

SUIT
Internet-Draft
Intended status: Standards Track
Expires: September 10, 2020

B. Moran
H. Tschofenig
Arm Limited
March 09, 2020

Strong Assertions of IoT Network Access Requirements
draft-moran-suit-mud-00

Abstract

The Manufacturer Usage Description (MUD) specification describes the access and network functionality required a device to properly function. The MUD description has to reflect the software running on the device and its configuration. Because of this, the most appropriate entity for describing device network access requirements is the same as the entity developing the software and its configuration.

A network presented with a MUD file by a device allows detection of misbehavior by the device software and configuration of access control.

This document defines a way to link a SUIT manifest to a MUD file offering a stronger binding between the two.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 10, 2020.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Architecture	3
4. Extensions to SUIT	4
5. Security Considerations	5
6. IANA Considerations	5
7. Normative References	5
Authors' Addresses	6

1. Introduction

Under [RFC8520], devices report a URL to a MUD manager in the network. RFC 8520 envisions different approaches for conveying the information from the device to the network such as:

- DHCP,
- IEEE802.1AB Link Layer Discovery Protocol (LLDP), and
- IEEE 802.1X whereby the URL to the MUD file would be contained in the certificate used in an EAP method.

The MUD manager then uses the the URL to fetch the MUD file, which contains access and network functionality required a device to properly function.

The MUD manager must trust the service from which the URL is fetched and to return an authentic copy of the MUD file. This concern may be mitigated using the optional signature reference in the MUD file. The MUD manager must also trust the device to report a correct URL. In case of DHCP and LLDP the URL is unprotected. When the URL to the MUD file is included in a certificate then it is authenticated and integrity protected. A certificate created for use with network access authentication is typically not signed by the entity that wrote the software and configured the device, which leads to conflation of local network access rights with rights to assert all network access requirements.

There is a need to bind the entity that creates the software and configuration to the MUD file because only that entity can attest the network access requirements of the device. This specification defines an extension to the SUIT manifest to include a MUD file (per reference or by value). When combining a manufacturer usage description with a manifest used for software/firmware updates (potentially augmented with attestation) then a network operator can get more confidence in the description of the access and network functionality required a device to properly function.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Architecture

The intended workflow is as follows:

- At the time of onboarding, devices report their manifest in use to the MUD manager.
- If the SUIT_MUD_container has been severed, the suit-reference-uri can be used to retrieve the complete manifest.
- The manifest authenticity is verified by the MUD manager, which enforces that the MUD file presented is also authentic and as intended by the device software vendor.

- Each time a device is updated, rebooted, or otherwise substantially changed, it will execute an attestation.
 - o Among other claims in the Entity Attestation Token (EAT) [I-D.ietf-rats-eat], the device will report its software digest(s), configuration digest(s), primary manifest URI, and primary manifest digest to the MUD manager.
 - o The MUD manager can then validate these attestation reports in order to check that the device is operating with the expected version of software and configuration.
 - o Since the manifest digest is reported, the MUD manager can look up the corresponding manifest.
- If the MUD manager does not already have a full copy of the manifest, it can be acquired using the reference URI.
- Once a full copy of the manifest is provided, the MUD manager can verify the device attestation report and apply any appropriate policy as described by the MUD file.

4. Extensions to SUIT

To enable strong assertions about the network access requirements that a device should have for a particular software/configuration pair, we include the ability to add MUD files to the SUIT manifest. However, there are also circumstances in which a device should allow the MUD to be changed without a firmware update. To enable this, we add a MUD url to SUIT along with a subject-key identifier, according to [RFC7093], mechanism 4 (the keyIdentifier is composed of the hash of the DER encoding of the SubjectPublicKeyInfo value).

The following CDDL describes the extension to the SUIT_Manifest structure:

```
? suit-manifest-mud => SUIT_Digest
```

The SUIT_Envelope is also amended:

```
? suit-manifest-mud => bstr .cbor SUIT_MUD_container
```

```
SUIT_MUD_container = {
  ? suit-mud-url => #6.32(tstr),
  ? suit-mud-ski => SUIT_Digest,
  ? suit-mud-file => bstr
}
```

The MUD file is included verbatim within the bstr. No limits are placed on the MUD file: it may be any RFC8520-compliant file.

5. Security Considerations

This specification links MUD files to other IETF technologies, particularly to SUIT manifests, for improving security protection and ease of use. By including MUD files (per reference or by value) in SUIT manifests an extra layer of protection has been created and synchronization risks can be minimized. If the MUD file and the software/firmware loaded onto the device gets out-of-sync a device may be firewalled and, with firewalling by networks in place, the device may stop functioning.

6. IANA Considerations

suit-manifest-mud must be added as an extension point to the SUIT manifest registry.

7. Normative References

[I-D.ietf-rats-eat]

Mandyam, G., Lundblade, L., Ballesteros, M., and J. O'Donoghue, "The Entity Attestation Token (EAT)", draft-ietf-rats-eat-03 (work in progress), February 2020.

[I-D.ietf-suit-manifest]

Moran, B., Tschofenig, H., and H. Birkholz, "A Concise Binary Object Representation (CBOR)-based Serialization Format for the Software Updates for Internet of Things (SUIT) Manifest", draft-ietf-suit-manifest-03 (work in progress), February 2020.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC7093] Turner, S., Kent, S., and J. Manger, "Additional Methods for Generating Key Identifiers Values", RFC 7093, DOI 10.17487/RFC7093, December 2013, <<https://www.rfc-editor.org/info/rfc7093>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

Authors' Addresses

Brendan Moran
Arm Limited

E-Mail: Brendan.Moran@arm.com

Hannes Tschofenig
Arm Limited

E-Mail: hannes.tschofenig@arm.com

SUIT
Internet-Draft
Intended status: Standards Track
Expires: November 26, 2021

B. Moran
H. Tschofenig
Arm Limited
May 25, 2021

Strong Assertions of IoT Network Access Requirements
draft-moran-suit-mud-02

Abstract

The Manufacturer Usage Description (MUD) specification describes the access and network functionality required a device to properly function. The MUD description has to reflect the software running on the device and its configuration. Because of this, the most appropriate entity for describing device network access requirements is the same as the entity developing the software and its configuration.

A network presented with a MUD file by a device allows detection of misbehavior by the device software and configuration of access control.

This document defines a way to link a SUIT manifest to a MUD file offering a stronger binding between the two.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 26, 2021.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Architecture	3
4. Extensions to SUIT	4
5. Security Considerations	5
6. IANA Considerations	5
7. Normative References	5
Authors' Addresses	6

1. Introduction

Under [RFC8520], devices report a URL to a MUD manager in the network. RFC 8520 envisions different approaches for conveying the information from the device to the network such as:

- DHCP,
- IEEE802.1AB Link Layer Discovery Protocol (LLDP), and
- IEEE 802.1X whereby the URL to the MUD file would be contained in the certificate used in an EAP method.

The MUD manager then uses the the URL to fetch the MUD file, which contains access and network functionality required a device to properly function.

The MUD manager must trust the service from which the URL is fetched and to return an authentic copy of the MUD file. This concern may be mitigated using the optional signature reference in the MUD file. The MUD manager must also trust the device to report a correct URL. In case of DHCP and LLDP the URL is unprotected. When the URL to the MUD file is included in a certificate then it is authenticated and integrity protected. A certificate created for use with network access authentication is typically not signed by the entity that wrote the software and configured the device, which leads to conflation of local network access rights with rights to assert all network access requirements.

There is a need to bind the entity that creates the software and configuration to the MUD file because only that entity can attest the network access requirements of the device. This specification defines an extension to the SUIT manifest to include a MUD file (per reference or by value). When combining a manufacturer usage description with a manifest used for software/firmware updates (potentially augmented with attestation) then a network operator can get more confidence in the description of the access and network functionality required a device to properly function.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Architecture

The intended workflow is as follows:

- At the time of onboarding, devices report their manifest in use to the MUD manager.
- If the SUIT_MUD_container has been severed, the suit-reference-uri can be used to retrieve the complete manifest.
- The manifest authenticity is verified by the MUD manager, which enforces that the MUD file presented is also authentic and as intended by the device software vendor.

- Each time a device is updated, rebooted, or otherwise substantially changed, it will execute an attestation.
 - o Among other claims in the Entity Attestation Token (EAT) [I-D.ietf-rats-eat], the device will report its software digest(s), configuration digest(s), primary manifest URI, and primary manifest digest to the MUD manager.
 - o The MUD manager can then validate these attestation reports in order to check that the device is operating with the expected version of software and configuration.
 - o Since the manifest digest is reported, the MUD manager can look up the corresponding manifest.
- If the MUD manager does not already have a full copy of the manifest, it can be acquired using the reference URI.
- Once a full copy of the manifest is provided, the MUD manager can verify the device attestation report and apply any appropriate policy as described by the MUD file.

4. Extensions to SUIT

To enable strong assertions about the network access requirements that a device should have for a particular software/configuration pair, we include the ability to add MUD files to the SUIT manifest. However, there are also circumstances in which a device should allow the MUD to be changed without a firmware update. To enable this, we add a MUD url to SUIT along with a subject-key identifier, according to [RFC7093], mechanism 4 (the keyIdentifier is composed of the hash of the DER encoding of the SubjectPublicKeyInfo value).

The following CDDL describes the extension to the SUIT_Manifest structure:

```
? suit-manifest-mud => SUIT_Digest
```

The SUIT_Envelope is also amended:

```
? suit-manifest-mud => bstr .cbor SUIT_MUD_container
```

```
SUIT_MUD_container = {  
  ? suit-mud-url => #6.32(tstr),  
  ? suit-mud-ski => SUIT_Digest,  
  ? suit-mud-file => bstr  
}
```

The MUD file is included verbatim within the bstr. No limits are placed on the MUD file: it may be any RFC8520-compliant file.

5. Security Considerations

This specification links MUD files to other IETF technologies, particularly to SUIT manifests, for improving security protection and ease of use. By including MUD files (per reference or by value) in SUIT manifests an extra layer of protection has been created and synchronization risks can be minimized. If the MUD file and the software/firmware loaded onto the device gets out-of-sync a device may be firewalled and, with firewalling by networks in place, the device may stop functioning.

6. IANA Considerations

suit-manifest-mud must be added as an extension point to the SUIT manifest registry.

7. Normative References

[I-D.ietf-rats-eat]

Mandyam, G., Lundblade, L., Ballesteros, M., and J. O'Donoghue, "The Entity Attestation Token (EAT)", draft-ietf-rats-eat-09 (work in progress), March 2021.

[I-D.ietf-suit-manifest]

Moran, B., Tschofenig, H., Birkholz, H., and K. Zandberg, "A Concise Binary Object Representation (CBOR)-based Serialization Format for the Software Updates for Internet of Things (SUIT) Manifest", draft-ietf-suit-manifest-12 (work in progress), February 2021.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.

[RFC7093] Turner, S., Kent, S., and J. Manger, "Additional Methods for Generating Key Identifiers Values", RFC 7093, DOI 10.17487/RFC7093, December 2013, <<https://www.rfc-editor.org/info/rfc7093>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

Authors' Addresses

Brendan Moran
Arm Limited

E-Mail: Brendan.Moran@arm.com

Hannes Tschofenig
Arm Limited

E-Mail: hannes.tschofenig@arm.com

anima Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 January 2021

M. Richardson
Sandelman Software Works
J. Yang
Huawei Technologies Co., Ltd.
27 July 2020

Security and Operational considerations for manufacturer installed keys
and anchors
draft-richardson-secdispatch-idevid-considerations-02

Abstract

This document provides a nomenclature to describe ways in which manufacturers secure private keys and public trust anchors in devices.

RFCEditor: please remove this paragraph. This work is occurring in <https://github.com/mcr/idevid-security-considerations>

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 28 January 2021.

Copyright Notice

Copyright (c) 2020 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Terminology	4
2.	Applicability Model	5
2.1.	A reference manufacturing/boot process	6
3.	Types of Trust Anchors	7
3.1.	Secured First Boot Trust Anchor	8
3.2.	Software Update Trust Anchor	8
3.3.	Trusted Application Manager anchor	8
3.4.	Public WebPKI anchors	9
3.5.	DNSSEC root	9
3.6.	What else?	9
4.	Types of Identities	9
4.1.	Manufacturer installed IDevID certificates	10
4.1.1.	Operational Considerations for Manufacturer IDevID Public Key Infrastructure	10
4.1.2.	Key Generation process	11
5.	Public Key Infrastructures (PKI)	14
5.1.	Number of levels of certification authorities	14
5.2.	Protection of CA private keys	15
5.3.	Supporting provisioned anchors in devices	16
6.	Evaluation Questions	16
6.1.	Integrity and Privacy of on-device data	16
6.2.	Integrity and Privacy of device identify infrastructure	17
6.3.	Integrity and Privacy of included trust anchors	17
7.	Privacy Considerations	18
8.	Security Considerations	18
9.	IANA Considerations	18
10.	Acknowledgements	18
11.	Changelog	18
12.	References	18
12.1.	Normative References	18
12.2.	Informative References	19
	Authors' Addresses	23

1. Introduction

An increasing number of protocols derive a significant part of their security by using trust anchors that are installed by manufacturers. Disclosure of the list of trust anchors does not usually cause a problem, but changing them in any way does. This includes adding, replacing or deleting anchors.

Many protocols also leverage manufacturer installed identities. These identities are usually in the form of [ieee802-1AR] Initial Device Identity certificates (IDevID). The identity has two components: a private key that must remain under the strict control of a trusted part of the device, and a public part (the certificate), which (ignoring, for the moment, personal privacy concerns) may be freely disclosed.

There also situations where identities which are tied up in the provision of symmetric shared secret. A common example is the SIM card ([_3GPP.51.011]), which now comes as a virtual SIM, but which is usually not provisioned at the factory. The provision of an initial, per-device default password also falls into the category of symmetric shared secret.

It is further not unusual for many devices (particularly smartphones) to also have one or more group identity keys. This is used in, for instance, in [fidotechnote] to make claims about being a particular model of phone (see [I-D.richardson-rats-usecases]). The keypair that does this is loaded into large batches of phones for privacy reasons.

The trust anchors are used for a variety of purposes. The following uses are specifically called out:

- * to validate the signature on a software update (as per [I-D.ietf-suit-architecture]).
- * to verify the end of a TLS Server Certificate, such as when setting up an HTTPS connection.
- * to verify the [RFC8366] format voucher that provides proof of an ownership change

Device identity keys are used when performing enrollment requests (in [I-D.ietf-anima-bootstrapping-keyinfra], and in some uses of [I-D.ietf-emu-eap-noob]). The device identity certificate is also used to sign Evidence by an Attesting Environment (see [I-D.ietf-rats-architecture]).

