                DNSSEC Signing Policy & Practice Statement Framework
                    draft-ietf-dnsop-dnssec-dps-framework-00

Abstract

   This document presents a framework to assist writers of DNSSEC
   Signing Policy and Practice Statements such as Regulatory Authorities
   and Registry Managers on both the TLD and secondary level, who is
   operating a DNS zone with Security Extensions (DNSSEC) implemented.

   In particular, the framework provides a comprehensive list of topics
   that potentially (at the writer's discretion) needs to be covered in
   a DNSSEC Signing Policy definition and Practice Statement.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

1.1. Background

The Domain Name System Security Extensions (DNSSEC, RFC 4033 [RFC4033]) is a set of IETF specifications for securing the Domain Name System. DNSSEC provides a way for software to validate that Domain Name System (DNS) data have not been modified during Internet transit.

The DNS was not originally designed with strong security mechanisms to provide integrity and authenticity of DNS data. Over the years, a number of vulnerabilities had been discovered that threaten the reliability and trustworthiness of the system. DNSSEC addresses these vulnerabilities by adding data origin authentication, data integrity verification and authenticated denial of existence capabilities to DNS by incorporating public key cryptography into the DNS hierarchy.

DNSSEC differs from the X.509 PKI in many ways. For instance, in DNSSEC there is no central control of assurance or trust levels. Each zone manager may have its own way of managing keys and operations, and there is no necessity to perform any coordination between different zones or levels in the DNS. The degree to which a relying party can trust the binding embodied in the DNSSEC chain is dependent on the weakest link in that chain. As an implication of this nature, the security of domains becomes more critical the higher up in the DNS hierarchy one gets.

To provide means for the relying parties to evaluate the trust and strength of the chain, registries may choose to publish DNSSEC Practice Statements (DPSs), comprising statements of critical security controls and procedures relevant to the relying parties.

Even though this document is heavily inspired by RFC 3647 [RFC3647] (and some content shamelessly copied), one significant difference is that the DPS framework is focused only on stating the security posture of the registry, rather than for all participants in the domain name system. The DNS is almost of a ubiquitous nature, and there exists no agreements with the relying (third) parties, which is basically everyone using the Internet.

1.2. Purpose

The purpose of this document is twofold. Firstly, the document aims to explain the concept of a DPS and describe the relationship between the DPS, the registry, the domain holders and the relying (third) parties. Second, this document aims to present a framework to

encourage and assist writers of DPSs in creating heterogenous and comparable policies and practices. In particular, the framework identifies the elements that may need to be considered in formulating a DPS. The purpose is not to define a particular Policy or Practice Statement, per se. Moreover, this document does not aim to provide legal advice or recommendations as to particular requirements or practices that should be contained within a DPS.

1.3. Scope

The scope of this document is limited to discussion of the topics that can be covered in a DPS and does not go into the specific details that could possibly be included in each topic. In particular, this document describes the types of information that should be considered for inclusion in a DPS.

This DPS framework should be viewed and used as a tool that could act as a checklist of factors that should be taken in to consideration prior to deploying DNSSEC, and an outline to create a document that states the actual DNSSEC operations. The framework is primarily aimed at organizations providing registry services, but may be used by high-value domain holders and serve as a check sheet for DNSSEC readiness at a high level.

This document assumes that the reader is familiar with the general concepts of DNS, DNSSEC and PKI.

2. Definitions

Policy Authority > needs explanation

3. Concepts

This section describes the concept of a DNSSEC Signing Policy and Practices Statement. Other related concepts are described as well.

3.1. DNSSEC

3.2. DPS

Most DNSSEC participants may not have the need to create a thorough and detailed statement of practices. For example, the registrant may itself be the relying party of its own zone and would already be aware of the nature and trustworthiness of its services. In other cases, a Registry may provide registration services providing only a very low level of assurances where the domain names being secured may

pose only marginal risks if compromised. In these cases, an
organization implementing DNS Security Extensions may only want to
use a registrant agreement. In such circumstances, that agreement
may serve as the only "statement of practices" used by one or more
registries within that Zone. Consequently, that agreement may also
be considered a DPS and can be entitled or subtitled as such.

A DPS should contain detailed information which is relevant to the
DNSSEC participants. Since the participants generally include the
Internet community, it should not contain such information which
could be considered to be sensitive details of a registry's
operation, but rather reference such processes and procedures, which
need not be public information.

The DPS does not automatically constitute a contract and do not
automatically bind DNSSEC participants as a contract would. In most
cases there exists no contractual agreement between the registry and
the relying party. Where a document serves the dual purpose of being
a registrant agreement and a DPS, the document is intended to be part
of a contract and constitutes a legally binding document to the
extent that a registrant agreement would ordinarily be considered as
such.

