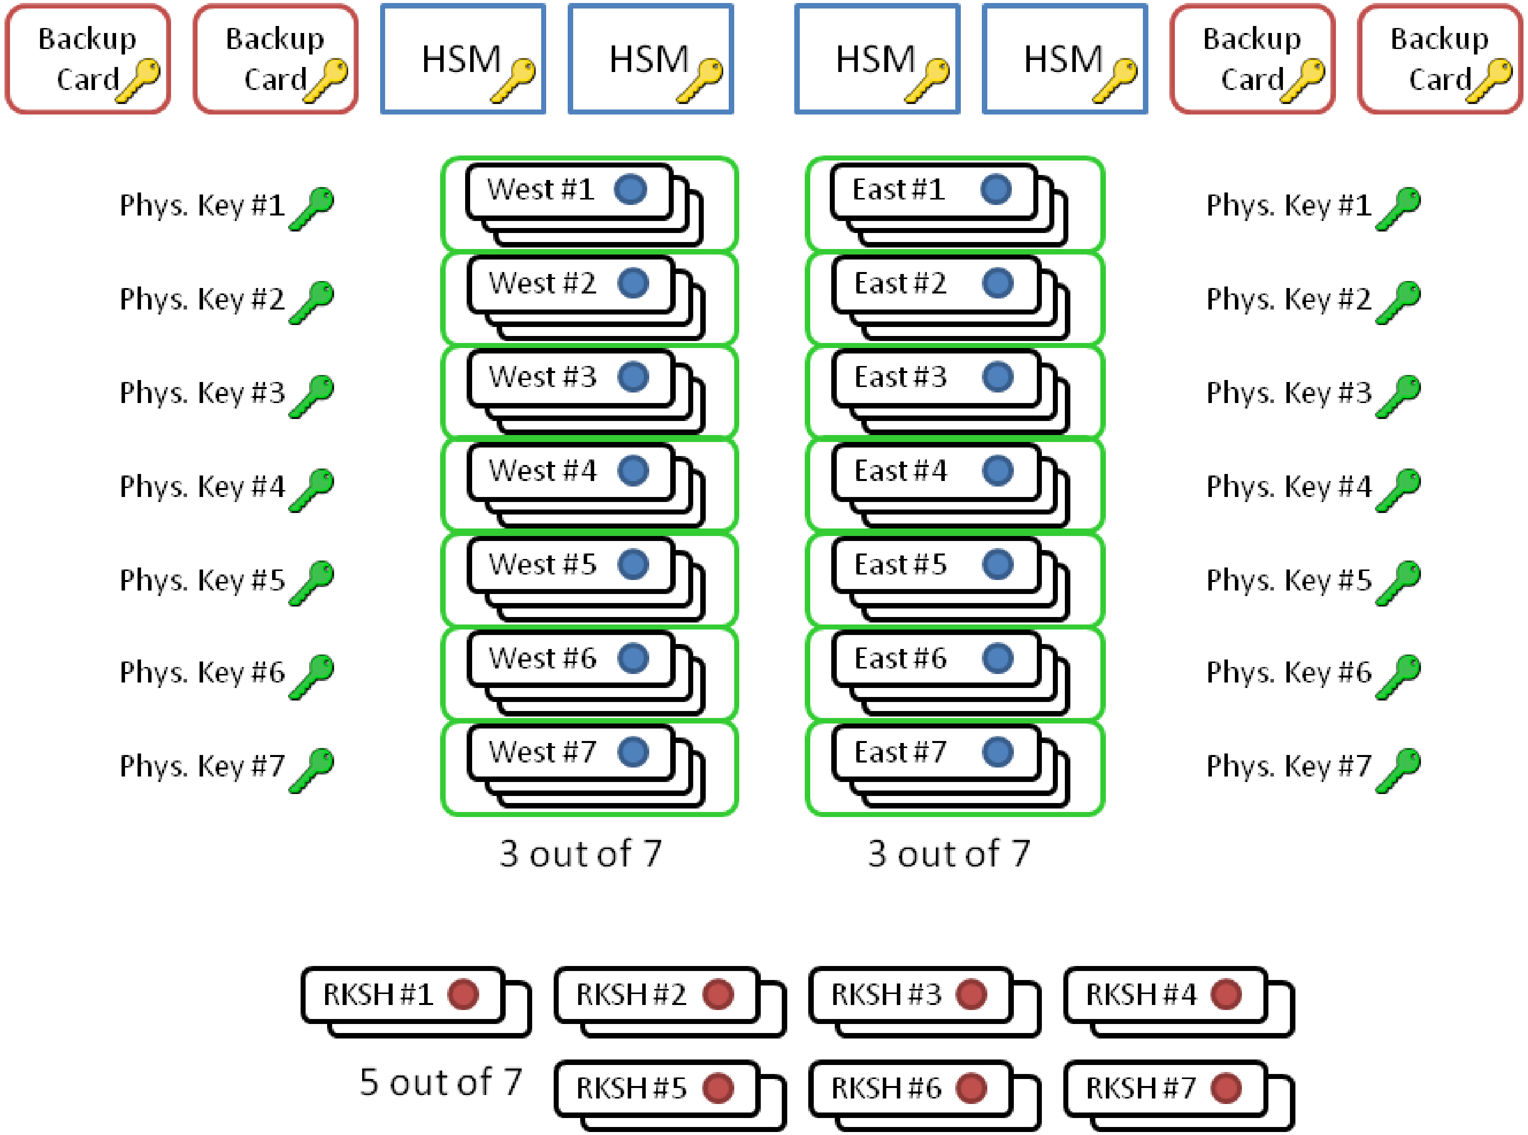

Access to Root KSK (5)

