

Network Security and Measurement

- Securing Names with DNSSEC -

Prof. Dr. Thomas Schmidt

<http://inet.haw-hamburg.de> | t.schmidt@haw-hamburg.de

Outline

1. The Attack Surface of the DNS
2. The Design of DNSSEC
3. DNSSEC Deployment
4. Orthogonal Approaches


Introduction to

THE ATTACK SURFACE OF THE DNS


Attacking Names



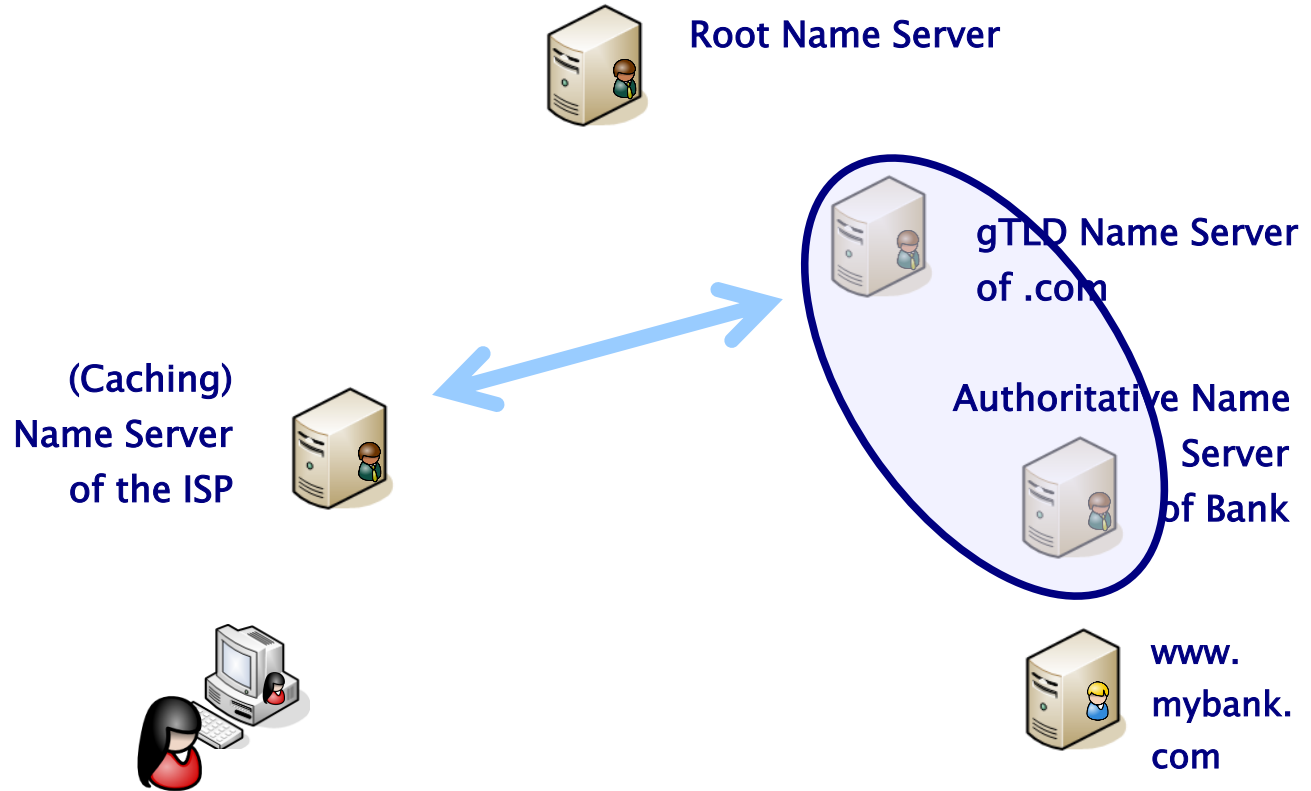
Attack Vectors (1)



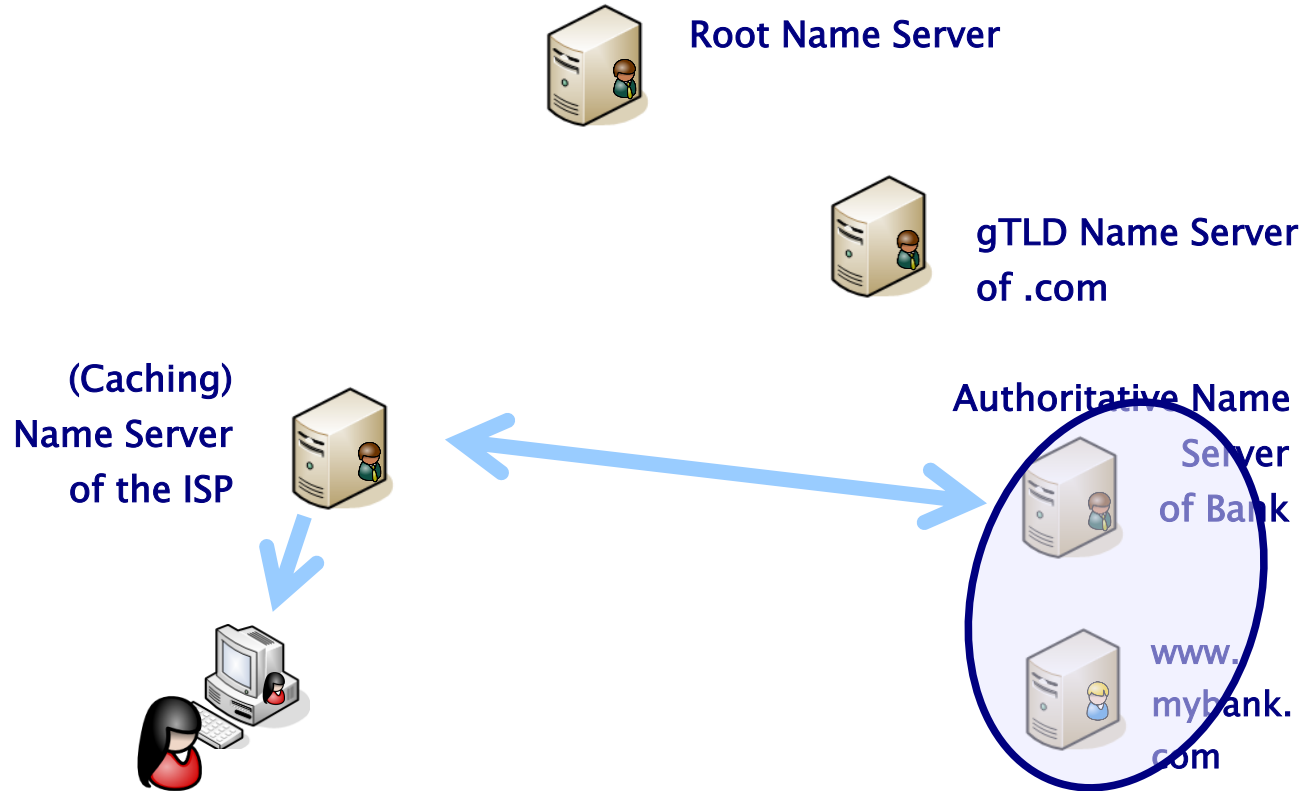
Attack Vectors (2)




Attack Vectors (3)



Attack Vectors (4)



Attack Surface of the DNS



RFC 973 on Trust in DNS, Jan. 1986



Paul Mockapetris

UDP checksums

Many versions of UNIX generate incorrect UDP checksums, and most ignore the checksum of incoming UDP datagrams. The typical symptom is that your UNIX domain code works fine with other UNIXes, but won't communicate with TOPS-20 or other systems. (JEEVES, the TOPS-20 server used for 3 of the 4 root servers, ignores datagrams with bad UDP checksums.)