These security artifacts are used to anchor other chains of information: an EAT Claim as to the version of software/firmware running on a device (XXX and [I-D.birkholz-suit-coswid-manifest]), an EAT claim about legitimate network activity (via [I-D.birkholz-rats-mud], or embedded in the IDevID in [RFC8520]). Known software versions lead directly to vendor/distributor signed Software Bill of Materials (SBOM), such as those described by [I-D.ietf-sacm-coswid] and the NTIA/SBOM work [ntiasbom] and CISQ/OMG SBOM work underway [cisqsbom].

In order to manage risks and assess vulnerabilities in a Supply Chain, it is necessary to determine a degree of trustworthiness in each device. A device may mislead audit systems as to its provenance, about its software load or even about what kind of device it is. (see [RFC7168] for a humorous example). In order to properly assess the security of a Supply Chain it is necessary to understand the kinds and severity of the threats which a device has been designed to resist. To do this, it is necessary to understand the ways in which the different trust anchors and identities are initially provisioned, are protected, and are updated.

To do this, this document details the different trust anchors (TrA) and identities (IDs) found in typical devices. The privacy and integrity of the TAs and IDs is often provided by a different, superior artifacts. This relationship is examined.

While many might desire to assign numerical values to different mitigation techniques in order to be able to rank them, this document does not attempt to do, as there are too many other (mostly human) factors that would come into play. Such an effort is more properly in the purvue of a formal ISO9001 process such as ISO14001.

1.1. Terminology

This document is not a standards track document, and it does not make use of formal requirements language.

This section will be expanded to include needed terminology as required.

The words Trust Anchor are contracted to TrA rather than TA, in order not to confuse with [I-D.ietf-teep-architecture]'s "Trusted Application".

This document defines a number of hyphenated terms, and they are summarized here:

device-generated : a private or symmetric key which is generated on

the device

infrastructure-generated : a private or symmetric key which is generated by some system, likely located at factory that built the device

mechanically-installed : when a key or certificate is programmed into flash by out-of-band mechanism like JTAG

mechanically-transferred : when a key or certificate is transferred into a system via private interface, such as serial console, JTAG managed mailbox, or other physically private interface

network-transferred : when a key or certificate is transferred into a system using a network interface which would be available after the device has shipped. This applies even if the network is physically attached using a bed-of-nails.

device/infrastructure-co-generated : when a private or symmetric key is derived from a secret previously synchronized between the silicon vendor and the factory using a common algorithm.

2. Applicability Model

There is a wide variety of devices to which this analysis can apply. (See [I-D.bormann-lwig-7228bis]) This document will use a J-group class C13 as a sample. This class is sufficiently large to experience complex issues among multiple CPUs, packages and operating systems, but at the same time, small enough that this class is often deployed in single-purpose IoT-like uses. Devices in this class often have Secure Enclaves (such as the "Grapeboard"), and can include silicon manufacturer controlled processors in the boot process (the Raspberry PI boots under control of the GPU).

Almost all larger systems (servers, laptops, desktops) include a Baseboard Management Controller (BMC), which ranges from a M-Group Class 3 MCU, to a J-Group Class 10 CPU (see, for instance [openbmc] which uses a Linux kernel and system inside the BMC). As the BMC usually has complete access to the main CPU's memory, I/O hardware and disk, the boot path security of such a system needs to be understood first as being about the security of the BMC.

2.1. A reference manufacturing/boot process

In order to provide for immutability and privacy of the critical TANs and IDs, many CPU manufacturers will provide for some kind of private memory area which is only accessible when the CPU is in certain privileged states. See the Terminology section of [I-D.ietf-teep-architecture], notably TEE, REE, and TAM, and also section 4, Architecture.

The private memory that is important is usually non-volatile and rather small. It may be located inside the CPU silicon die, or it may be located externally. If the memory is external, then it is usually encrypted by a hardware mechanism on the CPU, with only the key kept inside the CPU.

The entire mechanism may be external to the CPU in the form of a hardware-TPM module, or it may be entirely internal to the CPU in the form of a firmware-TPM. It may use a custom interface to the rest of the system, or it may implement the TPM 1.2 or TPM 2.0 specifications. Those details are important to performing a full evaluation, but do not matter that to this model (see initial-enclave-location below).

During the manufacturing process, once the components have been soldered to the board, the system is usually put through a system-level test. This is often done on as a "bed-of-nails" test [BedOfNails], where the board has key points attached mechanically to a test system. A [JTAG] process tests the System Under Test, and then initializes some firmware into the still empty flash storage. It is now common for a factory test image to be loaded first: this image will include code to initialize the private memory key described above, and will include a first-stage bootloader and some kind of (primitive) Trusted Application Manager (TAM). Embedded in the stage one bootloader will be a Trust Anchor that is able to verify the second-stage bootloader image.

After the system has undergone testing, the factory test image is erased, leaving the first-stage bootloader. One or more second-stage bootloader images is installed. The production image may be installed at that time, or if the second-stage bootloader is able to install it over the network, it may be done that way instead.

There are many variations of the above process, and this section is not attempting to be prescriptive, but to provide enough illustration to motivate subsequent terminology.

There process may be entirely automated, or it may be entirely driven by humans working in the factory.

Or a combination of the above.

These steps may all occur on an access-controlled assembly line, or the system boards may be shipped from one place to another (maybe another country) before undergoing testing.

Some systems are intended to be shipped in a tamper-proof state, but it is usually not desirable that bed-of-nails testing be possible without tampering, so the initialization process is usually done prior to rendering the system tamper-proof.

Quality control testing may be done prior to as well as after the application of tamper-proofing, as systems which do not pass inspection may be reworked to fix flaws, and this should ideally be impossible once the system has been made tamper-proof.

3. Types of Trust Anchors

Trust Anchors are fundamentally public keys. They are used to validate other digitally signed artifacts. Typically, these are chains of PKIX certificates leading to an End-Entity certificate (EE).

The chains are usually presented as part of an externally provided object, with the term "externally" to be understood as being as close as untrusted flash, to as far as objects retrieved over a network.

There is no requirement that there be any chain at all: the trust anchor can be used to validate a signature over a target object directly.

The trust anchors are often stored in the form of self-signed certificates. The self-signature does not offer any cryptographic assurance, but it does provide a form of error detection, providing verification against non-malicious forms of data corruption. If storage is at a premium (such as inside-CPU non-volatile storage) then only the public key itself need to be stored. For a 256-bit ECDSA key, this is 32-bytes of space.

When evaluating the degree of trust for each trust anchor there are four aspects that need to be determined:

- * can the trust anchor be replaced or modified?
- * can additional trust anchors be added?
- * can trust anchors be removed?

* how is the private key associated with the trust anchor stored?

The first three things are device specific properties of how the integrity of the trust anchor is maintained.

The fourth property has nothing to do with the device, but has to do with the reputation and care of the entity that maintains the private key.

Different anchors have different purposes.

These are:

3.1. Secured First Boot Trust Anchor

This anchor is part of the first-stage boot loader, and it is used to validate a second-stage bootloader which may be stored in external flash.

3.2. Software Update Trust Anchor

This anchor is used to validate the main application (or operating system) load for the device.

It can be stored in a number of places. First, it may be identical to the Secure Boot Trust Anchor.

Second, it may be stored in the second-stage bootloader, and therefore it's integrity is protected by the Secured First Boot Trust Anchor.

Third, it may be stored in the application code itself, where the application validates updates to the application directly (update in place), or via a double-buffer arrangement. The initial (factory) load of the application code initializes the trust arrangement.

In this situation the application code is not in a secured boot situation, as the second-stage bootloader does not validate the application/operating system before starting it, but it may still provide measured boot mechanism.

3.3. Trusted Application Manager anchor

This anchor is part of a [I-D.ietf-teep-architecture] Trusted Application Manager. Code which is signed by this anchor will be given execution privileges as described by the manifest which accompanies the code. This privilege may include updating anchors.

3.4. Public WebPKI anchors

These anchors are used to verify HTTPS certificates from web sites. These anchors are typically distributed as part of desktop browsers, and via desktop operating systems.

The exact set of these anchors is not precisely defined: it is usually determined by the browser vendor (e.g., Mozilla, Google, Apple, Safari, Microsoft), or the operating system vendor (e.g., Apple, Google, Microsoft, Ubuntu). In most cases these vendors look to the CA/Browser Forum ([CABFORUM]) for inclusion criteria.

3.5. DNSSEC root

This anchor is part of the DNS Security extensions. It provides an anchor for securing DNS lookups. Secure DNS lookups may be important in order to get access to software updates. This anchor is now scheduled to change approximately every 3 years, with the new key announced several years before it is used, making it possible to embed a keys that will be valid for up to five years.

This trust anchor is typically part of the application/operating system code and is usually updated by the manufacturer when they do updates. However, a system which is connected to the Internet may update the DNSSEC anchor itself through the mechanism described in [RFC5011].

There are concerns that there may be a chicken and egg situation for devices have remained in a powered off state (or disconnected from the Internet) for some period of years. That upon being reconnected, that the device would be unable to do DNSSEC validation. This failure would result in them being unable to obtain operating system updates that would then include the updates to the DNSSEC key.

3.6. What else?

TBD?

4. Types of Identities

Identities are installed during manufacturing time for a variety of purposes.

Identities require some private component. Asymmetric identities (e.g., RSA, ECDSA, EdDSA systems) require a co-responding public component, usually in the form of a certificate signed by a trusted third party.

The process of making this coordinated key pair and then installing it into the device is called identity provisioning.

4.1. Manufacturer installed IDevID certificates

[ieee802-1AR] defines a category of certificates that are to be installed by the manufacturer, which contain at the least, a device unique serial number.

A number of protocols depend upon this certificate.

- * [I-D.ietf-anima-bootstrapping-keyinfra] introduces a mechanism for new devices (called pledges) to be onboarded into a network without intervention from an expert operator. A number of derived protocols such as {{I-D.
- * [I-D.ietf-rats-architecture] depends upon a key provisioned into the Attesting Environment to sign Evidence.
- * [I-D.ietf-suit-architecture] may depend upon a key provisioned into the device in order to decrypt software updates.
- * TBD

4.1.1. Operational Considerations for Manufacturer IDevID Public Key Infrastructure

The manufacturer has the responsibility to provision a keypair into each device as part of the manufacturing process. There are a variety of mechanisms to accomplish this, which this document will overview.

There are three fundamental ways to generate IDevID certificates for devices:

1. generating a private key on the device, creating a Certificate Signing Request (or equivalent), and then returning a certificate to the device.
2. generating a private key outside the device, signing the certificate, and the installing both into the device.
3. deriving the private key from a previously installed secret seed, that is shared with only the manufacturer

There is a fourth situation where the IDevID is provided as part of a Trusted Platform Module (TPM), in which case the TPM vendor may be making the same tradeoffs.

The document [I-D.moskowitz-ecdsa-pki] provides some practical instructions on setting up a reference implementation for ECDSA keys using a three-tier mechanism.

This document recommends the use of ECDSA keys for the root and intermediate CAs, but there may be operational reasons why an RSA intermediate CA will be required for some legacy TPM equipment.

4.1.2. Key Generation process

4.1.2.1. On-device private key generation

Generating the key on-device has the advantage that the private key never leaves the device. The disadvantage is that the device may not have a verified random number generator. [factoringrsa] is an example of this scenario!

There are a number of options of how to get the public key securely from the device to the certification authority.

This transmission must be done in an integral manner, and must be securely associated with the assigned serial number. The serial number goes into the certificate, and the resulting certificate needs to be loaded into the manufacturer's asset database. This asset database needs to be shared with the MASA.

One way to do the transmission is during a factory Bed of Nails test (see [BedOfNails]) or Boundary Scan. When done via a physical connection like this, then this is referred to as a `_device-generated_ / _mechanically-transferred_`.

There are other ways that could be used where a certificate signing request is sent over a special network channel when the device is powered up in the factory. This is referred to as the `_device-generated_ / _network-transferred_` method.

Regardless of how the certificate signing request is sent from the device to the factory, and the certificate is returned to the device, a concern from production line managers is that the assembly line may have to wait for the certification authority to respond with the certificate.

After the key generation, the device needs to set a flag such that it no longer generates a new key, or will accept a new IDeVID via the factory connection. This may be a software setting, or could be as dramatic as blowing a fuse.

The risk is that if an attacker with physical access is able to put the device back into an unconfigured mode, then the attacker may be able to substitute a new certificate into the device. It is difficult to construct a rationale for doing this, unless the network initialization also permits an attacker to load or replace trust anchors at the same time.

Because the key is generated inside the device, it is assumed that the device can never be convinced to disclose the private key.

4.1.2.2. Off-device private key generation

Generating the key off-device has the advantage that the randomness of the private key can be better analyzed. As the private key is available to the manufacturing infrastructure, the authenticity of the public key is well known ahead of time.

If the device does not come with a serial number in silicon, then one should be assigned and placed into a certificate. The private key and certificate could be programmed into the device along with the initial bootloader firmware in a single step.

Aside from the change of origin for the randomness, a major advantage of this mechanism is that it can be done with a single write to the flash. The entire firmware of the device, including configuration of trust anchors and private keys can be loaded in a single write pass. Given some pipelining of the generation of the keys and the creation of certificates, it may be possible to install unique identities without taking any additional time.

The major downside to generating the private key off-device is that it could be seen by the manufacturing infrastructure. It could be compromised by humans in the factory, or the equipment could be compromised. The use of this method increases the value of attacking the manufacturing infrastructure.

If keys are generated by the manufacturing plant, and are immediately installed, but never stored, then the window in which an attacker can gain access to the private key is immensely reduced.

As in the previous case, the transfer may be done via physical interfaces such as bed-of-nails, giving the `_infrastructure-generated_ / _mechanically-transferred_` method.

There is also the possibility of having a `_infrastructure-generated_ / _network-transferred_` key. There is support for "server-generated" keys in [RFC7030], [I-D.gutmann-scep], and [RFC4210]. All methods strongly recommend encrypting the private key for transfer.

This is difficult to comply with as there is not yet any private key material in the device, so in many cases it will not be possible to encrypt the private key.

4.1.2.3. Key setup based on 256-bit secret seed

A hybrid of the previous two methods leverages a symmetric key that is often provided by a silicon vendor to OEM manufacturers.

Each CPU (or a Trusted Execution Environment [I-D.ietf-tee-architecture], or a TPM) is provisioned at fabrication time with a unique, secret seed, usually at least 256-bits in size.

This value is revealed to the OEM board manufacturer only via a secure channel. Upon first boot, the system (probably within a TEE, or within a TPM) will generate a key pair using the seed to initialize a Pseudo-Random-Number-Generator (PRNG). The OEM, in a separate system, will initialize the same PRNG and generate the same key pair. The OEM then derives the public key part, signs it and turns it into a certificate. The private part is then destroyed, ideally never stored or seen by anyone. The certificate (being public information) is placed into a database, in some cases it is loaded by the device as its IDevID certificate, in other cases, it is retrieved during the onboarding process based upon a unique serial number asserted by the device.

This method appears to have all of the downsides of the previous two methods: the device must correctly derive its own private key, and the OEM has access to the private key, making it also vulnerable. The secret seed must be created in a secure way and it must also be communicated securely.