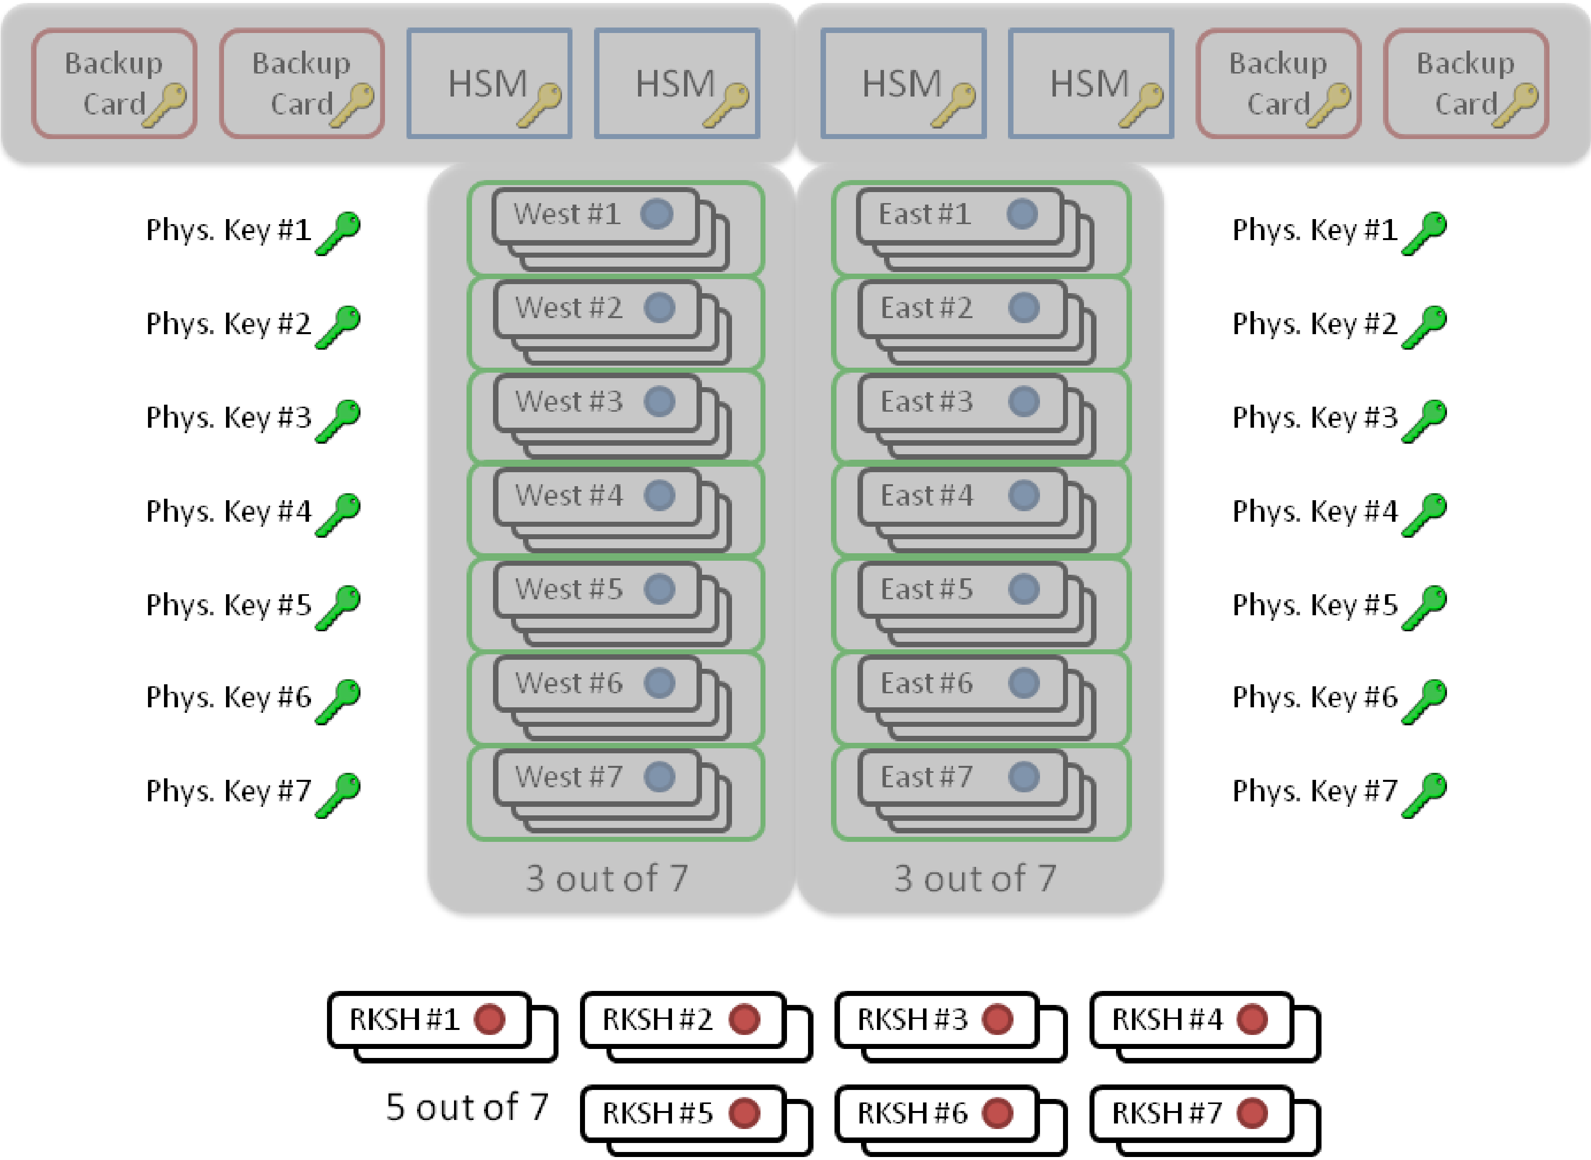

Root KSK

- RSA-2048/SHA-256: https://data.iana.org/root-anchors/ [25]
- Also signed by long-term ICANN keys for bootstrapping:
  1. X.509: RSA-2048/SHA-256, expires in 2029
  2. PGP: DSA-1024/SHA-1, key ID 0x0F6C91D2, no expiry date
- Rollover every 2-5 Years when appropriate (not scheduled)
- Private key owned by ICANN (stays in U.S.)
  - used every 3 months at KSK ceremony to sign new ZSK
- Offline operations, physical security
- HSM being used: AEP Keyper (€ 17,500 [26])
  - activated by 3 out of 7 smart cards