Making up data

There are lots of name servers which return RRs for the root servers with 99999999 or similar large values in the TTL. For example, some return RRs that suggest that ISIF is a root server. (It was months ago, but is no longer.)

One of the main ideas of the domain system is that everybody can get a chunk of the name space to manage as they choose. However, you aren't supposed to lie about other parts of the name space. Its OK to remember about other parts of the name space for caching or other purposes, but you are supposed to follow the TTL rules.

Now it may be that you put such records in your server or whatever to ensure a server of last resort. That's fine. But if you export these in answers to queries, you should be shot. These entries get put in caches and never die.

Suggested domain meta-rule:

If you must lie, lie only to yourself.

RFC 973 on Trust in DNS, Jan. 1986



Paul Mockapetris

RFC 973

January 1986

Domain System Changes and Observations

UDP checksums

Many versions of UNIX generate incorrect UDP checksums, and most ignore the checksum of incoming UDP datagrams. The typical symptom is that your UNIX domain code works fine with other UNIXes, but won't communicate with TOPS-20 or other systems. (JEEVES, the TOPS-20 server used for 3 of the 4 root servers, ignores datagrams with bad UDP checksums.)

Making up data

There are lots of name servers which return RRs for the root servers with 99999999 or similar large values in the TTL. For example, some return RRs that suggest that ISIF is a root server. (It was months ago, but is no longer.)

One of the main ideas of the domain system is that everybody can get a chunk of the name space to manage as they choose. However, you aren't supposed to lie about other parts of the name space. Its OK to remember about other parts of the name space for caching or other purposes, but you are supposed to follow the TTL rules.

Now it may be that you put such records in your server or whatever to ensure a server of last resort. That's fine. But if you export these in answers to queries, you should be shot. These entries get put in caches and never die.

Suggested domain meta-rule:

If you must lie, lie only to yourself.

RFC 973 on Trust in DNS, Jan. 1986



Paul Mockapetris

RFC 973

January 1986

Domain System Changes and Observations

UDP checksums

Many versions of UNIX generate incorrect UDP checksums, and most ignore the checksum of incoming UDP datagrams. The typical symptom is that your UNIX domain code works fine with other UNIXes, but won't communicate with TOPS-20 or other systems. (JEEVES, the TOPS-20 server used for 3 of the 4 root servers, ignores datagrams with bad UDP checksums.)

Making up data

There are lots of name servers which return RRs for the root servers with 99999999 or similar large values in the TTL. For example, some return RRs that suggest that ISIF is a root server. (It was months ago, but is no longer.)

One of the main ideas of the domain system is that everybody can get a chunk of the name space to manage as they choose. However, you aren't supposed to lie about other parts of the name space. Its OK to remember about other parts of the name space for caching or other purposes, but you are supposed to follow the TTL rules.

Now it may be that you put such records in your server or whatever to ensure a server of last resort. That's fine. But if you export these in answers to queries, you should be shot. These entries get put in caches and never die.

Suggested domain meta-rule:

If you must lie, lie only to yourself.

RFC 973 on Trust in DNS, Jan. 1986



Paul Mockapetris

Problem: This guideline does not comply anymore with malicious activities on the Internet.

RFC 973

January 1986

Domain System Changes and Observations

UDP checksums

Many versions of UNIX generate incorrect UDP checksums, and most ignore the checksum of incoming UDP datagrams. The typical symptom is that your UNIX domain code works fine with other UNIXes, but won't communicate with TOPS-20 or other systems. (JEEVES, the TOPS-20 server used for 3 of the 4 root servers, ignores datagrams with bad UDP checksums.)

Making up data

There are lots of name servers which return RRs for the root servers with 99999999 or similar large values in the TTL. For example, some return RRs that suggest that ISIF is a root server. (It was months ago, but is no longer.)


One of the main ideas of the domain system is that everybody can get a chunk of the name space to manage as they choose. However, you aren't supposed to lie about other parts of the name space. Its OK to remember about other parts of the name space for caching or other purposes, but you are supposed to follow the TTL rules.

Now it may be that you put such records in your server or whatever to ensure a server of last resort. That's fine. But if you export these in answers to queries, you should be shot. These entries get put in caches and never die.

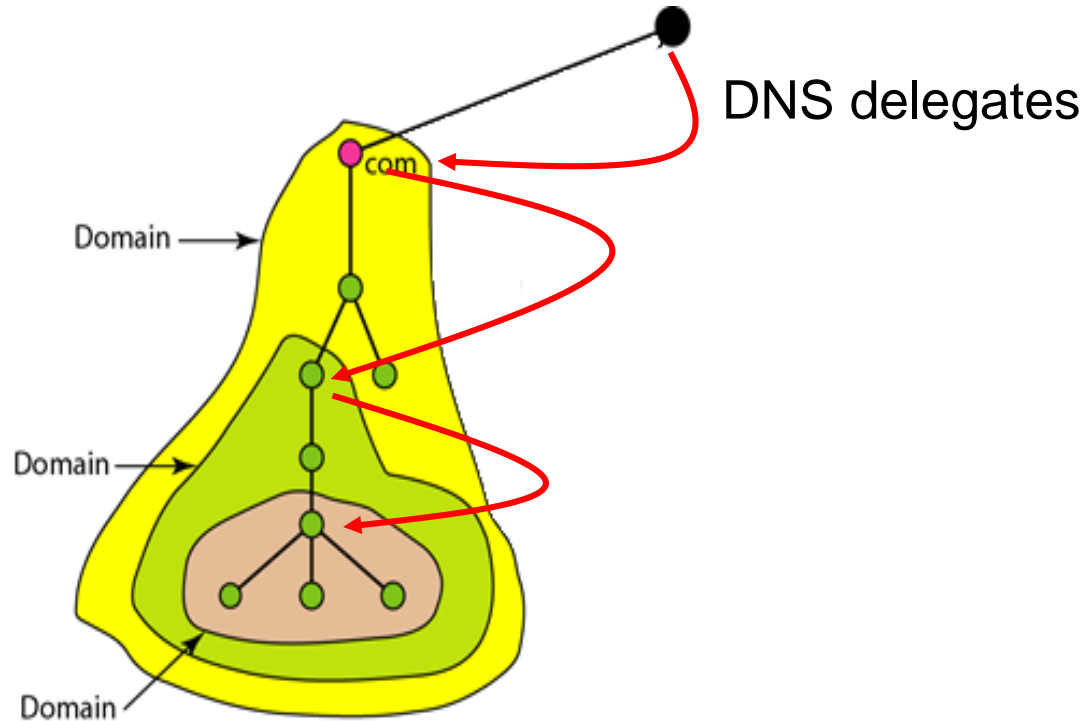
Suggested domain meta-rule:

If you must lie, lie only to yourself.


