

DNSSEC HOWTO, a tutorial in disguise

Olaf Kolkman

Revision 136

April 7, 2010

Download the most recent PDF release from:
<http://www.nlnetlabs.nl/dnssec-howto/dnssec-howto.pdf>

About This Document

This HOWTO is not a HOWTO.

In 2000-2001 this document started its life as an addendum to a DNSSEC course I organized at the RIPE NCC but in cause of time it has grown beyond the size of your typical HOWTO and became a (hopefully) comprehensive tutorial on the subject of DNSSEC and DNSSEC deployment.

This HOWTO is intended for those people who want to deploy DNSSEC and are seeking a document that lives between a typical high level description of the topic (see the excellent Surfnet White Paper on DNSSEC for that (<http://www.surfnet.nl/Documents/DNSSEC-web.pdf>)), the typical out of the box recipe, and an in depth description of the technology.

In this tutorial we touch upon the following topics:

Part I, intends to provide some background for those who want to deploy DNSSEC.

Part II, about the aspects of DNSSEC that deal with data security.

- Creating an island of security (Chapter 2, "Configuring a recursive name server to validate answers" and Chapter 3, "Securing a DNS zone") by configuring a recursive name server to validate the signed zones served by your organisations authoritative name servers. When you have learnt and implemented this, you can be sure that DNS data in your organisation is protected from change. Once you have created an island of security it is a small step to become part of a chain of trust.
- Delegating signing authority; building a chain of trust (Chapter 4, "Delegating of signing authority; becoming globally secure"). You will learn how to exchange keys with your parent and with your children.
- Chapter 5, "Rolling keys" covers maintaining keys and ensuring that during the rollover process clients will be able to maintain a consistent view of your DNS data.
- Part IV, covering aspects that deal with server to server security and transaction security.
- Chapter 9, "Securing zone transfers" is on the use of transaction security (TSIG) to provide authorisation and integrity for zone transfers.

Part III, describes a few tools that may turn out handy while figuring out what might have gone wrong.

The documentation is based on the so called DNSSEC-bis specifications that were finalised by the IETF DNSEXT working group in July 2004 and published in March 2005 as [3, 5, 4].

As of June 2009 the author is aware of the following open-source and or freeware implementations of the DNSSEC-bis specifications: BIND, Unbound and NSD. All our examples are based on BIND 9.6.1 and Unbound 1.3.0.

This document is not intended as an introduction to DNS. Basic knowledge of DNS and acronyms used is assumed. We have tried not to use jargon but when unavoidable we have tried to explain the meaning. If you want to know more about the topic of DNS in general then Paul Albitz and Cricket Lui's[2] or Ron Aitchinson's [1] text books provide an excellent introduction.

This document will be subject to change. Please regularly check http://www.nlnetlabs.nl/dnssec_howto/dnssec_howto.pdf for new versions.

Your corrections and additions are appreciated.

Contents

I	DNSSEC, the background	3
1	A motivation for DNSSEC	3
II	Securing DNS data	4
2	Configuring a recursive name server to validate answers	4
2.1	Introduction	4
2.2	Warning	4
2.3	Configuring the caching forwarder	5
2.3.1	Configuring a trust anchor	6
2.3.2	How and where to get these trust anchors?	6
2.3.3	Public repositories of trust-anchors	8
2.3.4	Testing	9
2.4	Finding trust-anchors	11
2.5	Maintaining Trust Anchors	11
2.6	Lookaside Validation	12
2.6.1	Configuring lookaside validation	12
2.7	Some Troubleshooting Tips	17
3	Securing a DNS zone	18
3.1	Introduction	18
3.2	Configuring authoritative servers	19
3.3	Creating key pairs	19
3.3.1	Key Maintenance Policy	19
3.3.1.1	Key- and zone-signing keys.	19
3.3.2	Creating the keys	20
3.4	Zone-signing	21

3.5	Caching forwarder configuration	29
3.6	Zone Re-Signing	29
3.7	Troubleshooting Signed Zones	29
3.8	Possible problems	31
3.9	Automating the signing process	31
4	Delegating of signing authority; becoming globally secure	32
4.1	Introduction	32
4.2	Practical steps	32
4.3	Possible problems	34
4.4	Registering with a DLV registry	34
5	Rolling keys	35
5.1	DNS traversal	35
5.2	"Pre-Publish" and "Double Signature" rollovers	36
5.3	Tools	37
5.4	ZSK rollover	37
5.4.1	ZSK preparation (production phase)	37
5.4.2	ZSK rollover (phase1)	39
5.4.3	ZSK Cleanup (phase2)	40
5.4.4	Modifying zone data during a rollover	40
5.5	Key-signing key rollovers	41
5.5.1	KSK preparation (production phase)	41
5.5.2	ZSK rollover (phase 1)	41
5.5.3	KSK cleanup (phase 2)	42
5.5.4	Multiple KSKs	42
III	Troubleshooting tools	44
6	Using drill for troubleshooting	44
7	Using dig for troubleshooting	44
8	DNSSEC tools	45
IV	Securing communication between Servers	48
9	Securing zone transfers	48
9.1	Introduction	48
9.2	Generating a TSIG key	49
9.2.1	Generating a TSIG secret with <code>dnssec-keygen</code>	49
9.2.2	Other ways to generate secrets	50
9.3	Configuring TSIG keys	50
9.4	Primary servers configuration of TSIG	51
9.5	Secondary servers configuration of TSIG	51
9.6	Securing the NOTIFY message too	52
9.7	Troubleshooting TSIG configuration	52
9.8	Possible problems	53
9.8.1	Timing problems	53
9.8.2	Multiple server directives	53

CONTENTS

V	Appendices	54
A	Deploying DNSSEC: the Milestones	54
A.1	Private Key policies and procedures	54
A.2	Public Key policies and procedures	54
A.3	Signing infrastructure	55
A.4	Server infrastructure	55
A.5	Monitoring tools	56
A.6	Serving a Secured zone	56
A.7	Requesting secure delegation	56
A.8	Secure Delegation Provisioning	56
B	BIND installation	57
C	Estimating zone size increase	58
D	Generating random numbers	58
E	Example DNSSEC key procedure	59
F	Perl's Net::DNS::SEC library	61

Part I

DNSSEC, the background

1 A motivation for DNSSEC

The Domain Name System (DNS) is one of these essential elements that allowed the Internet to evolve in what it is today.

It is used in almost every interaction that uses names in identifiers: Email, Web, SIP based Voice over IP, Web services, Spam filtering, Internet messaging, and many more. Yet, the DNS system was not designed with security in mind; over 3 decades ago the environment it was developed for (and in) just looked different. Regardless, given that the DNS is the largest distributed look-up system on the Internet, one could claim that the protocol designers were successful.

The fact that essential components in the DNS architecture, so called caching nameservers, are subject to types of attacks that allow the inception of false data has been known for almost 2 decades now.

As a result of such so called cache poisoning attacks, e-mails can be redirected and copied before they are delivered to their final destination, voice over IP calls can be tapped by third parties, and – given the circular dependency of the registration process on the DNS – SSL certificates may not be as protective as one would hope.

The DNS has become a utility that people depend on when moving about on the Internet. It is a core part of the Internet Infrastructure and trust in the DNS is necessary, albeit not sufficient, for trust in the Internet.

DNSSEC was designed to deal with cache poisoning and a set of other DNS vulnerabilities such as man in the middle attacks and data modification in authoritative servers. Its major objective is to provide the ability to validate the authenticity and integrity of DNS messages in such a way that tampering with the DNS information anywhere in the DNS system can be detected. This is the kind of protection the DNS desperately needs.

Unfortunately, because of the distributed nature of the DNS, DNSSEC needs to be deployed by a significant amount of DNS data providers before its utility becomes relevant. Custodians of the DNS infrastructure such as TLDs and the root system should provide a breeding ground on which DNSSEC can take off while ISPs and enterprise DNS administrators prepare their DNS infrastructure to validate signed data.

Obviously this is not going to be a project with immediate return on investment, it is a long term strategy to allow us to put similar trust in the Internet as we did 10 years ago.

By introducing DNSSEC in your environment you do not only protect yourself or the users of your data. But you also help in building a globally secure system that can be used to bootstrap trust relations in other protocols.

Part II

Securing DNS data

This part deals with securing data in zone files. We describe how to generate and manage keys, how to set up a recursive name server to validate signed zone data and how to sign and serve zones.

2 Configuring a recursive name server to validate answers

2.1 Introduction

We plan to configure a recursive name server to validate the data it receives. Users that use this recursive name server as their resolver will, then, only receive data that is either secure and validated or not secured in any way. As a result, secured data that fails validation will not find its way to the users¹. Having a validating recursive name server protects all those that use it as a forwarder against receiving spoofed DNS data.

Figure 1 illustrates how to configure the recursive DNS servers with a trusted key for "example.com" so that all the data served by the authoritative servers for "example.com" is validated before it is handed to the protected infrastructure that have the recursive servers configured as their forwarder (the name servers that usually are assigned through DHCP or configured in `/etc/resolv.conf`).

By configuring a public key for a specific zone, we tell the caching forwarder that all data coming from that zone should be signed with the corresponding private key. The zone acts as a secure entry point into the DNS tree and the key configured in the recursive name server acts as the start for a chain of trust. In an ideal situation you have only one key configured as a secure entry point: the key of the root zone.

We assume you have configured your BIND nameserver to be recursive only or that you use UNBOUND, which is recursive only.

We also assume that that name server in your organisation has been configured to run as an authoritative server for a secured zone called `example.net`. Notes on how to set up a secured zone can be found below in Chapter 3, "Securing a DNS zone"

2.2 Warning

Your recursive name server will treat the zones for which you configured trust anchors as being secured. If the zones for which you have configured trust anchors change their keys you will also have to reconfigure your trust anchors. Failure to do so will result in the data in these zones, or any child, being marked as bogus and therefore becoming invisible to users.

¹As a result users can trust the data they are dealing with, provided that the path between the validating recursive name server and the stub resolver can be trusted. On a shared network such as an IEEE802.11 network this is not the case.

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

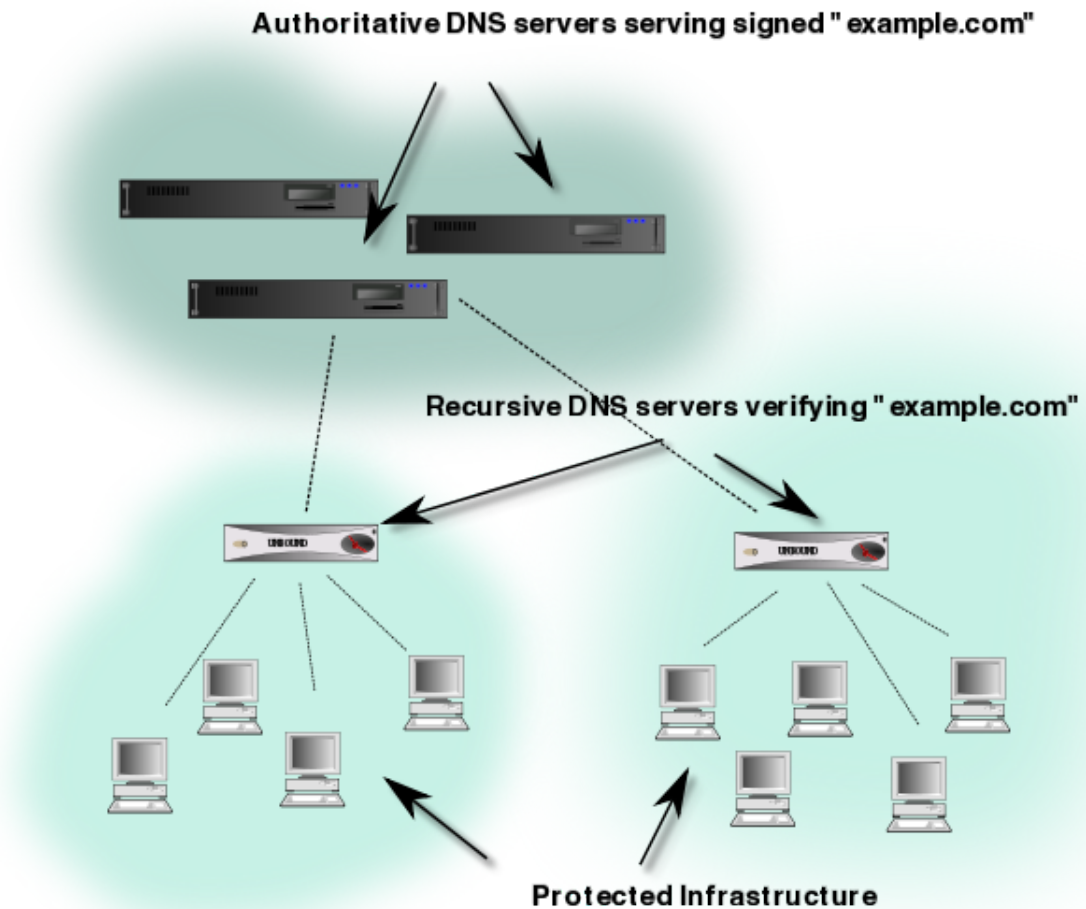


Figure 1: DNS environment

2.3 Configuring the caching forwarder

See AppendixB for information on compiling BIND with the correct switches to allow for DNSSEC. Do not forget enter the `dnssec-enable yes;` and `dnssec-validation yes;` statements in the `options` directive of your `named.conf`.²

As an alternative to bind you could use UNBOUND as a DNSSEC aware recursive nameserver. UNBOUND does not need any special configuration options except for the configuration of a trust-anchor to perform DNSSEC validation.

²the `dnssec-validation yes;` directive is only needed in the version 9.4 series. It will be the default for the 9.5 series of BIND

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

2.3.1 Configuring a trust anchor

A trust anchor is a public key that is configured as the entry point for a chain of authority. In the ideal case —where the root is signed and chains of trusts can be constructed through top-level domains to end-nodes — validating name servers would only need one of these trust anchors to be configured. During early deployment you will probably want to configure multiple trust anchors.

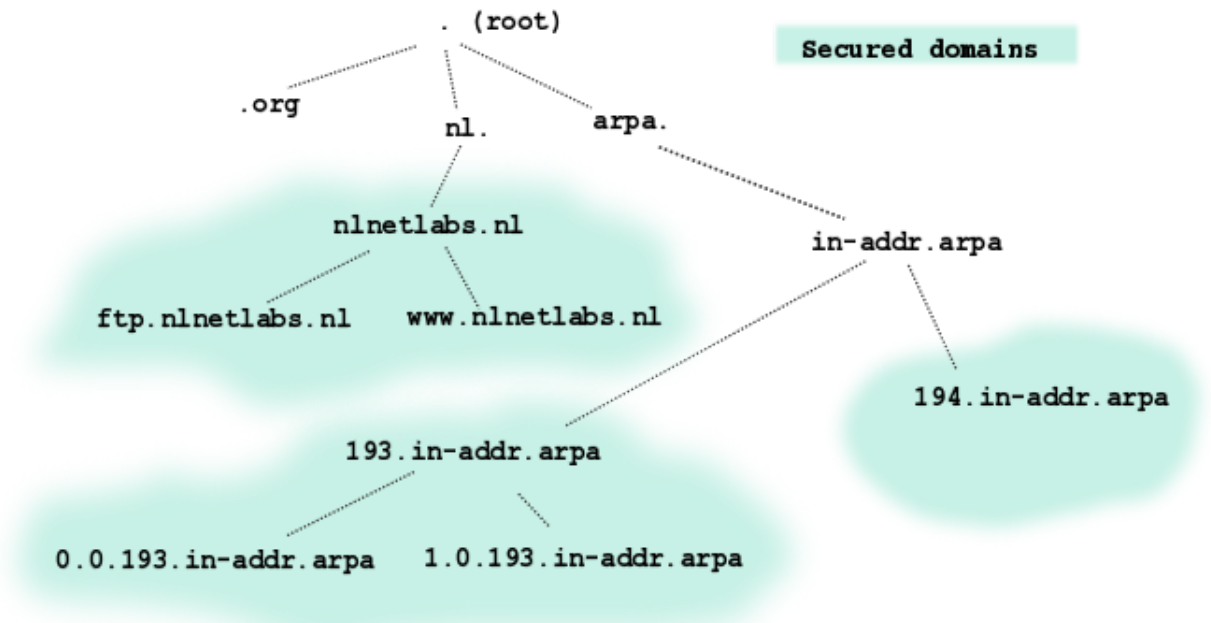


Figure 2: Trust anchors in the DNS tree

In Figure 2 we show a zone tree. In this tree, the domains `nlnetlabs.nl`, `194.in-addr.arpa`, `193.in-addr.arpa` and `0.0.193.in-addr.arpa` are assumed to be signed. It is also assumed that there is a secure delegation between `193.in-addr.arpa` and `0.0.193.in-addr.arpa`. In order to validate all these domains, the validating DNS client would have to configure trust anchors for `nlnetlabs.nl`, `194.in-addr.arpa` and `193.in-addr.arpa`.³

2.3.2 How and where to get these trust anchors?

To configure a trust-anchor you have to obtain the public key of the zone that you want to use as the start of the chain of authority that is to be followed when the data is validated. It is possible to get these straight from from the DNS, but there are two reasons why this may not be advisable.

³ Fetching and maintaining all these trust anchors in the absence of a signed root clearly does not scale and is one of the problems that early deployers will have to deal with. BIND 9.3.2 contains a so called "look-aside" validation option that may help with this trust anchor distribution issue. Also see section 2.6.

Firstly, you have to establish the authenticity of the key you are about to configure as your trust anchor. How you do this depends on what method the zone owner has made available for out of band validation of the key.

- You could do this by visiting the zone owners secure website and validate the key information. For instance, the RIPE NCC signs a number of reverse zones. They publish their public keys on through `<https://www.ripe.net/projects/disi/keys/>`
- You could give the zone owner a call if you personally know them.
- You could trust the key that is published on the bill you just received from the zone owner.
- You could just believe that your OS vendor did the validation on your behalf.

Secondly, you may have a choice of public keys, in which case you need to select the the proper "Secure Entry Point" key.

In DNSSEC a difference is made between key- and zone-signing keys. Key-signing keys exclusively sign the DNSKEY RR set at the apex, while zone-signing keys sign all RR sets in a zone. ⁴. Key-signing keys are often used as Secure Entry Points (SEP) keys. These SEP keys are the keys intended to be first used when building a chain of authority from a trust anchor to signed data. We advise a one-to-one mapping between SEP keys and key-signing keys. In practise key-signing keys have a lower rollover frequency than zone-signing keys so you should configure the SEP i.e. key-signing keys.

In addition to having the proper public key you should either be aware of the rollover policy of the zone owner, or that you have a tool that takes care of automated rollover. Failure to modify the trust anchor before the corresponding SEP key is rolled will result in validation failures.