Therefore, DPS's terms have a binding effect as contract terms only
if a separate document creates a contractual relationship between the
parties and where that document incorporates parts or all of the DPS
by reference.

3.3. Relationship between DNSSEC Signing Policy and Practice Statement

A DNSSEC Signing Policy and a DNSSEC Practice Statement address the
same set of topics that are of interest to the relying party in terms
of the degree to and purpose for which a signature should be trusted.
Their primary difference is in the focus of their provisions. A
Signing Policy sets forth the requirements and standards to be
implemented for a DNSSEC Signed Zone. In other words, the purpose of
the Signing Policy is to establish what participants must do. A
Practice statement, by contrast, states how a Registry and other
participants in a given Zone implement procedures and controls to
meet the requirements stated in the Policy. In other words, the
purpose of the Practice Statement is to disclose how the participants
perform their functions and implement controls.

An additional difference between a Signing Policy and a Practice
Statement relates the scope of coverage of the two kinds of
documents. Since a Signing Policy is a statement of requirements, it
best serves as the vehicle for communicating minimum operating
guidelines that must be met by complying parties. Thus, a Signing

Policy may apply to multiple Registries, multiple organizations, or
multiple zones. By contrast, a Practice Statement applies only to a
single Registry or single organization.

For example, a Regulatory Authority might define requirements in a
Signing Policy for DNSSEC operations for one or more Registries. The
Signing Policy will be a broad statement of the general requirements
for participants within the Registry's Zone.

A registry may be required to write its own Practice Statement to
support this Signing Policy by explaining how it meets the
requirements of the Signing Policy. Or, a Registry not governed by
any Signing Policy may choose to publish a Practice Statement to
provide transparency and gain community trust and acceptance.

An additional difference between a Signing Policy and a Practice
Statement concerns the level of detail of the provisions in each.
Although the level of detail may vary, a Practice Statement will
generally be more detailed than a Signing Policy. A Practice
Statement provides a detailed description of procedures and controls
in place to meet the Signing Policy requirements, while a Signing
Policy is more general.

The main differences between a Signing Policy and Practice Statement
can therefore be summarized as follows:

(a) Operation of a DNS Zone with DNSSEC may be governed by a Signing
     Policy, to establish requirements that state what the parties
     within the zone it must do. A single Registry or organization
     can use a Practice Statment to disclose how it meets the
     requirements of a Signing Policy or how it implements its
     practices and controls.
(b) A Signing Policy may facilitate interoperation of level of trust
     through several parts or levels in the DNS hierarchy. By
     contrast, a Practice Statement is a statement of a single
     Registry or organization.
(c) A Practice Statement is generally more detailed than a Signing
     Policy and specifies how the Registry meets the requirements
     specified in the one or more Signing Policies under which it
     operates DNSSEC.

3.4. Set of provisions

A set of provisions is a collection of Signing Policy and/or Practice
statements, spanning a range of standard topics for use in expressing
a Signing Policy or Practice Statement employing the approach
described in this framework by covering the topic appearing in
Section 5 below. They are also described in detail in Section 4

below.

A Signing Policy can be expressed as a single set of provisions.

A Practice Statement can be expressed as a single set of provisions
with each component addressing the requirements of one or more
Signing Policies, or, alternatively, as an organized collection of
sets of provisions. For example, a Practice Statement could be
expressed as a combination of the following:

(a) a list of Signing Policies supported by the DPS;
(b) for each Signing Policy in (a), a set of provisions that
     contains statements responding to that Signing Policy by filling
     in details not stipulated in that policy or expressly left to
     the discretion of the Signing Policy (in its Practice
     Statement); such statements serve to state how this particular
     Practice Statement implements the requirements of the particular
     Signing Policy; or
(c) a set of provisions that contains statements regarding the
     practices of the DNSSEC operations, regardless of Signing
     Policy.

The statements provided in (b) and (c) may augment or refine the
stipulations of an applicable Signing Policy, but generally must not
conflict with any of the stipulations of such Signing Policy. In
certain cases, however, a Policy Authority may permit exceptions to
the requirements in a Signing Policy, because certain compensating
controls of the registry are disclosed in its DPS that allow the
registry to provide assurances that are equivalent to the assurances
provided by registries that are in full compliance with the DPS.

This framework outlines the contents of a set of provisions, in terms
of eight primary components, as follows:

1. Introduction
2. Publication and Repositories
3. Operational Requirements
4. Facilities, Management, and Operational Controls
5. Technical Security Controls
6. Zone Signing
7. Compliance Audit
8. Legal Matters

Policy authorities can use this framework of eight primary components
to write a DNSSEC Signing Policy. Moreover, a Registry can use this
same framework to write a DNSSEC Practice Statement.