There are some advantages to the OEM however: the major one is that the problem of securely communicating with the device is outsourced to the silicon vendor. The private keys and certificates may be calculated by the OEM asynchronously to the manufacturing process, either done in batches in advance of actual manufacturing, or on demand when an IDevID is demanded. Doing the processing in this way permits the key derivation system to be completely disconnected from any network, and requires placing very little trust in the system assembly factory. Operational security such as often incorrectly presented fictionalized stories of a "mainframe" system to which only physical access is permitted begins to become realistic. That trust has been replaced with a heightened trust placed in the silicon (integrated circuit) fabrication facility.

The downsides of this method to the OEM are: they must be supplied by a trusted silicon fabrication system, which must communicate the set of secrets seeds to the OEM in batches, and they OEM must store and care for these keys very carefully. There are some operational advantages to keeping the secret seeds around in some form, as the same secret seed could be used for other things. There are some significant downsides to keeping that secret seed around.

5. Public Key Infrastructures (PKI)

[RFC5280] describes the format for certificates, and numerous mechanisms for doing enrollment have been defined (including: EST: [RFC7030], CMP: [RFC4210], SCEP [I-D.gutmann-scep]).

[RFC5280] provides mechanisms to deal with multi-level certification authorities, but it is not always clear what operating rules apply.

PKIs can suffer two kinds of failures: 1. disclosure of a private key. 2. loss of a private key.

A PKI which discloses one or more private certification authority keys is no longer secure. An attacker can create new identities, and forge certificates connecting existing identities to attacker controlled public/private keypairs. This can permit the attacker to impersonate the specific device.

If the PKI uses Certificate Revocation Lists (CRL)s, then an attacker can also revoke existing identities.

In the other direction, a PKI which loses access to a private key can no longer function. Existing identities may continue to function, unless CRLs or OCSP is in use, in which case, the inability to sign a fresh CRL or OCSP response will result in all identities becoming invalid.

This section details some nomenclature about the structure of certification authorities.

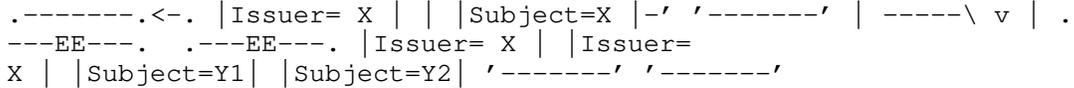
5.1. Number of levels of certification authorities

The certification authority (CA) starts with a Trust Anchor (TA). This is counted as the zeroth level of the authority.

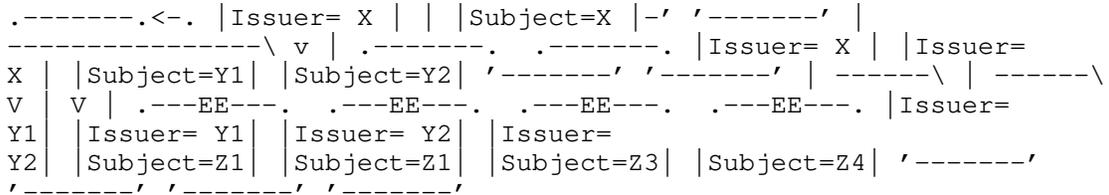
In the degenerate case of a self-signed certificate, then this a level zero PKI.

```
.-----<-. |Issuer= X | | |Subject=X |- ' '-----'
```

The Trust Anchor signs one or more certificates. When this first level authority trusts only End-Entity (EE) certificates, then this is a level one PKI.



When this first level authority signs intermediate certification authorities, and those certification authorities sign End-Entity certificates, then this is a level two PKI.



In general, when arranged as a tree, with the End-Entity certificates at the bottom, and the Trust Anchor at the top, then the level is where the deepest EE certificates are, counting from zero.

It is quite common to have a two-level PKI, where the root of the CA is stored in a Hardware Security Module, while the level one intermediate CA is available online.

5.2. Protection of CA private keys

The private key for the certification authorities must be protected from disclosure. The strongest protection is afforded by keeping them in a offline device, passing Certificate Signing Requests (CSR)s to the offline device by human process.

For examples of extreme measures, see [kskceremony]. There is however a wide spectrum of needs, as exemplified in [rootkeyceremony]. The SAS70 audit standard is usually used as a basis for the Ceremony, see [keyceremony2].

This is inconvenient, and may involve latencies of days, possibly even weeks to months if the offline device is kept in a locked environment that requires multiple keys to be present.

There is therefore a tension between protection and convenience. This is often accomplished by having some levels of the PKI be offline, and some levels of the PKI be online.

There is usually a need to maintain backup copies of the critical keys. It is often appropriate to use secret splitting technology such as Shamir Secret Sharing among a number of parties [shamir79] This mechanism can be setup such that some threshold k (less than the total n) of shares are needed in order to recover the secret.

5.3. Supporting provisioned anchors in devices

IDevID-type Identity (or Birth) Certificates which are provisioned into devices need to be signed by a certification authority maintained by the manufacturer. The manufacturer needs to maintain availability of this PKI.

Trust anchors which are provisioned in the devices will have corresponding private keys maintained by the manufacturer. The trust anchors will often anchor a PKI which is going to be used for a particular purpose. There will be End-Entity (EE) certificates of this PKI which will be used to sign particular artifacts (such as software updates), or communications protocols (such as TLS connections). The private key associated with these EE certificates are not stored in the device, but are maintained by the manufacturer. These need even more care than the private keys stored in the devices, as compromise of the software update key compromises all of the devices, not just a single device.

6. Evaluation Questions

This section recaps the set of questions that may need to be answered. This document does not assign valuation to the answers.

6.1. Integrity and Privacy of on-device data

`initial-enclave-location` : Is the location of the initial software trust anchor internal to the CPU package?

`initial-enclave-integrity-key` : If the first-stage bootloader is external to the CPU, and it is integrity protected, where is the key used to check the integrity?

`initial-enclave-privacy-key` : If the first-stage data is external to the CPU, is it encrypted?

`first-stage-initialization` : The number of people involved in the

first stage initialization. An entirely automated system would have a number zero. A factory with three 8 hour shifts might have a number that is a multiple of three. A system with humans involved may be subject to bribery attacks, while a system with no humans may be subject to attacks on the system which are hard to notice.

first-second-stage-gap : If a board is initialized with a first-stage bootloader in one location (factory), and then shipped to another location, there may situations where the device can not be locked down until the second step.

6.2. Integrity and Privacy of device identify infrastructure

For IDevID provisioning, which includes a private key and matching certificate installed into the device, the associated public key infrastructure that anchors this identity must be maintained by the manufacturer.

identity-pki-level : how deep are the IDevID certificates that are issued?

identity-time-limits-per-intermediate : how long is each intermediate CA maintained before before a new intermediate CA key is generated? There may be no time limit, only a device count limit.

identity-number-per-intermediate : how many identities are signed by a particular intermediate CA before it is retired? There may be no numeric limit, only a time limit.

identity-anchor-storage : how is the root CA key stored? How many people are needed to recover it?

6.3. Integrity and Privacy of included trust anchors

For each trust anchor (public key) stored in the device, there will be an associated PKI. For each of those PKI the following questions need to be answered.

pki-level : how deep is the EE that will be evaluated

pki-level-locked : (a boolean) is the level where the EE cert will be found locked by the device, or can levels be added or deleted by the PKI operator without code changes to the device.

pki-breadth : how many different EE certificates exist in this PKI

pki-lock-policy : can any EE certificate be used with this trust anchor to sign? Or, is there some kind of policy OID or Subject restriction? Are specific intermediate CAs needed that lead to the EE?

pki-anchor-storage: how is the root CA stored? How many people are needed to recover it?

7. Privacy Considerations

many yet to be detailed

8. Security Considerations

This entire document is a security considerations.

9. IANA Considerations

This document makes no IANA requests.

10. Acknowledgements

Robert Martin of MITRE provides some guidance about citing the SBOM efforts.

11. Changelog

12. References

12.1. Normative References

[BCP14] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, DOI 10.17487/RFC5280, May 2008, <<https://www.rfc-editor.org/info/rfc5280>>.

[I-D.moskowitz-ecdsa-pki]
Moskowitz, R., Birkholz, H., Xia, L., and M. Richardson, "Guide for building an ECC pki", Work in Progress, Internet-Draft, draft-moskowitz-ecdsa-pki-08, 14 February 2020, <<http://www.ietf.org/internet-drafts/draft-moskowitz-ecdsa-pki-08.txt>>.

[ieee802-1AR]

IEEE Standard, ., "IEEE 802.1AR Secure Device Identifier", 2009, <<http://standards.ieee.org/findstds/standard/802.1AR-2009.html>>.

12.2. Informative References

[I-D.richardson-anima-masa-considerations]

Richardson, M. and W. Pan, "Operational Considerations for Voucher infrastructure for BRSKI MASA", Work in Progress, Internet-Draft, draft-richardson-anima-masa-considerations-04, 9 June 2020, <<http://www.ietf.org/internet-drafts/draft-richardson-anima-masa-considerations-04.txt>>.

[I-D.ietf-anima-bootstrapping-keyinfra]

Pritikin, M., Richardson, M., Eckert, T., Behringer, M., and K. Watsen, "Bootstrapping Remote Secure Key Infrastructures (BRSKI)", Work in Progress, Internet-Draft, draft-ietf-anima-bootstrapping-keyinfra-41, 8 April 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-anima-bootstrapping-keyinfra-41.txt>>.

[RFC5011] StJohns, M., "Automated Updates of DNS Security (DNSSEC) Trust Anchors", STD 74, RFC 5011, DOI 10.17487/RFC5011, September 2007, <<https://www.rfc-editor.org/info/rfc5011>>.

[RFC8366] Watsen, K., Richardson, M., Pritikin, M., and T. Eckert, "A Voucher Artifact for Bootstrapping Protocols", RFC 8366, DOI 10.17487/RFC8366, May 2018, <<https://www.rfc-editor.org/info/rfc8366>>.

[RFC7030] Pritikin, M., Ed., Yee, P., Ed., and D. Harkins, Ed., "Enrollment over Secure Transport", RFC 7030, DOI 10.17487/RFC7030, October 2013, <<https://www.rfc-editor.org/info/rfc7030>>.

[I-D.gutmann-scep]

Gutmann, P., "Simple Certificate Enrolment Protocol", Work in Progress, Internet-Draft, draft-gutmann-scep-16, 27 March 2020, <<http://www.ietf.org/internet-drafts/draft-gutmann-scep-16.txt>>.

[RFC4210] Adams, C., Farrell, S., Kause, T., and T. Mononen, "Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)", RFC 4210, DOI 10.17487/RFC4210, September 2005, <<https://www.rfc-editor.org/info/rfc4210>>.

- [_3GPP.51.011]
3GPP, "Specification of the Subscriber Identity Module - Mobile Equipment (SIM-ME) interface", 3GPP TS 51.011 4.15.0, 15 June 2005, <<http://www.3gpp.org/ftp/Specs/html-info/51011.htm>>.
- [BedOfNails]
Wikipedia, ., "Bed of nails tester", July 2020, <https://en.wikipedia.org/wiki/In-circuit_test#Bed_of_nails_tester>.
- [pelionfcu]
ARM Pelion, "Factory provisioning overview", 28 June 2020, <<https://www.pelion.com/docs/device-management-provision/1.2/introduction/index.html>>.
- [factoringrsa]
"Factoring RSA keys from certified smart cards: Coppersmith in the wild", 16 September 2013, <<https://core.ac.uk/download/pdf/204886987.pdf>>.
- [RambusCryptoManager]
Qualcomm press release, "Qualcomm Licenses Rambus CryptoManager Key and Feature Management Security Solution", 2014, <<https://www.rambus.com/qualcomm-licenses-rambus-cryptomanager-key-and-feature-management-security-solution/>>.
- [kskceremony]
Verisign, "DNSSEC Practice Statement for the Root Zone ZSK Operator", 2017, <<https://www.iana.org/dnssec/dps/zsk-operator/dps-zsk-operator-v2.0.pdf>>.
- [rootkeyceremony]
Community, "Root Key Ceremony, Cryptography Wiki", April 2020, <https://cryptography.fandom.com/wiki/Root_Key_Ceremony>.
- [keyceremony2]
Digi-Sign, "SAS 70 Key Ceremony", April 2020, <<http://www.digi-sign.com/compliance/key%20ceremony/index>>.
- [shamir79] Shamir, A., "How to share a secret.", 1979, <<https://www.cs.jhu.edu/~sdoshi/crypto/papers/shamirturing.pdf>>.

- [nistsp800-57] NIST, "SP 800-57 Part 1 Rev. 4 Recommendation for Key Management, Part 1: General", 1 January 2016, <<https://csrc.nist.gov/publications/detail/sp/800-57-part-1/rev-4/final>>.
- [fidotechnote] FIDO Alliance, ., "FIDO TechNotes: The Truth about Attestation", July 2018, <<https://fidoalliance.org/fido-technotes-the-truth-about-attestation/>>.
- [ntiasbom] NTIA, ., "NTIA Software Component Transparency", n.d., <<https://www.ntia.doc.gov/SoftwareTransparency>>.
- [cisqsbom] CISQ/Object Management Group, ., "TOOL-TO-TOOL SOFTWARE BILL OF MATERIALS EXCHANGE", July 2020, <<https://www.it-cisq.org/software-bill-of-materials/index.htm>>.
- [openbmc] Linux Foundation/OpenBMC Group, ., "Defining a Standard Baseboard Management Controller Firmware Stack", July 2020, <<https://www.openbmc.org/>>.
- [JTAG] IEEE Standard, ., "1149.7-2009 - IEEE Standard for Reduced-Pin and Enhanced-Functionality Test Access Port and Boundary-Scan Architecture", 2009, <<https://ieeexplore.ieee.org/document/5412866>>.
- [rootkeyrollover] ICANN, ., "Proposal for Future Root Zone KSK Rollovers", 2019, <<https://www.icann.org/en/system/files/files/proposal-future-rz-ksk-rollovers-01nov19-en.pdf>>.
- [CABFORUM] CA/Browser Forum, ., "CA/Browser Forum Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates, v.1.2.2", October 2014, <<https://cabforum.org/wp-content/uploads/BRv1.2.2.pdf>>.
- [I-D.richardson-rats-usecases] Richardson, M., Wallace, C., and W. Pan, "Use cases for Remote Attestation common encodings", Work in Progress, Internet-Draft, draft-richardson-rats-usecases-07, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-richardson-rats-usecases-07.txt>>.

[I-D.ietf-suit-architecture]

Moran, B., Tschofenig, H., Brown, D., and M. Meriac, "A Firmware Update Architecture for Internet of Things", Work in Progress, Internet-Draft, draft-ietf-suit-architecture-11, 27 May 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-suit-architecture-11.txt>>.