Assume you have obtained the key-signing keys of `nlnetlabs.nl.`, `193.in-addr.arpa.`, and `195.in-addr.arpa.`. To configure those key as a trust anchor you will have tell your recursive nameserver to use those. For both BIND and UNBOUND you can follow the following procedure.

Create a separate file that contains the keys in a `trusted-keys` directive as shown in figure 3. The format is similar to the DNSKEY RR except that the "DNSKEY" label, the CLASS and the TTL, are omitted and quotes are placed around the name and the public key material.

You can use the include statements in both UNBOUND and BIND to read the files with the trusted keys directive into your configuration files.

2.3.3 Public repositories of trust-anchors

There are a few places where you can obtain trust-anchors.

The IANA ITAR (at <https://itar.iana.org>) is IANA's Interim Trust Anchor Repository to share the key material required to perform DNSSEC

⁴For an example see figure 3.4 where the DNSKEY RRset is signed with two keys with key id 17000 and 49656 and the other RRsets in the zone with the key with keyid 17000; the key with keyid 17000 is the zone-signing key and the key with key id 49656 is the key-singing key

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

```
// Trusted keys
// These are examples only, do not use in production

trusted-keys {
"nlnetlabs.nl." 3 5
  "AQPzzTWMz8qSWIQLfRnPckx2BiVmkVN6LPupO3mbz7Fh
  LSnm26n6iG9NLby97Ji453aWZY3M5/xJBSOS2vWtco2t
  8C0+xe01bc/d6ZTy32DHchpW6rDH1vp86L1+ha0tmwyy
  9QP7y2bVv5zSbFCrefk8qCUBgfHm9bHzMG1UBYtEIQ==";

"193.in-addr.arpa." 257 3 5
  "AwEAAc2RnCT1gjU22FbNC1baMQec77fq60z2H1CKscY1
  3idBZTp703ApMfAAFcMZQGkSmo8NP+47KqZJwG9ISLaT
  bUais3khgFVrf71IRzPJAMlXHsmOMmpq5xBORF66EDt/
  u2dau3qqz0fb/BrKcklGgnwBosgqaSPmWBQTuzJFqzi3
  4FQIt4xFHWYyt3B5qZ9h4dpUL96etvvx1N+z8t1Xjhlm
  Vauw1EPZnz2rmY6HEJFS2zjaI1FrDtY5/pooJjRWjobk
  RXL3iqjd5J/cmDikxjQCjwnbwBS+YvcwZCos4n9Xh2l5
  kf2k0cq9xCmZvEplfWJ9lWbVkfhpWam8qXXPN8E=";

"195.in-addr.arpa." 257 3 5
  "AwEAAaMn4k0rGaiHJBikvcf+mhPxzprL85Q40VA0hbRc
  a8FDDn6X1kuj95Nizy2vMr0y1MjIjo7a+GACGp6C/Rdj
  6nDimsRrUBr/G/dq+zBgg8qvRXWJZh+xzNCgkfv9gs1B
  eRlPnjXr1K/x5viTzQRDK3SYfHiCMVNxuYN+T7kniDLx
  QRUI/ASF3YxqNQ+Oo+T5L6nYt07uLeAUdxzToRdIHaeY
  iSnq52boA/3Yg6X8Kbo1uAUpeU4QDD7bOwq+obmaToLU
  m/FvNUKx019U2P2ItcsqRCHQut/RxK2pj8GGRDCCo1J
  5UAi7hiwP1eEWmbigbPnDQg++QDjegV39vTJQ2c=";
};
```

Figure 3: trust anchor configuration

verification of signed top-level domains, in lieu of a signed DNS root zone.

The tar publishes key information of TLDs, the exchange of information between IANA and the TLDs is the same as would normally take place for root-zone maintenance.

The UCLA secspider (at <https://secspider.cs.ucla.edu>) contains zones that have been submitted by users (via the on-line submission form), crawled from a large list of over 2.5 million zones, and walked (via NSEC walking). Their collection of keys in a format that is understood by Unbound and BIND is available at: <http://secspider.cs.ucla.edu/trust-anchors.conf>.

The IKS Jena TAR (at <https://www.iks-jena.de/leistungen/dnssec.php>)

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

is build from scanning the DNS and by open registration, it is made available through DLV.

2.3.4 Testing

As soon as a `trusted-key` has been configured, data from that zone or its sub zones will be validated by the caching forwarder. You can test this by querying your server⁵. If data is validated by the caching forwarder the `ad-bit` will be set by the name server (see the 'flags' in the following example).

```

; <<>> DiG 9.7.0b2 <<>> @192.168.2.204 example.net SOA +dnssec +multiline +retry=1
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 50415
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 2, AUTHORITY: 2, ADDITIONAL: 3

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
example.net.          IN SOA

;; ANSWER SECTION:
example.net.          100 IN SOA ns.example.net. olaf.nlnetlabs.nl. (
                        2002050501 ; serial
                        100      ; refresh (1 minute 40 seconds)
                        200      ; retry (3 minutes 20 seconds)
                        604800   ; expire (1 week)
                        100      ; minimum (1 minute 40 seconds)
                        )
example.net.          100 IN RRSIG SOA 5 2 100 20100507110046 (
                        20100407110046 17000 example.net.
                        0iHGv/u2i+D9YUDdW349dJPiDcNrL16ip+OpDALMb4b1
                        0xYzA2UG7GLdf0JH/jrnJ96F8Au101MuLhpkfbuccxgB
                        s7HH9BlqHwS085rJtluTiWyet/gopgdm1N4G49cHgmRY
                        djPGtUEz1s8ZSf+07I15v4JG8Dep0729KXUYoAY= )

;; AUTHORITY SECTION:
example.net.          100 IN NS ns.example.net.
example.net.          100 IN RRSIG NS 5 2 100 20100507110046 (
                        20100407110046 17000 example.net.
                        omQt5AIXfYtmqUDByq9QrXM2f9PuH0S5R/4iGrQzsw4/
                        Zp0w0eW+QrKrnlSQWpwUADsV10AFZL1SkDpX63508F17
                        hQ13v0XtoQi7qzMa6JqXvQYSjfkMk1Az3N5ZFE4RuaI
                        X8CaZePrkdHh9rFfTyuWJsC9hBkrpbQnlRyqv4U= )

;; ADDITIONAL SECTION:
ns.example.net.       100 IN A 192.168.2.203
ns.example.net.       100 IN RRSIG A 5 3 100 20100507110046 (
                        20100407110046 17000 example.net.
                        X8RgppDa83kHZeUN6b3ynA0sbFFUsY8cxS+384tY+eMP
                        FN07q4RCTPwpX+6jFN17lnSHWqJx+FPUIRF4kL4cxhsm
                        gydmuRA0gC+d081ERU3HxEizUBXu0SHYSNLAVw7Qmzbl
                        hb47NcGp/sod5tazrlabqCnqIrQ8uZaQL7F+fpI= )

;; Query time: 4 msec
;; SERVER: 192.168.2.204#53(192.168.2.204)
;; WHEN: Wed Apr 7 14:00:50 2010
;; MSG SIZE rcvd: 639

```

⁵We use the bind supplied `dig` tool, alternatively you can use NLnet Labs' `drill`, see section 6

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

It is important that you check that the validation is working correctly. This can be done by using the BIND log facilities on the machine that is configured as the validating recursive name server.

In BIND messages of a certain category can be logged to separate channels. The channels determine where the messages go and to what severity level they will need to be reported. The relevant category for DNSSEC validation is `dnssec`. In the example below the errors of the `dnssec` category are directed to the `dnssec_log` channel. In order to follow the validation process the channel has to log at least `severity debug 3`.

```
logging {
    channel dnssec_log {                // a DNSSEC log channel
        file "log/dnssec" size 20m;
        print-time yes;                // timestamp the entries
        print-category yes;           // add category name to entries
        print-severity yes;           // add severity level to entries
        severity debug 3;              // print debug message <= 3 t
    };

    category dnssec { dnssec_log; };
}
```

The output in the log file will look similar to the output below. The attempt for positive response validation shows how the validator tries to prove that the RR set is trusted by following the chain of trust to the appropriate secure entry point, your trusted-key statement. Chains of trust (see figure 4) start by the validation of a signature over a DNSKEY RRset, then these keys are used to validate the DS RRset that point to DNSKEY RRs in a child zone – which validates the DNSKEY RRs in the child zone –, or the DNSKEYs can be used to validate the data you have queried for. The log reflects the activity of the validator following the chain of trust.

```
validating @0x1008d7a00: example.net SOA: starting
validating @0x1008d7a00: example.net SOA: attempting positive response validation
validating @0x1008d8800: example.net DNSKEY: starting
validating @0x1008d8800: example.net DNSKEY: attempting positive response validation
validating @0x1008d8800: example.net DNSKEY: verify rdataset (keyid=49656): success
validating @0x1008d8800: example.net DNSKEY: signed by trusted key; marking as secure
validator @0x1008d8800: dns_validator_destroy
validating @0x1008d7a00: example.net SOA: in fetch_callback_validator
validating @0x1008d7a00: example.net SOA: keyset with trust 7
validating @0x1008d7a00: example.net SOA: resuming validate
validating @0x1008d7a00: example.net SOA: verify rdataset (keyid=17000): success
validating @0x1008d7a00: example.net SOA: marking as secure
validator @0x1008d7a00: dns_validator_destroy
```

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

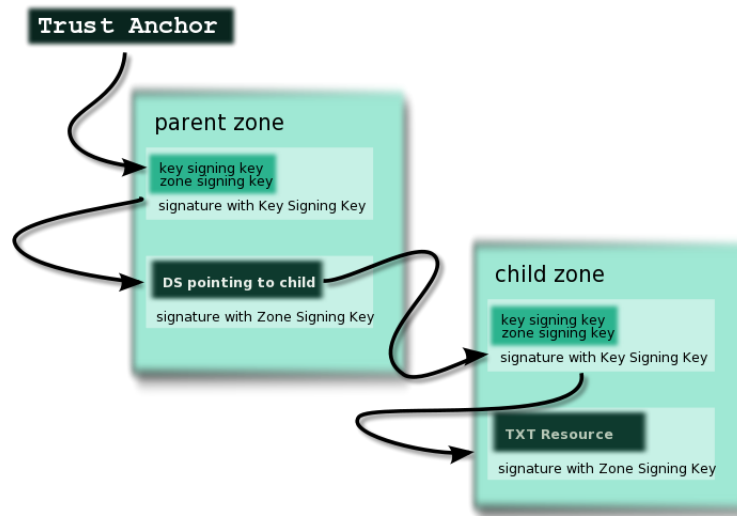


Figure 4: Chain of Trust

2.4 Finding trust-anchors

It is not trivial to find and maintain trust anchors. If you want to get started with validation of DNSSEC here are a few places where you can find more information.

- RIPE NCC maintains a set of keys on their secured website under <https://www.ripe.net/projects/disi/keys/index.html> (Note that this is a secured website, check the certificate).
- The Swedish TLD is signed. Its key can be found at <https://dnssec.nic.se/key.html>. (Also check the certificate of this site).
- The DNSSEC spider tries to locate secured zones and checks their status. You could use this site to find secured zones. See <http://secspider.cs.ucla.edu/islands.html>.

2.5 Maintaining Trust Anchors

Once trust anchors are configured you will need to make sure they are kept in sync with the key as published by the entity responsible for the zone to which it belongs.

Keeping in sync with trust anchors is incredibly important and while there is a standard to signal rollovers([16]) it is not yet widely deployed. Once a maintainer of a zone uses the RFC5011 mechanisms there are a number of tools to take care of the rollover, *e.g.* the NLnet Labs autotrust tool at <http://www.nlnetlabs.nl/projects/autotrust/>.

More attention to the maintenance of trust anchors will be paid in a future version of this HOWTO.

2.6 Lookaside Validation

Remember figure 2. If you would like to validate all these islands you will have to configure many trust-anchors, as in the example in figure 3.

In order to deal with this problem in absence of secure delegation from a small set of trust-anchors (ideally only 1, the root), BIND supports, as of version 9.3.2, a ⁶ mechanism called lookaside validation [19, 18].

In the lookaside validation a *DLV registry* will maintain all the trust-anchors you trust them to do the "Good Thing". The maintainers of zones that are secure register their trust-anchors with the *DLV registry* and (non-standard extensions) in BIND (as of 9.3.2) will allow you, operator of a validating nameserver, to make use of all trust anchors that are present in the DLV tree.

In the DLV scheme the trust anchors are published in a dedicated domain (`dlv.isc.org` in the figure 5). Whenever a validating resolver recognises that a zone is signed it will first try to validate it by assessing if it is within the island of trust configured by its local trust anchors. When the validated domain is not in a trusted island the resolver will lookup perform a lookup in the DLV domain and use the trust anchor from that zone if and when available.

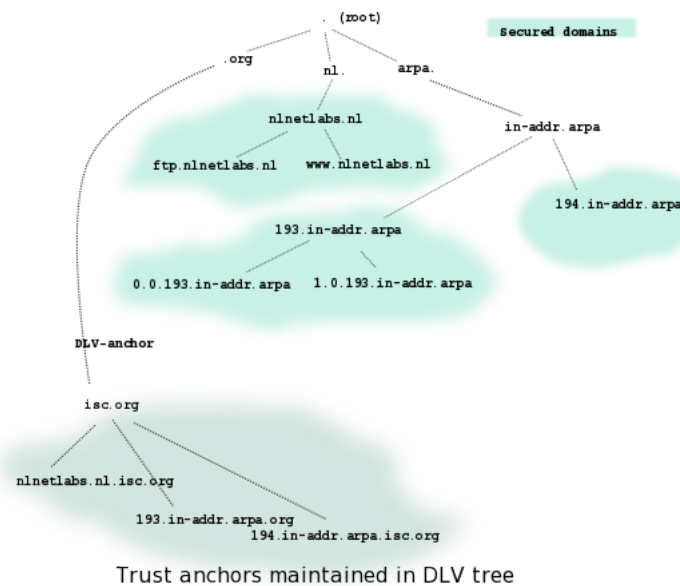


Figure 5: Trust anchors in the lookaside tree

⁶Authors opinion inserted here: I am not very enthusiastic about lookaside validation

Since anybody can, in theory, start a DLV registry it will very difficult to know who to trust and to get a full picture of which zones are secured and which not. On the other hand it is a mechanism to provide early deployment maximum benefit. At some point this section may disappear from the HOWTO.

At the moment of writing the author advises to use ISC's DLV registry.

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

2.6.1 Configuring lookaside validation

What follows is a generic description if you want to configure ISC's DLV as your authoritative lookaside domain you may want to read <http://www.isc.org/ops/dlv/>.

In the example below we assume that `dlv-registry.org` is the registry of our choice.

You have to perform two (additional) steps in order to turn on lookaside validation.

Configure a trust anchor for the DLV registry. You do this by defining a trust anchor for the island of trust defined by `dlv-registry.org` in `named.conf`. Obviously, this trust-anchor is not exclusive, any trust-anchor configured in your trusted-keys statement will have preference over the data in the DLV registry.

```
trusted-keys {
    //
    // this trust-anchor defines dlv-registry.org as a trusted island.
    //
    "dlv-registry.org." 257 3 5
        "AQXPX7B3JTdPPhMl ... u82ggY2BKPQ==";
    //
    // Other trust anchors below.
    //

    "nlnetlabs.nl." 3 5
        "AQPzzTWMz8qSWI ... zMG1UBYtEIQ==";

    "193.in-addr.arpa." 257 3 5
        "AwEAAc2RnCT1gj ... pWam8qXXPN8E=";

    "195.in-addr.arpa." 257 3 5
        "AwEAAaMN4kOrGai ... DjegV39vTJQ2c=";
};
```

Configuring how the DNS name space anchors in the DLV name space. By using the `dnssec-lookaside` statement in the `options` section of `named.conf`. The statement takes two arguments the first one is the domain in the DNS for which lookaside validation is to be applied. Usually this will be the full name space so the `."` (root) is configured. The second argument is the name of the trust-anchor where a lookup should be performed for a DLV record.

It is best to configure only one DLV trust-anchor.

```
options {
    // DNSSEC should be turned on, do not forget
    dnssec-enable yes;
    dnssec-validation yes;
```


2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

```
// This sets the dlvs registry "dlv-registry.org"
dnssec-lookaside "." trust-anchor "dlv-registry.org.";

// other options are skipped in this example
};
```

testing When you have your logging configured as described in section 2.3.4 *i.e.* you log errors of the `dnssec` category are directed to a channel that logs at least at severity debug 3, then your log output when querying for `example.net` SOA will be similar to what is shown below.

First, the amount of log-output to validate one query covers more than one page of fine print. On a production server this data for several validation sequences will be print mixed. It will be very hard to debug from logfiles on production servers if you have not first looked at what happens for a single query.

Second, the structure is that the validator first finds DNSSEC RRs, notices that those records are not secure according 'plain DNSSEC' and then moves to DLV validation.