Therefore, a Registry can establish a set of basic documents (with a

Signing Policy, Practice statement, and Registrant agreement) all
having the same structure and ordering of topics, thereby
facilitating comparisons and mappings among these documents and among
the corresponding documents of other Zones.

This basic framework may also be useful for the establishing of
agreements with registrars or outsourcing of certain services.

Drafters of DPSs are permitted to add additional levels of
subcomponents below the subcomponents described in Section 4 for the
purpose of meeting the needs of the drafter's particular
requirements.

4. Contents of a set of provisions

4.1. Introduction

This component identifies and introduces the set of provisions, and
indicates the types of entities and applications for which the
document (either the Signing Policy or the Practice Statement being
written) is targeted.

4.1.1. Overview

This subcomponent provides a general introduction to the document
being written. This subcomponent can also be used to provide a
synopsis of the community to which the Signing Policy or Practice
Statement applies.

4.1.2. Document Name and Identification

This subcomponent provides any applicable names or other identifiers
of the document.

4.1.3. Community and Applicability

This subcomponent addresses the stakeholders in DNSSEC along with the
expected roles and responsibilities. This includes but are not
limited to an entity signing the zone, an entity that relies on the
signed zone, other entities that have operational dependency on the
signed zone and an entity that entrusted the zone signing.

4.1.4. Specification Administration

This subcomponent includes the name and mailing address of the
organization that is responsible for drafting, registering,
maintaining, and updating of the DPS. It also includes the name,

electronic mail address, telephone number, and fax number of a
contact person. As an alternative to naming an actual person, the
document may name a title or role, an e-mail alias, and other
generalized contact information. In some cases, the organization may
state that its contact person, alone or in combination with others,
is available to answer questions about the document.

Moreover, when a formal or informal Policy Authority is responsible
for determining whether a Registry should be allowed to operate a
Zone, it may wish to approve the DPS of the Registry as being
suitable for the Policy Authority's Signing Policy. If so, this
subcomponent can include the name or title, electronic mail address
(or alias), telephone number, fax number, and other generalized
information of the entity in charge of making such a determination.
Finally, in this case, this subcomponent also includes the procedures
by which this determination is made.

4.2. Publication and Repositories

This component contains any applicable provisions regarding:

o An identification of the entity or entities that operate
   repositories within the community, such as a Registry;
o The responsibility of a registry to publish information regarding
   its practices, public keys, and the current status of such keys,
   which may include the responsibilities of making the DPS publicly
   available using various mechanisms and of identifying components,
   subcomponents, and elements of such documents that exist but are
   not made publicly available, for instance, security controls,
   clearance procedures, or business information due to their
   sensitivity;
o When information must be published and the frequency of
   publication; and
o Access control on published information objects.

4.3. Operational Requirements

This component describes the operational requirements when operating
DNSSEC.

4.3.1. Meaning of domain names

This section describes the meaning of names in child zones, if any.

4.3.2. Activation of DNSSEC for child zone

   This section describes how the child zone would be tied into the
   parent zone by incorporating DS record into the zone.

4.3.3. Identification and authentication of child zone manager

   This section will specify the methodology of identifying and
   authenticating the requester of the child zone to determine whether
   the request is valid or not.

4.3.4. Registration of delegation signer (DS) records

   This section describes how the delegation signer records are
   incorporated into the parent zone.

4.3.5. Method to prove possession of private key

   This section describes how the child zone proves the possession of
   the Key Signing Key to the parent zone when requesting a delegation
   signer record to be incorporated.

4.3.6. Removal of DS record

   This section will explain how, when and under which circumstances the
   DS record may be removed from the zone file.

4.4. Management, Operational, and Physical Controls

   This component describes non-technical security controls (that is,
   physical, procedural, and personnel controls) used by the Registry to
   securely perform the DNSSEC related functions such as key management,
   zone signing, key roll-over, zone distribution, auditing and
   archiving.

   These non-technical security controls are critical for trusting the
   signatures since lack of security may compromise DNSSEC operations
   resulting for example, in the creation of signatures with erroneous
   information or compromising the Key Signing Key and/or Zone Signing
   Key.

   Within each subcomponent, separate consideration will, in general,
   need to be given to each entity type.