How does the DNS Protect its Name Spaces?



How does the DNS Protect its Name Spaces?




How does the DNS Protect its Name Spaces?



DNS delegates
Higher-ranked server holds
Name Server (NS) Records

How does the DNS Protect its Name Spaces?




DNS delegates
Higher-ranked server holds
Name Server (NS) Records

NS Record:

```
example.com 3600 IN NS names.example.com
```

How does the DNS Protect its Name Spaces?



DNS delegates
Higher-ranked server holds
Name Server (NS) Records

NS Record:

```
example.com 3600 IN NS names.example.com
```

Glue Record:

```
names.example.com 3600 IN A 10.10.10.10
```

Core Technology

THE DESIGN OF DNSSEC

DNSSEC Design Objectives

Original Specification:
RFC 2535 (1999)

Current specifications:
RFCs 4033, 4034, 4035
+ updates

DNSSEC Design Objectives

Original Specification:
RFC 2535 (1999)

Current specifications:
RFCs 4033, 4034, 4035
+ updates

Goals

- Provide integrity (prevent spoofing) by
 - + Authenticating messages of name servers
 - + Authenticating resource records
- Proof of non-existence (prevent DoS against names)

DNSSEC Design Objectives

Original Specification:
RFC 2535 (1999)

Current specifications:
RFCs 4033, 4034, 4035
+ updates


Goals

- Provide integrity (prevent spoofing) by
 - + Authenticating messages of name servers
 - + Authenticating resource records
- Proof of non-existence (prevent DoS against names)

Non-Goals

- Confidentiality by hiding DNS data or requests
- Authorization of requests or requestors
- Protection against DDoS attacks (e.g., via traffic amplification)

DNSSEC Fundamentals



DNSSEC uses Public Key Cryptography

- Authenticate and verify Resource Record Sets (RRSets)
- Authenticate and verify zone delegations

Each Zone has key(s) for signing its **RRSets**

- Trust chain follows zone delegation
- Secured by Delegation Signer (**DS**) Records

Public Credentials are Stored in the DNS

DNSKEY

DNSSEC Resource
Record to store public
keys in the DNS

RRSIG

DNSSEC Resource
Record to store
signatures in the DNS

Signing Resource Record Sets

o RRset:

```
- www.opendnssec.se. 7200 IN A 192.168.10.3
                        -           A 10.0.0.3
                        -           A 172.25.215.
```

o DNSKEY RDATA:

```
- opendnssec.se. 3600 IN DNSKEY 256 3 5
                        AQOvhvXXU61Pr8sCwELcqqq1g4JJCALG4C9EtraBKVd+vG
                        IF/unwigfLOAO3nHp/cgGrG6gJYe8OWKYNgq3kDChN
```

o RRSIG RDATA:

```
- opendnssec.se. 3600 IN RRSIG A 5 2 3600
                        20050611144523 20050511144523 3112 opendnssec.se.
                        VJ+8ijXvbrTLeoAiEk/qMrdudRnYZM1VlqhNvhYuAcYKe2X/jqYfMfj
                        fSUrmhPo+0/GOZj66DJubZPmNSYXw==
```

Signing Resource Record Sets

o RRset:

```
- www.opendnssec.se. 7200 IN A 192.168.10.3
  -                               A 10.0.0.3
  -                               A 172.25.215.
```

o DNSKEY RDATA:

```
- opendnssec.se. 3600 IN DNSKEY 256 3 5
  AQOvhvXXU61Pr8sCwELcqqq1g4JJCALG4C9EtraBKVd+vG
  IF/unwigfLOAO3nHp/cgGrG6gJYe8OWKYNgq3kDChN
```

Flags, Protocol, Algorithm

o RRSIG RDATA:

```
- opendnssec.se. 3600 IN RRSIG A 5 2 3600
  20050611144523 20050511144523 3112 opendnssec.se.
  VJ+8ijXvbrTLeoAiEk/qMrdudRnYZM1VlqhNvhYuAcYKe2X/jqYfMfj
  fSUrmhPo+0/GOZj66DJubZPmNSYXw==
```

Signing Resource Record Sets

o RRset:

```

- www.opendnssec.se. 7200 IN A 192.168.10.3
                        -      A 10.0.0.3
                        -      A 172.25.215.
  
```

o DNSKEY RDATA:

```

- opendnssec.se. 3600 IN DNSKEY 256 3 5
                                AQOvhvXXU61Pr8sCwELcqqq1g4JJCALG4C9EtraBKVd+vG
                                IF/unwigfLOAO3nHp/cgGrG6gJYe8OWKYNgq3kDChN
  
```

Flags, Protocol, Algorithm

o RRSIG RDATA:

```

- opendnssec.se. 3600 IN RRSIG A 5 2 3600
                                20050611144523 20050511144523 3112 opendnssec.se.
                                VJ+8ijXvbrTLeoAiEk/qMrdudRnYZM1VlqhNvhYuAcYKe2X/jqYfMfj
                                fSUrmhPo+0/GOZj66DJubZPmNSYXw==
  
```

Type covered, Algorithm, # of labels covered, orig. TTL

Signing Resource Record Sets

o RRset:

```

- www.opendnssec.se. 7200   IN  A   192.168.10.3
  -                   -      A   10.0.0.3
  -                   -      A   172.25.215.
  
```

o DNSKEY RDATA:

```

- opendnssec.se. 3600 IN DNSKEY 256 3 5
  AQOvhvXXU61Pr8sCwELcqqq1g4JJCALG4C9EtraBKVd+vG
  IF/unwigfLOAO3nHp/cgGrG6gJYe8OWKYNgq3kDChN
  
```

Flags, Protocol, Algorithm

o RRSIG RDATA:

```

- opendnssec.se. 3600 IN RRSIG A 5 2 3600
  20050611144523 20050511144523 3112 opendnssec.se
  VJ+8ijXvbrTLeoAIEk/qmVrdudRnYZMTviqNvvhYuAcYKe2X/jqYfMfj
  fSUrmhPo+0/GOZj66DJubZPmNSYXw==
  
```

Type covered, Algorithm, # of labels covered, orig. TTL

Signature time range, key tag, signer's name

Delegation Signer (DS) Record

Handle for building the chain of trust along names

A DS record is the hash of the DNSKEY published at the parent zone to delegate trust to the child zone

Delegation Signer (DS) Record

Handle for building the chain of trust along names

A DS record is the hash of the DNSKEY published at the parent zone to delegate trust to the child zone

Example (name, types, key-tag, algorithm, digest-type, digest):

```
opendnssec.se. IN DS 27295 5 1  
5AEF372D65BC594A7AF5E0E77CDDA55E0C  
43A56A
```

Delegation Signer (DS) Record

The DS records are signed by the parent

DS MUST NOT be in the child zone!