Third, small chains of trust are build, from the DLV trust-anchor, via DNSKEY RRs to the signatures over the data. Try to follow these trust anchors in the example output so it will be easier to identify them in production logs.

```
validating @0x1008d0a00: . NS: starting
validating @0x1008d0a00: . NS: looking for DLV
validating @0x1008d0a00: . NS: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008d0a00: . NS: looking for DLV dlv-registry.org
validating @0x1008d0a00: . NS: DLV lookup: wait
validating @0x1008ca000: dlv-registry.org DLV: starting
validating @0x1008ca000: dlv-registry.org DLV: attempting negative response validation
  validating @0x1008da800: dlv-registry.org SOA: starting
    validating @0x1008da800: dlv-registry.org SOA: attempting positive response validation
  validating @0x1008db600: dlv-registry.org DNSKEY: starting
    validating @0x1008db600: dlv-registry.org DNSKEY: attempting positive response validation
      validating @0x1008db600: dlv-registry.org DNSKEY: verify rdataset (keyid=8916): success
        validating @0x1008db600: dlv-registry.org DNSKEY: signed by trusted key; marking as secure
          validator @0x1008db600: dns_validator_destroy
            validating @0x1008da800: dlv-registry.org SOA: in fetch_callback_validator
              validating @0x1008da800: dlv-registry.org SOA: keyset with trust 7
                validating @0x1008da800: dlv-registry.org SOA: resuming validate
                  validating @0x1008da800: dlv-registry.org SOA: verify rdataset (keyid=27467): success
                    validating @0x1008da800: dlv-registry.org SOA: marking as secure
                      validator @0x1008da800: dns_validator_destroy
            validating @0x1008ca000: dlv-registry.org DLV: in authvalidated
          validating @0x1008ca000: dlv-registry.org DLV: resuming nsecvalidate
            validating @0x1008da800: dlv-registry.org NSEC: starting
              validating @0x1008da800: dlv-registry.org NSEC: attempting positive response validation
                validating @0x1008da800: dlv-registry.org NSEC: keyset with trust 7
                  validating @0x1008da800: dlv-registry.org NSEC: verify rdataset (keyid=27467): success
                    validating @0x1008da800: dlv-registry.org NSEC: marking as secure
                      validator @0x1008da800: dns_validator_destroy
            validating @0x1008ca000: dlv-registry.org DLV: in authvalidated
          validating @0x1008ca000: dlv-registry.org DLV: looking for relevant nsec
        validating @0x1008ca000: dlv-registry.org DLV: nsec proves name exists (owner) data=0
      validating @0x1008ca000: dlv-registry.org DLV: resuming nsecvalidate
```

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

```

validating @0x1008ca00: dlv-registry.org DLV: nonexistence proof(s) found
validator @0x1008ca00: dns_validator_destroy
validating @0x1008d0a00: . NS: in dlvfetched: ncache nxrrset
validating @0x1008d0a00: . NS: DLV not found
validating @0x1008d0a00: . NS: marking as answer
validator @0x1008d0a00: dns_validator_destroy
validating @0x1008d0a00: example.net SOA: starting
validating @0x1008d0a00: example.net SOA: looking for DLV
validating @0x1008d0a00: example.net SOA: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008d0a00: example.net SOA: looking for DLV example.net.dlv-registry.org
validating @0x1008d0a00: example.net SOA: DNS_R_COVERINGNSEC
validating @0x1008d0a00: example.net SOA: covering nsec: not in range
validating @0x1008d0a00: example.net SOA: DLV lookup: wait
validating @0x1008da800: example.net.dlv-registry.org DLV: starting
validating @0x1008da800: example.net.dlv-registry.org DLV: attempting positive response validation
validating @0x1008da800: example.net.dlv-registry.org DLV: keyset with trust 7
validating @0x1008da800: example.net.dlv-registry.org DLV: verify rdataset (keyid=27467): success
validating @0x1008da800: example.net.dlv-registry.org DLV: marking as secure
validator @0x1008da800: dns_validator_destroy
validating @0x1008d0a00: example.net SOA: in dlvfetched: success
validating @0x1008d0a00: example.net SOA: DLV example.net found
validating @0x1008d0a00: example.net SOA: dlv_validator_start
validating @0x1008d0a00: example.net SOA: restarting using DLV
validating @0x1008d0a00: example.net SOA: attempting positive response validation
validating @0x1008da800: example.net DNSKEY: starting
validating @0x1008da800: example.net DNSKEY: looking for DLV
validating @0x1008da800: example.net DNSKEY: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008da800: example.net DNSKEY: looking for DLV example.net.dlv-registry.org
validating @0x1008da800: example.net DNSKEY: DLV example.net found
validating @0x1008da800: example.net DNSKEY: dlv_validator_start
validating @0x1008da800: example.net DNSKEY: restarting using DLV
validating @0x1008da800: example.net DNSKEY: attempting positive response validation
validating @0x1008da800: example.net DNSKEY: dlv_validatezonekey
validating @0x1008da800: example.net DNSKEY: Found matching DLV record: checking for signature
validating @0x1008da800: example.net DNSKEY: verify rdataset (keyid=17000): RRSIG failed to verify
validating @0x1008da800: example.net DNSKEY: verify rdataset (keyid=49656): success
validating @0x1008da800: example.net DNSKEY: marking as secure
validator @0x1008da800: dns_validator_destroy
validating @0x1008d0a00: example.net SOA: in fetch_callback_validator
validating @0x1008d0a00: example.net SOA: keyset with trust 7
validating @0x1008d0a00: example.net SOA: resuming validate
validating @0x1008d0a00: example.net SOA: verify rdataset (keyid=17000): success
validating @0x1008d0a00: example.net SOA: marking as secure
validator @0x1008d0a00: dns_validator_destroy

```

When using lookaside validation assessing the log output in case of corrupted zone data is a challenge. Below is the output of the validator when it tries to figure out if a query that returns a corrupted result is valid or not. The conclusion is reached in the last few lines.

```

validating @0x1008cfa00: . NS: starting
validating @0x1008cfa00: . NS: looking for DLV
validating @0x1008cfa00: . NS: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008cfa00: . NS: looking for DLV dlv-registry.org
validating @0x1008cfa00: . NS: DLV lookup: wait
validating @0x1008d7800: corrupt.example.net A: starting
validating @0x1008d7800: corrupt.example.net A: looking for DLV
validating @0x1008d7800: corrupt.example.net A: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008d7800: corrupt.example.net A: looking for DLV corrupt.example.net.dlv-registry.org
validating @0x1008d7800: corrupt.example.net A: DNS_R_COVERINGNSEC
validating @0x1008d7800: corrupt.example.net A: covering nsec: trust 1
validating @0x1008d7800: corrupt.example.net A: DLV lookup: wait
validating @0x1008dc600: dlv-registry.org DLV: starting

```

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

```

validating @0x1008dc600: dlv-registry.org DLV: attempting negative response validation
validating @0x1008dd400: dlv-registry.org SOA: starting
validating @0x1008dd400: dlv-registry.org SOA: attempting positive response validation
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: starting
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: attempting negative response validation
validating @0x1008df000: dlv-registry.org SOA: starting
validating @0x1008df000: dlv-registry.org SOA: attempting positive response validation
validating @0x1008dfe00: dlv-registry.org DNSKEY: starting
validating @0x1008dfe00: dlv-registry.org DNSKEY: attempting positive response validation
validating @0x1008dfe00: dlv-registry.org DNSKEY: verify rdataset (keyid=8916): success
validating @0x1008dfe00: dlv-registry.org DNSKEY: signed by trusted key; marking as secure
validator @0x1008dfe00: dns_validator_destroy
validating @0x1008dd400: dlv-registry.org SOA: in fetch_callback_validator
validating @0x1008dd400: dlv-registry.org SOA: keyset with trust 7
validating @0x1008dd400: dlv-registry.org SOA: resuming validate
validating @0x1008dd400: dlv-registry.org SOA: verify rdataset (keyid=27467): success
validating @0x1008dd400: dlv-registry.org SOA: marking as secure
validator @0x1008dd400: dns_validator_destroy
validating @0x1008dc600: dlv-registry.org DLV: in authvalidated
validating @0x1008dc600: dlv-registry.org DLV: resuming nsecvalidate
validating @0x1008df000: dlv-registry.org SOA: in fetch_callback_validator
validating @0x1008df000: dlv-registry.org SOA: keyset with trust 7
validating @0x1008df000: dlv-registry.org SOA: resuming validate
validating @0x1008df000: dlv-registry.org SOA: verify rdataset (keyid=27467): success
validating @0x1008df000: dlv-registry.org SOA: marking as secure
validator @0x1008df000: dns_validator_destroy
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: in authvalidated
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: resuming nsecvalidate
validating @0x1008df000: example.net.dlv-registry.org NSEC: starting
validating @0x1008df000: example.net.dlv-registry.org NSEC: attempting positive response validation
validating @0x1008df000: example.net.dlv-registry.org NSEC: keyset with trust 7
validating @0x1008df000: example.net.dlv-registry.org NSEC: verify rdataset (keyid=27467): success
validating @0x1008df000: example.net.dlv-registry.org NSEC: marking as secure
validator @0x1008df000: dns_validator_destroy
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: in authvalidated
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: looking for relevant nsec
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: nsec range ok
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: resuming nsecvalidate
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: in checkwildcard: *.example.net.dlv-registry.org
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: looking for relevant nsec
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: nsec range ok
validating @0x1008de200: corrupt.example.net.dlv-registry.org DLV: nonexistence proof(s) found
validator @0x1008de200: dns_validator_destroy
validating @0x1008d7800: corrupt.example.net A: in dlvfetched: ncache nxdomain
validating @0x1008d7800: corrupt.example.net A: looking for DLV example.net.dlv-registry.org
validating @0x1008d7800: corrupt.example.net A: DLV lookup: wait
validating @0x1008dd400: dlv-registry.org NSEC: starting
validating @0x1008dd400: dlv-registry.org NSEC: attempting positive response validation
validating @0x1008dd400: dlv-registry.org NSEC: keyset with trust 7
validating @0x1008dd400: dlv-registry.org NSEC: verify rdataset (keyid=27467): success
validating @0x1008dd400: dlv-registry.org NSEC: marking as secure
validator @0x1008dd400: dns_validator_destroy
validating @0x1008dc600: dlv-registry.org DLV: in authvalidated
validating @0x1008dc600: dlv-registry.org DLV: looking for relevant nsec
validating @0x1008dc600: dlv-registry.org DLV: nsec proves name exists (owner) data=0
validating @0x1008dc600: dlv-registry.org DLV: resuming nsecvalidate
validating @0x1008dc600: dlv-registry.org DLV: nonexistence proof(s) found
validator @0x1008dc600: dns_validator_destroy
validating @0x1008cfa00: . NS: in dlvfetched: ncache nxrrset
validating @0x1008cfa00: . NS: DLV not found
validating @0x1008cfa00: . NS: marking as answer
validator @0x1008cfa00: dns_validator_destroy
validating @0x1008cfa00: example.net.dlv-registry.org DLV: starting
validating @0x1008cfa00: example.net.dlv-registry.org DLV: attempting positive response validation
validating @0x1008cfa00: example.net.dlv-registry.org DLV: keyset with trust 7

```

2 CONFIGURING A RECURSIVE NAME SERVER TO VALIDATE ANSWERS

```

validating @0x1008cfa00: example.net.dlv-registry.org DLV: verify rdataset (keyid=27467): success
validating @0x1008cfa00: example.net.dlv-registry.org DLV: marking as secure
validator @0x1008cfa00: dns_validator_destroy
validating @0x1008d7800: corrupt.example.net A: in dlvfetched: success
validating @0x1008d7800: corrupt.example.net A: DLV example.net found
validating @0x1008d7800: corrupt.example.net A: dlv_validator_start
validating @0x1008d7800: corrupt.example.net A: restarting using DLV
validating @0x1008d7800: corrupt.example.net A: attempting positive response validation
validating @0x1008cfa00: example.net DNSKEY: starting
validating @0x1008cfa00: example.net DNSKEY: looking for DLV
validating @0x1008cfa00: example.net DNSKEY: plain DNSSEC returns unsecure (.): looking for DLV
validating @0x1008cfa00: example.net DNSKEY: looking for DLV example.net.dlv-registry.org
validating @0x1008cfa00: example.net DNSKEY: DLV example.net found
validating @0x1008cfa00: example.net DNSKEY: dlv_validator_start
validating @0x1008cfa00: example.net DNSKEY: restarting using DLV
validating @0x1008cfa00: example.net DNSKEY: attempting positive response validation
validating @0x1008cfa00: example.net DNSKEY: dlv_validatezonekey
validating @0x1008cfa00: example.net DNSKEY: Found matching DLV record: checking for signature
validating @0x1008cfa00: example.net DNSKEY: verify rdataset (keyid=17000): RRSIG failed to verify
validating @0x1008cfa00: example.net DNSKEY: verify rdataset (keyid=49656): success
validating @0x1008cfa00: example.net DNSKEY: marking as secure
validator @0x1008cfa00: dns_validator_destroy
validating @0x1008d7800: corrupt.example.net A: in fetch_callback_validator
validating @0x1008d7800: corrupt.example.net A: keyset with trust 7
validating @0x1008d7800: corrupt.example.net A: resuming validate
validating @0x1008d7800: corrupt.example.net A: verify rdataset (keyid=17000): RRSIG failed to verify
validating @0x1008d7800: corrupt.example.net A: failed to verify rdataset
validating @0x1008d7800: corrupt.example.net A: verify failure: RRSIG failed to verify
validating @0x1008d7800: corrupt.example.net A: no valid signature found
validator @0x1008d7800: dns_validator_destroy

```

2.7 Some Troubleshooting Tips

Suppose that you have configured a trust anchor and you are experiencing problems. For instance, your nameserver returns "SERVFAIL" for particular queries. Well, "SERVFAIL" is the default return code that a validating nameserver returns when it flags data as being bogus. Bogus data can be caused by two things. Either you are under attack or you are experiencing a configuration error either by the operator of one of the zones in the chain of trust or by the operator of the validating recursive nameserver.

In addition to looking at the logs there are a number of tools at your disposal (see III). To assess if problems occur because of misconfiguration, or bugs, in your validating nameserver, or because of problems with the signed zones you will need a troubleshooting strategy.

One of the approaches you could take is to first use `drill` (see 6) or `dig` (see 7) to perform a 'sigchase' or a 'trace' with a key copied to your local file system, circumventing your validating recursive nameserver. In that way you will be able to check if the chain of trust can actually be built from the data. Make sure you use the correct trust-anchor when tracing data.

When you have verified that the chain of trust can be built from the data in the DNS it is time to troubleshoot the validating nameserver. This is easy when you have access to the log files but may be more troublesome if you don't. You could use `dig` to query the validating nameservers with and without the `+cd` flag. That flag sets a bit in the query that instructs the nameserver not to perform

validation. When the individual pieces in the chain of trust (`drill` returned those when using the `trace` option) you may be able to find inconsistencies that indicate that an expired trust anchor has been configured. You start by querying the DNSKEY RRset for which you assume there is a trust-anchor considered and work your way down. Or, alternatively, you query for the data you were looking for, use the data in the RRSIG RR to find for which DNSKEY RR to query, then query for a DS RR at the same name and work your way up (similar to the 'sigchase' in `drill`).

While troubleshooting there are a number of failures that are easily introduced. Take the following example.

In his ISP column [9, 8, 7], Geoff Huston documented his experiences as an early deployer. He blindly configured a trust-anchor for the "nlnetlabs.nl" zone. While this zone was signed it was done in an experimental setup whereby not all servers for the zone were configured with the same version of the protocol. In this case one of the servers would not provide RRSIGs with the answer, something which may give rise to re-occurring but hard to predict failures.

There are two things to learn from this: Never blindly configure trust anchors in validating resolvers and make sure that when you serve zones all your servers conform to the DNSSEC protocol specification.

Another failure is that one of the RRsets in the chain of trust has expired signatures. Check this by looking at the date fields in the RRSIG RRs.

More problematic to find may be a rollover, where DNSKEY RRs or RRSIG RRs have been removed to early so that there is an inconsistency between data in cache and data that needs to be validated (also see section 5). Using the `+cd` to with `dig` and looking at the TTLs might help to distinguish if you are trying to validate RRSIGs for which there are no DNSKEYs available (or vice verse).

3 Securing a DNS zone

3.1 Introduction

If a zone has been signed and its key has been configured in a validating recursive name server we usually refer to it as being an "island of security". It apparently does not have a secured parent and stands alone in the sea of other unsecured domains. Usually creating an "island of security" is the first step to becoming part of the secure DNS. The "island of security" will remain "insecure" for resolvers that have no trust anchor configured for the domain.

If a zone owner decides to create "an island of security" she will sign her zones and distribute the "secure entry points" to the system administrators that want to validate her zone data. Once the island of security has been set up the island can become part of the secure tree by exchanging the "secure entry point" with the parent.

After creation of the key-pairs used for signing and validation we want to sign the zone data for our own organisation (e.g. `example.net.`) and configure the caching forwarders on our organisations network to validate data against the public key of our organisation.

In the text below, we assume that your organisation's domain names are maintained in one zone. If domain name administration is delegated to sub-zones, see section Chapter 4, "Delegating of signing authority; becoming globally secure".

Signing the zone data is the task of the zone administrator, configuring the caching forwarder is a task of system administrators.

The examples are based on the example zone in section Figure 7.

3.2 Configuring authoritative servers

All the authoritative servers will need to be configured to deal with the DNSSEC protocol. How this is done for BIND is explained in AppendixB. The essential steps are compiling bind with openssl and enabling dnssec through the use of the `dnssec-enable yes;` directive in the `options` section of `named.conf`.

That is all there is to it.

3.3 Creating key pairs

3.3.1 Key Maintenance Policy

Before generating keys, you will need to think about your key maintenance policy. Such policy should address

- What will be the sizes of your keys?
- Will you separate the key- and zone-signing keys functionality?
- How often will you roll the keys?
- How will system administrators that intend to use your zone as a trust anchor get hold of the appropriate public key and what mechanism will you provide to enable them to validate the authenticity of your public key?
- How will you signal a key rollover or how can you make sure that all interested parties are aware of a rollover?

Some of these issues may be easy to address. For example, your organisation may have established mechanisms to distribute the public keys, there may be obvious ways to publish an upcoming rollover such as the possibility of publishing the event in a corporate newspaper. Alternatively, it may be possible to notify all relevant parties by mail when a corporate X.509 hierarchy is available for e-mail validation.

3.3.1.1 Key- and zone-signing keys. The author thinks it is good practise to use zone-signing keys and key-signing keys (also see Chapter 5, "Rolling keys"). The key-signing keys are usually the first keys from your zone that are used to build a chain of authority to the data that needs to be validated. Therefore, these keys are often called a secure entry point key (or SEP key). These SEP keys are the ones that you should exchange with your parents or that validating resolvers configure as their trust anchors.

Throughout this document we assume that you use separate key- and zone-signing keys and that the key-signing keys are exclusively used as secure entry

3 SECURING A DNS ZONE

point keys and can be identified by the SEP[11] bit in the flag field; the flag field is odd.

3.3.2 Creating the keys

`dnssec-keygen` is the tool that we use to generate key pairs. The arguments that we have to provide `dnssec-keygen` are shown in Figure 6.

The output can be found in two files. The name of the files contain relevant information:

Kdomain_name+algorithm_id+key_id.extension

The *domain_name* is the name specified on the command line. It is used by other BIND DNSSEC tools, if you use a different name from the domain name, you might confuse those tools. The *algorithm_id* identifies the algorithm used: 1 for RSAMD5, 3 for DSA, 5 for RSASHA1 and 54 for HMAC-MD5 (TSIG only). The *key_id* is an identifier for the key material. This *key_id* is used by the RRSIG Resource Record. The *extension* is either `key` or `private`, the first is the public key, the second is the private key.

We create an RSASHA1 zone-signing key pair for `example.net`:

```
# dnssec-keygen -r/dev/random -a RSASHA1 -b 1024 -n ZONE example.net
Kexample.net.+005+17000
```

Because of the considerations in Section 3.3.1 you will also need to create SEP keys. Create keys with the SEP bit set by specifying the `-f KSK` flag with `dnssec-keygen`.

```
# dnssec-keygen -r/dev/random -f KSK -a RSASHA1 -b 1280 -n ZONE example.net
Kexample.net.+005+49656
```

Lets have a look at the content of these files⁷

```
cat Kexample.net.+005+17000.key
example.net. IN DNSKEY 256 3 5 (
    AQPI4+0M1V055RS2Hqv+8w8V20Dh+SQmFzHQtZMuzLH3UxWEOGmG5Gfj
    ijandJeAZTKLpERXB6RfHHTGG81D3IO1azWN6DiVFEVzgr0otAdDonfY
    +oEsRw== )
```

The public key (`.key` extension) is exactly as it appears in your zone file. Note, that the TTL value is not specified. This key has a "flag" value of 256. Since this value is an even number, the key is not marked as a SEP key and should be used for zone-signing.