4.4.1. Physical Controls

   In this subcomponent, the physical controls on the facility housing
   the entity systems are described. Topics addressed may include:

o Site location and construction, such as the construction
   requirements for high-security areas and the use of locked rooms,
   cages, safes, and cabinets;
o Physical access, i.e., mechanisms to control access from one area
   of the facility to another or access into high-security zones,
   such as locating Registry operations in a secure computer room
   monitored by guards, cameras or security alarms and requiring
   movement from zone to zone to be accomplished using a token,
   biometric readers, and/or access control lists;
o Power and air conditioning;
o Water exposures;
o Fire prevention and protection;
o Media storage, for example, requiring the storage of backup media
   in a separate location that is physically secure and protected
   from fire, smoke, particle and water damage;
o Waste disposal; and
o Off-site backup.

4.4.2. Procedural Controls

In this subcomponent, requirements for recognizing trusted roles are
described, together with the responsibilities for each role.
Examples of trusted roles include system administrators, security
officers, and system auditors.

For each task identified, the number of individuals required to
perform the task (n out m rule if applicable) should be stated for
each role. Identification and authentication requirements for each
role may also be defined.

This component also includes the separation of duties in terms of the
roles that cannot be performed by the same individuals.

4.4.3. Personnel Controls

This subcomponent addresses the following:

o Qualifications, experience, and clearances that personnel must
   have as a condition of filling trusted roles or other important
   roles. Examples include credentials, job experiences, and
   official government clearances that candidates for these positions
   must have before being hired;
o Background checks and clearance procedures that are required in
   connection with the hiring of personnel filling trusted roles or
   perhaps other important roles; such roles may require a check of
   their criminal records, references, and additional clearances that
   a participant undertakes after a decision has been made to hire a

      particular person;
   o Training requirements and training procedures for each role
      following the hiring of personnel;
   o Any retraining period and retraining procedures for each role
      after completion of initial training;
   o Frequency and sequence for job rotation among various roles;
   o Sanctions against personnel for unauthorized actions, unauthorized
      use of authority, and unauthorized use of entity systems for the
      purpose of imposing accountability on a participant's personnel;
   o Controls on personnel that are independent contractors rather than
      employees of the entity; examples include:

      * Bonding requirements on contract personnel;
      * Contractual requirements including indemnification for damages
         due to the actions of the contractor personnel;
      * Auditing and monitoring of contractor personnel; and
      * Other controls on contracting personnel.

   o Documentation to be supplied to personnel during initial training,
      retraining, or otherwise.

4.4.4. Audit Logging Procedures

   This subcomponent is used to describe event logging and audit
   systems, implemented for the purpose of maintaining a secure
   environment. Elements include the following:

   o Types of events recorded, such as attempts to access the system,
      and requests made to the system;
   o Frequency with which audit logs are processed or archived, for
      example, weekly, following an alarm or anomalous event, or when
      ever the audit log is n% full;
   o Period for which audit logs are kept;
   o Protection of audit logs:

      * Who can view audit logs, for example only the audit
         administrator;
      * Protection against modification of audit logs, for instance a
         requirement that no one may modify or delete the audit records
         or that only an audit administrator may delete an audit file as
         part of rotating the audit file; and
      * Protection against deletion of audit logs.

   o Audit log back up procedures;
   o Whether the audit log accumulation system is internal or external
      to the entity;

o Whether the subject who caused an audit event to occur is notified
   of the audit action; and
o Vulnerability assessments, for example, where audit data is run
   through a tool that identifies potential attempts to breach the
   security of the system.

4.4.5. Compromise and Disaster Recovery

   This subcomponent describes requirements relating to notification and
   recovery procedures in the event of compromise or disaster. Each of
   the following may need to be addressed separately:

   o Identification or listing of the applicable incident and
      compromise reporting and handling procedures.
   o The recovery procedures used if computing resources, software,
      and/or data are corrupted or suspected to be corrupted. These
      procedures describe how a secure environment is re-established,
      whether the Key Signing Key or Zone Signing key requires a roll
      over, how to assess the damage and carry out the root cause
      analysis.
   o The recovery procedures used if the Key Signing Key or Zone
      Signing Key is compromised. These procedures describe how a
      secure environment is re- established, how the keys are rolled
      over, how the new Trust Anchor is provided to the users and how
      new zone file is published.
   o The entity's capabilities to ensure business continuity following
      a natural or other disaster. Such capabilities may include the
      availability of a disaster recovery site at which operations may
      be recovered. They may also include procedures for securing its
      facility during the period of time following a natural or other
      disaster and before a secure environment is re-established, either
      at the original site or at a disaster recovery site. For example,
      procedures to protect against theft of sensitive materials from an
      earthquake-damaged site.

4.4.6. Entity termination

   This subcomponent describes requirements relating to procedures for
   termination, termination notification and transition of
   responsibilities of a Registry. The major purpose is to ensure that
   the transition process will be transparent to the relying party and
   will not affect the service.