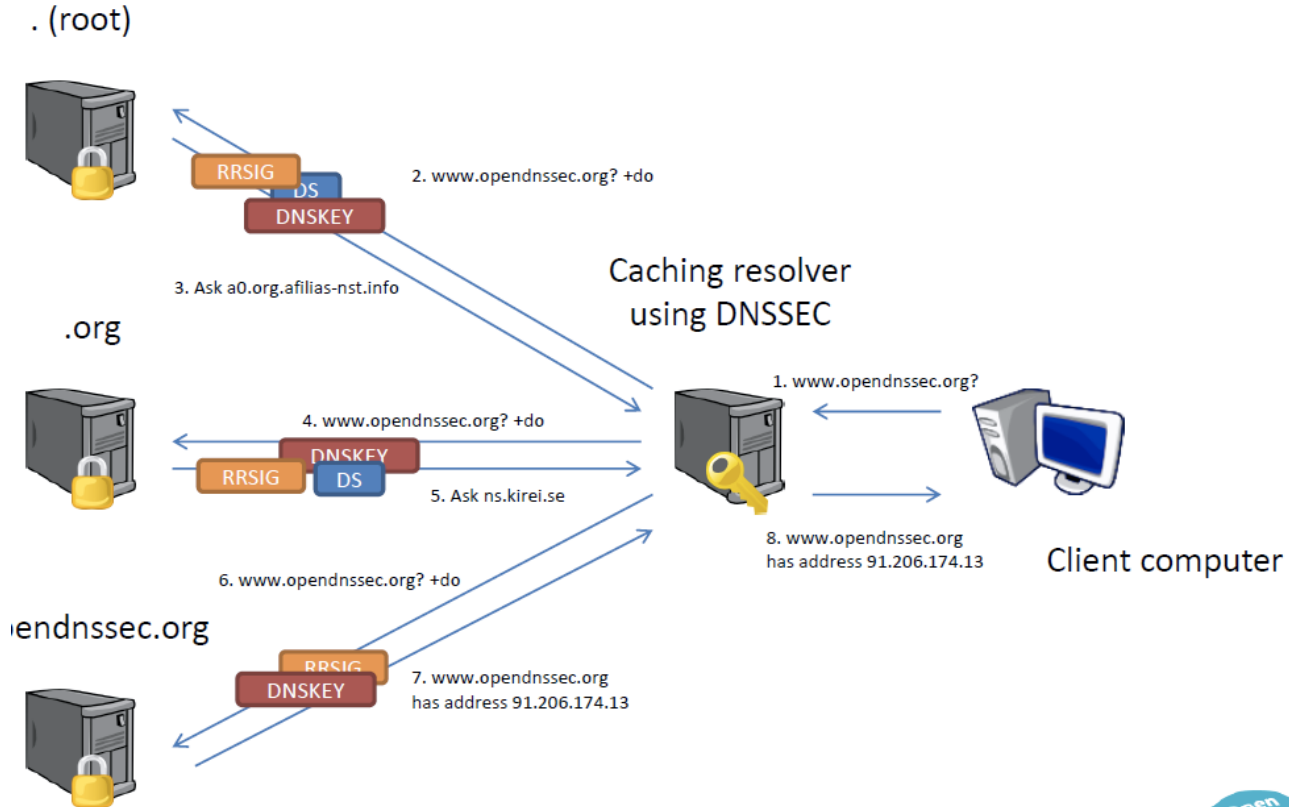
Handle for building the chain of trust along names

A DS record is the hash of the DNSKEY published at the parent zone to delegate trust to the child zone

Example (name, types, key-tag, algorithm, digest-type, digest):

```
opendnssec.se. IN DS 27295 5 1  
5AEF372D65BC594A7AF5E0E77CDDA55E0C  
43A56A
```

Resolving DNSSEC



DNSSEC Cryptography

Caveat:
Keys may
be cached

Problem: Keys in trust chain locked with parents
– changes are difficult ...

Solution: Two keys

– The **Key Signing Key (KSK)** for trust establishment

– The **Zone Signing Key (ZSK)** for signing RRs

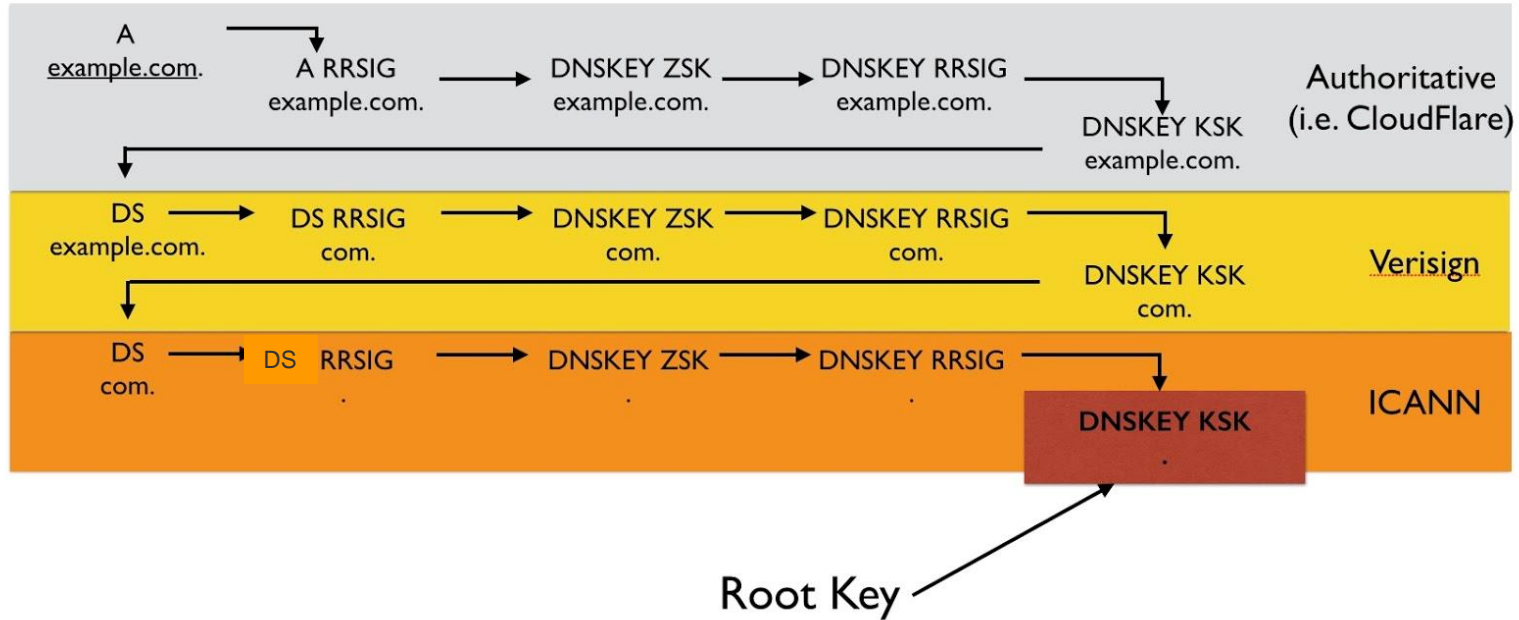
KSK signs the ZSK, it maybe offline for protection