⁷ We slightly edited the output for readability. We printed the base64 material over several lines and introduced the brackets so that this is a legal multi-line representation of a RR.

The private key (`.private` extension) contains all the parameters that make an RSASHA1 private key. The private key of a RSA key contains different parameters to DSA. Here is the private key (with base64 material truncated):

```
cat Kexample.net.+005+17000.private
Private-key-format: v1.2
Algorithm: 5 (RSASHA1)
Modulus: yOPtDNVd0eUUth6r/vMPFdtA4fkkJhcx0LWTLsyx91MVhNBphu...
PublicExponent: Aw==
PrivateExponent: he1Iszjo0UNjJBRyqdfdfY+eAlqYYGWTl4HkMyd3L+j...
Prime1: +X0kNW1JrepBnVw5o9fDUyWAT5zqxKt0YR4vJZ19991tLZAmD04...
Prime2: ziIX5qfpZGBuzfd847TqtDfYcWv5UfUrPAIa/11g31eUUNERmsB...
Exponent1: p1NteP0Gc/GBE5LRF+Us4hkANRNHLcei6210w75T+p0eHmAZ...
Exponent2: iWwP7xqbmEBJ3qT97SNHIs/logf7i/jHfVa8qj5ALDpi4Ith...
Coefficient: rmmgD9P7/ywQJ4F0epdGq0UoQZmqrPQsraDTD8vkU1wLju...
```

This private key should be kept secure⁸ i.e. the file permissions should be set so that the zone administrator will be able to access them when a zone needs to be signed. The BIND tools will, by default,⁹ look for the keys in the directory where signing is performed (see Section 3.4), that might not be the most secure place on your OS.

3.4 Zone-signing

When you create key pairs, you should include them in your zone file. Refer to the example in Figure 7, where we use the `$include` directive to include the keys. We increase the serial number in the SOA record before signing.

In the example below we will use the RSASHA1 type keys for zone and key-signing keys.

Once the key is included in the zone file we are ready to sign the zone using the `dnssec-signzone` tool (see Figure 8 for all the arguments). We use the `-o` flag to specify the origin of the zone; by default the origin is deduced from the zone file name.

With the `'-k key_name'` we specify which key is to be used as the key-signing key. That key will only sign the DNSKEY RR set in the apex of the zone. The keys that come as arguments at the end of the command are used to sign all the RR data for which the zone is authoritative. If you do not specify the keys, BIND will use the ones for which the public keys are included in the zone and use the SEP flag to distinguish between key- and zone-signing keys.

In practise you would not want to rely on the default, since in key rollover scenarios you will have a public key in your zone file but you would not want to use that for zone-signing (in order to avoid double signatures and therefore longer signature generation times and more resource consumption on your name server). Below is the command issued to sign a zone with the 49656 key as key-signing key and the 17000 key as zone-signing key.

⁸ At the RIPE NCC we are working on a dedicated signing server that has SSH based access control. Based on which key is used to login, a dedicated shell is opened; A zone maintenance shell allows signing of zones; A key maintenance shell for key maintenance. Only system administrators can access the key-material itself. A beta version of this tool is available on RIPE NCC's website <http://www.ripe.net/disi/dnssec_maint_tool/>.

⁹ It is possible to fully specify the path to the keys.

3 SECURING A DNS ZONE

```

/usr/local/sbin/dnssec-signzone \
-o example.net \
-k Kexample.net.+005+49656 \
db.example.net \
Kexample.net.+005+17000.key

```

The signed zone file is reproduced in figure 3.4 . Note that the apex DNSKEY RRset is the only RRset with two signatures, made with the zone- and key-signing keys. The other RRsets are only signed with the zone-signing keys.

The signing process completed the following:

- Sorted the zone in 'canonical' order¹⁰.
- Inserted NSEC records for every label.
- Added the key-id as a comment to each DNSKEY-record.
- Signed the DNSKEY RR set with two keys; the key-signing key and the zone-signing key.
- Signed the other RRs with the zone-signing key.
- It created two files, `dsset-example.net` and `keyset-example.net`. These two files are relevant when building a chain of trust. Per default the files are created in the 'current directory' i.e. the directory in which you ran the `dnssec-signzone` command, but when specifying the `-d`, with its directory, then the files end up there.

The signatures were created with a default life time of 30 days from the moment of signing. Once signatures have expired data can not be validated and your zone will be marked 'bogus'. Therefore you will have to re-sign your zone within 30 days. Zone re-signing is discussed below.

The signed zone is stored in `db.example.net.signed`, make sure you have configured `named` to use this file to serve the zones from.

¹⁰ A specific ordering as defined by the DNSSEC protocol.

3 SECURING A DNS ZONE

Usage:

dnssec-keygen [options] name

Version: 9.7.0b2

name: owner of the key

Options:

- K <directory>: write keys into directory
- a <algorithm>:
 - RSA | RSAMD5 | DSA | RSASHA1 | NSEC3RSASHA1 | NSEC3DSA |
 - RSASHA256 | RSASHA512 |
 - DH | HMAC-MD5 | HMAC-SHA1 | HMAC-SHA224 | HMAC-SHA256 |
 - HMAC-SHA384 | HMAC-SHA512(default: RSASHA1, or NSEC3RSASHA1 if using -3)
- 3: use NSEC3-capable algorithm
- b <key size in bits>:

RSAMD5:	[512..4096]
RSASHA1:	[512..4096]
NSEC3RSASHA1:	[512..4096]
RSASHA256:	[512..4096]
RSASHA512:	[1024..4096]
DH:	[128..4096]
DSA:	[512..1024] and divisible by 64
NSEC3DSA:	[512..1024] and divisible by 64
HMAC-MD5:	[1..512]
HMAC-SHA1:	[1..160]
HMAC-SHA224:	[1..224]
HMAC-SHA256:	[1..256]
HMAC-SHA384:	[1..384]
HMAC-SHA512:	[1..512]

(if using the default algorithm, key size defaults to 2048 for KSK, or 1024 for all others)
- n <nametype>: ZONE | HOST | ENTITY | USER | OTHER
(DNSKEY generation defaults to ZONE)
- c <class>: (default: IN)
- d <digest bits> (0 => max, default)
- E <engine name>
- e: use large exponent (RSAMD5/RSASHA1 only)
- f <keyflag>: KSK | REVOKE
- g <generator>: use specified generator (DH only)
- p <protocol>: (default: 3 [dnssec])
- s <strength>: strength value this key signs DNS records with (default: 0)
- T <rrrtype>: DNSKEY | KEY (default: DNSKEY; use KEY for SIG(0))
- t <type>: AUTHCONF | NOAUTHCONF | NOAUTH | NOCONF (default: AUTHCONF)
- r <randomdev>: a file containing random data
- h: print usage and exit
- m <memory debugging mode>:
 - usage | trace | record | size | mctx
- v <level>: set verbosity level (0 - 10)

Date options:

- P date/[+-]offset: set key publication date (default: now)
- A date/[+-]offset: set key activation date (default: now)
- R date/[+-]offset: set key revocation date
- I date/[+-]offset: set key inactivation date
- D date/[+-]offset: set key deletion date
- G: generate key only; do not set -P or -A
- C: generate a backward-compatible key, omitting all dates

Output:

K<name>+<alg>+<id>.key, K<name>+<alg>+<id>.private

3 SECURING A DNS ZONE

```
; example.net zone
;
$TTL 100
$ORIGIN example.net.
@           100      IN      SOA     ns.example.net. (
                                olaf.nlnetlabs.nl.
                                2002050501
                                100
                                200
                                604800
                                100
                                )

ns.example.net.      NS           ns.example.net.
                    A           192.168.2.203

a                   A           192.168.2.1
b                   A           192.168.2.2

*                   A           192.168.2.10
b.a                 A           192.168.2.11

; These are the keys that need to be published in the DNSKEY RRset
;
$include Kexample.net.+005+17000.key      ; ZSK
$include Kexample.net.+005+49656.key      ; KSK
```

Figure 7: example.net example zone

3 SECURING A DNS ZONE

Usage:

```
dnssec-signzone [options] zonefile [keys]
```

Version: 9.7.0b2

Options: (default value in parenthesis)

```
-S:          smart signing: automatically finds key files
             for the zone and determines how they are to be used
-K directory:
             directory to find key files (.)
-d directory:
             directory to find dsset files (.)
-g:          generate dsset file, and/or include DS records
             from child zones' dsset files
-s [YYYYMMDDHHMMSS|+offset]:
             RRSIG start time - absolute|offset (now - 1 hour)
-e [YYYYMMDDHHMMSS|+offset|"now"+offset]:
             RRSIG end time - absolute|from start|from now (now + 30 days)
-i interval:
             cycle interval - resign if < interval from end ( (end-start)/4 )
-j jitter:
             randomize signature end time up to jitter seconds
-v debuglevel (0)
-o origin:
             zone origin (name of zonefile)
-f outfile:
             file the signed zone is written in (zonefile + .signed)
-I format:
             file format of input zonefile (text)
-O format:
             file format of signed zone file (text)
-N format:
             soa serial format of signed zone file (keep)
-r randomdev:
             a file containing random data
-a:          verify generated signatures
-c class (IN)
-E engine:
             name of an OpenSSL engine to use
-p:          use pseudorandom data (faster but less secure)
-P:          disable post-sign verification
-T TTL:     TTL for newly added DNSKEYs      -t:          print statistics
-u:          update or replace an existing NSEC/NSEC3 chain
-x:          sign DNSKEY record with KSKs only, not ZSKs
-z:          sign all records with KSKs
-C:          generate a keyset file, for compatibility
             with older versions of dnssec-signzone -g
-n ncpus (number of cpus present)
-k key_signing_key
-l lookasidezone
-3 NSEC3 salt
-H NSEC3 iterations (10)
-A NSEC3 optout
```

Signing Keys: (default: all zone keys that have private keys)
keyfile (Kname+alg+tag)

3 SECURING A DNS ZONE

```

; File written on Wed Apr 7 13:50:10 2010
; dnssec\_signzone version 9.7.0b2
example.net.          100      IN SOA  ns.example.net. olaf.nlnetlabs.nl. (
                        2002050501 ; serial
                        100        ; refresh (1 minute 40 seconds)
                        200        ; retry (3 minutes 20 seconds)
                        604800    ; expire (1 week)
                        100        ; minimum (1 minute 40 seconds)
                        )
100      RRSIG  SOA 5 2 100 20100507105010 (
                        20100407105010 17000 example.net.
                        T38xkzDBM8ZniNRYcNdgd2RR+wnNY2lQeXWT
                        TybmPpdgzgH4miFRMY2YlH1Hane6L5uSu9z2
                        s2SQaBldEz2sSt3XC3D3FMwBrwkaLTuUxs5g
                        BnavL8UsoEbbu00rtLRSLfJZGG7nqOuudLx
                        rLpZQBj+jwvse5eeAR8SniW3LI= )
100      NS     ns.example.net.
100      RRSIG  NS 5 2 100 20100507105010 (
                        20100407105010 17000 example.net.
                        sxiC2j6CHuhEeGa6v4PlaW+i3vFTRlU4S+zm
                        1yoeYEpamtcl+gT+8UBRQxGxpuopxQcSKRJJ
                        ugnQFvHvsoN3C1eiEQDSu40c9rwibppQMaus
                        2piNovQD3Hc/osf8f10o8xl/jik90YHYCADs
                        g1Qp30Cp1h8q1tj0MzxhLsjczVY= )
100      NSEC   *.example.net. NS SOA RRSIG NSEC DNSKEY
100      RRSIG  NSEC 5 2 100 20100507105010 (
                        20100407105010 17000 example.net.
                        LVyBGAnXmX8apqyZLTf2AqebfcNRPmknubwf
                        z1WL3B5Dc8rPcdB83S3bBnEPfccMQjTt6+Sn
                        vzcKpbI4Jkf+9LiKw/OVzKxjf+JYiMIwSN0j
                        tyqVGQlh4veOhkYImv0+1W+PCKTSCWqyutUc
                        lgOEjKmVukdZBEiJ5kG/bc/ROGO= )
100      DNSKEY 256 3 5 (
                        AQP4+0M1V055RS2Hqv+8w8V20Dh+SQmFzHQ
                        tZMuzLH3UxWEOGmG5GfjiJandJeAZTKLpERX
                        B6RfHTHGG81D3IO1azWN6DiVFEVzgrOotAdd
                        onfYF8gUT03ZnRcXlkJk41h12N0fq6rkODaF
                        nfMHCpI3WZ/MJqe+9hLJtis+oEsRw==
                        ) ; key id = 17000
100      DNSKEY 257 3 5 (
                        AQ0zgs4qea+ImJ10CworkabHqFvPKybVT7b
                        nDIkJ2HvXWslbwNWJ660x3N6ftpCTc9wWBMw
                        5+x0h7ilTwFPruma2gURwEywZaMG9ipILOxm
                        K04a5I+8R2QTH4BM0WaIKnv5jCHose/l9LL3
                        Y8MApsjP6gOWNM8b9aVTjBFnf0xEF7sOSBBB
                        E4G2/og5Fr+H8DYaotqgJ3nrzRfYA0gSXwwb
                        ) ; key id = 49656
100      RRSIG  DNSKEY 5 2 100 20100507105010 (

```

Figure 9: Example of a signed zone file

3 SECURING A DNS ZONE

```

20100407105010 17000 example.net.
nDu/CirOLW+MwDcnWwkgwZxJ7mnZKQghwKVk
i+OIiUqpixMjgQVjTxM4erCvCv13G4A4A1Ki
Ist+UF2iT+t3quGMWLMrtXAPEXqSSF53NQjA
Grj7XvP7MwE4541DLulvBLfWqlh6XJU2UPmI
jVrEi2uzlRz6fG+23IDOFHz6YuM= )
100 RRSIG DNSKEY 5 2 100 20100507105010 (
20100407105010 49656 example.net.
S5tubUZjvfvxD8qtfhv3yAT/LYarYQcsIgIP
Mccq91F114CqguAPLFWC2C4tLY6DUBuIwwFr
qpuIGt/nuR10DuIGHCsQ6Tc69CH/FYpSNWeB
whovrm5sZmaz6quJ5nvVcJ0x1XMG1kJ0166G
8q13RqRDbLMPqxLEZ4KQxSyPsvHCEHo0USKV
feUJuE3UZ+27pIwGNLj1dBRR4d4K1SjVug== )
*.example.net. 100 IN A 192.168.2.10
100 RRSIG A 5 2 100 20100507105010 (
20100407105010 17000 example.net.
hJuQxTvVGOD4cfPvLIspjqy1e7ONBT9CUsVX
CLafNnc3tUH10bdGQRRCMZS9eomN7dEaBU4Y
+1sRQq1p5J4tdxvkssyf3YrRnHhI84Y0sOmv
dnRLz92/Zt0CJ+qEmmNgNsBz11f+HwUe135W
l6H5aNWwydVY3Hb7HCXoqBdfEVI= )
100 NSEC a.example.net. A RRSIG NSEC
100 RRSIG NSEC 5 2 100 20100507105010 (
20100407105010 17000 example.net.
ceSC9NHxTCe9gi6Rqdg9kegHOC0aX+X8Mqbo
Z9R/zcyN8RJHbWTUAAI2pFa3NXPZVVTsk3jg
xJSvjgqaICus70Fo4M3s8qob2eC/p0e3vd8l
DI136rWXEPuzzppTA03yp1BFEbtYyIjbUiVk
Hu+38AwwpY2JX1CULINKLHFwbRo= )
a.example.net. 100 IN A 192.168.2.1
100 RRSIG A 5 3 100 20100507105010 (
20100407105010 17000 example.net.
dL9oxGCv3EWLXv94UZCctZqMBmzCjusfY+UW
Od2ampbUC2i3OpPiZB8KrWVWdkFqYtYrQKgT
TztQQChVYxbDey2JEhzfMYPBe2McGA7UPiz0
kboE1wTBpSwxpuZ/2yMqPdJ0W82xkfSbE7xh
AGjmU5qVA/Q69JZhLjF1mRcXMzo= )
100 NSEC b.a.example.net. A RRSIG NSEC
100 RRSIG NSEC 5 3 100 20100507105010 (
20100407105010 17000 example.net.
P3iT7v++wnKtd21jczq+yK3kB1r3GwZtAXwV
8vnnPIjR90FkmYfRyww00aoYKFjXXjZ79KNP
J3RsELb0kt13Ks3UGq1M6fmgxAVc68HcBcCC
xJRIHua4/80UE2d+YeASC6Fbw1wh1sr64oWo
q1ZKq30XrQIOji6FBpktyAFRzno= )
b.a.example.net. 100 IN A 192.168.2.11
100 RRSIG A 5 4 100 20100507105010 (
20100407105010 17000 example.net.
c+2Cw15aGoWbBmJsXv9HN4scEihcpGCKg8KM

```

Figure 9: continued

3 SECURING A DNS ZONE

```

RtU+SJvR/C1lKAVlKXCh5vHvRdrYadFy1xF6
jyA5HhHaY2wXk//ISiiAAEoBEyiYPN/Tzft4
o0mNPr1xHckfTlH03TUdG7r13VA= )
100      NSEC      b.example.net. A RRSIG NSEC
100      RRSIG     NSEC 5 4 100 20100507105010 (
20100407105010 17000 example.net.
AvnXboVvEIItxk/cAOK5YvznGhL1xv2IuYDXp
ScXArSfCk0gZoctpSf5P6RZDr5AqUK1uTVqb
v00RScn6dxRwZMEL+d1w09D3SCWGNz7x0yR8
5xEpV/1e4HCGsScjU4P6QKT9eC4GqGSEozrp
I7pqIFm2YFmWZq0qCbSZx+T2D4g= )
b.example.net.      100      IN A      192.168.2.2
100      RRSIG     A 5 3 100 20100507105010 (
20100407105010 17000 example.net.
CPCBrh5fNvofisFHKR9a0EQXtBCdYfyKc8pd
+pAUjn4Fi6phuytensezVE6p09AgvUgzT5NK9
fsEUp2tJEdKIXqCvnfEkzfdn5zXZ09VouAtt
7WpaDaRo6m6B9R9/ir/xpJ2RHlZ8APGgmIZ/
VhZQRAYf604cPcmjjuSUK67uFfc= )
100      NSEC      ns.example.net. A RRSIG NSEC
100      RRSIG     NSEC 5 3 100 20100507105010 (
20100407105010 17000 example.net.
Wql8otn6W/o6HXLdbVMhT5HsY1Z6cGpdZPVt
alatZyD5y8EU919jLD4UQaqN1YpFnieTxnNe
armGmD0oHBwf280uxCV9LioJbTreNva5SRzD
D/NnGveXhMCUwhEfI9hZfv07HHTFRFGVD3wL
vycZGrDaoI5qYJgR6J+0cEdbst8= )
ns.example.net.    100      IN A      192.168.2.203
100      RRSIG     A 5 3 100 20100507105010 (
20100407105010 17000 example.net.
kNDc37s1NhNf7+0+HPjPbDLVRP9DY04PURbG
fDOW+/jEt/DustUFT7XMfLOdAZou7SWHW2Hf
3CJCuTnbQhWWdvR/muLMGJTPt0XKryL/wcHu
6SeljdQQZ4BigHNtwIvdiE+fq8U66oq4/Vey
P6ZPgRsyLhAoSxeqW4InJv4SJNc= )
100      NSEC      example.net. A RRSIG NSEC
100      RRSIG     NSEC 5 3 100 20100507105010 (
20100407105010 17000 example.net.
gk+BpOLgggTMMvAVjdMfEU59Ee1jyD7CYo/c
zaq9Typc+ZoKVWeN8bMMgd4/LE1Nh2pk306W
/TqYHEudQJraceN3Q/GCtv9JCVNS9h88jPP7
Sdmyng9IQBCUi0oa0d10hP0xg0JmX0aGw2+N
lKJ9JeaT7yhw67vrufEJ9lxyWN8= )

```

Figure 9: continued

3.5 Caching forwarder configuration

Now the DNS servers publish signed data, we need to configure the 'clients' to validate it. The clients in this context are recursive name servers. Just configure your recursive name server with the public SEP key generated for the zone. This is described in Section 2.3 (above), also see Figure 3.