Root ZSK

- RSA-1024, SHA-256
- Rollover every 3 months
- Private key owned by VeriSign (stays in U.S.)
  - used twice daily to sign root zone
- Semi-automatic operations [27]
  - $\geq$2 trusted persons or $\geq$1 trusted person and an automated process
- HSM attached to production network
  - activated by 3 out of 16 smart cards
- Root zone signatures valid $\leq$10 days

DNSCurve

- Alternative concept to secure DNS [28] [29]
- Elliptic curve cryptography (ECC) instead of RSA
- Far less complex
- No new DNS resource records
- Keys are associated to nameservers, not zones
- Secures link between one authoritative server and one resolver
    - Unlike DNSSEC no end-to-end security
- Online cryptography instead of pre-generated signatures

Secure Path of DNSCurve

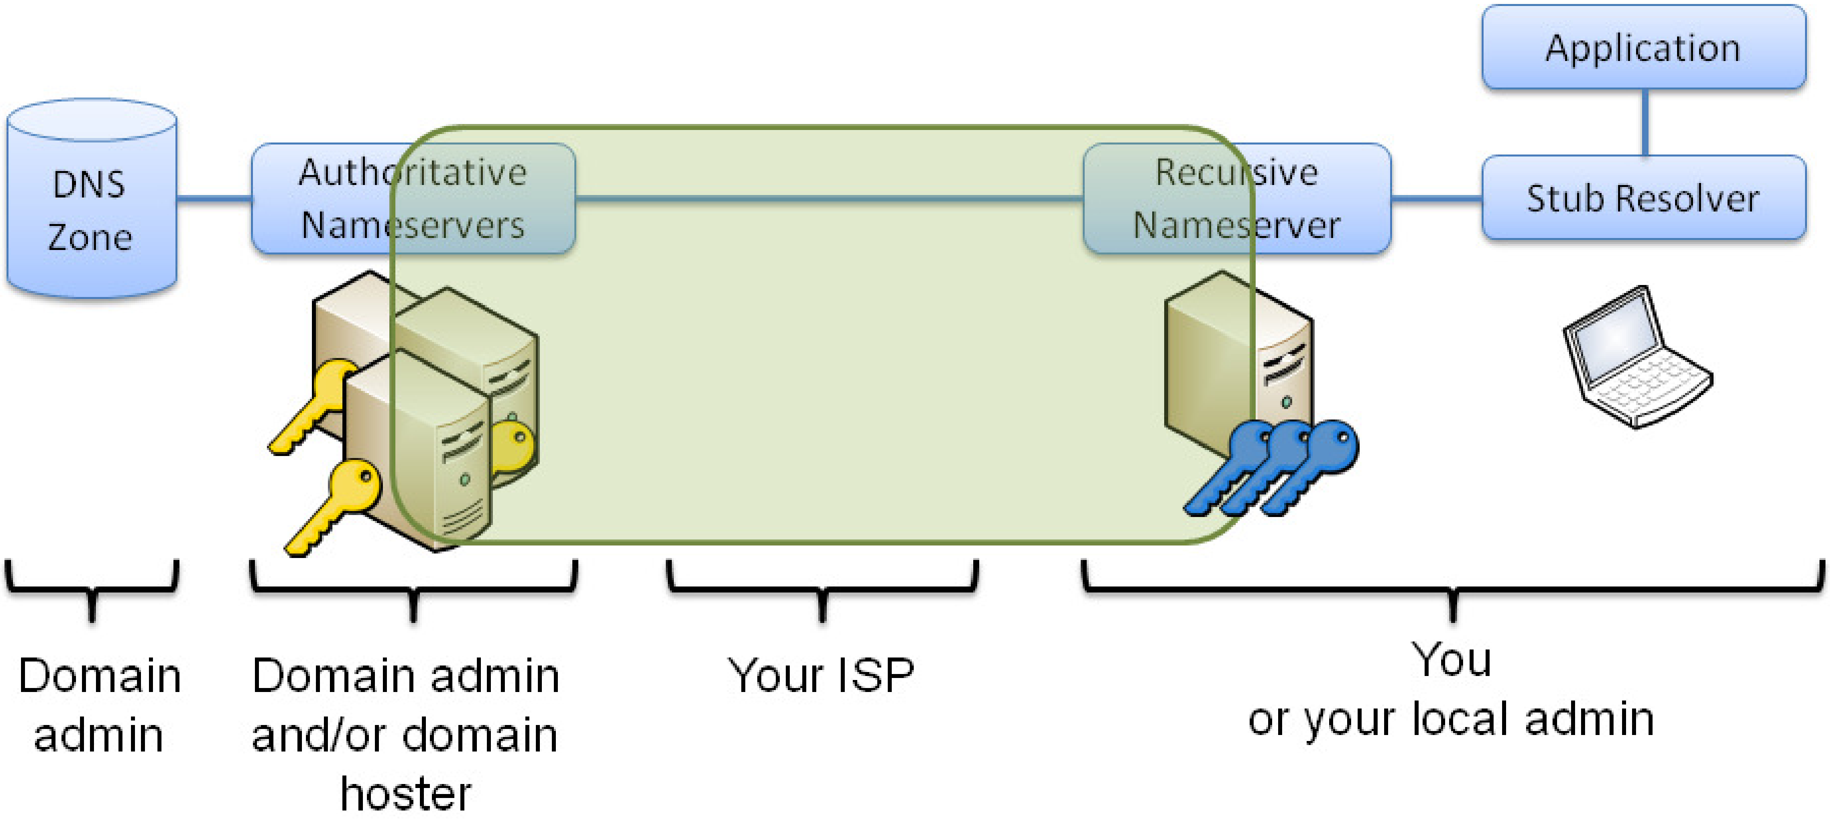

DNSCurve Messages

- New custom message format over UDP/53
  - also TXT tunneling for compatibility with strict firewalls
  - tunnel packets may be >512 bytes but EDNS is not used
- Put legacy DNS message into crypto box
- Each packet contains a nonce and is unique
  - replay attacks not possible
  - no expiration of signatures
  - system time doesn't need to be correct
  - NXDOMAIN secure without NSEC or other data
- Bonus: crypto boxes are encrypted
  - but: watch nameserver address, server name in TLS handshake etc.

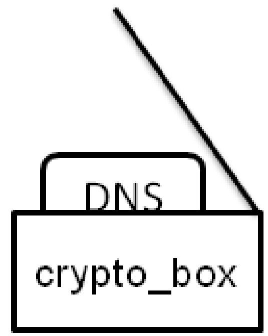

DNSCurve Cryptography

- Networking and Cryptography Lib (NaCl)
- ECC Curve25519 for Diffie-Hellman key exchange
  - 255 Bit public keys (in general faster than RSA)
  - shared key between resolver/nameserver can be cached and reused
  - other cryptographic operations are symmetric key
- Client: public key included in query
- Server: public key encoded as server name in parent zone

```
example.net.   IN  NS uz5wmnnvkbdd29t79yzg9fr2s2rx[...].example.net.
```

  - no extra resource record needed
  - secure if parent uses also DNSCurve

DNSCurve in Root Zone?