[I-D.ietf-emu-eap-noob]

Aura, T. and M. Sethi, "Nimble out-of-band authentication for EAP (EAP-NOOB)", Work in Progress, Internet-Draft, draft-ietf-emu-eap-noob-02, 12 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-emu-eap-noob-02.txt>>.

[I-D.ietf-rats-architecture]

Birkholz, H., Thaler, D., Richardson, M., Smith, N., and W. Pan, "Remote Attestation Procedures Architecture", Work in Progress, Internet-Draft, draft-ietf-rats-architecture-05, 10 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-rats-architecture-05.txt>>.

[I-D.birkholz-suit-coswid-manifest]

Birkholz, H., "A SUIT Manifest Extension for Concise Software Identifiers", Work in Progress, Internet-Draft, draft-birkholz-suit-coswid-manifest-00, 17 July 2018, <<http://www.ietf.org/internet-drafts/draft-birkholz-suit-coswid-manifest-00.txt>>.

[I-D.birkholz-rats-mud]

Birkholz, H., "MUD-Based RATS Resources Discovery", Work in Progress, Internet-Draft, draft-birkholz-rats-mud-00, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-birkholz-rats-mud-00.txt>>.

[RFC8520] Lear, E., Droms, R., and D. Romascanu, "Manufacturer Usage Description Specification", RFC 8520, DOI 10.17487/RFC8520, March 2019, <<https://www.rfc-editor.org/info/rfc8520>>.

[I-D.ietf-sacm-coswid]

Birkholz, H., Fitzgerald-McKay, J., Schmidt, C., and D. Waltermire, "Concise Software Identification Tags", Work in Progress, Internet-Draft, draft-ietf-sacm-coswid-15, 1 May 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-sacm-coswid-15.txt>>.

[RFC7168] Nazar, I., "The Hyper Text Coffee Pot Control Protocol for Tea Efflux Appliances (HTCPCP-TEA)", RFC 7168, DOI 10.17487/RFC7168, April 2014, <<https://www.rfc-editor.org/info/rfc7168>>.

[I-D.bormann-lwig-7228bis]
Bormann, C., Ersue, M., Keranen, A., and C. Gomez, "Terminology for Constrained-Node Networks", Work in Progress, Internet-Draft, draft-bormann-lwig-7228bis-06, 9 March 2020, <<http://www.ietf.org/internet-drafts/draft-bormann-lwig-7228bis-06.txt>>.

[I-D.ietf-teep-architecture]
Pei, M., Tschofenig, H., Thaler, D., and D. Wheeler, "Trusted Execution Environment Provisioning (TEEP) Architecture", Work in Progress, Internet-Draft, draft-ietf-teep-architecture-12, 13 July 2020, <<http://www.ietf.org/internet-drafts/draft-ietf-teep-architecture-12.txt>>.

Authors' Addresses

Michael Richardson
Sandelman Software Works

Email: mcr+ietf@sandelman.ca

Jie Yang
Huawei Technologies Co., Ltd.

Email: jay.yang@huawei.com

Independent Stream
Internet-Draft
Intended status: Informational
Expires: 7 January 2024

M. Schanzenbach
Fraunhofer AISEC
C. Grothoff
Berner Fachhochschule
B. Fix
GNUnet e.V.
6 July 2023

The GNU Name System
draft-schanzen-gns-28

Abstract

This document contains the GNU Name System (GNS) technical specification. GNS is a decentralized and censorship-resistant domain name resolution protocol that provides a privacy-enhancing alternative to the Domain Name System (DNS) protocols.

This document defines the normative wire format of resource records, resolution processes, cryptographic routines and security considerations for use by implementers.

This specification was developed outside the IETF and does not have IETF consensus. It is published here to inform readers about the function of GNS, guide future GNS implementations, and ensure interoperability among implementations including with the pre-existing GNUnet implementation.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 7 January 2024.

Copyright Notice

Copyright (c) 2023 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Notation	4
2.	Terminology	4
3.	Overview	7
3.1.	Names and Zones	7
3.2.	Publishing Binding Information	8
3.3.	Resolving Names	9
4.	Zones	10
4.1.	Zone Top-Level Domain	12
4.2.	Zone Revocation	13
5.	Resource Records	17
5.1.	Zone Delegation Records	19
5.1.1.	PKEY	20
5.1.2.	EDKEY	23
5.2.	Redirection Records	27
5.2.1.	REDIRECT	27
5.2.2.	GNS2DNS	27
5.3.	Auxiliary Records	28
5.3.1.	LEHO	28
5.3.2.	NICK	29
5.3.3.	BOX	30
6.	Record Encoding for Remote Storage	31
6.1.	The Storage Key	33
6.2.	Plaintext Record Data (RDATA)	34
6.3.	The Resource Records Block	35
7.	Name Resolution	37
7.1.	Start Zones	38
7.2.	Recursion	39
7.3.	Record Processing	40
7.3.1.	REDIRECT	41
7.3.2.	GNS2DNS	41
7.3.3.	BOX	42

7.3.4.	Zone Delegation Records	43
7.3.5.	NICK	43
8.	Internationalization and Character Encoding	44
9.	Security and Privacy Considerations	44
9.1.	Availability	44
9.2.	Agility	45
9.3.	Cryptography	45
9.4.	Abuse Mitigation	46
9.5.	Zone Management	47
9.6.	DHTs as Remote Storage	48
9.7.	Revocations	48
9.8.	Zone Privacy	49
9.9.	Zone Governance	49
9.10.	Namespace Ambiguity	50
10.	IANA Considerations	51
10.1.	GNS Record Types Registry	51
10.2.	.alt Subdomains Registry	52
11.	IANA Considerations	53
12.	Implementation and Deployment Status	53
13.	Acknowledgements	54
14.	Normative References	54
15.	Informative References	57
Appendix A.	Usage and Migration	59
A.1.	Zone Dissemination	59
A.2.	Start Zone Configuration	60
A.3.	Globally Unique Names and the Web	61
A.4.	Migration Paths	62
Appendix B.	Example flows	63
B.1.	AAAA Example Resolution	63
B.2.	REDIRECT Example Resolution	64
B.3.	GNS2DNS Example Resolution	65
Appendix C.	Base32GNS	66
Appendix D.	Test Vectors	67
D.1.	Base32GNS en-/decoding	67
D.2.	Record sets	68
D.3.	Zone revocation	81
Authors' Addresses	84

1. Introduction

This specification describes the GNU Name System (GNS), a censorship-resistant, privacy-preserving and decentralized domain name resolution protocol. GNS cryptographically secures the binding of names to arbitrary tokens, enabling it to double in some respects as an alternative to some of today's public key infrastructures.

In the terminology of the Domain Name System (DNS) [RFC1035], GNS roughly follows the idea of a local root zone deployment (see [RFC8806]), with the difference that the design encourages alternative roots and does not expect all deployments to use the same or any specific root zone. In the GNS reference implementation, users can autonomously and freely delegate control of names to zones through their local configurations. GNS expects each user to be in control of their setup. By following Section 9.10 guidelines, users should manage to avoid any confusion as to how names are resolved.

Name resolution and zone dissemination is based on the principle of a petname system where users can assign local names to zones. The GNS has its roots in ideas from the Simple Distributed Security Infrastructure [SDSI], enabling the decentralized mapping of secure identifiers to memorable names. A first academic description of the cryptographic ideas behind GNS can be found in [GNS].

This document defines the normative wire format of resource records, resolution processes, cryptographic routines and security considerations for use by implementers.

This specification was developed outside the IETF and does not have IETF consensus. It is published here to guide implementers of GNS and to ensure interoperability among implementations.

1.1. Requirements Notation

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

Apex Label This type of label is used to publish resource records in a zone that can be resolved without providing a specific label. It is the GNS method to provide what is the "zone apex" in DNS [RFC4033]. The apex label is represented using the character U+0040 ("@" without the quotes).

Application A component which uses a GNS implementation to resolve names into records and processes its contents.

Blinded Zone Key Key derived from a zone key and a label. The zone key and any blinded zone key derived from it are unlinkable without knowledge of the specific label used for the derivation.

Extension Label This type of label is used to refer to the authoritative zone that the record is in. The primary use for the extension label is in redirections where the redirection target is defined relative to the authoritative zone of the redirection record (see Section 5.2). The extension label is represented using the character U+002B ("+" without the quotes).

Label Separator Labels in a name are separated using the label separator U+002E ("." without the quotes). In GNS, with the exceptions of zone Top-Level Domains (see below) and boxed records (see Section 5.3.3), every label separator in a name indicates delegation to another zone.

Label A GNS label is a label as defined in [RFC8499]. Labels are UTF-8 strings in Unicode Normalization Form C (NFC) [Unicode-UAX15]. The apex label and the extension label have special purposes in the resolution protocol which are defined in the rest of the document. Zone administrators MAY disallow certain labels that might be easily confused with other labels through registration policies (see also Section 9.4).

Name A name in GNS is a domain name as defined in [RFC8499]: Names are UTF-8 [RFC3629] strings consisting of an ordered list of labels concatenated with a label separator. Names are resolved starting from the rightmost label. GNS does not impose length restrictions on names or labels. However, applications MAY ensure that name and label lengths are compatible with DNS and in particular IDNA [RFC5890]. In the spirit of [RFC5895], applications MAY preprocess names and labels to ensure compatibility with DNS or support specific user expectations, for example according to [Unicode-UTS46]. A GNS name may be indistinguishable from a DNS name and care must be taken by applications and implementors when handling GNS names (see Section 9.10). In order to avoid misinterpretation of example domains with (reserved) DNS domains this draft uses the suffix ".gns.alt" in examples which is also registered in the GANA ".alt Subdomains" registry [GANA] (see also [I-D.ietf-dnsop-alt-tld]).

Resolver The component of a GNS implementation which provides the recursive name resolution logic defined in Section 7.

Resource Record A GNS resource record is the information associated with a label in a GNS zone. A GNS resource record contains information as defined by its resource record type.

Start Zone In order to resolve any given GNS name an initial start

zone must be determined for this name. The start zone can be explicitly defined as part of the name using a zone Top-Level Domain (zTLD). Otherwise, it is determined through a local suffix-to-zone mapping (see Section 7.1).

Top-Level Domain The rightmost part of a GNS name is a GNS Top-Level Domain (TLD). A GNS TLD can consist of one or more labels. Unlike DNS Top-Level Domains (defined in [RFC8499]), GNS does not expect all users to use the same global root zone. Instead, with the exception of Zone Top-Level Domains (see Section 4.1), GNS TLDs are typically part of the configuration of the local resolver (see Section 7.1), and might thus not be globally unique.

Zone A GNS zone contains authoritative information (resource records). A zone is uniquely identified by its zone key. Unlike DNS zones, a GNS zone does not need to have a SOA record under the apex label.

Zone Key A key which uniquely identifies a zone. It is usually a public key of an asymmetric key pair. However, the established technical term "public key" is misleading, as in GNS a zone key may be a shared secret that should not be disclosed to unauthorized parties.

Zone Key Derivation Function The zone key derivation function (ZKDF) blinds a zone key using a label.

Zone Master The component of a GNS implementation which provides local zone management and publication as defined in Section 6.

Zone Owner The holder of the secret (typically a private key) that (together with a label and a value to sign) allows the creation of zone signatures that can be validated against the respective blinded zone key.

Zone Top-Level Domain A GNS Zone Top-Level Domain (zTLD) is a sequence of GNS labels at the end of a GNS name which encodes a zone type and zone key of a zone (see Section 4.1). Due to the statistical uniqueness of zone keys, zTLDs are also globally unique. A zTLD label sequence can only be distinguished from ordinary TLD label sequences by attempting to decode the labels into a zone type and zone key.

Zone Type The type of a GNS zone determines the cipher system and binary encoding format of the zone key, blinded zone keys, and cryptographic signatures.

3. Overview

GNS exhibits the three properties that are commonly used to describe a petname system:

1. Global names through the concept of zone top-level domains (zTLDs): As zones can be uniquely identified by their zone key and are statistically unique, zTLDs are globally unique mappings to zones. Consequently, GNS domain names with a zTLD suffix are also globally unique. Names with zTLDs suffixes are not human-readable.
2. Memorable petnames for zones: Users can configure local, human-readable references to zones. Such petnames serve as zTLD monikers which provide convenient names for zones to the local operator. The petnames may also be published as suggestions for other users searching for good label to use when referencing the respective zone.
3. A secure mapping from names to records: GNS allows zone owners to map labels to resource records or to delegate authority of names in the subdomain induced by a label to other zones. Zone owners may choose to publish this information to make it available to other users. Mappings are encrypted and signed using keys derived from the respective label before being published to remote storage. When names are resolved, signatures on resource records including delegations are verified by the recursive resolver.

In the remainder of this document, the "implementer" refers to the developer building a GNS implementation including the resolver, zone master, and supporting configuration such as start zones (see Section 7.1).

3.1. Names and Zones

It follows from the above that GNS does not support names which are simultaneously global, secure and human-readable. Instead, names are either global and not human-readable or not globally unique and human-readable. An example for a global name pointing to the record "example" in a zone is:

```
example.000G006K2TJNMD9VTCYRX7BRVV3HAEPS15E6NHDXKPJA1KAJJEG9AFF884
```

Now consider the case where a user locally configured the petname "pet.gns.alt" for the zone with the "example" record of the name above. The name "example.pet.gns.alt" would then point to the same record as the globally unique name above, but name resolution would only work on the local system where the "pet.gns.alt" petname is configured.

The delegation of petnames and subsequent resolution of delegation builds on ideas from the Simple Distributed Security Infrastructure [SDSI]. In GNS, any user can create and manage any number of zones (see Section 4) if their system provides a zone master implementation. For each zone, the zone type determines the respective set of cryptographic operations and the wire formats for encrypted data, public keys and signatures. A zone can be populated with mappings from labels to resource records (see Section 5) by its owner. A label can be mapped to a delegation record which results in the corresponding subdomain being delegated to another zone. Circular delegations are explicitly allowed, including delegating a subdomain to its immediate parent zone. In order to support (legacy) applications as well as to facilitate the use of petnames, GNS defines auxiliary record types in addition to supporting existing DNS records.

3.2. Publishing Binding Information

Zone contents are encrypted and signed before being published in a remote key-value storage (see Section 6) as illustrated in Figure 1. In this process, unique zone identification is hidden from the network through the use of key blinding. Key blinding allows the creation of signatures for zone contents using a blinded public/private key pair. This blinding is realized using a deterministic key derivation from the original zone key and corresponding private key using record label values as inputs from which blinding factors are derived. Specifically, the zone owner can derive blinded private keys for each record set published under a label, and a resolver can derive the corresponding blinded public keys. It is expected that GNS implementations use decentralized remote storages such as distributed hash tables (DHT) in order to facilitate availability within a network without the need for dedicated infrastructure. Specification of such a distributed or decentralized storage is out of scope of this document, but possible existing implementations include those based on [RFC7363], [Kademlia] or [R5N].