3.6 Zone Re-Signing

When the signatures in your zone are due to expire or if you have added new records to your zone, you will have to re-sign your zone. There are two ways to re-sign your zone data. You may choose either option depending on levels of automation, the zone size and the frequency with RRSIG RRs are generated.

- You can regenerate the signed zone from the unsigned zone file. The signer will need to sort the zone again, generate all the NSEC records and generate all RRSIG records.
If you generate your zone file from a back-end database this is probably the preferred method.
- You can add the new records to the already signed zone file and then run that zone file through the signer. The BIND signer will insert new records and associate NSECs in the already sorted zone file and will only sign new records and records for which the signatures are reaching the end of their validity period.

You should build tools to maintain your signed zones e.g.: using `cron`, `perl` and `make`. (also see Appendix F)

3.7 Troubleshooting Signed Zones

You can check the format of your `named.conf` using the `named-checkconf` program. The `named-checkzone` program can be used to check zone files. These programs use the same routines to parse the configuration and zone files as `named` but only check syntax.

One can use `dig` and a name server configured with a `trusted-key` to validate ones setup. If data cannot be cryptographically validated, the forwarder will return with a `SERVFAIL` status. You can test this by intentionally corrupting a resource record in the signed zone file. This is typical output of `dig` when querying for corrupted data ¹¹:

```

; <<>> DiG 9.7.0b2 <<>> @192.168.2.204 corrupt.example.net A +dnssec +multiline +retry=1
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: SERVFAIL, id: 35409
;; flags: qr rd ra; QUERY: 1, ANSWER: 0, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:

```

¹¹ We corrupted the data by modifying `rdata` in the signed zone-file to generate this example

3 SECURING A DNS ZONE

```

; EDNS: version: 0, flags: do; udp: 4096
;; QUESTION SECTION:
;corrupt.example.net.   IN A

;; Query time: 5 msec
;; SERVER: 192.168.2.204#53(192.168.2.204)
;; WHEN: Wed Apr 7 14:00:56 2010
;; MSG SIZE rcvd: 48

```

Note that a caching forwarder will not do cryptographic validation of the zones for which it is authoritative. Therefore, if your caching forwarder is a primary or secondary server for a particular zone you will always get an answer as it is assumed that data from disk is secure.

Further troubleshooting needs to be done on a server configured as a validating recursive name server. Below is an example of the log output of the validating nameserver when we queried for corrupted data.

```

validating @0x1008d5a00: corrupt.example.net A: starting
validating @0x1008d5a00: corrupt.example.net A: attempting positive response validation
validating @0x1008d6800: example.net DNSKEY: starting
validating @0x1008d6800: example.net DNSKEY: attempting positive response validation
validating @0x1008d6800: example.net DNSKEY: verify rdataset (keyid=49656): success
validating @0x1008d6800: example.net DNSKEY: signed by trusted key; marking as secure
validator @0x1008d6800: dns_validator_destroy
validating @0x1008d5a00: corrupt.example.net A: in fetch_callback_validator
validating @0x1008d5a00: corrupt.example.net A: keyset with trust 7
validating @0x1008d5a00: corrupt.example.net A: resuming validate
validating @0x1008d5a00: corrupt.example.net A: verify rdataset (keyid=17000): RRSIG failed to verify
validating @0x1008d5a00: corrupt.example.net A: failed to verify rdataset
validating @0x1008d5a00: corrupt.example.net A: verify failure: RRSIG failed to verify
validating @0x1008d5a00: corrupt.example.net A: no valid signature found
validator @0x1008d5a00: dns_validator_destroy

```

(This output was generated by using the `dnssec` category to a logging channel with severity `debug 3`; configured.)

Similarly one can use the logs produced by `unbound` for troubleshooting. When setting `verbosity: 3` then the log files are very verbose but also tell us precisely what went wrong. Like in the excerpt of a log file below.

```

info: resolving (init part 3): <example.net. DNSKEY
info: processQueryTargets: <example.net. DNSKEY
info: sending query: <example.net. DNSKEY
debug: sending to target: <example.net.>
debug: cache memory msg=298385 rrset=306734 infra=11248
debug: iterator[module 1] operate: extstate:module_wait_reply
info: iterator operate: query <example.net. DNSKEY
info: response for <example.net. DNSKEY
info: reply from <example.net.>
info: query response was
info: finishing processing for <example.net. DNSKEY
debug: validator[module 0] operate: extstate:module_wait_module
info: validator operate: query <example.net. DNSKEY
info: validated DNSKEY <example.net. DNSKEY
debug: validator[module 0] operate: extstate:module_wait_subquery
info: validator operate: query <corrupt.example.net. A
debug: verify: signature
info: validator: response has failed ANSWER rrset: <corrupt.example.net. A
info: Validate: message contains bad
debug: cache memory msg=298750 rrset=307821 infra=11248

```

3.8 Possible problems

SOA serial If you forget to increase the serial number before re-signing your zone, secondary servers may not pick up the new signatures. This may cause some of the authoritative servers to time out so some resolvers will be able to validate your signature while others will not.

'Zone-signing key' rollover If a zone administrator makes a distinction between zone and key-signing keys then the rollover of a zone-signing key will not involve any action from the administrators of the validators. If a key-signing key is to be changed care should be taken that all resolvers in the organisation have been supplied with a new trusted-key.

If the zone is only locally secured (i.e. is not part of a chain of trust) then the rollover of a key-signing key is relatively simple. Remember that to validate data there has to be at least one signature that can be validate with the trusted-keys in resolvers. For a limited time you use two key-signing keys to sign your zone: the old and new key. During that time you start reconfiguring the resolvers in your organisation with new trusted-keys. Once all resolvers have the new key configured in their trusted-key> statement, the zones should be signed with the new key only.

Also see Chapter 5, "Rolling keys".

Slave server problems Slave servers will need to run code that is DNSSEC enabled if one of the authoritative servers for a zone is not DNSSEC aware. Problems may arise for the DNS client that tries to fetch data from those DNSSEC oblivious servers.

The load on all your name servers will increase. Zone files, memory and bandwidth consumption will grow. Factors 2-5 are not uncommon; See "Hints and tips" for some numbers.

3.9 Automating the signing process

In June 2009 I am aware of two open source tools for the key maintenance and zone signing. Both are still under development.

Open-DNSSEC (at <http://www.opendnssec.org/>) was created as an open-source turn-key solution for DNSSEC. It secures zone data just before it is published in an authoritative name server.

OpenDNSSEC takes in unsigned zones, adds the signatures and other records for DNSSEC and passes it on to the authoritative name servers for that zone.

DNS is complicated, and so is digital signing; their combination in DNSSEC is of course complex as well. The idea of OpenDNSSEC is to handle such difficulties, to relieve the administrator of them after a one-time effort for setting it up.

4 DELEGATING OF SIGNING AUTHORITY; BECOMING GLOBALLY SECURE

The storage of keys is done through a PKCS #11 standard interface. To deploy OpenDNSSEC, an implementation of this interface is needed, for example a software library, an HSM or perhaps a simpler token.

DNSSHIM (at <http://registro.br/dnsshim/index-EN.html>) an open-source software that implements the Domain Name System (DNS) protocol for the Internet. Its main feature is to work as a Hidden Master nameserver providing information only to authoritative slave servers. Furthermore, DNSSHIM has the ability to manage, store and automatically resign zones using the security extension DNSSEC.

More about open-dnssec in a future version of the HOWTO.

4 Delegating of signing authority; becoming globally secure

This section is subject to change as the tools needed are currently being modified/developed.

4.1 Introduction

We have covered how to deploy DNSSEC in a single zone. We now want to build a chain of trust, so that once a client has securely obtained a public key high in the DNS hierarchy, it can follow the chain to validate data in your or your children's zones.

During the validation process a resolver will start from a configured trust anchor. It will use that to validate the keys set at the apex of the zone. Once the key-set has been validated the keys in that key-set can be used to validate any other data in a zone, such as A, AAAA and PTR resource records. In order to trust a child zone the validator will follow a pointer, stored in the DS resource record, that points to a key in the child's key-set that will be used to validate the keys in that zone. That DS RR is signed by the parents zone-signing-key and points to the child's key-signing key (figure 4).

4.2 Practical steps

Below we will describe how to set up a zone that is globally secure based on the parental signature over the DS record pointing to the child key-signing key.

In the example we use `net` as parent and `example.net` as child. At the start of the process we assume that the parent zone is already locally secure but has not secured the delegation yet. This means that the parent has no DS RR for `example.net`. and that resolvers following the chain of trust via `net`. will treat the `example.net`. zone as variably insecure. The `example.net`. zone is assumed not to be secure. Much of the procedure will be similar to Chapter 3, "Securing a DNS zone", but, since key-sets are used, some details are different.

Our goal is to publish a parent zone with a DS RR. The DS RR is related to the key signing key as generated by the child (the DS RR contains a cryptographic

4 DELEGATING OF SIGNING AUTHORITY; BECOMING GLOBALLY SECURE

hash over data in the DNSKEY RR). Therefore, the child needs to send some information to the parent. To ease the process BIND introduces key-sets and ds-sets.

A key-set is a small file, with the same syntax as a zone file, that contains one or more key-signing keys of the child. The ds-set is also a similar file but this file contains the DS RR that is to be included in the parent zone. These files are created when signing a zone as described in Section 3.4. For for the `example.net` zone they will be called `dsset-example.net` and `keyset-example.net`.

There are many imaginable ways to get the key-set to the parent. For instance

- the child sends a mail (cryptographically signed, to allow for integrity and authentication checks) with either the ds-set or the key-set.
- the parent can fetch the appropriate key from the child's DNS and create a key-set file itself. This is done by putting the key material in a file called `keyset-child-domainname`.
- a web based registration system interface is used to acquire the key-set.

In an operational environment it is extremely important that the authenticity and the integrity of the DNSKEY is established. The zone administrator¹² of the parent will need to validate that the key came from the zone administrator of the child zone. If possible, this should be confirmed by an out-of-DNS mechanism. The parent could use a customers database to validate if the key was actually sent by the zone administrator. If a wrong key is signed the child zone will be vulnerable for attacks; signing the wrong key breaks DNSSEC.

The parent stores the key-sets in the directory where zone files are stored, or, when you want to maintain some file system hygiene, in a directory that is to be specified with the `-d` flag of `dnssec-signzone`. The `signzone` tool will automatically generate (or include) the appropriate DS records if the `-g` flag is provided and a `keyset-child-domainname` (or the ds-set) is found. Although the key-set generated by the child contains signatures the RRSIG RRs do not need to be available in the `keyset-child-domain` file at the parent, the `sign` tool will not perform signature validation.

Below is an example of how to invoke the command:

```
dnssec-signzone -r /dev/random -g -d /registry/tld-zone/child-keys/ \
-o tld -f tld.signed db.tld
```

An alternative method of including DS RRs into ones zone is by concatenating to, or, including the ds-sets in the zone file.

```
cat /registry/tld-zone/child-dssets/dsset-* >> tld
dnssec-signzone -r /dev/random -o tld -f tld.signed db.tld
```

When the parent signs its own zone and uses the `-d` flag with `dnssec-signzone` its own ds- and key-set will end up in the specified directory, that can be quite confusing.

¹² The person who is responsible for publishing the zone data

5 ROLLING KEYS

4.3 Possible problems

Public Key Algorithm To be globally secure you need to use at least one key of an algorithm that is mandatory to implement. Mandatory to implement are RSA/SHA1 and DSA keys. We recommend the use of RSA/SHA1 keys only.

Parent indicating child security It is important that a DNSKEY is published in the DNS before the parent includes a signed DS RR for that key.

If the parent includes a DS RR while the child has not yet published the key then the child will go 'bad'; By not having a DS RR for the child, the parent indicates the child to be insecure.

As a parent you should always validate that the child publishes a signed DNSKEY before including a DS RR.

4.4 Registering with a DLV registry

If your parent is not yet secure you could consider gesturing with a *DLV registry* so that 3rd parties can still make use of the security your domain provides (for client side configuration see section 2.6)

BIND's `dnssec-signzone` contains an option to create a file that contains the data relevant to the *DLV registry*. Suppose that your favourite *DLV registry* is anchored under `dlv-registry.org` then signing with the `-l <dlv-registry-anchor>` option will create a `dlvset` file.

For example:

```
/usr/local/sbin/dnssec-signzone \
-l dlv-registry.org \
-o example.net \
-k Kexample.net.+005+49656 \
db.example.net \
Kexample.net.+005+17000.key
```

will create a file called `dlvset-example.net.` that contains the following information:

```
example.net.dlv-registry.org. IN DLV 49656 5 1 3850EFB913AEC66275BCA53221587D445702397E
```

This author suggests you use the ISC lookaside registration service. See <http://www.isc.org/index.pl?/ops/dlv/>

5 Rolling keys

A rollover is the process in which one key in a zone is replaced by another key. Since keys have a limited lifetime they need to be changed occasionally. Care needs to be taken that existing chains of trust are not broken during the rollover.

The rollover is defined by the moment keys generated with the "new" private key are first introduced into the zone. The key pair may have been generated well in advance and the public key may also have been made public well in advance.

If the rollover is planned we refer to it as scheduled rollover. If the rollover is the result of a (suspected) compromise or loss of private key it is called an unscheduled or emergency key rollover.

There are two types of scheduled key rollovers. The rollovers of key-signing keys and the rollovers of zone-signing keys.

Although the DNSSEC protocol does not make a distinction between zone- and key-signing keys we strongly advise you to make this distinction as it provides a clear separation between the keys that can be rolled without external interaction (the zone-signing keys) and the keys that need external interaction (the key-signing keys). You should use the `-f KSK` flag with `dnssec-keygen` when creating key-signing keys so that you can always make a distinction between key- and zone-signing keys by looking at the so-called flag field in the DNSKEY resource record. Its flag-field will be odd (257 mostly) when you deal with a key-signing, or SEP, key.

5.1 DNS traversal

Whenever data in a zone file is replaced by other data, it will need to propagate through the DNS before DNS clients actually see it. In a non-DNSSEC environment this may hardly ever be noticed, but when operating DNSSEC allowing data to traverse through the DNS is critical.

DNS data with its associated signatures and the public key with which this data is validated travel through the DNS independently. This also implies that the public keys and the signatures are independently cached and therefore expire from caches at different times. As a consequence it can happen that an RRSIG is validated with a DNSKEY from a cache and that the RRSIG and DNSKEY come from different versions of the zone; i.e. the public key relates to a key that is older than the signature. The reverse, where the signatures are older than the public keys that are used for validation can also happen.

As a zone administrator you have to be aware of this behaviour and take into account that your signatures will need to validate with any future or previous version of your key-set. [10] describes the details which differ for zone-signing and key-signing key rollovers. There are two approaches for this. The "pre-publish" and the "double signature" rollover.

First let us take a closer look at how data traverses through the DNS. See Figure 10 for reference.

At t_0 new data replaces data from a previous version of the zone file. The data is published on the authoritative master (or primary server). It will take some time (which we refer to as zone synchronisation time) before the new version of the zone is picked up by all authoritative servers. In the worst case scenario, a change to a slave server will not be able to reach the master server and the zone will expire. So the maximum value of the zone synchronisation time will be the value of the SOA expiration parameter.

5 ROLLING KEYS

Assume that at some time (τ_1) between publication of the new zone on the master server (τ_0) and the time the new zone is picked up by a slave server (τ_2) a query for the data is done by a recursive caching name server. That recursive server will return the *old* data to any of its clients for the time that is set by the TTL value on the *old* RRset. Only after τ_4 , will the recursive server go back and query for new data picking up the new records.

Note that the τ_4 does not only depend on $\tau_1 + \text{TTL}$ but is also upper bound by the signature expiration time of the signature on the old RRset.




Figure 10: DNS Data Propagation

5.2 "Pre-Publish" and "Double Signature" rollovers

During a pre-publish rollover the public key is introduced in the DNSKEY RRset well before RRSIGs are made with the private part of the key. The "new" public keys are then available in caches when the RRSIGs over the data show up on the authoritative name servers and caching name servers can use cached DNSKEY RRs to validate the new data.

During a double signature rollover the new key pair is introduced and signatures are generated with both the new and the old key. Both public keys are published in the DNS. After the period it takes for this data to propagate through the DNS, the old key is removed and only the new key is published and used for signing.

5.3 Tools

To properly maintain 'state' you will need an operational note book. For each of your zone there will be multiple KSKs and ZSKs and these keys all have a 'state'. The situation may become very confusing. Below we give an overview of the operations using an "operational note book" At the RIPE NCC a tool has been developed that replaces the "operational note book" and that links to

the signing operations. This tool is available through: http://www.ripe.net/disi/dnssec_maint_tool/.

Sparta has developed a daemon and a control tool, `rollerd` and `rollctl` respectively. `Rollerd` automates key rollovers. That is, it automates the steps necessary to change over from one Zone Signing Key (ZSK) to the next using the Pre-Publish Method of key rollover. It can also automate the less frequent Key Signing Key (KSK) change over using the Double Signature Method of key rollover. See RFC 4641[10] for a descriptions of these key rollover methods.

OpenDNSSEC (www.opendnssec.org) is an initiative to develop a comprehensive system for key-maintenance and signing tool that can maintain signed zone on the basis of pre-defined zone signing and key-maintenance policies. The first prototypes are, at the moment this document is written, expected mid 2009.

5.4 ZSK rollover

During a Zone-signing key (ZSK) rollover we use a "pre-publish" scheme.

5.4.1 ZSK preparation (production phase)

Use the trivial `example.com` zone (Figure 11) as an example. The zone is stored in `db.example.com`.