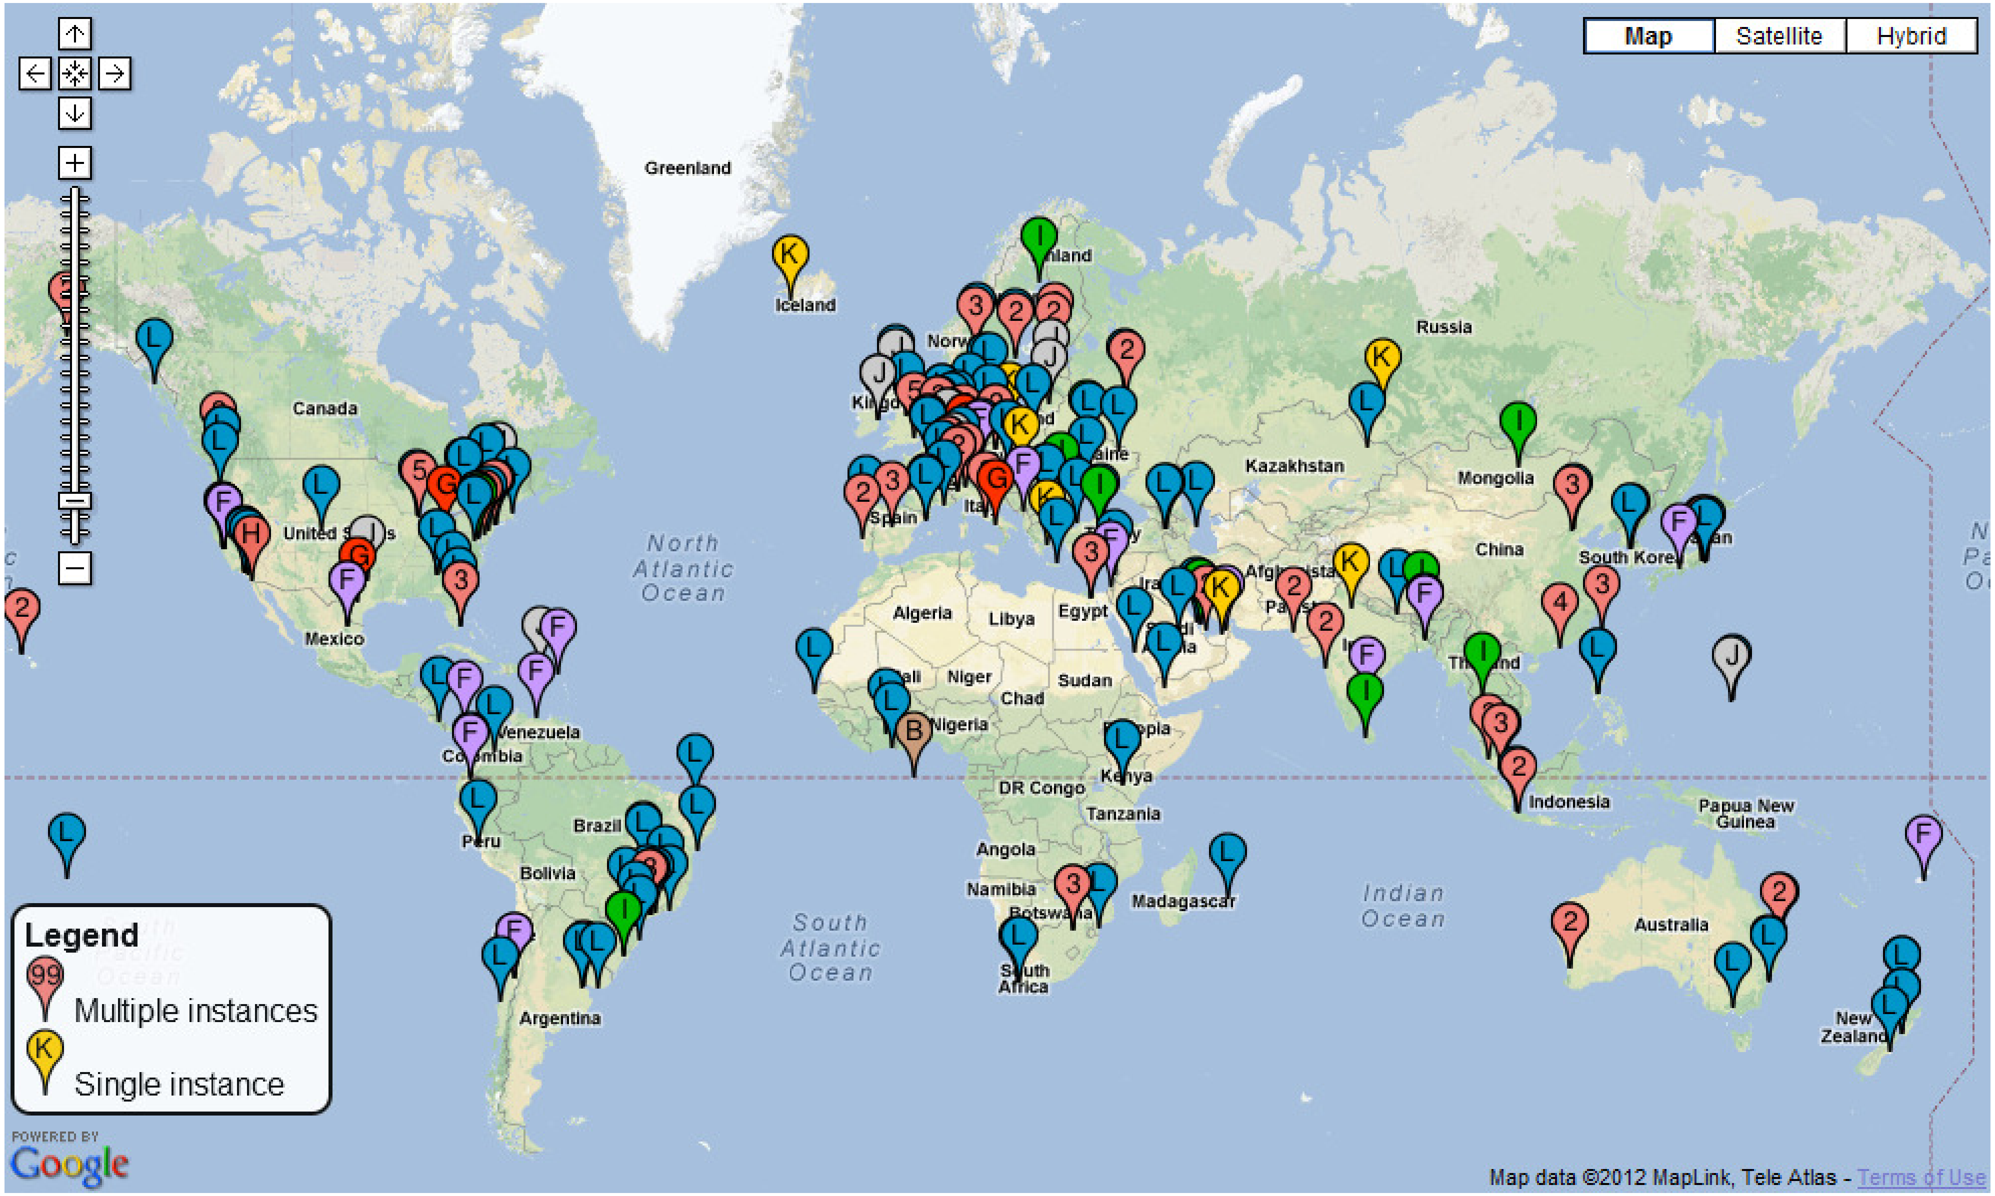

[30]

Deployment and Implications

- Private key must be online on nameserver
  - not feasible for root and top-level
- CPU exhaustion attack on authoritative servers → impact?
- Response size increases slightly
  - amplification factor comparable to legacy DNS
- No multi-hop caching → impact on TLD nameservers?
- DNSCurve happily carries DNSSEC-signed data
- Bummer: can't securely get *uz5* key from DNSSEC-signed parent
  - DNSSEC signs in delegations only DS records, not server names
- How to securely retrieve DNSCurve public key?

Namecoin

- Peer-to-peer-based naming system [31]
  - namespace controlled by majority, not centralized instance
- Bitcoin fork with all basic currency functions [32]
  - miners generate namecoins by solving hash puzzles
  - users send namecoins to each other, signed with ECDSA
  - all transactions are publicly shared by all users
- Transactions to store and update name data
  - in general arbitrary name/value data (255+1023 bytes)
  - primary use case is DNS-like data
  - small namecoin fee for each transaction
- Names expire if not refreshed within 250d

Resolving .bit Names

- Domain names are under virtual .bit TLD
  - not assigned in ICANN root (also not applied for as new gTLD)
- All users in Namecoin P2P network share a copy of all names
  - Namecoin ensures integrity → local secure name lookup
- How can outsiders resolve .bit names? (e.g. mobile devices)
- Point domain search suffix to Namecoin DNS gateway
  - **bad**, some guy on the Internet will get your NXDOMAIN queries
- Use public Namecoin DNS gateway as resolver
  - **worse**, some guy on the Internet will get all your DNS queries
- No secure .bit lookup for outsiders
  - and incompatible with DNSSEC: root says there is no .bit