– Changes involve third parties

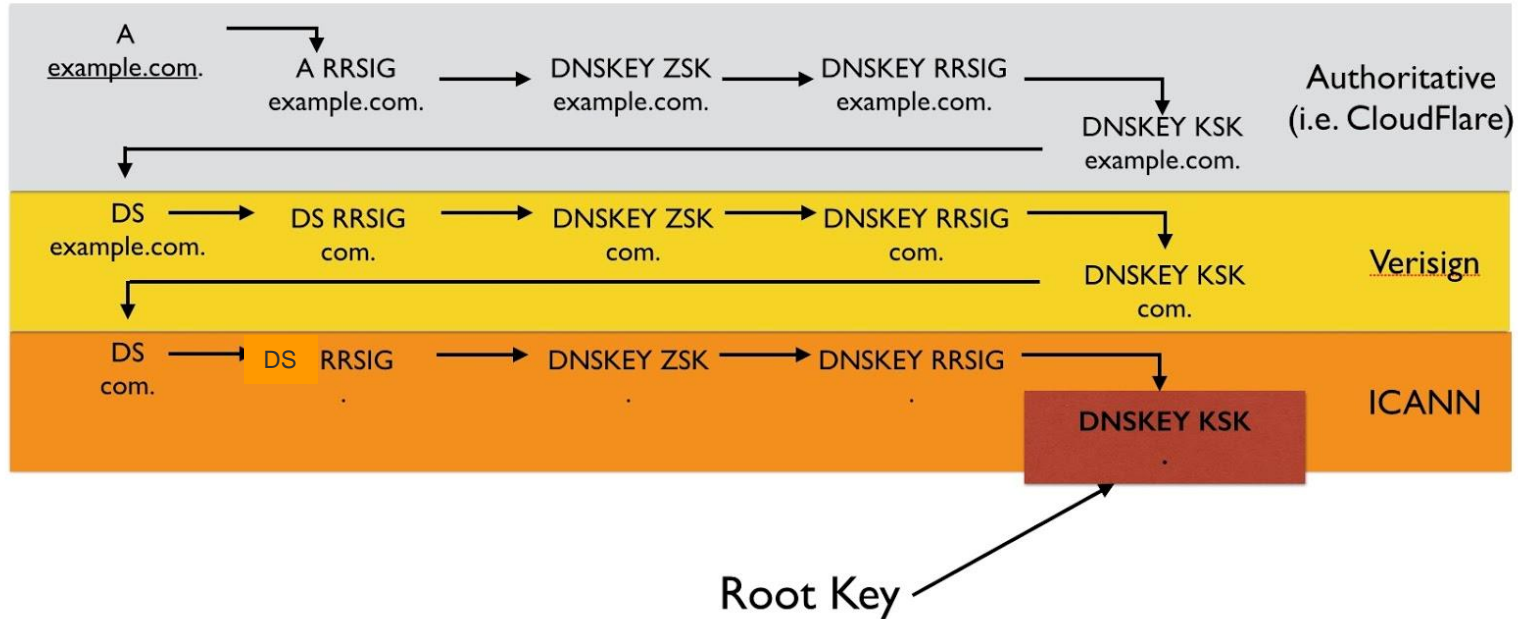
ZSK signs daily DNS changes, needed 'on disk'

– Changes without third parties

DNSSEC Trust CHAIN



DNSSEC Trust CHAIN



The root of trust is the KSK DNSKEY for the DNS root.
This key is universally known and published.

Walking the Chain of Trust

Locally configured

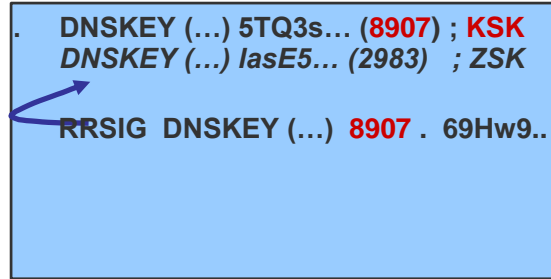
Trusted key: . **8907**

\$ORIGIN .

1

2

```
. DNSKEY (...) 5TQ3s... (8907) ; KSK  
DNSKEY (...) lasE5... (2983) ; ZSK  
RRSIG DNSKEY (...) 8907 . 69Hw9..
```



Walking the Chain of Trust

Locally configured

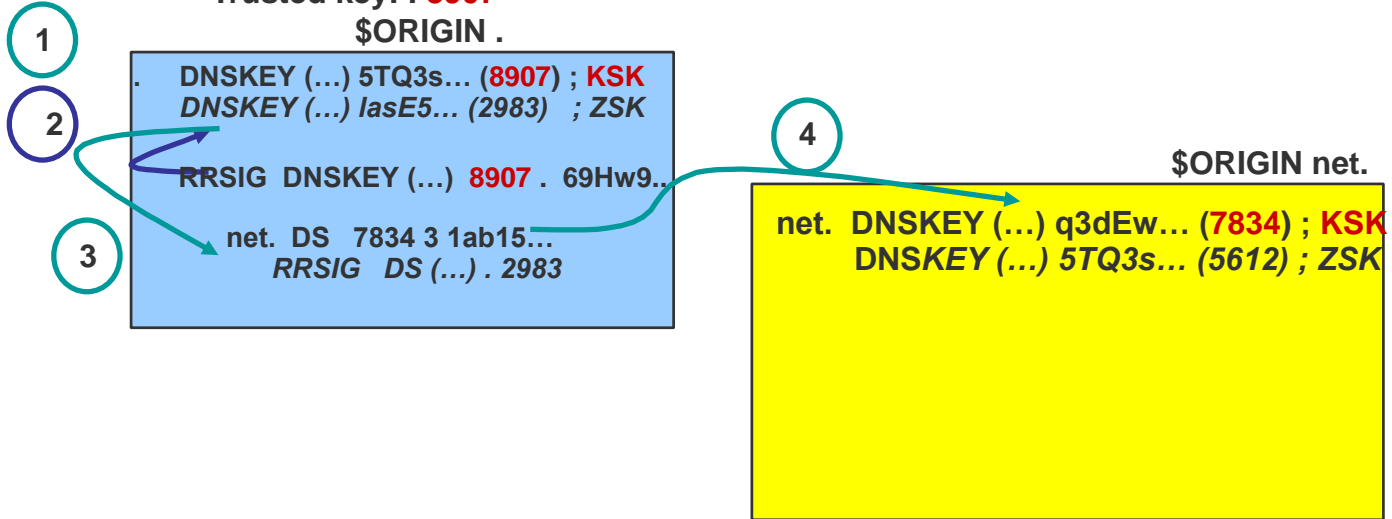
Trusted key: . **8907**

\$ORIGIN .