Assuming that we first start to publish `example.com` we generate two ZSK keys and one KSK key.

```
dnssec-keygen -a RSASHA1 -b 1024 -n ZONE example.com
Kexample.com.+005+63935
dnssec-keygen -a RSASHA1 -b 1024 -n ZONE example.com
Kexample.com.+005+64700
dnssec-keygen -f KSK -a RSASHA1 -b 1024 -n ZONE example.com
Kexample.com.+005+54915
```

In the operational note book we note that key 63935 will be used as the *active* and key 64700 as the *passive* ZSK. Both keys will be available through the key-set but only the active key is used for signing.

After we generated the keys we include them in the zone by adding the following include statements to `db.example.com`

```
;; ZSKs
$include Kexample.com.+005+63935.key
$include Kexample.com.+005+64700.key

;; KSKs
$include Kexample.com.+005+54915.key
```

5 ROLLING KEYS

```

; example.net zone
;
$TTL 100
$ORIGIN example.net.
@           100      IN      SOA     ns.example.net. (
                                olaf.nlnetlabs.nl.
                                2002050501
                                100
                                200
                                604800
                                100
                                )

                                NS      ns.example.net.
ns.example.net.  A      192.168.2.203

a              A      192.168.2.1
b              A      192.168.2.2

*              A      192.168.2.10
b.a           A      192.168.2.11

```

Figure 11: Trivial example.com

Then sign the zone. Since we do not want to use `dnssec-signzone`'s default behaviour, (which is to use all available keys for signing), we have to fully specify which keys to use on the command line. Since you will have to do this frequently the operational note book will come in handy.

```

dnssec-signzone -k Kexample.com.+005+54915.key -o example.com \
                db.example.com Kexample.com.+005+63935

```

Note that we supplied the KSK as an argument with the `-k` switch and only the active key ZSK as a signing key.

5.4.2 ZSK rollover (phase1)

Note down the signature expiration of the DNSKEY RR as it is now available in the DNS. This value can be used as an upper limit for the duration of this phase. It is the value of `t4` in Figure 10. In the DNSKEY RR set below the signature expiration time is August 21, 2004 around 11:35 UTC. If *all*¹³ the

¹³ including the last parameter in the SOA that is used for negative caching

5 ROLLING KEYS

TTLs in your zone are not higher than for example 600, you would not have to wait that long. You would have to wait until you see the new zone published in all authoritative servers and an additional 10 minutes. The signature expiration is an upper bound.

```

100      DNSKEY 256 3 5 (
          AQPQyhg865V4zkFZN+FiCLAZPWwAf5I43pW
          Ucu0iejT92AVu0eH0kbH5YiHV97r+QjAdZ7K
          W7W+bvbqKBR5P4QMVNm8zCs5Trb90cOY0+bb
          LYZG3aG69wUfF1pjvmFV5zUSRHCLMEzXb5NS
          XdazgdhhuM07L2e2EfJGp5qijtRwpQ==
          ) ; key id = 63935
100      DNSKEY 256 3 5 (
          AQPWrMsW0sGSTD7iE9ou+s7886WeSLIq/l/J
          CgqwAn7j1ECGAAN6cSHV5jWvovcWFthapWdG
          DpC1uL48AcWtVWkRABGjU8Q16CAy0EcZ+24V
          4cul+VluBt1YjuNfUlye+k5V+lmkjXBQ3Qdf
          E8/owjsdx9mTkeQC4qiFjUxWXT14DQ==
          ) ; key id = 64700
100      DNSKEY 257 3 5 (
          APhZQ29Xg6ONLgR+qdJENZpk1U+WQF0abmp
          Ni3Ce0YrR+bd01Q/2WDI6BbWCLdIb9Yf1Raj
          hmyb+AmzmjNzhw8VjcY9Sr2zIcG50ctuZ80g
          t7fcGrCbEM9fIDIKdDRlf+SY80nGEMi6sI4m
          bZ4zoh+nWfNrTxQR5hHv074uSAvZyQ==
          ) ; key id = 54915
100      RRSIG  DNSKEY 5 2 100 20040821114554 (
          20040722114554 54915 example.com.
          gcnf3rf+D6izv9A//16u+Jx/LDVinLtcpkWR
          yxDV5goS2SnoLfyEryqbSAyKbh4redyQCjSW
          /HZXFBoPYrAy8fqaY1AfjVP+q9zJPvysUOp+
          2T6mm8/9pcZoGXw1wPjPUAz+AF0oJnoaWo7t
          764xvZc47kAI1pTORTizV2BofcU= )
100      RRSIG  DNSKEY 5 2 100 20040821114554 (
          20040722114554 63935 example.com.
          T7gRcEZkxEl5iGJdCzSu470g9ydM05Uggvcz
          A9jETiTuRbttyYua7qDZ0jNrzt4GVZ6s/UBw
          tbGCqyMU/sVvaulP4h8oerX44bw5eP/mluLY
          T9rwm2jBI1rZSPDdGdp8lJ2vvrXASYSF2FXg
  
```

At the moment of the rollover you have to make your current passive key (64700) active and your current active key (63935) passive. Also make a note that this key is to be removed from the keyset in the next phase of the rollover.

Increase the SOA serial number and re-sign the zone using the new active key.

```

dnssec-signzone -k Kexample.com.+005+54915.key -o \
example.com db.example.com \
Kexample.com.+005+64700
  
```

5 ROLLING KEYS

Publish this zone in the DNS and make sure it remains published long enough to propagate through the DNS.

5.4.3 ZSK Cleanup (phase2)

After the data has propagated though the DNS, you have to replace the passive ZSK (63935) by a new passive key.

Start with generating the passive ZSK.

```
dnssec-keygen -a RSASHA1 -b 1024 -n ZONE example.com
Kexample.com.+005+01844
```

Add the new passive key (01844) into the zone file. Remove the old passive key (63935) from the zone file.

```
;; ZSKs
$include Kexample.com.+005+64700.key
$include Kexample.com.+005+01844.key

;; KSKs
$include Kexample.com.+005+54915.key
```

Increase the SOA serial and re-sign using the same active key as in phase1.

```
dnssec-signzone -k Kexample.com.+005+54915.key -o example.com db.example.com \
Kexample.com.+005+64700
```

After publishing your zone you are back in the "production" phase. You should not proceed in a new rollover until the current DNSKEY RRset has had a chance to propagate through the system

You can now delete the 63935 key. We suggest you move the key pair to a separate directory or make a backup.

5.4.4 Modifying zone data during a rollover

You can at any time modify zone data other than the data in the key-set. As long as you use the suitable active ZSK for signing.

5.5 Key-signing key rollovers

During a key-signing key (KSK) rollover we use a "double signature" scheme.

5 ROLLING KEYS

5.5.1 KSK preparation (production phase)

We again use the trivial `example.com` zone (Figure 11) as an example. The zone is stored in `db.example.com`. It contains a active and a passive ZSK (63935 and 64700 respectively) and a KSK (54915). The include statements are the same as Section 5.4.1:

```
;; ZSKs
$include Kexample.com.+005+63935.key
$include Kexample.com.+005+64700.key

;; KSKs
$include Kexample.com.+005+54915.key
```

and the command to sign the zone is also the same.

```
dnssec-signzone -k Kexample.com.+005+54915.key -o example.com \
db.example.com Kexample.com.+005+63935
```

5.5.2 ZSK rollover (phase 1)

We start the rollover by generating a new KSK

```
dnssec-keygen -f KSK -a RSASHA1 -b 1024 -n ZONE example.com
Kexample.com.+005+06456
```

Insert the new KSK into the zone file:

```
;; ZSKs
$include Kexample.com.+005+63935.key
$include Kexample.com.+005+64700.key

;; KSKs
$include Kexample.com.+005+54915.key
$include Kexample.com.+005+06456.key
```

Sign the zone with both KSKs and the active ZSK.

```
dnssec-signzone -k Kexample.com.+005+54915.key \
-k Kexample.com.+005+06456.key -o example.com \
db.example.com \ Kexample.com.+005+64700
```


You have now introduced the new KSK key.

Since you are rolling a KSK you will have to upload this key to your parent or have to configure it into your trust anchors (see Section 2.3). The public key you will have to upload and configure is the new one with key-id 06456.

If your parent has a DS RR pointing to your old key, it will take time before that DS RR has expired from caches. The upper limit on the `t4` parameter is the signature expiration time in the DS RR pointing to the old KSK (54915).

`dnssec-signzone` provides two files that will help you during this process. `dsset-example.com.` and `keyset-example.com..` The "dsset" file contains the DS RRs that relate to the KSKs in your zone file and the "keyset" file contains the KSKs published in your zone file. Remember that since you are replacing keys only one of these entries (06456) will need to be sent/appear at your parent.

5.5.3 KSK cleanup (phase 2)

Once you are satisfied that all trust anchors are updated and the parental DS RR has travelled through the DNS you can remove the old key from the set of includes:

```
;; ZSKs
$include Kexample.com.+005+63935.key
$include Kexample.com.+005+64700.key

;; KSKs
$include Kexample.com.+005+06456.key
```

Sign the zone with the new KSK and the active ZSK.

```
dnssec-signzone -k Kexample.com.+005+06456.key \
-o example.com db.example.com \
Kexample.com.+005+64700
```

From this moment on you are in the production phase again.

5.5.4 Multiple KSKs

This algorithm also applies if you have multiple KSKs. The steps are:

- generate and include the new KSK in the zone;
- sign the zone with all KSKs; wait for propagation;
- remove one of the KSKs and sign with all left over KSKs.

Part III

Troubleshooting tools

This part describes a few practical trouble shooting tools that may help to understand what goes wrong, if something goes wrong.

6 Using drill for troubleshooting

Both `dig`, from the BIND distribution, as `drill` can be used for trouble shooting DNSSEC set ups. See section 7 for more information about `dig`, we will first discuss `drill`.

`drill` is part of the `ldns` library available from <http://www.nlnetlabs.nl/ldns/>. Installation instructions are also available on that page. (It is as simple as: `./configure ; make ; make install`).

Drill's `-T` and `-S` switches are particularly helpful when troubleshooting DNSSEC setups. Using `drill` with the `-T` follows the chain of trust from the root to the leaves and indicates the security status (see figure 13). With the `-S` flag `drill` will chase the signatures from the leave-node back to the root, looking for the relevant records (see figure 14). When using the `-T` or `-S` flag you will have to specify a file that contains a trust-anchor in RR format *i.e.* just as in the files generated by `dnssec-keygen` (see page 20).

The `ldns` library does not only come with `drill`. You will find a few useful utilities in its `examples` directory. Among others there are:

```
ldns-key2ds  Creates a DS record from a DNSKEY record
ldns-keyfetcher  Fetches DNSSEC public keys for zones
ldns-keygen  Generate private/pubkey key pair for DNSSEC.
ldns-signzone  Signs a zone file according to DNSSECbis.
ldns-walk  'Walks' a DNSSEC zone
```

7 Using dig for troubleshooting

`dig` has a few switches that come in useful when troubleshooting DNSSEC setups.

```
+multiline Structures the output of dig so that it is easily readable. As a
  bonus the keyid will be printed as a comment behind DNSKEY RRs.
+cd Sets the "checking disabled" bit on the query. You would typically use this
  when your validating recursive name server reports a SERVFAIL and you
  need to establish if the is due to DNSSEC marking this data as "bad".
+dnssec Forces the server being queried to include the DNSSEC related data.
  Use in combination with the +cd to establish if data from a zone is signed
  at all or if you want to determine if the validity intervals on the signatures
  are correct.
+trace Traces delegation chain. This option may be helpful if you trying to
  figure out where the delegation points are.
```

`+sigchase` Traces the the signature chain. You will also need to have a `./trusted-keys.keys` or `/etc/trusted-keys.keys` available that contains trusted key entries.

8 DNSSEC tools

A number of open-source DNSSEC tools can be found at www.dnssec-tools.org. The site contains a number of generic maintenance tools for zone and key administration, some DNSSEC applications (mozilla and spam assassin plug-ins), and a few troubleshooting tools.

One of these tools has is `dnspktflow` tool that visualizes DNS 'streams' in a graph. This tool may help, in combination with the tools above, to create a bit of insight of what is going on. For instance, in figure 15 the packets that are send and received during the trace in figure 14 are shown.

bin/drill version 1.5.1 (ldns version 1.5.1)
Written by NLnet Labs.

Copyright (c) 2004-2008 NLnet Labs.
Licensed under the revised BSD license.
There is NO warranty; not even for MERCHANTABILITY or FITNESS
FOR A PARTICULAR PURPOSE.

Usage: bin/drill name [@server] [type] [class]
 <name> can be a domain name or an IP address (-x lookups)
 <type> defaults to A
 <class> defaults to IN

arguments may be placed in random order

Options:

-D enable DNSSEC (DO bit)
 -T trace from the root down to <name>
 -S chase signature(s) from <name> to a know key [*]
 -V <number> verbosity (0-5)
 -Q quiet mode (overrules -V)

-f file read packet from file and send it
 -i file read packet from file and print it
 -w file write answer packet to file
 -q file write query packet to file
 -h show this help
 -v show version

Query options:

-4 stay on ip4
 -6 stay on ip6
 -a fallback to EDNS0 and TCP if the answer is truncated
 -b <bufsize> use <bufsize> as the buffer size (defaults to 512 b)
 -c <file> use file for recursive nameserver configuration (/etc/resolv.conf)
 -k <file> specify a file that contains a trusted DNSSEC key [**]
 used to verify any signatures in the current answer
 -o <mnemonic> set flags to: [QR|qr] [AA|aa] [TC|tc] [RD|rd] [CD|cd] [RA|ra] [AD|ad]
 lowercase: unset bit, uppercase: set bit
 -p <port> use <port> as remote port number
 -s show the DS RR for each key in a packet
 -u send the query with udp (the default)
 -x do a reverse lookup

when doing a secure trace:

-r <file> use file as root servers hint file
 -t send the query with tcp (connected)
 -d <domain> use domain as the start point for the trace
 -y <name:key[:algo]> specify named base64 tsig key, and optional an
 algorithm (defaults to hmac-md5.sig-alg.reg.int)
 -z don't randomize the nameservers before use

[*] = enables/implies DNSSEC

[**] = can be given more than once

ldns-team@nlnetlabs.nl | <http://www.nlnetlabs.nl/ldns/>

Figure 12: drill arguments

```
;; Domain: .
[T] . 100 IN DNSKEY 256 3 5 ;{id = 63380 (zsk), size = 1024b}
. 100 IN DNSKEY 257 3 5 ;{id = 63276 (ksk), size = 1280b}
Checking if signing key is trusted:
New key: . 100 IN DNSKEY 256 3 5 AQPQyahT00aR/Pi6p+EG+JldP9wQJurw7iB8Mz+De45akis0/VeprAd0s771A5Nxxz+42XP+VcGMqSIhUV1x1WIXD14ZHBG8mKocfNtOmn0Jh
Trusted key: . 3600 IN DNSKEY 257 3 5 AQ0v6tbkmW+1IrU3q2SrG+36bmJfFs1+u7XaMgytcTftdBoYBquXqYVmYt4EtI9soyZBG6jS5CtkrzWnk721/ZFgYatkrSq80WQ
Trusted key: . 100 IN DNSKEY 256 3 5 AQPQyahT00aR/Pi6p+EG+JldP9wQJurw7iB8Mz+De45akis0/VeprAd0s771A5Nxxz+42XP+VcGMqSIhUV1x1WIXD14ZHBG8mKocf
Key is now trusted!
Trusted key: . 100 IN DNSKEY 257 3 5 AQ0v6tbkmW+1IrU3q2SrG+36bmJfFs1+u7XaMgytcTftdBoYBquXqYVmYt4EtI9soyZBG6jS5CtkrzWnk721/ZFgYatkrSq80WQ
[T] net. 100 IN DS 13467 5 2 ec9b094786b82c46aa3054c7352b59904b697119d515b4ea536ec3dd9a10ed81
net. 100 IN DS 13467 5 1 de01426e08ddb9186502ccc1081390cd7da0e178
;; Domain: net.
[T] net. 100 IN DNSKEY 256 3 5 ;{id = 62972 (zsk), size = 1024b}
net. 100 IN DNSKEY 257 3 5 ;{id = 13467 (ksk), size = 1280b}
Checking if signing key is trusted:
New key: net. 100 IN DNSKEY 256 3 5 AQPVP6Jea1Lo/dU19UCmumJqR36Mx6ecjXxdwTByhyn//3S/NgDTgYbTzvFq+dmjch1ccuCIWS9PPB9tmQ5grQ2YMwee2VLTrtRnOPX2gm8
Trusted key: . 3600 IN DNSKEY 257 3 5 AQ0v6tbkmW+1IrU3q2SrG+36bmJfFs1+u7XaMgytcTftdBoYBquXqYVmYt4EtI9soyZBG6jS5CtkrzWnk721/ZFgYatkrSq80WQ
Trusted key: . 100 IN DNSKEY 256 3 5 AQPQyahT00aR/Pi6p+EG+JldP9wQJurw7iB8Mz+De45akis0/VeprAd0s771A5Nxxz+42XP+VcGMqSIhUV1x1WIXD14ZHBG8mKocf
Trusted key: . 100 IN DNSKEY 257 3 5 AQ0v6tbkmW+1IrU3q2SrG+36bmJfFs1+u7XaMgytcTftdBoYBquXqYVmYt4EtI9soyZBG6jS5CtkrzWnk721/ZFgYatkrSq80WQ
Trusted key: net. 100 IN DNSKEY 256 3 5 AQPVP6Jea1Lo/dU19UCmumJqR36Mx6ecjXxdwTByhyn//3S/NgDTgYbTzvFq+dmjch1ccuCIWS9PPB9tmQ5grQ2YMwee
Key is now trusted!
Trusted key: net. 100 IN DNSKEY 257 3 5 AQ0sAH77fsZc53KoA9QfrXABEoXa5ZIBmUmJuXvnp36/LEgfn4ItTn0GKMHesgIV/qopgs+4dZxSIRAjGF92y3Mp3+79
[T] example.net. 100 IN DS 49656 5 1 3850efb913aec66275bca53221587d445702397e
example.net. 100 IN DS 49656 5 2 9e06b299abe811d699e077fff990ff5a1b496c914deb22697ba22a1da31f0a6e
;; Domain: example.net.
[T] example.net. 100 IN DNSKEY 256 3 5 ;{id = 17000 (zsk), size = 1024b}
example.net. 100 IN DNSKEY 257 3 5 ;{id = 49656 (ksk), size = 1280b}
[T] example.net. 100 IN SOA ns.example.net. olaf.nlnetlabs.nl. 2002050501 100 200 604800 100
;;[S] self sig OK; [B] bogus; [T] trusted
```

Figure 13: Output of `drill -T -k <root.ksk> example.net SOA`

```
;; Chasing: example.net. SOA

DNSSEC Trust tree:
example.net. (SOA)
|---example.net. (DNSKEY keytag: 17000)
|---example.net. (DNSKEY keytag: 49656)
|---example.net. (DS keytag: 49656)
|---net. (DNSKEY keytag: 62972)
|---net. (DNSKEY keytag: 13467)
|---net. (DS keytag: 13467)
|---. (DNSKEY keytag: 63380)
|---. (DNSKEY keytag: 63276)
;; Chase successful
```

Figure 14: Output of `drill -S -k <root.ksk> example.net SOA`

Figure 15: Example of `dnspktflow` output

Part IV

Securing communication between Servers

This part considers transaction security issues, which are typically not considered to be part of DNSSEC and can be deployed completely independently. This part focuses on securing the transactions between authoritative servers, but the same techniques can be used to secure dynamic updates.