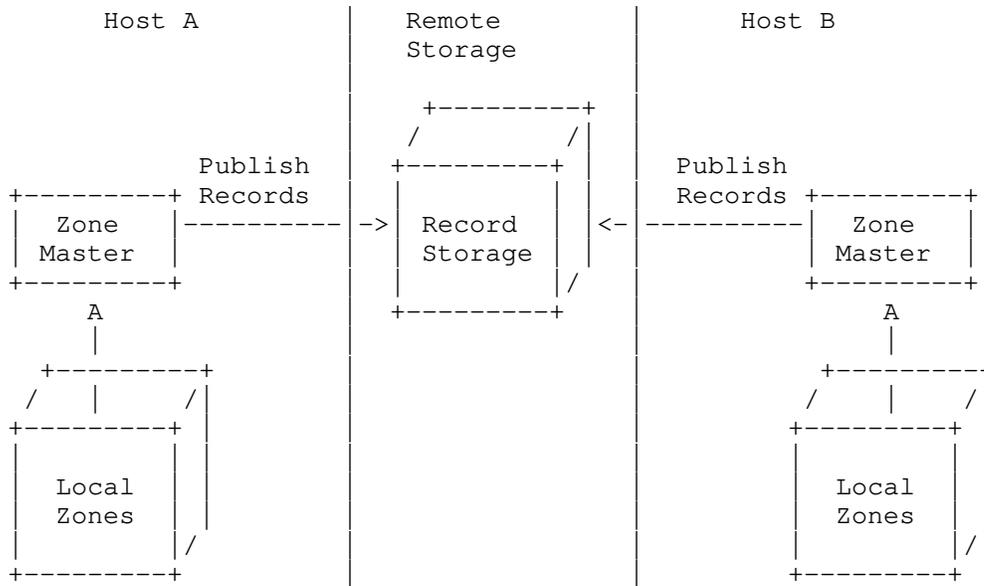


Figure 1: An example diagram of two hosts publishing GNS zones.

A zone master implementation SHOULD be provided as part of a GNS implementation to enable users to create and manage zones. If this functionality is not implemented, names can still be resolved if zone keys for the initial step in the name resolution have been configured (see Section 7) or if the names end with a zTLD suffix.

3.3. Resolving Names

Applications use the resolver to lookup GNS names. Starting from a configurable start zone, names are resolved by following zone delegations recursively as illustrated in Figure 2. For each label in a name, the recursive GNS resolver fetches the respective record set from the storage layer (see Section 7). Without knowledge of the label values and the zone keys, the different derived keys are unlinkable both to the original zone key and to each other. This prevents zone enumeration (except via expensive online brute force attacks): To confirm affiliation of a query or the corresponding encrypted record set with a specific zone requires knowledge of both the zone key and the label, neither of which are disclosed to remote storage by the protocol. At the same time, the blinded zone key and digital signatures associated with each encrypted record set allow resolvers and oblivious remote storage to verify the integrity of the published information without disclosing anything about the originating zone or the record sets.

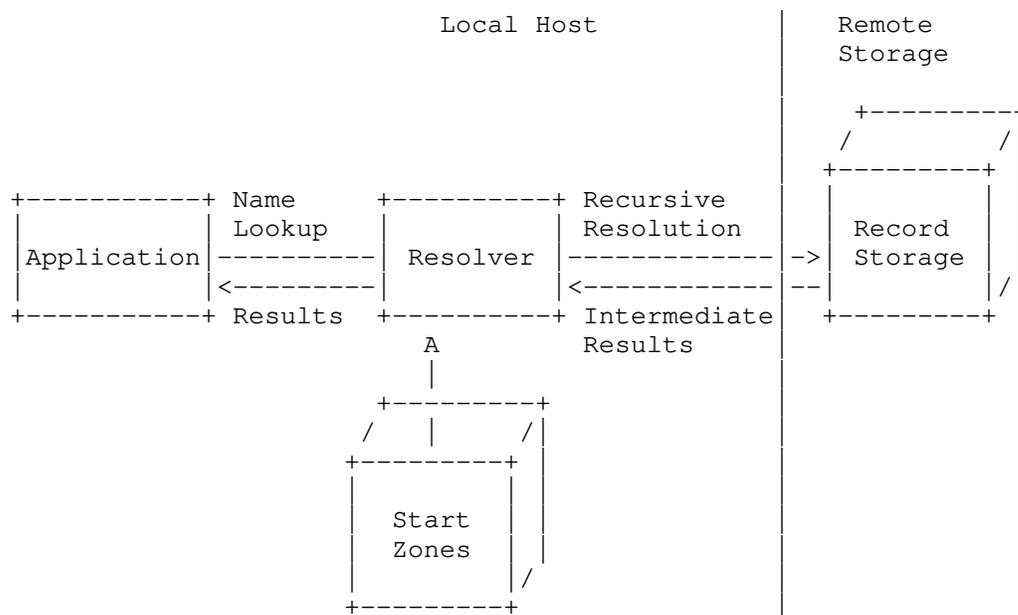


Figure 2: High-level view of the GNS resolution process.

4. Zones

A zone in GNS is uniquely identified by its zone type and zone key. Each zone can be referenced by its zone Top-Level Domain (zTLD) string (see Section 4.1) which encodes the zone type and zone key. A zone type (ztype) is a unique 32-bit number which corresponds to a resource record type number identifying a delegation record type in the GANA "GNS Record Types" registry [GANA]. The ztype is a unique identifier for the set cryptographic functions of the zone and the format of the delegation record type. Any ztype registration MUST define the following set of cryptographic functions:

KeyGen() -> d, zk is a function to generate a new private key d and the corresponding public zone key zk.

ZKDF(zk, label) -> zk' is a zone key derivation function which blinds a zone key zk using a label. zk and zk' must be unlinkable. Furthermore, blinding zk with different values for the label must result in different, unlinkable zk' values.

S-Encrypt(zk, label, expiration, plaintext) -> ciphertext is a symmetric encryption function which encrypts the plaintext to derive ciphertext based on key material derived from the zone key zk, a label and an expiration timestamp. In order to leverage

performance-enhancing caching features of certain underlying storages, in particular DHTs, a deterministic encryption scheme is recommended.

S-Decrypt(zk,label,expiration,ciphertext) -> plaintext is a symmetric decryption function which decrypts the ciphertext into plaintext based on key material derived from the zone key, a label, and an expiration timestamp.

Sign(d,message) -> signature is a function to sign a message using the private key d, yielding an unforgeable cryptographic signature. In order to leverage performance-enhancing caching features of certain underlying storages, in particular DHTs, a deterministic signature scheme is recommended.

Verify(zk,message,signature) -> boolean is a function to verify the signature was created using the private key d corresponding to the zone key zk where d,zk := Keygen(). The function returns a boolean value of "TRUE" if the signature is valid, and otherwise "FALSE".

SignDerived(d,label,message) -> signature is a function to sign a message (typically encrypted record data) that can be verified using the derived zone key zk' := ZKDF(zk,label). In order to leverage performance-enhancing caching features of certain underlying storages, in particular DHTs, a deterministic signature scheme is recommended.

VerifyDerived(zk,label,message,signature) -> boolean is function to verify the signature using the derived zone key zk' := ZKDF(zk,label). The function returns a boolean value of "TRUE" if the signature is valid, and otherwise "FALSE".

The cryptographic functions of the default ztypes are specified with their corresponding delegation records in Section 5.1. In order to support cryptographic agility, additional ztypes MAY be defined in the future which replace or update the default ztypes defined in this document. All ztypes MUST be registered as dedicated zone delegation record types in the GANA "GNS Record Types" registry (see [GANA]). When defining new record types the cryptographic security considerations of this document apply, in particular Section 9.3.

4.1. Zone Top-Level Domain

A zone Top-Level Domain (zTLD) is a string which encodes the zone type and zone key into a domain name suffix. A zTLD is used as a globally unique references to a zone in the process of name resolution. It is created by encoding a binary concatenation of the zone type and zone key (see Figure 3). The used encoding is a variation of the Crockford Base32 encoding [CrockfordB32] called Base32GNS. The encoding and decoding symbols for Base32GNS including this modification are defined in Figure 30. The functions for encoding and decoding based on this table are called Base32GNS-Encode and Base32GNS-Decode, respectively.

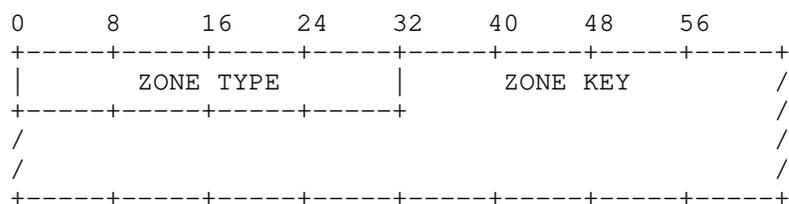


Figure 3: The binary representation of the zTLD

The ZONE TYPE must be encoded in network byte order. The format of the ZONE KEY depends entirely on the ZONE TYPE.

Consequently, a zTLD is encoded and decoded as follows:

```
zTLD := Base32GNS-Encode(ztype||zkey)
ztype||zkey := Base32GNS-Decode(zTLD)
```

where "||" is the concatenation operator.

The zTLD can be used as-is as a rightmost label in a GNS name. If an application wants to ensure DNS compatibility of the name, it MAY also represent the zTLD as follows: If the zTLD is less than or equal to 63 characters, it can be used as a zTLD as-is. If the zTLD is longer than 63 characters, the zTLD is divided into smaller labels separated by the label separator. Here, the most significant bytes of the "ztype||zkey" concatenation must be contained in the rightmost label of the resulting string and the least significant bytes in the leftmost label of the resulting string. This allows the resolver to determine the ztype and zTLD length from the rightmost label and to subsequently determine how many labels the zTLD should span. A GNS implementation MUST support the division of zTLDs in DNS compatible label lengths. For example, assuming a zTLD of 130 characters, the division is:

zTLD[126..129].zTLD[63..125].zTLD[0..62]

4.2. Zone Revocation

In order to revoke a zone key, a signed revocation message MUST be published. This message MUST be signed using the private key of the zone. The revocation message is broadcast to the network. The specification of the broadcast mechanism is out of scope for this document. A possible broadcast mechanism for efficient flooding in a distributed network is implemented in [GNUnet]. Alternatively, revocation messages could also be distributed via a distributed ledger or a trusted central server. To prevent flooding attacks, the revocation message MUST contain a proof of work (PoW). The revocation message including the PoW MAY be calculated ahead of time to support timely revocation.

For all occurrences below, "Argon2id" is the Password-based Key Derivation Function as defined in [RFC9106]. For the PoW calculations the algorithm is instantiated with the following parameters:

S The salt. Fixed 16-byte string: "GnsRevocationPow".

t Number of iterations: 3

m Memory size in KiB: 1024

T Output length of hash in bytes: 64

p Parallelization parameter: 1

v Algorithm version: 0x13

y Algorithm type (Argon2id): 2

X Unused

K Unused

Figure 4 illustrates the format of the data "P" on which the PoW is calculated.

22.

EPOCH A single epoch. Its value is fixed at 365 days in microseconds.

The revocation message wire format is illustrated in Figure 5.

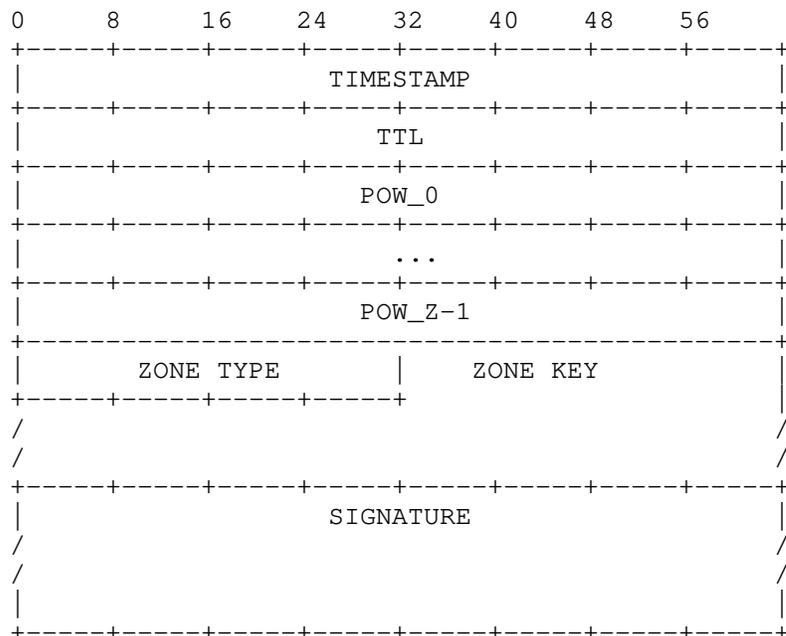


Figure 5: The Revocation Message Wire Format.

TIMESTAMP denotes the absolute 64-bit date when the revocation was computed. In microseconds since midnight (0 hour), January 1, 1970 UTC in network byte order. This is the same value as the time stamp used in the individual PoW calculations.

TTL denotes the relative 64-bit time to live of the record in microseconds in network byte order. The field SHOULD be set to $EPOCH * 1.1$. Given an average number of leading zeros D' , then the field value MAY be increased up to $(D' - D + 1) * EPOCH * 1.1$. Validators MAY reject messages with lower or higher values when received.

POW_i The values calculated as part of the PoW, in network byte order. Each POW_i MUST be unique in the set of POW values. To facilitate fast verification of uniqueness, the POW values must be given in strictly monotonically increasing order in the message.

1. The signature MUST be verified against the zone key.
2. The set of POW values MUST NOT contain duplicates which MUST be checked by verifying that the values are strictly monotonically increasing.
3. The average number of leading zeroes D' resulting from the provided POW values MUST be greater than or equal to D . Implementers MUST NOT use an integer data type to calculate or represent D' .

The TTL field in the revocation message is informational. A revocation MAY be discarded without checking the POW values or the signature if the TTL (in combination with TIMESTAMP) indicates that the revocation has already expired. The actual validity period of the revocation MUST be determined by examining the leading zeroes in the POW values.

The validity period of the revocation is calculated as $(D' - D + 1) * EPOCH * 1.1$. The EPOCH is extended by 10% in order to deal with unsynchronized clocks. The validity period added on top of the TIMESTAMP yields the expiration date. If the current time is after the expiration date, the revocation is considered stale.

Verified revocations MUST be stored locally. The implementation MAY discard stale revocations and evict them from the local store at any time.

Implementations MUST broadcast received revocations if they are valid and not stale. Should the calculated validity period differ from the TTL field value, the calculated value MUST be used as TTL field value when forwarding the revocation message. Systems might disagree on the current time, so implementations MAY use stale but otherwise valid revocations but SHOULD NOT broadcast them. Forwarded stale revocations MAY be discarded.

Any locally stored revocation MUST be considered during delegation record processing (see Section 7.3.4).