4.5. Technical Security Controls

This component is used to define the security measures taken by the
Registry to protect its cryptographic keys and activation data (e.g.,
PINs, passwords, or manually-held key shares). This component may
also be used to impose constraints on repositories, child zone
operators, and other participants to protect their private keys,
activation data for their private keys, and critical security
parameters. Secure key management is critical to ensure that all
secret and private keys and activation data are protected and used
only by authorized personnel.

This component also describes other technical security controls used
by the Registry to perform securely the functions of key generation,
authentication, registration, auditing, and archiving. Technical
controls include life-cycle security controls (including software
development environment security, trusted software development
methodology) and operational security controls.

This component can also be used to define other technical security
controls on repositories, authoritative name servers, registrants,
and other participants.

4.5.1. Key Pair Generation and Installation

Key pair generation and installation need to be considered for the
Registry. The following questions potentially need to be answered:

1. Who generates the Registry's public, private key pair?
    Furthermore, how is the key generation performed? Is the key
    generation performed by hardware or software?
2. How is the private key installed in all parts of the key
    management system?
3. How is the Registry's public key provided securely to potential
    relying parties?
4. What are the key sizes and algorithm?
5. Who generates the public key parameters, and is the quality of
    the parameters checked during key generation?
6. For what purposes may the key be used, or for what purposes
    should usage of the key be restricted?

4.5.2. Private key protection and Cryptographic Module Engineering
       Controls

Requirements for private key protection and cryptographic modules
need to be considered for key generation and creation of signatures.
The following questions potentially need to be answered:

1. What standards, if any, are required for the cryptographic
    module used to generate the keys? A cryptographic module can be
    composed of hardware, software, firmware, or any combination of
    them. For example, are the zones signatures required to be
    generated using modules compliant with the US FIPS 140-2? If
    so, what is the required FIPS 140-2 level of the module? Are
    there any other engineering or other controls relating to a
    cryptographic module, such as the identification of the
    cryptographic module boundary, input/output, roles and services,
    finite state machine, physical security, software security,
    operating system security, algorithm compliance, electromagnetic
    compatibility, and self tests.

2. Is the private key under n out of m multi-person control?(7) If
    yes, provide n and m (two person control is a special case of n
    out of m, where n = m = 2)?

3. Is the private key escrowed?(8) If so, who is the escrow agent,
    what form is the key escrowed in (examples include plaintext,
    encrypted, split key), and what are the security controls on the
    escrow system?

4. Is the private key backed up? If so, who is the backup agent,
    what form is the key backed up in (examples include plaintext,
    encrypted, split key), and what are the security controls on the
    backup system?

5. Is the private key archived? If so, who is the archival agent,
    what form is the key archived in (examples include plaintext,
    encrypted, split key), and what are the security controls on the
    archival system?

6. Under what circumstances, if any, can a private key be
    transferred into or from a cryptographic module? Who is
    permitted to perform such a transfer operation? In what form is
    the private key during the transfer (i.e., plaintext, encrypted,
    or split key)?

7. How is the private key stored in the module (i.e., plaintext,
    encrypted, or split key)?

8. Who can activate (use) the private key? What actions must be
    performed to activate the private key (e.g., login, power on,
    supply PIN, insert token/key, automatic, etc.)? Once the key is
    activated, is the key active for an indefinite period, active
    for one time, or active for a defined time period?

9. Who can deactivate the private key and how? Examples of methods
    of deactivating private keys include logging out, turning the
    power off, removing the token/key, automatic deactivation, and
    time expiration.

10. Who can destroy the private key and how? Examples of methods of
    destroying private keys include token surrender, token
    destruction, and overwriting the key.

11. Provide the capabilities of the cryptographic module in the
     following areas: identification of the cryptographic module
     boundary, input/output, roles and services, finite state
     machine, physical security, software security, operating system
     security, algorithm compliance, electromagnetic compatibility,
     and self tests. Capability may be expressed through reference
     to compliance with a standard such as U.S. FIPS 140-1,
     associated level, and rating.

4.5.3. Other Aspects of Key Pair Management

Other aspects of key management need to be considered for the
Registry and other participants. For each of these types of
entities, the following questions potentially need to be answered:

1. Is the public key archived? If so, who is the archival agent and
    what are the security controls on the archival system?
2. What is the operational period of the keys. What are the usage
    periods, or active lifetimes, for the subscriber's key pair?