9 Securing zone transfers

9.1 Introduction

The communication between hosts can be secured (authenticated and encrypted) using a scheme based on symmetric cryptography. By sharing a key the administrators of two servers can be sure that DNS data is only being exchanged between those two boxes and that the data has not been tampered with in transit.

The most known mechanism used to enable this is referred to as TSIG[17] and is based on a shared secret. The shared secret is used to sign the content of each DNS packet. The signature can be used for both authentication and for integrity checking of the data. In order to prevent a malicious third party retransmitting captured data (replay attack) a time stamp is included in the data. The TSIG mechanism can also be used to prevent unauthorised zone transfers; only owners of the secret key are able to do a zone transfer¹⁴. We will describe how primary server `ns.foo.example` and secondary server `ns.example.com` need to be configured to enable TSIG for zone transfers.

To configure TSIG perform the following steps:

- Synchronise clocks.
- Create and distribute a shared secret, the TSIG key.
- At the primary server, create an access list specifying which keys are allowed to transfer.
- At the secondary server, specify which keys to use when contacting which primary servers

The first item is a prerequisite for DNSSEC. If you do DNSSEC you should be in sync with the rest of the world: Use NTP. Time zones can be confusing. Use `date -u` to validate if your machine has the proper UTC time.

TSIG configuration is a task for system administrators.

¹⁴ The data in the DNS is public data; disabling zone transfers does not guarantee your DNS data will become 'invisible'

9.2 Generating a TSIG key

There are various ways to create a shared secret.

9.2.1 Generating a TSIG secret with `dnssec-keygen`

`dnssec-keygen` is the tool used to generate a base64 encoded random number that will be used as the `secret`. The arguments that we have to provide `dnssec-keygen` to generate a TSIG key are (also see Figure 6):

```
dnssec-keygen -a hmac-md5 -b 256 -n HOST ns.foo.example.ns.example.com.
```

The command produces two files¹⁵. The name of the files contain relevant information:

```
Kdomain_name+algorithm_id+key_id.extension
```

The `domain_name` is the name specified as the name of the key. The name specified here does not need to be a name that you can query in the DNS but should be a name can be encoded as a domain name. The convention is to concatenate the DNS names of the two servers.

One can use `dnssec-keygen` to generate a truly random secret or use a passphrase - we describe both methods in Section 9.2. In this particular case `ns.foo.example.` and `ns.example.com`. The `algorithm_id` identifies the algorithm used: 5 for HMAC-MD5 (1 and 3 are for RSA and DSA respectively, see Appendix A.1. in [5]). The `key_id` is an identifier for the key material, it is not of relevance for symmetric keys. The `extension` is either `key` or `private`, the first is the public key and the second is the private key.

The format of these files differs a bit but they contain exactly the same information; a base64 encoded random number that you are going to use as a shared secret. Do not be misled by the extensions `private` and `key`, both files should be kept secure. Since the secret material is copied to the configuration files and these files are not used in production you should actually consider deleting them.

Note that the `-n HOST` and the `name` are not used for the generation of the base64 encoded random number. It is a convention to use the unique domain name label used to identify the key as the `name`.

```
# dnssec-keygen -r /dev/random -a HMAC-MD5 -b 128 -n HOST \
  ns.foo.example.ns.example.com.
Kns.foo.example.ns.example.com.+157+12274

# cat Kns.foo.example.ns.example.com.+157+12274.key
ns.foo.example.ns.example.com. IN DNSKEY 512 3 157 gQQqMJA/LGHwJa8vtD7u6w==

# cat Kns.foo.example.ns.example.com.+157+12274.private
Private-key-format: v1.2
Algorithm: 157 (HMAC_MD5)
```

¹⁵ It is a feature of the `dnssec-keygen` program to always creates two files, even when it is generating symmetric keys.

9 SECURING ZONE TRANSFERS

Key: gQ0qMJA/LGHwJa8vtD7u6w==

The base64 encoded random number is the thing you need to extract from either of these files (*i.e.* gQ0qMJA/LGHwJa8vtD7u6w==) it specifies the **secret** in the **key** statement:

```
key ns.foo.example.ns.example.com.{
    algorithm hmac-md5;
    secret "gQ0qMJA/LGHwJa8vtD7u6w==";
};
```

This key definition should be included in both primary and secondary name server configuration files and should be exactly the same on both sides (in this example `ns.foo.example.ns.example.com.` is used on both name servers). It is recommended to generate a secret for each different party, with which you are involved and you will need to maintain as many secrets as zones for which you have secondaries.

9.2.2 Other ways to generate secrets

The `dnssec-keygen` command provides you with a truly random bit sequence. It might be difficult to communicate the secret to your colleague running a secondary server on the other side of the world. In those cases you may want to choose to fall back to a pass-phrase that can be communicated over the telephone.

You can use any base64 encoder to convert the pass-phrase to a valid string in the key-definition.

```
# echo "Crypto Rules" | mmencode
Q3J5cHRvIFJ1bGVzCg==
```

If `mmencode` is not available maybe this perl script can assist you.

```
#!/usr/bin/perl
use MIME::Base64;
print encode_base64("@ARGV") ;
```

Actually any string that can be base64 decoded will do, for example `ThisIsAValidBase64String` can also be used as **secret**.

9.3 Configuring TSIG keys

To secure a zone transfer, the primary server and the secondary server administrators have to configure a TSIG key in `named.conf`. The TSIG key consists of a secret and a hashing algorithm and are identified by domain names. We

9 SECURING ZONE TRANSFERS

recommend that you maintain the list of secret keys in a separate file which is readable by root only and included in the `named.conf` file (e.g. by `include /var/named/shared.keys`).

The key statement looks like:

```
key ns.foo.example.ns.example.com. {
    algorithm hmac-md5
    secret "gQ0qMJA/LGHwJa8vtD7u6w==";
};
```

This statement needs to be exactly the same for the two parties involved.

9.4 Primary servers configuration of TSIG

Both the primary and secondary server should have shared secret configured by using the key statement in a file included in `named.conf` (see above).

The primary server can now use the key in what BIND calls an [3] `address_match_list`. These lists appear in the `allow-notify`, `allow-query`, `allow-transfer` and `allow-recursion` statements which controls access to the server. (Also see section 6.1.1 and 6.2.14.3 of the on-line BIND documentation).

Relevant at this point is the `allow-transfer` in the `zone` statement. Using the key generated above, the primary server for `foo.example` would have the following statement in `named.conf`:

```
zone "foo.example" {
    type master;
    file db.foo.example.signed;
    \\ allow transfer only from secondary server that has
    \\ key ns.foo.example.ns.example.com.
    allow-transfer { key ns.foo.example.ns.example.com. ; };
    notify yes;
};
```

9.5 Secondary servers configuration of TSIG

Both the primary and secondary server should have shared secret configured by using the key statement in `named.conf` (see above).

The `server` definition in `named.conf` is used to instruct the name server to use a specific key when contacting another name server.

```
\\ secondary for foo.example.
\\ primary server ns.foo.example is on 10.1.1.2
server 10.1.1.2 {
```

```

    keys { ns.foo.example.ns.example.com. };
};

```

9.6 Securing the NOTIFY message too

The setup above will provide signatures for the zone transfer from the primary to the secondary. Since the session is initiated by the secondary server¹⁶ it is the secondary server that sets up the secure link. Therefore, the secondary server has the `server` definition in its `named.conf`. Alternatively you can secure the traffic to the secondary server that was initiated by the primary server. Think of the NOTIFY messages sent to the secondary server when the zone content changed. That traffic will be TSIG signed as soon as you add a `server` with the secondary's IP address in the primary's `named.conf`. You can use the same key as for the zone transfer.

Once the primary server has configured its server to use TSIG to sign the NOTIFY messages the secondary server can use the key in the `allow-notify` access control list.

9.7 Troubleshooting TSIG configuration

You can check the format of your `named.conf` using the `named-checkconf` program. This program reads the configuration file using the same routines as `named` itself.

To troubleshoot your configuration, you have the log file and `dig` at your disposal.

Before adding the `allow-transfer {key ns.foo.example.ns.example.com.};` you should be able to transfer the domain from any machine. `dig @ns.foo.example foo.example AXFR` should be successful. After key configuration the same command should fail and give you output similar to:

```

; <<>> DiG 9.2.0rc1 <<>> @ns.foo.example foo.example AXFR
; global options printcmd
; Transfer failed.

```

You can test if the key is configured correctly in two ways.

Method 1 Ask the zone administrator to increase the SOA serial and to have the zone reloaded on the primary server. The secondary server should pick up the changes.

The log file of the secondary server will have entries similar to:

```

... general: info: zone foo.example/IN: transferred serial 2001082801
... xfer-in: info: transfer of 'foo.example/IN' from 10.1.1.2\#53: end of transfer

```

¹⁶ Remember the zone transfer is over TCP

9 SECURING ZONE TRANSFERS

Method 2 Use `dig` to test the key by using the `-k` flag.

```
dig @ns.foo.example -k Kns.foo.example.ns.example.com.+157+12274.key \
    foo.example AXFR
```

Alternatively you can use the `-y` switch and specify the key-name and the secret ¹⁷ with the `-y` switch.

```
dig @ns.foo.example \
    -y ns.foo.example.ns.example.com.:gQ0qMJA/LGHwJa8vtD7u6w== \
    foo.example AXFR
```

If the key did not match the log file of the primary server against which you tried this, will have entries similar to the following.

```
... security: error: client 10.1.1.6#1379: zone transfer 'foo.example.com/IN' denied
```

9.8 Possible problems

9.8.1 Timing problems

Machines that are involved in a TSIG signed transaction need to have their clocks synchronised to within a few ¹⁸ minutes. Use 'NTP' to synchronise the machines and make sure the time zones are correctly configured. A wrong time-zone configuration can lead to hard to spot problems; use `date -u` to check what your machine thinks is the 'UTC' time.

9.8.2 Multiple server directives

TSIG is a mechanism to protect communication on a per machine basis. Having multiple server directives for the same server or multiple keys in one server directive will lead to unexpected results.

¹⁷ Take care when using secrets on the command line of multi-user systems: on most Unix systems command line arguments are visible on the output of `ps` or via the `/proc` file system

¹⁸ BIND has 5 minutes hard coded

Part V

Appendices

A Deploying DNSSEC: the Milestones

In order to deploy DNSSEC in existing infrastructure a number of steps will need to be made. Try taking the following milestones, and their interdependencies into account when developing a project plan.

A.1 Private Key policies and procedures

prerequisites none

description The policies and procedures concerning the private key handling are a prerequisite for designing your architecture. There are a number of issues that need thought, documentation and management signoff.

- The use of key signing and zone signing keys (KSK and ZSKs)
- Generation mechanism for keys
 - Who generates the keys
 - Where are the keys generated e.g. is there special hardware involved
 - How are the keys stored after generation
 - Technical consideration: what is the source of randomness.
 - What key lengths are used
- Who has access to the keys, is this different for KSKs and ZSKs
- Where and how are the keys stored in the production environment
- How long will the keys be used

Some of these decisions may influence the complexity of your operation. In practice the considerations will need to be based on a risk analysis *i.e.* what will happen if the key would be compromised.

A.2 Public Key policies and procedures

prerequisites

description The policies concerning the public key could be published publicly so that all users that use your key as a trust-anchor know what they are up to.

These policies are closely related to your private key policy and should document

A DEPLOYING DNSSEC: THE MILESTONES

- If and how a distinction can be made between zone signing keys (ZSKs) and key signing keys (KSK).
- Signing frequencies of KSK and ZSKs.
- Rollover frequencies of KSK and ZSKs.
- Rollover techniques used (Further considerations see [10] and [15]).
- How the keys are published and which off-band mechanism(s) for validation are offered.
- How changes to the policy are announced in a way such these announcements can be validated.
- How will emergencies be communicate and how can people validate such messages.

In appendix E we have reproduced an example public key policy statement inspired on the one in use by the RIPE NCC.

A.3 Signing infrastructure

prerequisites A.2

description The signing infrastructure is all the infrastructure needed to turn unsigned records into signed records. The architecture of this setup depends on how your provisioning system generates data for the DNS and how the data is put into the DNS. If you use proprietary systems without zone-files you may not be able to depend on standard tools.

The design also depends on how access to the private keys (ZSK and KSKs) is arranged.

A.4 Server infrastructure

prerequisites External, software needs to support same features.

description All the servers, primary and secondary zones, need to be able to support the DNSSEC protocol. If one of the servers does not support DNSSEC there will be failures. If you choose to deploy NSEC3 all the servers will need to support NSEC3

There could be concerns with respect to growth of the size of the zonefiles as kept in memory or to increased traffic or CPU loads. In practice the memory concerns may be the most problematic. But the increase in memory can be predicted, see C, [12], or [13].

For well provisioned server infrastructure the increase in CPU or network traffic should not be a problem.

A DEPLOYING DNSSEC: THE MILESTONES

A.5 Monitoring tools

prerequisites A.2

description Tools are essential to make sure that what is served is correct, that signatures will not expire and that your service levels are maintained.

When adding DNSSEC you will have to adapt or create monitoring tools, also see ??.

A.6 Serving a Secured zone

prerequisites A.4 A.5 A.2 A.2

description At this moment you are ready to load the first signed zone into your nameservers. It is also the moment you can publish your trust anchors.

A.7 Requesting secure delegation

prerequisites A.6

description Once your zone is served you can request a secure delegation with your parent, or in absence of a signed parent approach ISC for an entry in their DLV registry.

A.8 Secure Delegation Provisioning

prerequisites A.6

description For the provisioning of secure delegations you will have to perform two subtasks.

You will have to provide a method for the zone-owners to securely enter the key signing key into your provisioning system. It is good to realize that the method you use for exchanging NS RRs has the same security requirements as the method for exchanging the DS information.

The details of the system depend a lot on how your registry, and possibly registrars, are set up. For registry-registrar interaction it may be good to note that EPP[6] supports DNSSEC.

You are advised to devise a method to prevent your children to enter secure entry-points for which there are no keys in the DNS *i.e.* you should try to prevent security lameness.

The second subtask is that the DS RRs should be pulled from your provisioning system into the nameservers. The details of this process are also dependent on your organization, but the process used for NS records can probably be cloned.

B BIND INSTALLATION

B BIND installation

There are two open-source reference implementations of DNSSEC for authoritative servers known to the author: BIND and NSD (<http://www.nlnetlabs.nl/nsd>). BIND is currently the only open-source recursive name server known to do DNSSEC validation.

DNSSEC is available as of BIND 9.3.0. The latest versions of BIND can be found on ISC's ftp server. DNSSEC support is only compiled if the openssl library is configured during compilation.

Make sure you fetch the latest version of BIND (take care of the patch level of the release indicated by `-Pnumber`) and verify the checksum.

configure with the `--with-openssl` flag.

If you want to have the "sigchase" capability (see Section 7) compiled into dig you will have to set the `STF_CDEFINES` variable to the `-DDIG_SIGCHASE=1`

Check the output of config to confirm that openssl was found. For example:

```
cd /usr/local/src
tar -xzf bind-9.4.1-P1.tar.gz
openssl sha1 bind-9.4.1-P1.tar.gz
...
cd bind-9.4.1-P1
./configure --prefix=/usr/local --with-openssl=/sw/
...
Checking whether byte ordering is bigendian... yes
checking for OpenSSL library... using openssl from /sw/lib and /sw/include
checking whether linking with OpenSSL works... yes
...
```

Please note that BIND 9.4.1 does not have DNSSEC enabled by default. Therefore you have to use the `dnssec-enable` and the `dnssec-validation` directives in the `options` section of `named.conf`.

```
options {
// turn on dnssec awareness
dnssec-enable yes;
dnssec-validation yes;
};
```

C Estimating zone size increase

When planning to sign zones you have to consider that zone-signing will increase your zone file size and the amount of memory used in the authoritative name servers. We have performed some measurements where we took a number of zone files, signed them and loaded them on a name server.

We started with the 1.8 thousand zones that the RIPE NCC serves on their authoritative servers. For a number of these zones the RIPE NCC is the primary server but for the largest part these are zones for which the RIPE NCC is secondary. The zones can roughly be split into two classes; "end-node" zones and "delegation" zones. In end-node-zones, the data for most names in the zone is authoritative (containing e.g. A, AAAA or PTR for most names). Delegation zones contain mostly delegations (NS) records, typically these are Top-Level Domains and "/16 reverse delegation" domains.

We signed the zone files with a 1024 bit RSASHA1 zone-signing key. During the signing NSEC RRs with corresponding RRSIGs are added and all RRsets in the zone are signed. Since a delegation NS RR is not an authoritative piece of data, no signature is created.

Typically for an "end-node" zone one NSEC and two RRSIGs are introduced into the zone, while for a delegation-type zone only one of each type of security record is introduced. In Figure 16 we plotted the zone file size increase as a function of the number of NSEC records. The number of NSEC records correlates with the number of domain names in a zone. In the figure you can clearly see a bi-modal distribution. One for the "end-node" type of zones and one for the delegation type of zone.

We fitted two linear relations to this data and found that for a delegation type of zone the size increase is 350 bytes per owner name while for an "end-node" zone the increase is 672 bytes per owner name.

In Figure 17, we plotted the relation between increase in core size versus the zone file size increase due to signing. The relation is linear and the slope is roughly 0.73. The core size increase is roughly 200 and 500 bytes for delegation type and end-node type zones respectively.

You can use these parameters for approximate size calculations. Results may vary depending on the size and the algorithm of the key you use, the version of BIND ¹⁹ and the content of the zone.