5. Resource Records

A GNS implementation SHOULD provide a mechanism to create and manage local zones as well as a persistence mechanism (such as a local database) for resource records. A new local zone is established by selecting a zone type and creating a zone key pair. If this mechanism is not implemented, no zones can be published in the storage (see Section 6) and name resolution is limited to non-local start zones (see Section 7.1).

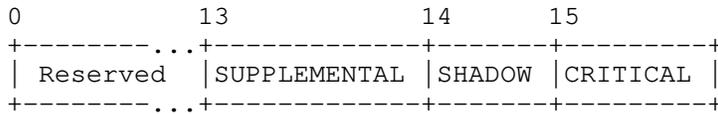


Figure 8: The Resource Record Flag Wire Format.

CRITICAL If this flag is set, it indicates that processing is critical. Implementations that do not support the record type or are otherwise unable to process the record **MUST** abort resolution upon encountering the record in the resolution process.

SHADOW If this flag is set, this record **MUST** be ignored by resolvers unless all (other) records of the same record type have expired. Used to allow zone publishers to facilitate good performance when records change by allowing them to put future values of records into the storage. This way, future values can propagate and can be cached before the transition becomes active.

SUPPLEMENTAL This is a supplemental record. It is provided in addition to the other records. This flag indicates that this record is not explicitly managed alongside the other records under the respective name but might be useful for the application.

5.1. Zone Delegation Records

This section defines the initial set of zone delegation record types. Any implementation **SHOULD** support all zone types defined here and **MAY** support any number of additional delegation records defined in the GANA "GNS Record Types" registry (see [GANA]). Not supporting some zone types will result in resolution failures in case the respective zone type is encountered. This can be a valid choice if some zone delegation record types have been determined to be cryptographically insecure. Zone delegation records **MUST NOT** be stored and published under the apex label. A zone delegation record type value is the same as the respective ztype value. The ztype defines the cryptographic primitives for the zone that is being delegated to. A zone delegation record payload contains the public key of the zone to delegate to. A zone delegation record **MUST** have the **CRITICAL** flag set and **MUST** be the only non-supplemental record under a label. There **MAY** be inactive records of the same type which have the **SHADOW** flag set in order to facilitate smooth key rollovers.

In the following, "||" is the concatenation operator of two byte strings. The algorithm specification uses character strings such as GNS labels or constant values. When used in concatenations or as input to functions the null-terminator of the character strings **MUST NOT** be included.

5.1.1. PKEY

In GNS, a delegation of a label to a zone of type "PKEY" is represented through a PKEY record. The PKEY DATA entry wire format can be found in Figure 9.

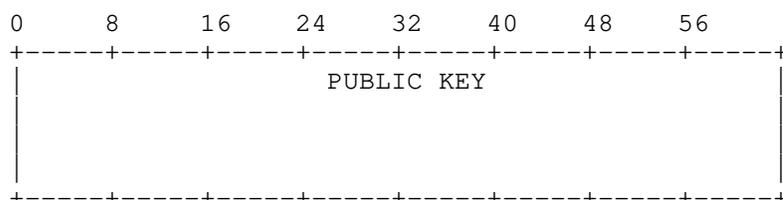


Figure 9: The PKEY Wire Format.

PUBLIC KEY A 256-bit Ed25519 public key.

For PKEY zones the zone key material is derived using the curve parameters of the twisted Edwards representation of Curve25519 [RFC7748] (the reasoning behind choosing this curve can be found in Section 9.3) with the ECDSA scheme [RFC6979]. The following naming convention is used for the cryptographic primitives of PKEY zones:

d is a 256-bit Ed25519 private key (private scalar).

zk is the Ed25519 public zone key corresponding to *d*.

p is the prime of edwards25519 as defined in [RFC7748], i.e. $2^{255} - 19$.

G is the group generator ($X(P), Y(P)$). With $X(P), Y(P)$ of edwards25519 as defined in [RFC7748].

L is the order of the prime-order subgroup of edwards25519 in [RFC7748].

KeyGen() The generation of the private scalar *d* and the curve point $zk := d * G$ (where *G* is the group generator of the elliptic curve) as defined in Section 2.2. of [RFC6979] represents the KeyGen() function.

The zone type and zone key of a PKEY are 4 + 32 bytes in length. This means that a zTLD will always fit into a single label and does not need any further conversion. Given a label, the output *zk'* of the ZKDF(*zk*, label) function is calculated as follows for PKEY zones:

```
ZKDF(zk,label):  
  PRK_h := HKDF-Extract ("key-derivation", zk)  
  h := HKDF-Expand (PRK_h, label || "gns", 512 / 8)  
  zk' := (h mod L) * zk  
  return zk'
```

The PKEY cryptosystem uses a hash-based key derivation function (HKDF) as defined in [RFC5869], using SHA-512 [RFC6234] for the extraction phase and SHA-256 [RFC6234] for the expansion phase. PRK_h is key material retrieved using an HKDF using the string "key-derivation" as salt and the zone key as initial keying material. h is the 512-bit HKDF expansion result and must be interpreted in network byte order. The expansion information input is a concatenation of the label and the string "gns". The multiplication of zk with h is a point multiplication, while the multiplication of d with h is a scalar multiplication.

The Sign() and Verify() functions for PKEY zones are implemented using 512-bit ECDSA deterministic signatures as specified in [RFC6979]. The same functions can be used for derived keys:

```
SignDerived(d,label,message):  
  zk := d * G  
  PRK_h := HKDF-Extract ("key-derivation", zk)  
  h := HKDF-Expand (PRK_h, label || "gns", 512 / 8)  
  d' := (h * d) mod L  
  return Sign(d',message)
```

A signature (R,S) is valid if the following holds:

```
VerifyDerived(zk,label,message,signature):  
  zk' := ZKDF(zk,label)  
  return Verify(zk',message,signature)
```

The S-Encrypt() and S-Decrypt() functions use AES in counter mode as defined in [MODES] (CTR-AES-256):

```

S-Encrypt(zk,label,expiration,plaintext):
  PRK_k := HKDF-Extract ("gns-aes-ctx-key", zk)
  PRK_n := HKDF-Extract ("gns-aes-ctx-iv", zk)
  K := HKDF-Expand (PRK_k, label, 256 / 8)
  NONCE := HKDF-Expand (PRK_n, label, 32 / 8)
  IV := NONCE || expiration || 0x00000000000000000001
  return CTR-AES256(K, IV, plaintext)

```

```

S-Decrypt(zk,label,expiration,ciphertext):
  PRK_k := HKDF-Extract ("gns-aes-ctx-key", zk)
  PRK_n := HKDF-Extract ("gns-aes-ctx-iv", zk)
  K := HKDF-Expand (PRK_k, label, 256 / 8)
  NONCE := HKDF-Expand (PRK_n, label, 32 / 8)
  IV := NONCE || expiration || 0x00000000000000000001
  return CTR-AES256(K, IV, ciphertext)

```

The key K and counter IV are derived from the record label and the zone key zk using a hash-based key derivation function (HKDF) as defined in [RFC5869]. SHA-512 [RFC6234] is used for the extraction phase and SHA-256 [RFC6234] for the expansion phase. The output keying material is 32 bytes (256 bits) for the symmetric key and 4 bytes (32 bits) for the nonce. The symmetric key K is a 256-bit AES [RFC3826] key.

The nonce is combined with a 64-bit initialization vector and a 32-bit block counter as defined in [RFC3686]. The block counter begins with the value of 1, and it is incremented to generate subsequent portions of the key stream. The block counter is a 32-bit integer value in network byte order. The initialization vector is the expiration time of the resource record block in network byte order. The resulting counter (IV) wire format can be found in Figure 10.

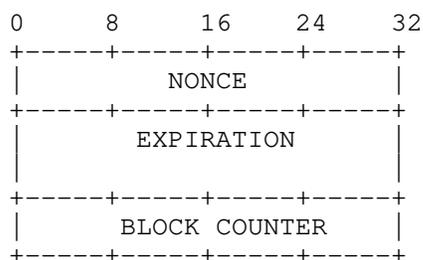


Figure 10: The Block Counter Wire Format.

5.1.2. EDKEY

In GNS, a delegation of a label to a zone of type "EDKEY" is represented through a EDKEY record. The EDKEY DATA entry wire format is illustrated in Figure 11.

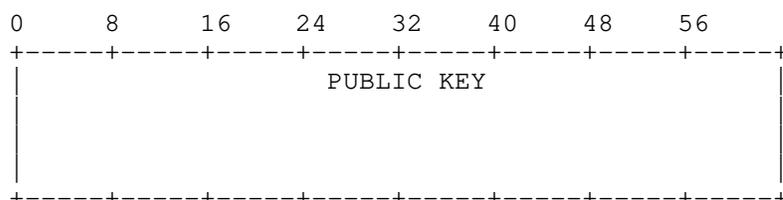


Figure 11: The EDKEY DATA Wire Format.

PUBLIC KEY A 256-bit EdDSA zone key.

For EDKEY zones the zone key material is derived using the curve parameters of the twisted edwards representation of Curve25519 [RFC7748] (a.k.a. Ed25519) with the Ed25519 scheme [ed25519] as specified in [RFC8032]. The following naming convention is used for the cryptographic primitives of EDKEY zones:

d is a 256-bit EdDSA private key.

a is an integer derived from **d** using the SHA-512 hash function as defined in [RFC8032].

zk is the EdDSA public key corresponding to **d**. It is defined as the curve point $a \cdot G$ where G is the group generator of the elliptic curve as defined in [RFC8032].

p is the prime of edwards25519 as defined in [RFC8032], i.e. $2^{255} - 19$.

G is the group generator $(X(P), Y(P))$. With $X(P), Y(P)$ of edwards25519 as defined in [RFC8032].

L is the order of the prime-order subgroup of edwards25519 in [RFC8032].

KeyGen() The generation of the private key **d** and the associated public key $zk := a \cdot G$ where G is the group generator of the elliptic curve and **a** is an integer derived from **d** using the SHA-512 hash function as defined in Section 5.1.5 of [RFC8032] represents the KeyGen() function.

The zone type and zone key of an EDKEY are 4 + 32 bytes in length. This means that a zTLD will always fit into a single label and does not need any further conversion.

The "EDKEY" ZKDF instantiation is based on [Tor224]. The calculation of a is defined in Section 5.1.5 of [RFC8032]. Given a label, the output of the ZKDF function is calculated as follows:

```
ZKDF(zk,label):
  /* Calculate the blinding factor */
  PRK_h := HKDF-Extract ("key-derivation", zk)
  h := HKDF-Expand (PRK_h, label || "gns", 512 / 8)
  /* Ensure that h == h mod L */
  h[31] &= 7

  zk' := h * zk
  return zk'
```

Implementers SHOULD employ a constant time scalar multiplication for the constructions above to protect against timing attacks. Otherwise, timing attacks could leak private key material if an attacker can predict when a system starts the publication process.

The EDKEY cryptosystem uses a hash-based key derivation function (HKDF) as defined in [RFC5869], using SHA-512 [RFC6234] for the extraction phase and HMAC-SHA256 [RFC6234] for the expansion phase. PRK_h is key material retrieved using an HKDF using the string "key-derivation" as salt and the zone key as initial keying material. The blinding factor h is the 512-bit HKDF expansion result. The expansion information input is a concatenation of the label and the string "gns". The result of the HKDF must be clamped and interpreted in network byte order. a is the 256-bit integer corresponding to the 256-bit private key d . The multiplication of zk with h is a point multiplication, while the division and multiplication of a and a_1 with the co-factor are integer operations.

The Sign(d ,message) and Verify(zk ,message,signature) procedures MUST be implemented as defined in [RFC8032].

Signatures for EDKEY zones use a derived private scalar d' which is not compliant with [RFC8032]. As the corresponding private key to the derived private scalar is not known, it is not possible to deterministically derive the signature part R according to [RFC8032]. Instead, signatures MUST be generated as follows for any given message and private zone key: A nonce is calculated from the highest 32 bytes of the expansion of the private key d and the blinding factor h . The nonce is then hashed with the message to r . This way, the full derivation path is included in the calculation of the R value of the signature, ensuring that it is never reused for two different derivation paths or messages.

```

SignDerived(d,label,message):
  /* Key expansion */
  dh := SHA-512 (d)
  /* EdDSA clamping */
  a := dh[0..31]
  a[0] &= 248
  a[31] &= 127
  a[31] |= 64
  /* Calculate zk corresponding to d */
  zk := a * G

  /* Calculate blinding factor */
  PRK_h := HKDF-Extract ("key-derivation", zk)
  h := HKDF-Expand (PRK_h, label || "gns", 512 / 8)
  /* Ensure that h == h mod L */
  h[31] &= 7

  zk' := h * zk
  a1 := a >> 3
  a2 := (h * a1) mod L
  d' := a2 << 3
  nonce := SHA-256 (dh[32..63] || h)
  r := SHA-512 (nonce || message)
  R := r * G
  S := r + SHA-512(R || zk' || message) * d' mod L
  return (R,S)

```

A signature (R,S) is valid if the following holds:

```

VerifyDerived(zk,label,message,signature):
  zk' := ZKDF(zk,label)
  (R,S) := signature
  return S * G == R + SHA-512(R, zk', message) * zk'

```

The `S-Encrypt()` and `S-Decrypt()` functions use XSalsa20 as defined in [XSalsa20] (XSalsa20-Poly1305):

5.2. Redirection Records

Redirect records are used to redirect resolution. Any implementation SHOULD support all redirection record types defined here and MAY support any number of additional redirection records defined in the GANA "GNS Record Types" registry [GANA]. Redirection records MUST have the CRITICAL flag set. Not supporting some record types can result in resolution failures. This can be a valid choice if some redirection record types have been determined to be insecure, or if an application has reasons to not support redirection to DNS for reasons such as complexity or security. Redirection records MUST NOT be stored and published under the apex label.

5.2.1. REDIRECT

A REDIRECT record is the GNS equivalent of a CNAME record in DNS. A REDIRECT record MUST be the only non-supplemental record under a label. There MAY be inactive records of the same type which have the SHADOW flag set in order to facilitate smooth changes of redirection targets. No other records are allowed. Details on processing of this record is defined in Section 7.3.1. A REDIRECT DATA entry is illustrated in Figure 13.

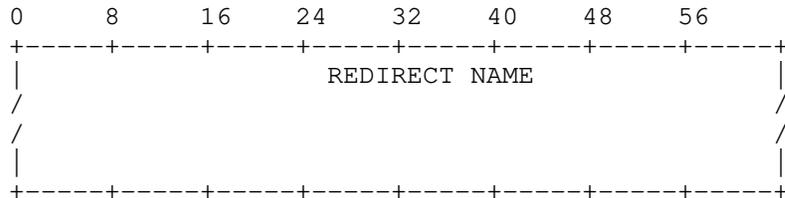


Figure 13: The REDIRECT DATA Wire Format.