4.5.4. Activation data

Activation data refers to data values other than whole private keys
that are required to operate private keys or cryptographic modules
containing private keys, such as a PIN, passphrase, or portions of a
private key used in a key-splitting scheme. Protection of activation
data prevents unauthorized use of the private key, and potentially
needs to be considered for the Registry. Such consideration
potentially needs to address the entire life-cycle of the activation
data from generation through archival and destruction. For each of
the entity types, all of the questions listed in 4.5.1 through 4.5.3
potentially need to be answered with respect to activation data
rather than with respect to keys.

4.5.5. Computer Security Controls

This subcomponent is used to describe computer security controls such
as: use of the trusted computing base concept, discretionary access
control, labels, mandatory access controls, object re-use, audit,
identification and authentication, trusted path, security testing,
and penetration testing. Product assurance may also be addressed.

A computer security rating for computer systems may be required. The
rating could be based, for example, on the Trusted System Evaluation
Criteria (TCSEC), Canadian Trusted Products Evaluation Criteria,
European Information Technology Security Evaluation Criteria (ITSEC),

or the Common Criteria for Information Technology Security
Evaluation, ISO/IEC 15408:1999. This subcomponent may also address
requirements for product evaluation analysis, testing, profiling,
product certification, and/or product accreditation related activity
undertaken.

4.5.6. Network Security Controls

This subcomponent addresses network security related controls,
including firewalls.

4.5.7. Timestamping

This subcomponent addresses requirements or practices relating to the
use of timestamps on various data. It may also discuss whether or
not the time-stamping application must use a trusted time source.

4.5.8. Life Cycle Technical Controls

This subcomponent addresses system development controls and security
management controls.

System development controls include development environment security,
development personnel security, configuration management security
during product maintenance, software engineering practices, software
development methodology, modularity, layering, use of failsafe design
and implementation techniques (e.g., defensive programming) and
development facility security.

Security management controls include execution of tools and
procedures to ensure that the operational systems and networks adhere
to configured security. These tools and procedures include checking
the integrity of the security software, firmware, and hardware to
ensure their correct operation.

This subcomponent can also address life-cycle security ratings based,
for example, on the Trusted Software Development Methodology (TSDM)
level IV and V, independent life-cycle security controls audit, and
the Software Engineering Institute's Capability Maturity Model (SEI-
CMM).

4.6. Zone Signing

This component covers all aspects of zone signing including, the
cryptographic specification surrounding the Key Signing Key and Zone
Signing Key, methodology and signing scheme for key roll-over and the
actual zone signing. Child zones and other relying parties may
depend on the information in this section to understand the expected

data in the signed zone and determine their own behavior. In
addition, this section will be used to state the compliance to the
cryptographic and operational requirements pertaining to zone signing
if applicable.

4.6.1. Key lengths and algorithms

This subcomponent describes the key generation algorithm and the key
length used to create the Key Signing Key and the Zone Signing Key.

4.6.2. Authenticated denial of existence

Authenticated denial of existence refers to the usage of NSEC (RFC
4034 [RFC4034]), NSEC3 (RFC 5155 [RFC5155])or any other record
defined in the future that is used to authenticate the denial of
existence of the resource record.

4.6.3. Signature format

This subcomponent is used to describe the signing method used for the
zone signing.

4.6.4. Zone signing key roll-over

This subcomponent explains the Zone signing key roll-over scheme.

4.6.5. Key signing key roll-over

This subcomponent addresses the Key signing key roll-over scheme.

4.6.6. Signature life-time and re-signing frequency

This subcomponent identifies the life-cycle of the Resource Record
Signature (RRSIG) record.

4.6.7. Verification of zone signing key set

This subsection addresses the controls around the keyset signing
process performed by the Key Signing Key. The procedures surrounding
KSK management may be different from those of the ZSK, hence it may
be necessary to authenticate the data signed by the KSK.

4.6.8. Verification of resource records

This subsection addresses the controls around the verification of the
resource records in order to validate and authenticate the data to be
signed.

4.6.9. Resource records time-to-live

   This subcomponent specifies the time-to-live (TTL) for each DNSSEC
   related resource record such as DNSKEY, NSEC/NSEC3, DS and RRSIG.

4.7. Compliance Audit

   The ideal and the only way to prove the statements in the DNSSEC
   Signing Policy or Practices Statement is to conduct an audit. This
   component describes the outline of how the audit is conducted at the
   registry.

4.7.1. Frequency of entity compliance audit

   This subcomponent describes the frequency the compliance audit. An
   audit could be considered as a health check of the service therefore
   it is ideal to have an audit at least once a year to know the current
   status.

4.7.2. Identity/qualifications of auditor

   This subcomponent addresses what is the qualifications for the
   auditor. For instance it may be an auditor from a specific
   association or an auditor that has a certain certifications.