D Generating random numbers

The generation of keys and, for the DSA algorithm, the generation of signatures requires random numbers. You should take care that the random number generator produces "genuine" random numbers. The quality of random numbers generated in software is debatable. This also applies to /dev/random devices. These extract "randomness" from hardware response times. You should ensure that your operating system produces a flow of good random numbers. For a

¹⁹ These measurements were done with a snapshot version of BIND 9.3

E EXAMPLE DNSSEC KEY PROCEDURE

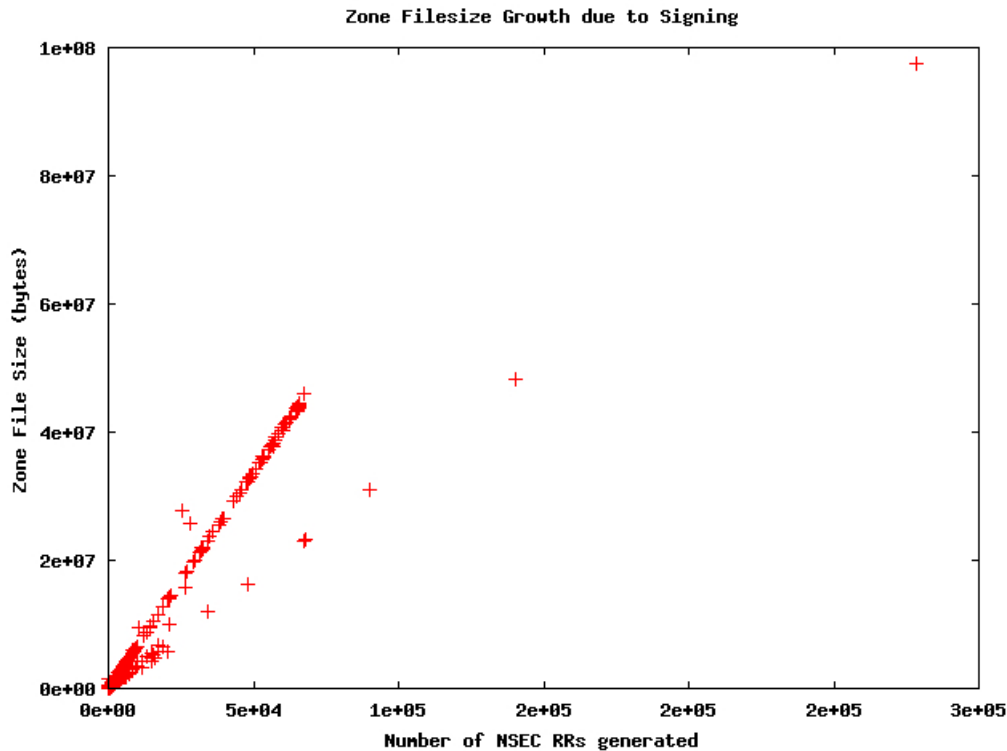


Figure 16: Zone size vs number of owner names

machine that does not have any external sources of "randomness", this may be tricky to achieve and cause your key generator or signer to block and wait for entropy ("randomness").

One relatively simple tool to test "randomness" of data streams is `ent` from Fourmilab <<http://www.fourmilab.ch/random/>>. Alternatively you could use tools from NIST <<http://csrc.nist.gov/rng/>>. "Good" measurement result from `ent` or the NIST tools should not be taken as a proof that your random number generator is perfect. There could be systematic effects that are hard to find using this particular tool ²⁰.

Relatively cheap sources of random data are USB crypto tokens. For more information about these tokens and random number generation see the Openfortress website <<http://openfortress.org/cryptodoc/random/>>.

E Example DNSSEC key procedure

This is an example key policy is based on, so not a verbatim copy of, the key policy as developed for RIPE NCC.

²⁰ For instance the predictable behavior of hardware during the boot of an OS may cause bit streams generated shortly after two boot sequences to be correlated

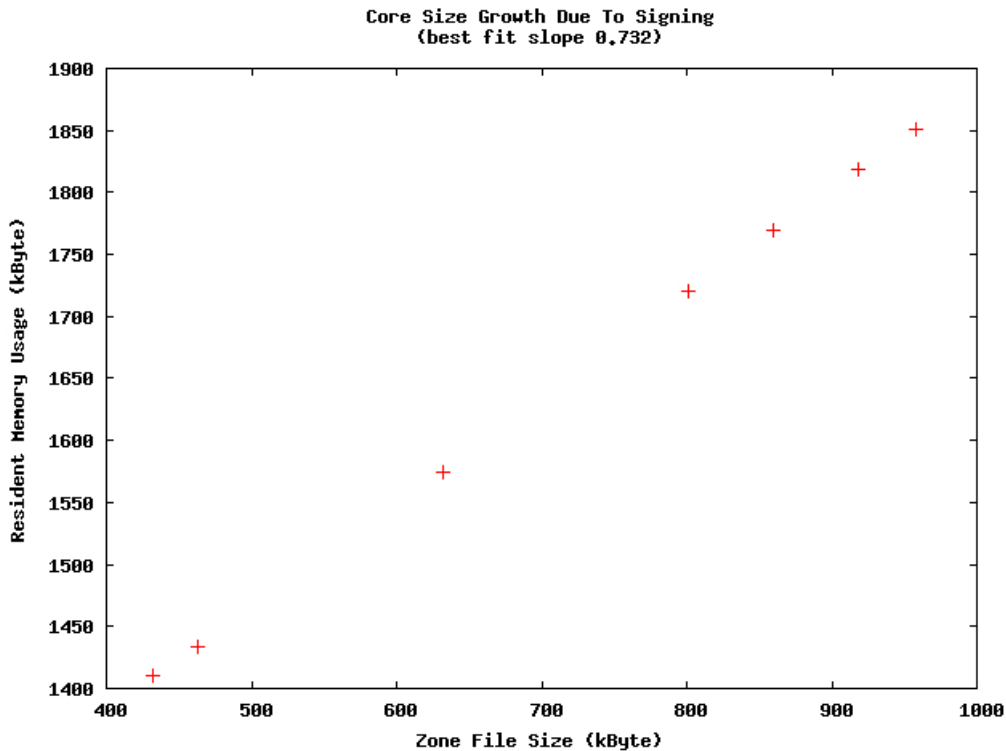


Figure 17: Core size vs zone file size

- This procedure²¹ applies to each zone that the is signed.
- Each zone with at least one Zone Signing Key (ZSK). A ZSK is zone specific.
- The ZSK will be published in the DNSKEY Resource Record (RR) set and signed with a Key Signing Key (KSK).
- The KSKs will have a SEP flag set so that they can be distinguished from the ZSKs in the DNSKEY RR set.
- The ZSK may be rolled without making any announcement. The 'pre-publish rollover scheme' as published in RFC4641 [10] is used. This will avoid breaks in the chain of trust.
- During the first two years of deployment, the KSK of each signed zone will be rolled twice each year. The rollover scheme that we will follow is the 'double signature scheme' published in RFC4661[10]. There will be an overlap of three months to allow zone administrators to configure their new key.
 - o At t=0 KSK1 signs the keyset. At t=3months KSK1 and KSK2 sign the keyset. DNS clients are expected to configure KSK2 during the three months that follow. At t=6months only KSK2 signs the keyset until (at

²¹Based on the procedure published at the RIPE NCC

t=9 months) KSK3 is introduced and a new rollover starts. o All zones at the RIPE NCC will roll their KSK simultaneously. Signatures are valid for one month. However, after announcing the change, the signing validity period may be changed to the shortest operationally possible period. Also see RFC4641[10] section 4.4.4. The ZSK will be an RSA/SHA1 key of 1200 bits (e RFC4641[10] section 3.5) The KSK will be an RSA/SHA1 key of 2048 bits.

- The KSKs to be used as 'trust-anchors' for our zones are published on a secure website in the format used in the 'trusted-keys' statement in BIND9 named configuration files.
- The KSK will also be published in the ISC DLV registry, but only until the root is signed.
- Any changes to this procedure and other announcements will be signed with our PGP key and published on our secure website and a dedicated mailing list.

F Perl's Net::DNS::SEC library

If you want to build tools to maintain your DNSSEC zones you may want to have a look at the `Net::DNS::SEC` library available on (CPAN <<http://search.cpan.org/dist/Net-DNS-SEC/>>). Using this extension to the `Net::DNS` library it is fairly easy to write scripts such as the one below that validate that the signature over a SOA will not expire within the next 24 hours.

```
#!/usr/local/bin/perl -T -Wall
#

# checkexpire.pl

# Example script that queries an authoritative server for a SOA
# record and verifies that the signatures over the record are still
# valid and will not expire in the next 24 hours.

# This annotated and somewhat verbose script is written for
# demonstration purposes only hence some possible error conditions
# are not tested for.

use strict;
use Net::DNS::SEC;
use Time::Local;

# The domain and its master server.
my $domain="secret-wg.org";
my $authoritative_server="ns.secret-wg.org";

# Setting up the resolver (use perldoc Net::DNS::Resolver for the
```

F PERL'S NET::DNS::SEC LIBRARY

```

# documentation of this class and its methods)
my $res = Net::DNS::Resolver->new();

# Query the default resolver to find out what the address is of the
# authoritative server.
my $answerpacket_auth_server= $res->query($authoritative_server,"A");

# Digest the packet see perldoc Net::DNS::Packet and perldoc
# Net::DNS::RR::A We ignore error checking. The first RR in the answer
# section is assumed to be the A RR for ns.secret-wg.org.

my $auth_address=($answerpacket_auth_server->answer)[0]->address;

# Set up the resolver object so it queries the authoritative server.
$res->nameserver( $auth_address );

# Set up the resolver so that it talks DNSSEC
$res->dnssec(1);

# Send the query for the soa to the authoritative nameserver.
my $packet=$res->send($domain,"SOA");

# Digest the answer section, realizing there may be more than one
# RRSIG (per definition there is always one SOA RR.

my $soa;
my @soasig;
foreach my $rr ( $packet->answer ){
    if ($rr->type eq "SOA"){
        $soa=$rr;
        next;
    }
    if ($rr->type eq "RRSIG"){
        push @soasig,$rr;
        next;
    }
}

die "NO SOA RR found" unless $soa;
die "NO RRSIGs over the SOA found" unless @soasig;
print @soasig ." signatures found\n";

# Fetch the keys that belong to this zone (DNSKEYs live, like the SOA
# at the apex.)

my @keyrr;
$packet=$res->send($domain,"DNSKEY");
foreach my $rr ( $packet->answer ){
    if ($rr->type eq "DNSKEY"){

```

```

        push @keyrr,$rr;
        next;
    }
}

die "NO DNSKEYS found for $domain" unless @keyrr;

# Now loop over each signature, fetch the public part of the key with
# which the signature was made, validate the signature and do the date
# comparisson.

# See perldoc Net::DNS::RR::RRSIG for the methods to access the RRSIGs
# internals

SIGLOOP: foreach my $sig ( @soasig ){
    print "Checking signature made with key ".$sig->keytag ."\n";
    # verify the signature.
    # first select the key with the proper keytag from the key set.
    my $keyfound=0;
    KEYLOOP: foreach my $key (@keyrr){
        next KEYLOOP if ($key->keytag != $sig->keytag);
        $keyfound=$key;
        last KEYLOOP;
    }
    print "WARNING: NO public key found to validate:\n " .
        $sig->string."\n" unless $keyfound;

    # Do the actual validation.
    if (! $sig->verify([ $soa ],$keyfound)){
        # The signature did not validate. Say why.
        print "WARN: Signature made with " . $sig->keytag . " failed to verify:\n".
            $sig->vrfyerrstr;
    }else{

        # The signature validated.
        # Lets verify if we have more than 24 hours before expiration.

        $sig->sigexpiration =~ /(\d{4})(\d{2})(\d{2})(\d{2})(\d{2})(\d{2})/;
        my $expiration=timegm ($6, $5, $4, $3, $2-1, $1-1900);
        my $hourstogo=($expiration-time())/3600;
        print "WARNING: Signature made with ".$sig->tag. "will expire within ".
            $hourstogo . " hours\n" if $hourstogo <24;
    }
}
}

```


F PERL'S NET::DNS::SEC LIBRARY

```
#####  
# $Id: expire.pl 21 2004-10-11 14:52:09Z olaf $  
#####
```

References

- [1] Ron Aitchison. *Pro DNS and BIND*. Apress, 2005.
- [2] Paul Albitz and Cricket Liu. *DNS and BIND, 4th Edition*. O'Reilly, 4 edition, April 2001.
- [3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *DNS Security Introduction and Requirements*. RFC 4033 (Proposed Standard), March 2005. <http://www.ietf.org/rfc/rfc4033.txt>.
- [4] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Protocol Modifications for the DNS Security Extensions*. RFC 4035 (Proposed Standard), March 2005. <http://www.ietf.org/rfc/rfc4035.txt>, (Updated by RFC 4470).
- [5] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Resource Records for the DNS Security Extensions*. RFC 4034 (Proposed Standard), March 2005. <http://www.ietf.org/rfc/rfc4034.txt>, (Updated by RFC 4470).
- [6] S. Hollenbeck. *Domain Name System (DNS) Security Extensions Mapping for the Extensible Provisioning Protocol (EPP)*. RFC 4310 (Proposed Standard), December 2005. <http://www.ietf.org/rfc/rfc4310.txt>.
- [7] Geoff Huston. *DNSSEC The Opinion*. The ISOC ISP column, November 2006. <http://ispcolumn.isoc.org/2006-10/dnssec3.html>.
- [8] Geoff Huston. *DNSSEC The Practice*. The ISOC ISP column, September 2006. <http://ispcolumn.isoc.org/2006-09/dnssec2.html>.
- [9] Geoff Huston. *DNSSEC The Theory*. The ISOC ISP column, August 2006. <http://ispcolumn.isoc.org/2006-08/dnssec.html>.
- [10] O. Kolkman and R. Gieben. *DNSSEC Operational Practices*. RFC 4641 (Informational), September 2006. <http://www.ietf.org/rfc/rfc4641.txt>.
- [11] O. Kolkman, J. Schlyter, and E. Lewis. *Domain Name System KEY (DNSKEY) Resource Record (RR) Secure Entry Point (SEP) Flag*. RFC 3757 (Proposed Standard), April 2004. <http://www.ietf.org/rfc/rfc3757.txt>, (Obsoleted by RFCs 4033, 4034, 4035).
- [12] Olaf Kolkman. *Measuring the resource requirements of DNSSEC*. RIPE NCC web pages. <http://www.ripe.net/ripe/docs/ripe-352.html>.
- [13] NLnet Labs. *NSD Memory Size Estimate*, 2008. <http://www.nlnetlabs.nl/projects/nsd/nsd-memsize.html>.
- [14] Evi Nemeth. *Securing the DNS*. ;login:, pages 21–31, November 2000.

REFERENCES

- [15] M. StJohns. *Automated Updates of DNSSEC Trust Anchors*, November 2006. <http://www.ietf.org/internet-drafts/draft-ietf-dnsext-trustupdate-time%rs-05.txt>, (Internet Drafts are subject to change and have a limited lifetime).
- [16] M. StJohns. *Automated Updates of DNS Security (DNSSEC) Trust Anchors*. RFC 5011 (Proposed Standard), September 2007. <http://www.ietf.org/rfc/rfc5011.txt>.
- [17] P. Vixie, O. Gudmundsson, D. Eastlake 3rd, and B. Wellington. *Secret Key Transaction Authentication for DNS (TSIG)*. RFC 2845 (Proposed Standard), May 2000. <http://www.ietf.org/rfc/rfc2845.txt>, (Updated by RFC 3645).
- [18] Paul Vixie. *Preventing Child Neglect in DNSSECbis Using Lookaside Validation (DLV)*. *IEICE-Transactions on Communications*, E88-B, number 4:21–31, April 2005.
- [19] Paul Vixy and Mark Andrews. *DNSSEC Lookaside Validation (DLV)*, April 2006. <http://www.isc.org/pubs/tn/isc-tn-2006-1.html>.

Acknowledgements

The RIPE NCC for their initial investment of my resources in this document.

There are numerous people who helped compiling these notes, either by helping me to understand DNSSEC or by giving feedback on earlier versions of this document. Special thanks go to Sebastian Anding, Rossen Antonov, Roy Arends, Adrian Bedford, Emma Bretherick, Daniel Diaz, Miek Gieben, Geoff Huston, Daniel Karrenberg, Marc Lampo, Ed Lewis, Cricket Liu, Rick van Rein, Andrew Ruthven, Jakob Shlyter, Görran Uddeborg, Paul Vixie, and Wouter Wijngaards

who provided feedback on this or earlier versions of this document Also, thanks to the participants of a DNSSEC workshop at InERLab in November 2006 for being guinea-pigs for version 1.8. A workshop organized by USC/ISI on operational testing of DS in Washington DC has provided a substantial amount of material for version 1.3 of the document. Finally, a 'login;' article by Evi Nemeth[14], the text in the BIND book and the various presentations by Edward Lewis have been the examples on which I based version 1.1 of this document.

Colophon

This document will be subject to change. Please regularly check for new versions. <http://www.nlnetlabs.nl/dnssec_howto/>. Your corrections and additions are appreciated.

If you have questions, remarks or contributions please contact dnssec-howto-editor@nlnetlabs.nl

The sourcetext of this document has shortly been in Docbook format. As of version 1.8 it is again authored as tex. I have much more experience and control over the output with TeX. As of version 1.8 some of the examples and log-outputs are maintained with some shell scripts and `make`.

The source text and the "DNS infrastructure" needed to create example output is all under subversion version control.

This PDF version of the howto has been created with `pdf2latex` the html version has been created with `tex4ht`. The tex source, or a snapshot of the subversion repository, is available on request (dnssec-howto-editor@nlnetlabs.nl).

Document History

Revision 136 July 4, 2009

The version numbering has changed to be based on subversion versions.

136saw the addition of part I and appendix A . Some sections were rewritten and more information about TARs and tools was added.

This release is a snapshot, there are sections that need more work such as the section on automatic key maintenance and automated rollover.

Also some errors were corrected and some formatting improved.

Revision 1.8 Public release, January 2006

The document was brought under subversion control. That changed the revision numbering from dotted notation to continuous revision numbers. The publication versions are now manually remained. The document source has been back-ported from docbook to latex. Section 4 was rewritten and names and addresses were made consistent. Most of the examples are now generated automatically in order to maintain consistency after software upgrades.

Some minor editorial corrections were performed post publication. The subversion Id is relevant.

Revision 1.7 Public release, April 2005.

Paragraph added to clarify the TSIG signing of NOTIFY messages. Minor editorial fixes. References were added.

Revisions 1.6 Public release, December 2004

REFERENCES

Minor editorial fixes.

Revisions 1.5 Public release.

Revisions 1.4.4.1 - 1.4.4.9 Several snapshots that were not publically available. The document source was ported from TeX to docbook source and was updated to reflect experiences with keymanagement. Large chunks were rewritten and we added a number of figures. This version of the documentation is based on the DNSSEC bis specification and bind9.3.0beta implementation.

Version 1.3, Oct 2002 The first modification of the document. First experiences with DS have been incooperated and the document has been rewritten to be useful as a more generic HOWTO and introduction to \dnssec operations.

Version 1.2 was not published.

Version 1.1 Was compiled for a DNSSEC tutorial in Prague, October 8, 2000. The document was a set of notes to be used in a workshop setup.

COPYRIGHT

Copyright © 2006, 2007, 2008, 2009 NLnet Labs

Copyright © 2002, 2003, 2004, 2005 RIPE NCC

Legal Notice

This document and the information contained herein is provided on an **as is** basis and **NLnet Labs and RIPE NCC disclaim all warranties, express or implied, including but not limited to any warranty that the use of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.**

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works.