REDIRECT NAME The name to continue with. The value of a redirect record can be a regular name, or a relative name. Relative GNS names are indicated by an extension label (U+002B, "+") as rightmost label. The string is UTF-8 encoded and 0-terminated.

5.2.2. GNS2DNS

A GNS2DNS record delegates resolution to DNS. The resource record contains a DNS name for the resolver to continue with in DNS followed by a DNS server. Both names are in the format defined in [RFC1034] for DNS names. There MAY be multiple GNS2DNS records under a label. There MAY also be DNSSEC DS records or any other records used to secure the connection with the DNS servers under the same label. There MAY be inactive records of the same type(s) which have the SHADOW flag set in order to facilitate smooth changes of redirection

targets. No other non-supplemental record types are allowed in the same record set. A GNS2DNS DATA entry is illustrated in Figure 14.

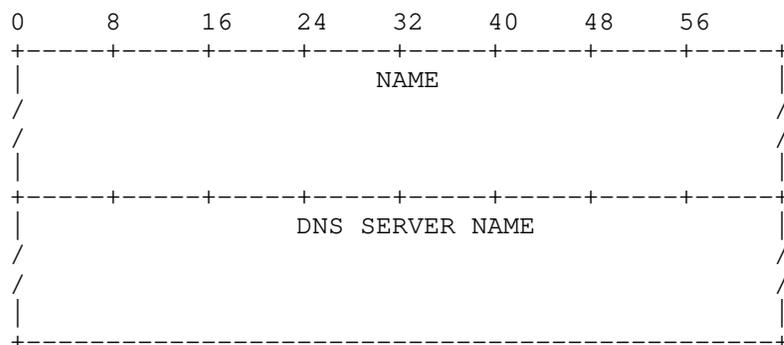


Figure 14: The GNS2DNS DATA Wire Format.

NAME The name to continue with in DNS. The value is UTF-8 encoded and 0-terminated.

DNS SERVER NAME The DNS server to use. This value can be an IPv4 address in dotted-decimal form or an IPv6 address in colon-hexadecimal form or a DNS name. It can also be a relative GNS name ending with a "+" as the rightmost label. The implementation **MUST** check the string syntactically for an IP address in the respective notation before checking for a relative GNS name. If all three checks fail, the name **MUST** be treated as a DNS name. The value is UTF-8 encoded and 0-terminated.

NOTE: If an application uses DNS names obtained from GNS2DNS records in a DNS request they **MUST** first be converted to an IDNA compliant representation [RFC5890].

5.3. Auxiliary Records

This section defines the initial set of auxiliary GNS record types. Any implementation **SHOULD** be able to process the specified record types according to Section 7.3.

5.3.1. LEHO

This record is used to provide a hint for Legacy Hostnames: Applications can use the GNS to lookup IPv4 or IPv6 addresses of internet services. However, sometimes connecting to such services does not only require the knowledge of an address and port, but also requires the canonical DNS name of the service to be transmitted over the transport protocol. In GNS, legacy host name records provide

applications the DNS name that is required to establish a connection to such a service. The most common use case is HTTP virtual hosting and TLS Server Name Indication [RFC6066], where a DNS name must be supplied in the HTTP "Host"-header and the TLS handshake, respectively. Using a GNS name in those cases might not work as it might not be globally unique. Furthermore, even if uniqueness is not an issue, the legacy service might not even be aware of GNS.

A LEHO resource record is expected to be found together in a single resource record with an IPv4 or IPv6 address. A LEHO DATA entry is illustrated in Figure 15.

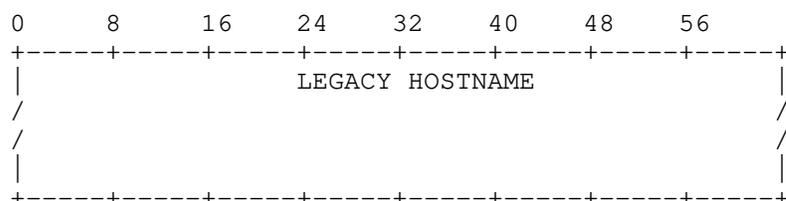


Figure 15: The LEHO DATA Wire Format.

LEGACY HOSTNAME A UTF-8 string (which is not 0-terminated) representing the legacy hostname.

NOTE: If an application uses a LEHO value in an HTTP request header (e.g. "Host:" header) it MUST be converted to an IDNA compliant representation [RFC5890].

5.3.2. NICK

Nickname records can be used by zone administrators to publish a label that a zone prefers to have used when it is referred to. This is a suggestion to other zones what label to use when creating a delegation record (Section 5.1) containing this zone key. This record SHOULD only be stored locally under the apex label "@" but MAY be returned with record sets under any label as a supplemental record. Section 7.3.5 details how a resolver must process supplemental and non-supplemental NICK records. A NICK DATA entry is illustrated in Figure 16.

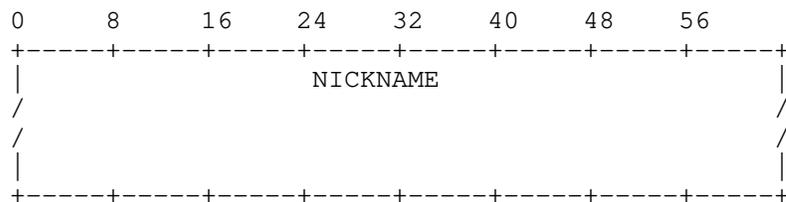


Figure 16: The NICK DATA Wire Format.

NICKNAME A UTF-8 string (which is not 0-terminated) representing the preferred label of the zone. This string MUST be a valid GNS label.

5.3.3. BOX

GNS lookups are expected to return all of the required useful information in one record set. This avoids unnecessary additional lookups and cryptographically ties together information that belongs together, making it impossible for an adversarial storage to provide partial answers that might omit information critical for security.

This general strategy is incompatible with the special labels used by DNS for SRV and TLSA records. Thus, GNS defines the BOX record format to box up SRV and TLSA records and include them in the record set of the label they are associated with. For example, a TLSA record for "_https._tcp.example.org" will be stored in the record set of "example.org" as a BOX record with service (SVC) 443 (https) and protocol (PROTO) 6 (tcp) and record TYPE "TLSA". For reference, see also [RFC2782]. A BOX DATA entry is illustrated in Figure 17.

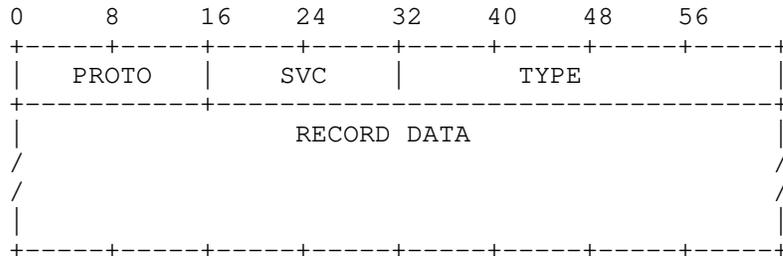


Figure 17: The BOX DATA Wire Format.

PROTO the 16-bit protocol number in network byte order. Values below 2^8 are reserved for 8-bit Internet Protocol numbers allocated by IANA [RFC5237] (e.g. 6 for TCP). Values above 2^8 are allocated by the GANA "Overlay Protocols" registry [GANAL].

SVC the 16-bit service value of the boxed record in network byte order. In case of TCP and UDP it is the port number.

TYPE is the 32-bit record type of the boxed record in network byte order.

RECORD DATA is a variable length field containing the "DATA" format

of TYPE as defined for the respective TYPE. Thus, for TYPE values below 2^{16} , the format is the same as the respective record type's binary format in DNS.

6. Record Encoding for Remote Storage

Any API which allows storing a block under a 512-bit key and retrieving one or more blocks from a key can be used by an implementation for remote storage. To be useful, the API MUST permit storing at least 176 byte blocks to be able to support the defined zone delegation record encodings, and SHOULD allow at least 1024 byte blocks. In the following, it is assumed that an implementation realizes two procedures on top of a storage:

```
PUT(key,block)
GET(key) -> block
```

A GNS implementation publishes blocks in accordance to the properties and recommendations of the underlying remote storage. This can include a periodic refresh operation to preserve the availability of published blocks.

There is no mechanism to explicitly delete individual blocks from remote storage. However, blocks include an EXPIRATION field which guides remote storage implementations to decide when to delete blocks. Given multiple blocks for the same key, remote storage implementations SHOULD try to preserve and return the block with the largest EXPIRATION value.

All resource records from the same zone sharing the same label are encrypted and published together in a single resource records block (RRBLOCK) in the remote storage under a key q as illustrated in Figure 18. A GNS implementation MUST NOT include expired resource records in blocks. An implementation MUST use the PUT storage procedure when record sets change to update the zone contents. Implementations MUST ensure that the EXPIRATION fields of RRBLOCKS increases strictly monotonically for every change, even if the smallest expiration time of records in the block does not.

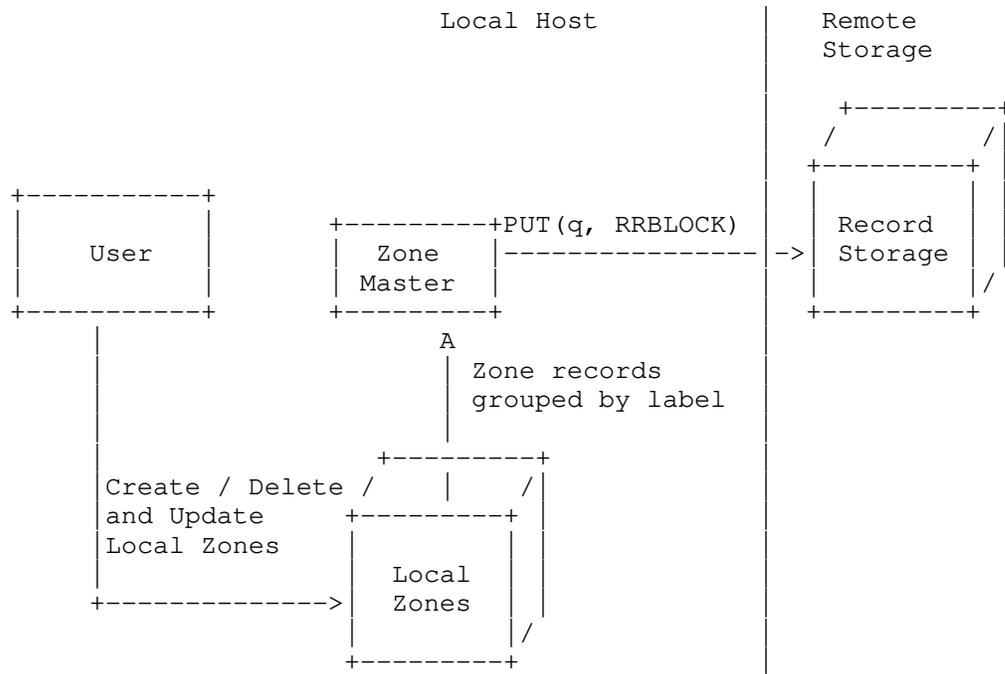


Figure 18: Management and publication of local zones in the distributed storage.

The storage key derivation and records block creation is specified in the following sections and illustrated in Figure 19.

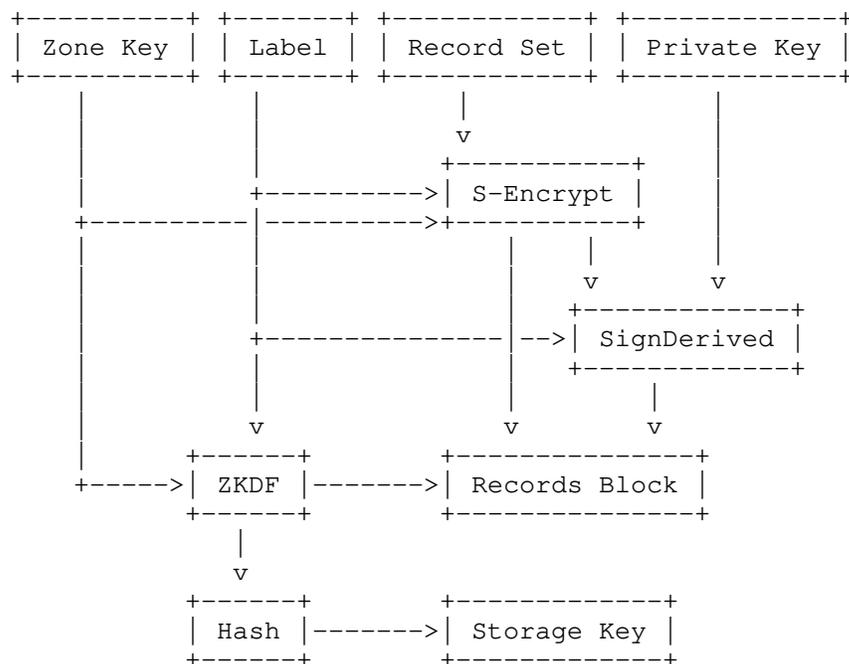


Figure 19: Storage key and records block creation overview.

6.1. The Storage Key

The storage key is derived from the zone key and the respective label of the contained records. The required knowledge of both zone key and label in combination with the similarly derived symmetric secret keys and blinded zone keys ensures query privacy (see [RFC8324], Section 3.5).

Given a label, the storage key q is derived as follows:

$$q := \text{SHA-512}(\text{ZKDF}(\text{zk}, \text{label}))$$

label is a UTF-8 string under which the resource records are published.

zk is the zone key.

q Is the 512-bit storage key under which the resource records block is published. It is the SHA-512 hash [RFC6234] over the derived zone key.

6.2. Plaintext Record Data (RDATA)

GNS records from a zone are grouped by their labels such that all records under the same label published together as a single block in the storage. Such grouped record sets MAY be paired with supplemental records.

Record data (RDATA) is the format used to encode such a group of GNS records. The binary format of RDATA is illustrated in Figure 20.

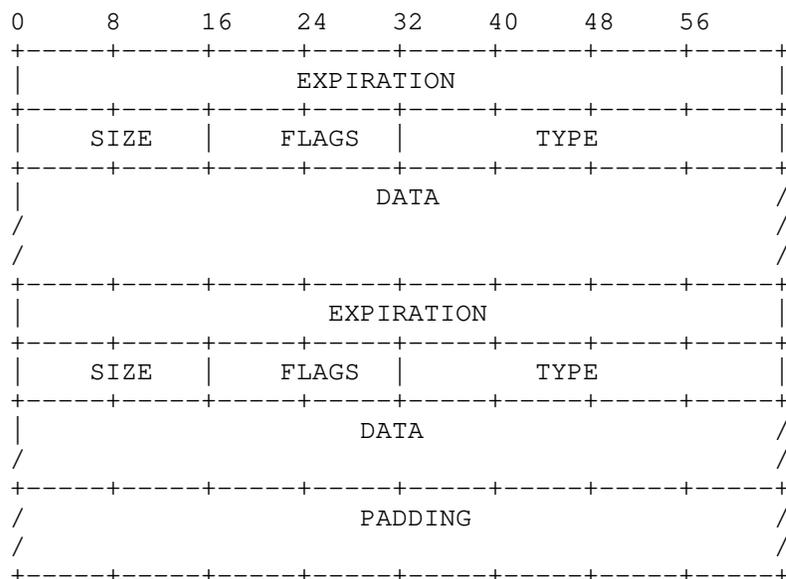


Figure 20: The RDATA Wire Format.