4.7.3. Auditor's relationship to audited party

   This subcomponent is used to clarify the relationship between the
   auditor and the entity being audited. This becomes important if
   there is any requirements or guidlines for selection of the auditor.

4.7.4. Topics covered by audit

   Topics covered by audit refers to the scope of the audit. Since the
   DNSSEC Signing Policy and Practices Statement is the document to be
   audited against, it is ideal to set the scope to the scope of the
   DPS. However, the scope may be narrowed down or expanded as needed
   for example in case there is not enough resources to conduct a full
   audit, some portion under development and not ready for the audit.

4.7.5. Actions taken as a result of deficiency

   This subcomponent specifies the action taken in order to correct the
   discrepancy. This could be the remediation process for the audit
   findings or any other action to correct the descrepancy with the
   DNSSEC Signing Policy or Practices Statement.

[4.7.6](#). **Communication of results**

[4.8](#). **Legal Matters**

   This component covers legal matters. Sections [9.1](#) and [9.2](#) of the
   framework discuss the business issues of fees to be charged for
   various services and the financial responsibility of participants to
   maintain resources for ongoing operations and for paying judgments or
   settlements in response to claims asserted against them. The
   remaining sections are generally concerned with legal topics.

   With respect to many of the legal subcomponents within this
   component, a DPS drafter may choose to include in the document terms
   and conditions that apply directly to registrants or relying parties.
   For instance, a Registry may set forth limitations of liability that
   apply to registrants and relying parties. The inclusion of terms and
   conditions is likely to be appropriate where the DPS is itself a
   contract or part of a contract.

   In other cases, however, the DPS is not a contract or part of a
   contract; instead, it is configured so that its terms and conditions
   are applied to the parties by separate documents, which may include
   associated agreements, such as subscriber or relying party
   agreements. In that event, a DPS drafter may write a Polict so as to
   require that certain legal terms and conditions appear (or not
   appear) in such associated agreements. For example, a Signing Policy
   might include a subcomponent stating that a certain limitation of
   liability term must appear in a Registrys registrant agreements.
   Another example is a Signing Policy that contains a subcomponent
   prohibiting the use of a subscriber or relying party agreement
   containing a limitation upon Registry liability inconsistent with the
   provisions of the Signing Policy. A DPS drafter may use legal
   subcomponents to disclose that certain terms and conditions appear in
   associated subscriber, relying party, or other agreements in use by
   the Registry. A DPS might explain, for instance, that the Registry
   writing it uses an associated subscriber or relying party agreement
   that applies a particular provision for limiting liability.

[4.8.1](#). **Fees**

   This subcomponent contains any applicable provisions regarding fees
   charged by the Registry for DNSSEC or services related to DNSSEC.

[4.8.2](#). **Financial responsibility**

   This subcomponent contains requirements or disclosures relating to
   the resources available to the Registry, and to remain solvent and
   pay damages in the event they are liable to pay a judgment or

settlement in connection with a claim arising out of such operations.
Such provisions include:

o A statement that the participant maintains a certain amount of
   insurance coverage for its liabilities to other participants;
o A statement that a participant has access to other resources to
   support operations and pay damages for potential liability, which
   may be couched in terms of a minimum level of assets necessary to
   operate and cover contingencies that might occur, and a right
   under an agreement to an indemnity under certain circumstances;
   and
o A statement that a participant has a program that offers first-
   party insurance or warranty protection to other participants in
   connection with their use of the Registry services.

4.8.3. Confidentiality of business information

This subcomponent contains provisions relating to the treatment of
confidential business information. Specifically, this subcomponent
addresses:

o The scope of what is considered confidential information;
o The types of information that are considered to be outside the
   scope of confidential information; and
o The responsibilities of participants that receive confidential
   information to secure it from compromise, and refrain from using
   it or disclosing it to third parties.

4.8.4. Privacy of personal information

This subcomponent relates to the protection that participants,
particularly the Registry, may be required to afford to personally
identifiable private information of registrants and other
participants. Specifically, this subcomponent addresses the
following, to the extent pertinent under applicable law:

o The designation and disclosure of the applicable privacy plan that
   applies to a participant's activities, if required by applicable
   law or policy;
o Information that is or is not considered private within the
   Registry;
o Any responsibility of participants that receive private
   information to secure it, and refrain from using it and from
   disclosing it to third parties;

   o Any requirements as to notices to, or consent from individuals
      regarding use or disclosure of private information; and

   o Any circumstances under which a participant is entitled or
      required to disclose private information pursuant to judicial,
      administrative process in a private or governmental proceeding, or
      in any legal proceeding.