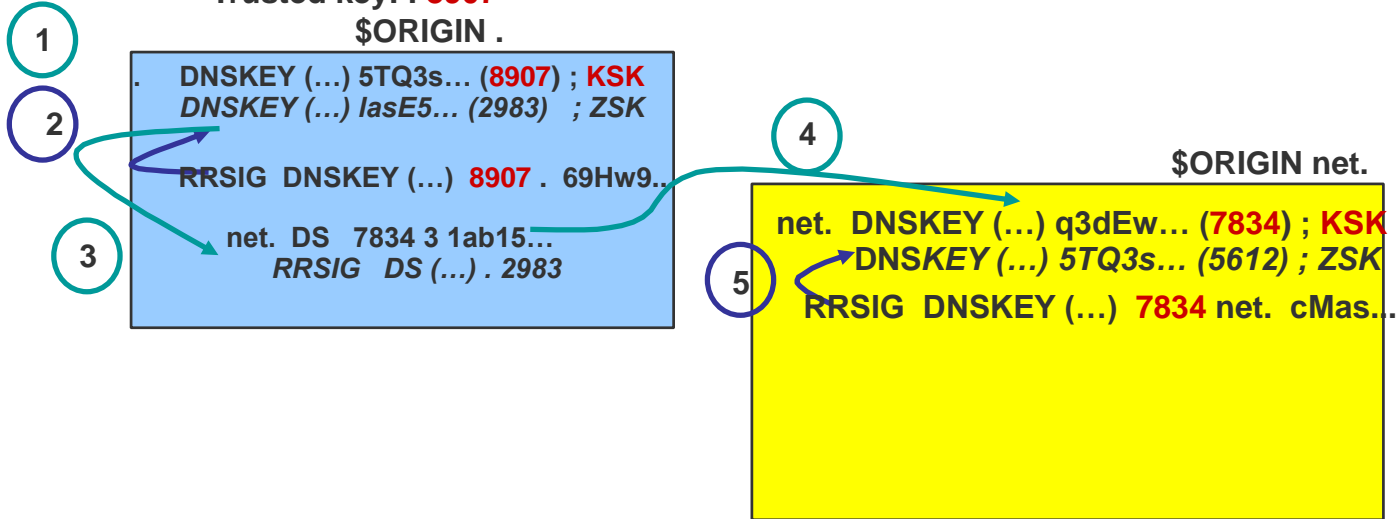
Walking the Chain of Trust

Locally configured
 Trusted key: . 8907
 \$ORIGIN .



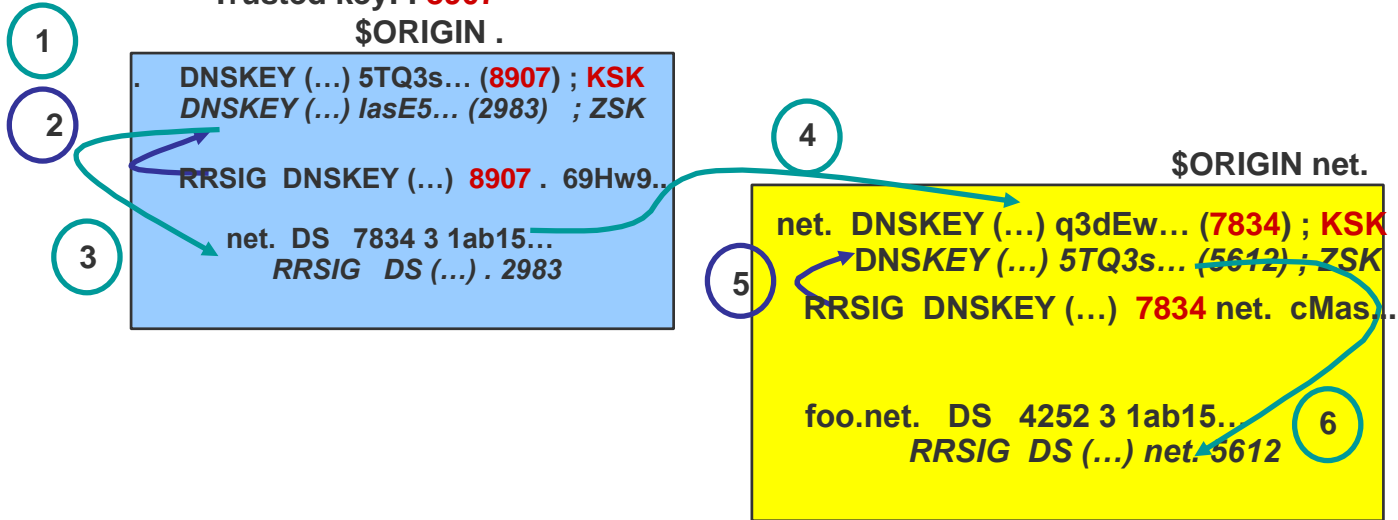
Walking the Chain of Trust

Locally configured
 Trusted key: . 8907
 \$ORIGIN .