Zooko's Triangle

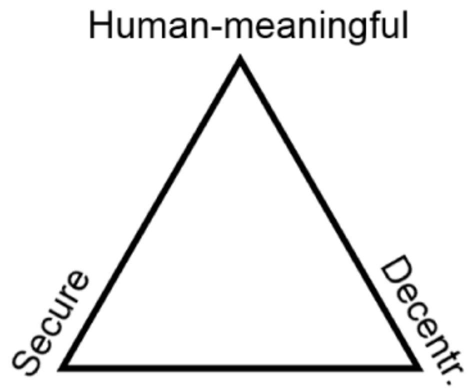

- Desirable properties of a naming system:
  1. secure (i.e. ensures integrity)
  2. decentralized
  3. human-meaningful
- Claim: any naming system can fulfill at most two of them [33]
- DNSSEC: secure with human-meaningful names
  - Not decentralized, instead hierarchical with powerful root
- Namecoin: decentralized with human-meaningful names
  - also secure if you participate in the P2P system
  - but what about scalability and efficiency?

Secure Name Resolution?

[34]

References

[1] DNSCurve: The nsec3walker tool, 2011-01-03
[2] ICANN: TLD DNSSEC Report, 2012-12-26
[3] Registro.br: Domínios Registrados por DPN, 2012-12-26
[4] VeriSign: Domains Secured with DNSSEC, 2012-12-26
[5] CZ.NIC: Statistics, 2012-12-25
[6] PowerDNS: Total number of DNSSEC delegations in the .NL zone, 2012-12-01
[7] SIDN: Statistics, 2012-12-01
[8] .SE: Domain Growth per Type, 2012-12-26
[9] RFC 3514: The Security Flag in the IPv4 Header, 2003-04-01
[10] RIPE NCC: Status for k.root-servers.net, 2012-08-09
[11] Comcast DNS: Analysis of NASA.GOV Validation Failure, 2012-01-24
[12] Simon McCalla: DNSSEC incident report, 2010-09-24
[13] Keith Cowing (NASA Watch): Comcast Blocks Customer Access to NASA.gov, 2012-01-18

References (2)

[14] P. Vixie, V. Schryver: DNS Response Rate Limiting (DNS RRL), 2012-06

[15] Ondřej Caletka: Wildcard domains DNSSEC resolver test

[16] Red Hat Bugzilla: Bug 824219

[17] Anonymous: The Collateral Damage of Internet Censorship by DNS Injection, 2012-07-03

[18] P. Eckersleyer & J. Burns: Is the SSLiverse a Safe Place?, 2010

[19] RFC 6698: The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA, 2012-08

[20] RFC 5011: Automated Updates of DNS Security (DNSSEC) Trust Anchors

[21] Image credit: Microsoft Bing Maps

[22] Image credit: Terremark Inc.

[23] Image credit: Kim Davies, KSK Ceremony 1, 2010-06-16

[24] Image credit: ICANN, http://data.iana.org/ksk-ceremony/

References (3)

[25] Fingerprint of root KSK as of 2012-12-26:
“. IN DS 19036 8 2 49AAC11D7B6F6446702E54A1607371607A1A41855200FD2CE1CDDE32F24E8FB5”

[26] Jakob Schlyter: Hardware Security Modules

[27] T. Okubo et al.: DNSSEC Practice Statement for the Root Zone ZSK operator, 2010-05-28

[28] http://dnscurve.org/

[29] Matthew Dempsky: DNSCurve: Link-Level Security for the Domain Name System, 2010-02-26

[30] Image credit: http://root-servers.org & Google Maps, 2012-12-27

[31] http://dot-bit.org

[32] Matthäus Wander: How Bitcoin Works, 2011-06-29

[33] Zooko Wilcox-O'Hearn: Names: Decentralized, Secure, Human-Meaningful: Choose Two, 2003-09-22

[34] Image credit: Sven Wolter, Wikimedia Commons