4.8.5. Limitations of liability

   This subcomponent can include limitations of liability in a DPS or
   limitations that appear or must appear in an agreement associated
   with the DPS, such as a subscriber or relying party agreement. These
   limitations may fall into one of two categories: limitations on the
   elements of damages recoverable and limitations on the amount of
   damages recoverable, also known as liability caps. Often, contracts
   contain clauses preventing the recovery of elements of damages such
   as incidental and consequential damages, and sometimes punitive
   damages. Frequently, contracts contain clauses that limit the
   possible recovery of one party or the other to an amount certain or
   to an amount corresponding to a benchmark, such as the amount a
   vendor was paid under the contract.

4.8.6. Term and termination

   This subcomponent can include the time period in which a DPS remains
   in force and the circumstances under which the document, portions of
   the document, or its applicability to a particular participant can be
   terminated. In addition or alternatively, the DPS may include
   requirements that certain term and termination clauses appear in
   agreements, such as subscriber or relying party agreements. In
   particular, such terms can include:

   o The term of a document or agreement, that is, when the document
      becomes effective and when it expires if it is not terminated
      earlier.

   o Termination provisions stating circumstances under which the
      document, certain portions of it, or its application to a
      particular participant ceases to remain in effect.

   o Any consequences of termination of the document. For example,
      certain provisions of an agreement may survive its termination and
      remain in force. Examples include acknowledgements of
      intellectual property rights and confidentiality provisions.
      Also, termination may trigger a responsibility of parties to
      return confidential information to the party that disclosed it.

[5](#). **Security Considerations**

[6](#). **Outline of a set of provisions**

```
    1.  INTRODUCTION
      1.1.  Overview
      1.2.  Document name and identification
      1.3.  Community and Applicability
        1.3.1.  Registry
        1.3.2.  Registrar
        1.3.3.  Registrant
        1.3.4.  Relying Party
        1.3.5   Auditor
        1.3.4.  Applicability
      1.4.  Specification Administration
        1.4.1.  Specification administration organization
        1.4.2.  Contact Information
        1.4.3.  Specification change procedures
    2.  PUBLICATION AND REPOSITORIES
      2.1.  Repositories
      2.2.  Publication of key signing keys
      2.3.  Access controls on repositories
    3.  OPERATIONAL REQUIREMENTS
      3.1.  Meaning of domain names
      3.2.  Activation of DNSSEC for child zone
      3.3.  Identification and authentication of child zone manager
      3.4.  Registration of delegation signer (DS) records
      3.5.  Method to prove possession of private key
      3.6.  Removal of DS record
        3.6.1.  Who can request removal
        3.6.2.  Procedure for removal request
        3.6.3.  Emergency removal request
    4.  FACILITY, MANAGEMENT AND OPERATIONAL CONTROLS
      4.1.  Physical Controls
        4.1.1.  Site location and construction
        4.1.2.  Physical access
        4.1.3.  Power and air conditioning
        4.1.4.  Water exposures
        4.1.5.  Fire prevention and protection
        4.1.6.  Media storage
        4.1.7.  Waste disposal
        4.1.8.  Off-site backup
      4.2.  Procedural Controls
        4.2.1.  Trusted roles
        4.2.2.  Number of persons required per task
        4.2.3.  Identification and authentication for each role
```

7. Acknowledgements

   The authors gratefully acknowledges, in no particular order, the
   contributions of the following persons:
      Richard Lamb
      Jakob Schlyter

8. References

8.1. Normative References

   [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S.
              Rose, "DNS Security Introduction and Requirements",
              RFC 4033, March 2005.

8.2. Informative References

   [RFC3647] Chokhani, S., Ford, W., Sabett, R., Merrill, C., and S.
              Wu, "Internet X.509 Public Key Infrastructure Certificate
              Policy and Certification Practices Framework", RFC 3647,
              November 2003.

   [RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S.
              Rose, "Resource Records for the DNS Security Extensions",
              RFC 4034, March 2005.

   [RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS
              Security (DNSSEC) Hashed Authenticated Denial of
              Existence", RFC 5155, March 2008.

Authors' Addresses

   Fredrik Ljunggren
   Kirei AB
   P.O. Box 53204
   Goteborg SE-400 16
   Sweden

   Email: fredrik@kirei.se

Anne-Marie Eklund-Lowinder
.SE (The Internet Infrastructure Foundation)
P.O. Box 7399
Stockholm SE-103 91
Sweden

Email: amel@iis.se

Tomofumi Okubo
VeriSign Inc.
21345 Ridgetop Circle
Dulles, VA 20166-6503
USA

Email: tookubo@verisign.com