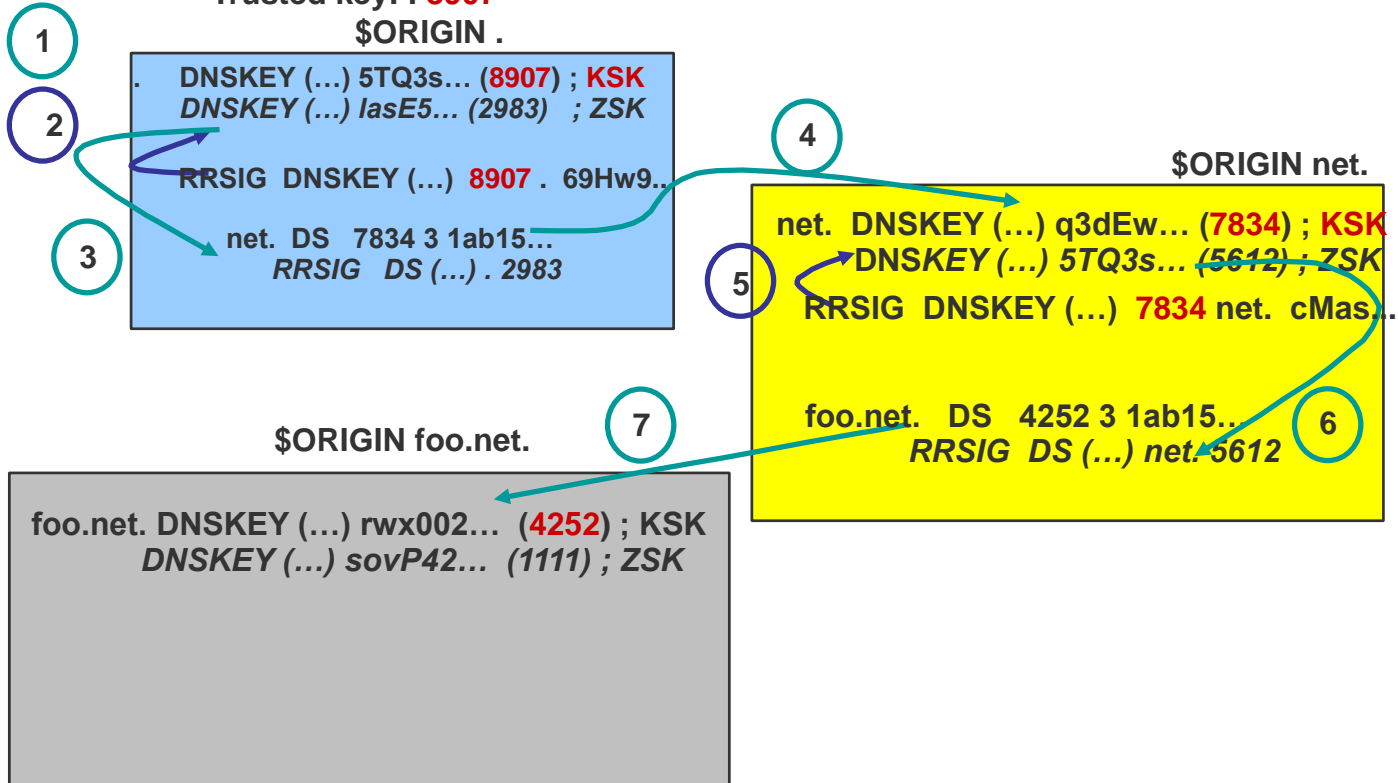
Walking the Chain of Trust

Locally configured
Trusted key: . 8907
\$ORIGIN .



Walking the Chain of Trust

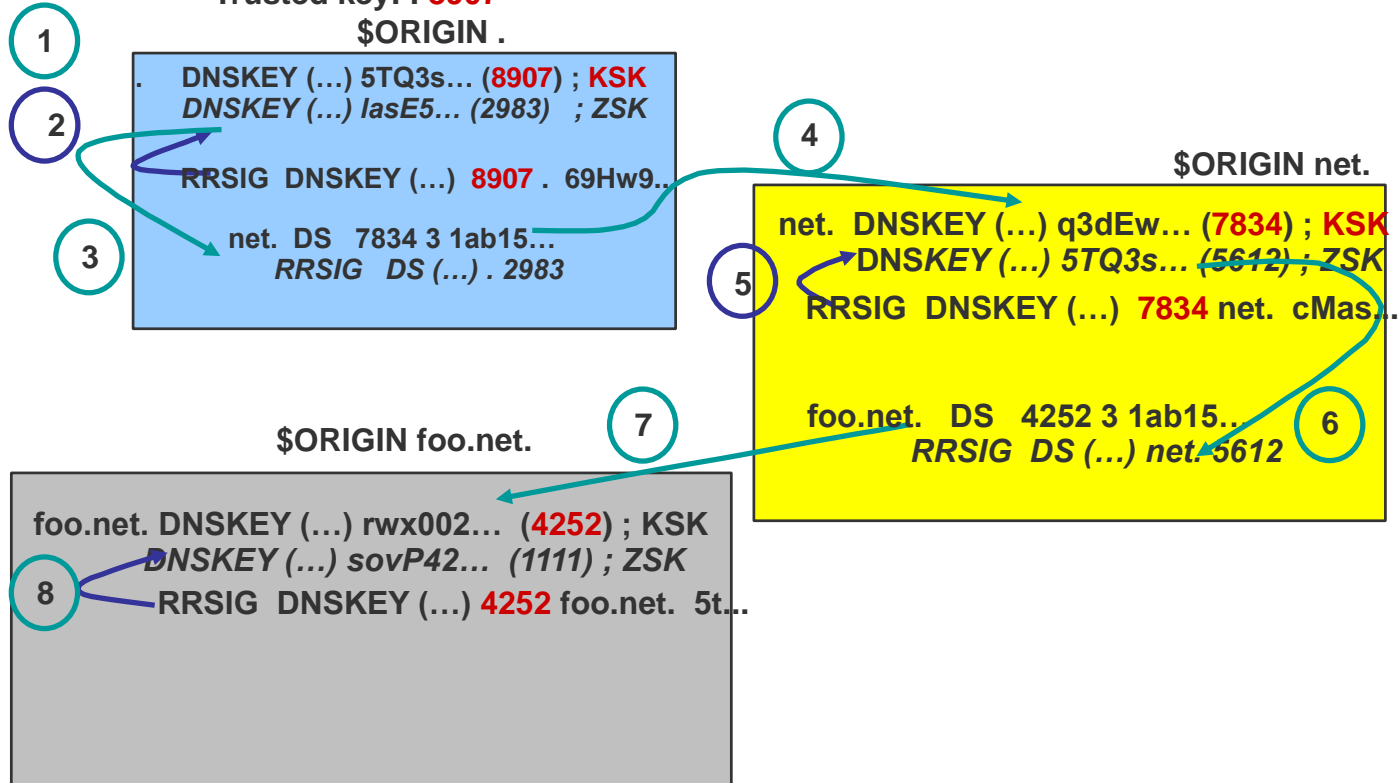
Locally configured
Trusted key: . 8907
\$ORIGIN .



Source: [NLnet Labs](#)

Walking the Chain of Trust

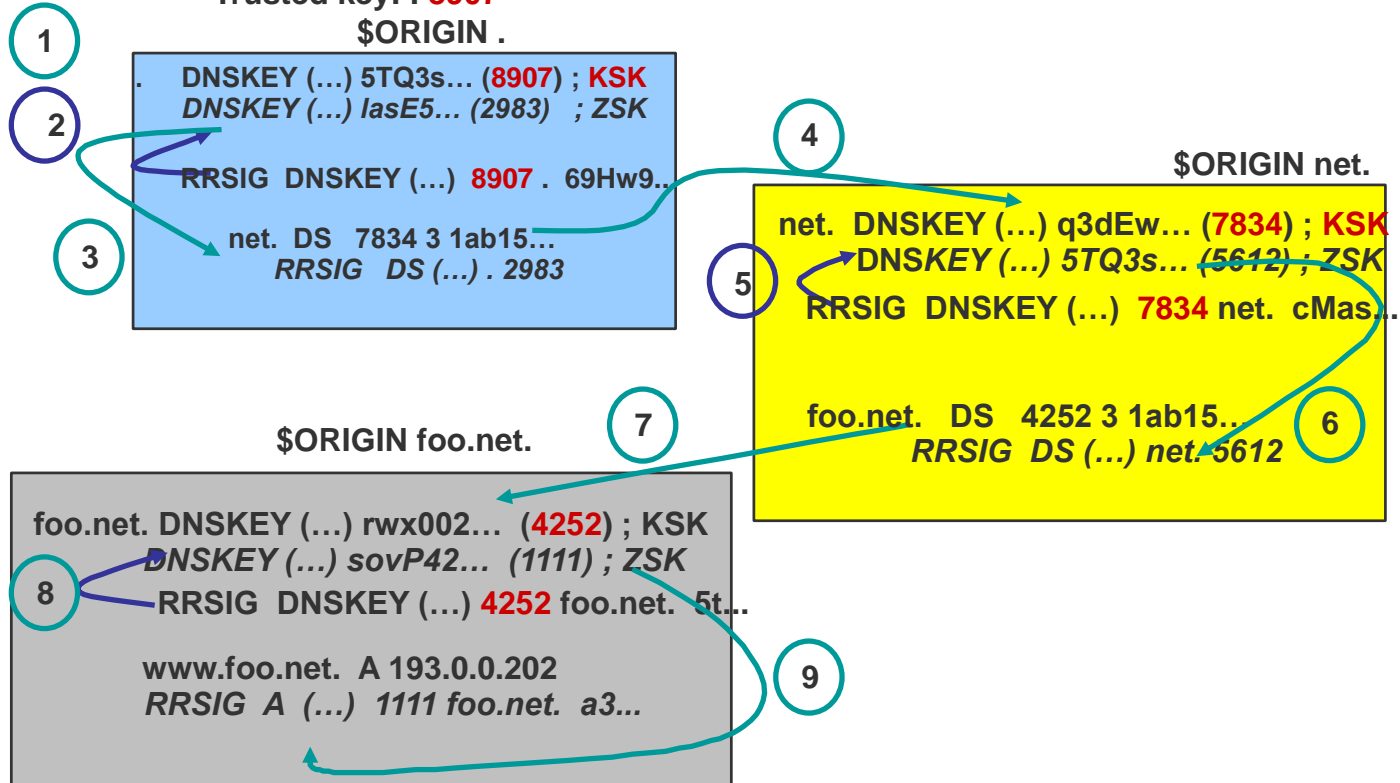
Locally configured
Trusted key: . 8907
\$ORIGIN .



Source: [NLnet Labs](#)

Walking the Chain of Trust

Locally configured
Trusted key: . 8907
\$ORIGIN .



Source: [NLnet Labs](#)

Summary on Verifying the Chain of Trust

Summary on Verifying the Chain of Trust

Data in zone can be trusted if signed by a
Zone-Signing-Key

Summary on Verifying the Chain of Trust

Data in zone can be trusted if signed by a Zone-Signing-Key

Zone-Signing-Keys can be trusted if signed by a Key-Signing-Key

Summary on Verifying the Chain of Trust

Data in zone can be trusted if signed by a Zone-Signing-Key

Zone-Signing-Keys can be trusted if signed by a Key-Signing-Key

Key-Signing-Key can be trusted if pointed to by trusted DS record (from parent)

Summary on Verifying the Chain of Trust

Data in zone can be trusted if signed by a Zone-Signing-Key

Zone-Signing-Keys can be trusted if signed by a Key-Signing-Key

Key-Signing-Key can be trusted if pointed to by trusted DS record (from parent)

DS record can be trusted if signed by the parents Zone-Signing-Key

Summary on Verifying the Chain of Trust

Secure entry point:

DS or DNSKEY
exchanged out-of-
band, locally stored

Data in zone can be trusted if signed by a
Zone-Signing-Key

Zone-Signing-Keys can be trusted if
signed by a Key-Signing-Key

Key-Signing-Key can be trusted if pointed
to by trusted DS record (from parent)

DS record can be trusted if signed by the
parents Zone-Signing-Key

Provide Proof of Non-existent Names

NSEC

- Points to the next label
(domain name) in the zone
- Enables zone walk (“get next”)
 - Zone walk often unwanted

Provide Proof of Non-existent Names

NSEC

NSEC3

- Points to the next label
(domain name) in the zone
- Enables zone walk (“get next”)
 - Zone walk often unwanted

Provide Proof of Non-existent Names

NSEC

- Points to the next label
(domain name) in the zone
- Enables zone walk (“get next”)
 - Zone walk often unwanted

NSEC3

- Prevents ‘walking in the clear’
- Translates into hashes
(linked list of hashed names)
 - Non-existence of hash proves
non-existence of name

Provide Proof of Non-existent Names

NSEC

- Points to the next label
(domain name) in the zone
- Enables zone walk (“get next”)
 - Zone walk often unwanted

NSEC3

- Prevents ‘walking in the clear’
- Translates into hashes
(linked list of hashed names)
 - Non-existence of hash proves
non-existence of name

Create new RRs: NSEC, NSEC3 and NSEC3PARAM

How to ...

DNSSEC DEPLOYMENT

Deployment Options for Clients

Full DNSSEC Resolver

- Fully DNSSEC compliant
- Performs DNSSEC validation on its own


Stub resolvers

- Client completely trusts local DNS server (e.g., from the ISP)
- Client decides autonomously about unauthenticated data
 - DNS query includes DO bit (DNSSEC OK Bit): Enforce the server to perform validation
 - DNS server performs DNSSEC validation and answers with AD flag (Authenticated Data) or error

DNSSEC Deployment

Country Top Level Domains

ccTLD DNSSEC Status on 2020-09-14



www.internetsociety.org/deploy360/dnssec/maps

DNSSEC enabled zones

18.10.2020: 7,852,116

See: secspider.net

New Developments

ORTHOGONAL APPROACHES

DNS over (D)TLS (DoT) – RFCs 7858, 8094

Privacy extension between DNS client and recursive resolver – the ‘Last Mile’

Encrypts and authenticates transport, not DNS data

Servers use UDP/TCP port 853

Deployment initiative to provide DoT services:
E.g., Quad9 – 9.9.9.9 (anycast)

DNS over HTTP(s) (DoH) - RFC 8484

Recent counter approach to DNSSEC (10/18)

- Web-centric Over-the-Top service (OTT)
 - Easy to run, independent of providers
 - Can be activated in browsers
- Different trust model: Trust the (central) DoH server instead of the DNS data
- Full privacy on the net
 - queries TLS-encrypted
- Rapid deployment: selected centralized servers by Google, Cloudflare, ...

Problems with OTT DoH

- DNS records remain unauthenticated by RFC 8484
- Centralized approach – no more distributed caching
- Querier visible to DoH server *and* beyond: Modern DNS resolvers use EDNS subnet field (RFC7871)
- Invisible to local providers – hinders debugging and performance optimization

Summary

DNSSEC is a major building block for securing the Internet infrastructure

- It provides Integrity and Authenticity for DNS Resource Records
- It builds trust along the name delegation chain

Deployment is slowly progressing

DoT adds privacy extensions to the last mile