EXPIRATION, SIZE, TYPE, FLAGS and DATA These fields were defined in the resource record format in Section 5.

PADDING When serializing records into RDATA, a GNS implementation MUST ensure that the size of the RDATA is a power of two using the padding field. The field MUST be set to zero and MUST be ignored on receipt. As a special exception, record sets with (only) a zone delegation record type are never padded.

6.3. The Resource Records Block

The resource records grouped in an RDATA are encrypted using the S-Encrypt() function defined by the zone type of the zone to which the resource records belong and prefixed with meta data into a resource record block (RRBLOCK) for remote storage. The GNS RRBLOCK wire format is illustrated in Figure 21.

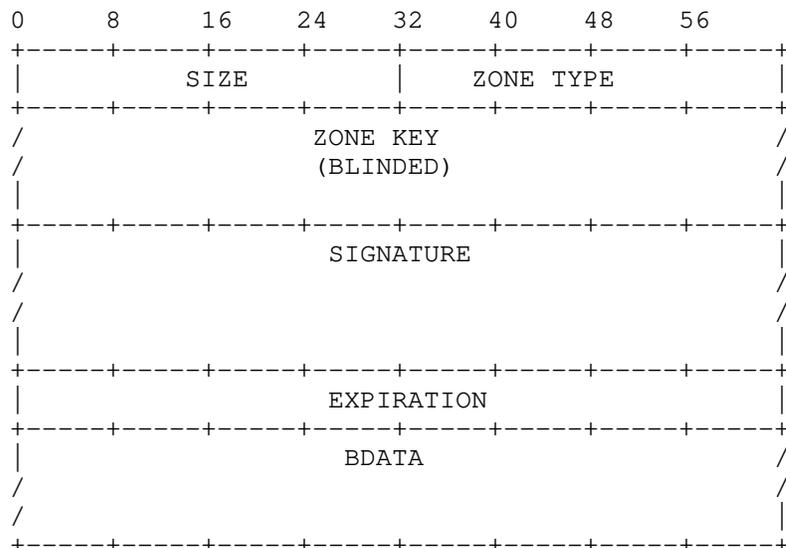


Figure 21: The RRBLOCK Wire Format.

SIZE A 32-bit value containing the length of the block in bytes in network byte order. Despite the message format's use of a 32-bit value, implementations MAY refuse to publish blocks beyond a certain size significantly below the theoretical block size limit of 4 GB.

ZONE TYPE is the 32-bit ztype in network byte order.

ZONE KEY (BLINDED) is the blinded zone key "ZKDF(zk, label)" to be used to verify SIGNATURE. The length and format of the blinded public key depends on the ztype.

SIGNATURE The signature is computed over the EXPIRATION and BDATA fields as detailed in Figure 22. The length and format of the signature depends on the ztype. The signature is created using the SignDerived() function of the cryptosystem of the zone (see Section 4).

BDATA Field as defined in the RRBLOCK message above.

7. Name Resolution

Names in GNS are resolved by recursively querying the record storage. Recursive in this context means that a resolver does not provide intermediate results for a query to the application. Instead, it MUST respond to a resolution request with either the requested resource record or an error message in case resolution fails. Figure 23 illustrates how an application requests the lookup of a GNS name (1). The application MAY provide a desired record type to the resolver. Subsequently, a Start Zone is determined (2) and the recursive resolution process started. This is where the desired record type is used to guide processing. For example, if a zone delegation record type is requested, the resolution of the apex label in that zone must be skipped, as the desired record is already found. Details on how the resolution process is initiated and each iterative result (3a,3b) in the resolution is processed are provided in the sections below. The results of the lookup are eventually returned to the application (4). The implementation MUST NOT filter the returned resource record sets according to the desired record type. Filtering of record sets is typically done by the application.

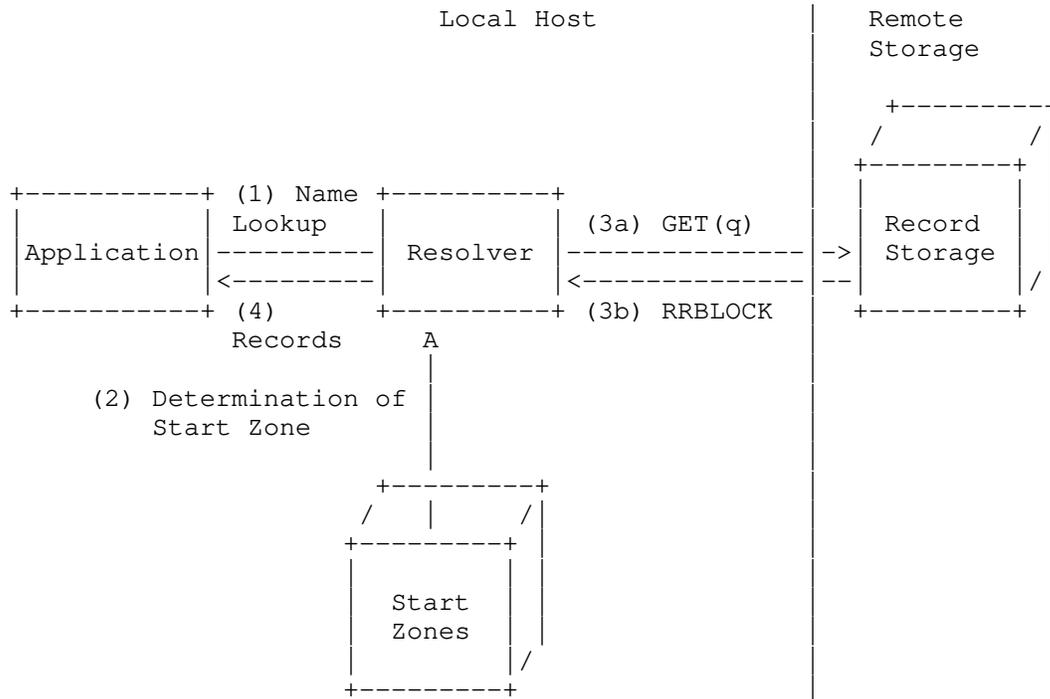


Figure 23: The recursive GNS resolution process.

7.1. Start Zones

The resolution of a GNS name starts by identifying the start zone suffix. Once the start zone suffix is identified, recursive resolution of the remainder of the name is initiated (see Section 7.2). There are two types of start zone suffixes: zTLDs and local suffix-to-zone mappings. The choice of available suffix-to-zone mappings is at the sole discretion of the local system administrator or user. This property addresses the issue of a single hierarchy with a centrally controlled root and the related issue of distribution and management of root servers in DNS (see [RFC8324], Sections 3.10 and 3.12).

For names ending with a zTLD the start zone is explicitly given in the suffix of the name to resolve. In order to ensure uniqueness of names with zTLDs any implementation MUST use the given zone as start zone. An implementation MUST first try to interpret the rightmost label of the given name as the beginning of a zTLD (see Section 4.1). If the rightmost label cannot be (partially) decoded or if it does not indicate a supported ztype, the name is treated as a normal name and start zone discovery MUST continue with finding a local suffix-to-zone mapping. If a valid ztype can be found in the rightmost label, the implementation MUST try to synthesize and decode the zTLD to retrieve the start zone key according to Section 4.1. If the zTLD cannot be synthesized or decoded, the resolution of the name fails and an error is returned to the application. Otherwise, the zone key MUST be used as the start zone:

```
Example name: www.example.<zTLD>
=> Start zone: zk of type ztype
=> Name to resolve from start zone: www.example
```

For names not ending with a zTLD the resolver MUST determine the start zone through a local suffix-to-zone mapping. Suffix-to-zone mappings MUST be configurable through a local configuration file or database by the user or system administrator. A suffix MAY consist of multiple GNS labels concatenated with a label separator. If multiple suffixes match the name to resolve, the longest matching suffix MUST be used. The suffix length of two results MUST NOT be equal. This indicates a misconfiguration and the implementation MUST return an error. The following is a non-normative example mapping of start zones:

```
Example name: www.example.xyz.gns.alt
Local suffix mappings:
xyz.gns.alt = zTLD0 := Base32GNS(ztype0 || zk0)
example.xyz.gns.alt = zTLD1 := Base32GNS(ztype1 || zk1)
example.com.gns.alt = zTLD2 := Base32GNS(ztype2 || zk2)
...
=> Start zone: zk1
=> Name to resolve from start zone: www
```

The process given above MAY be supplemented with other mechanisms if the particular application requires a different process. If no start zone can be discovered, resolution MUST fail and an error MUST be returned to the application.

7.2. Recursion

In each step of the recursive name resolution, there is an authoritative zone `zk` and a name to resolve. The name MAY be empty. If the name is empty, it is interpreted as the apex label `"@"`. Initially, the authoritative zone is the start zone.

From here, the following steps are recursively executed, in order:

1. Extract the right-most label from the name to look up.
2. Calculate `q` using the label and `zk` as defined in Section 6.1.
3. Perform a storage query `GET(q)` to retrieve the RRBLOCK.
4. Check that (a) the block is not expired, (b) the SHA-512 hash of the derived authoritative zone key `zk'` from the RRBLOCK matches the query `q`, and (c) the signature is valid. If any of these tests fail, the RRBLOCK MUST be ignored and, if applicable, the storage lookup `GET(q)` MUST continue to look for other RRBLOCKS.
5. Obtain the RDATA by decrypting the BDATA contained in the RRBLOCK using `S-Decrypt()` as defined by the zone type, effectively inverting the process described in Section 6.3.

Once a well-formed block has been decrypted, the records from RDATA are subjected to record processing.

7.3. Record Processing

In record processing, only the valid records obtained are considered. To filter records by validity, the resolver MUST at least check the expiration time and the FLAGS field of the respective record. Specifically, the resolver MUST disregard expired records. Furthermore, SHADOW and SUPPLEMENTAL flags can also exclude records from being considered. If the resolver encounters a record with the CRITICAL flag set and does not support the record type the resolution MUST be aborted and an error MUST be returned. The information that the critical record could not be processed SHOULD be returned in the error description. The implementation MAY choose not to return the reason for the failure, merely complicating troubleshooting for the user.

The next steps depend on the context of the name that is being resolved:

- * Case 1: If the filtered record set consists of a single REDIRECT record, the remainder of the name is prepended to the REDIRECT data and the recursion is started again from the resulting name. Details are described in Section 7.3.1.
- * Case 2: If the filtered record set consists exclusively of one or more GNS2DNS records resolution continues with DNS. Details are described in Section 7.3.2.
- * Case 3: If the remainder of the name to be resolved is of the format "_SERVICE._PROTO" and the record set contains one or more matching BOX records, the records in the BOX records are the final result and the recursion is concluded as described in Section 7.3.3.
- * Case 4: If the current record set consist of a single delegation record, resolution of the remainder of the name is delegated to the target zone as described in Section 7.3.4.
- * Case 5: If the remainder of the name to resolve is empty the record set is the final result. If any NICK records are in the final result set, they MUST first be processed according to Section 7.3.5. Otherwise, the record result set is directly returned as the final result.
- * Finally, if none of the above is applicable, resolution fails and the resolver MUST return an empty record set.

7.3.1. REDIRECT

If the remaining name is empty and the desired record type is REDIRECT, in which case the resolution concludes with the REDIRECT record. If the rightmost label of the redirect name is the extension label (U+002B, "+"), resolution continues in GNS with the new name in the current zone. Otherwise, the resulting name is resolved via the default operating system name resolution process. This can in turn trigger a GNS name resolution process depending on the system configuration. In case resolution continues in DNS, the name MUST first be converted to an IDNA compliant representation [RFC5890].

In order to prevent infinite loops, the resolver MUST implement loop detection or limit the number of recursive resolution steps. The loop detection MUST be effective even if a REDIRECT found in GNS triggers subsequent GNS lookups via the default operating system name resolution process.

7.3.2. GNS2DNS

When a resolver encounters one or more GNS2DNS records and the remaining name is empty and the desired record type is GNS2DNS, the GNS2DNS records are returned.

Otherwise, it is expected that the resolver first resolves the IP addresses of the specified DNS name servers. The DNS name MUST be converted to an IDNA compliant representation [RFC5890] for resolution in DNS. GNS2DNS records MAY contain numeric IPv4 or IPv6 addresses, allowing the resolver to skip this step. The DNS server names might themselves be names in GNS or DNS. If the rightmost label of the DNS server name is the extension label (U+002B, "+"), the rest of the name is to be interpreted relative to the zone of the GNS2DNS record. If the DNS server name ends in a label representation of a zone key, the DNS server name is to be resolved against the GNS zone zk.

Multiple GNS2DNS records can be stored under the same label, in which case the resolver MUST try all of them. The resolver MAY try them in any order or even in parallel. If multiple GNS2DNS records are present, the DNS name MUST be identical for all of them. Otherwise, it is not clear which name the resolver is supposed to follow. If different DNS names are present the resolution fails and an appropriate error is SHOULD be returned to the application.

If there are DNSSEC DS records or any other records used to secure the connection with the DNS servers stored under the label, the DNS resolver SHOULD use them to secure the connection with the DNS server.

Once the IP addresses of the DNS servers have been determined, the DNS name from the GNS2DNS record is appended to the remainder of the name to be resolved, and resolved by querying the DNS name server(s). The synthesized name has to be converted to an IDNA compliant representation [RFC5890] for resolution in DNS. If such a conversion is not possible, the resolution MUST be aborted and an error MUST be returned. The information that the critical record could not be processed SHOULD be returned in the error description. The implementation MAY choose not to return the reason for the failure, merely complicating troubleshooting for the user.

As the DNS servers specified are possibly authoritative DNS servers, the GNS resolver MUST support recursive DNS resolution and MUST NOT delegate this to the authoritative DNS servers. The first successful recursive name resolution result is returned to the application. In addition, the resolver SHOULD return the queried DNS name as a supplemental LEHO record (see Section 5.3.1) with a relative expiration time of one hour.

Once the transition from GNS into DNS is made through a GNS2DNS record, there is no "going back". The (possibly recursive) resolution of the DNS name MUST NOT delegate back into GNS and should only follow the DNS specifications. For example, names contained in DNS CNAME records MUST NOT be interpreted by resolvers that support both DNS and GNS as GNS names.

GNS resolvers SHOULD offer a configuration option to disable DNS processing to avoid information leakage and provide a consistent security profile for all name resolutions. Such resolvers would return an empty record set upon encountering a GNS2DNS record during the recursion. However, if GNS2DNS records are encountered in the record set for the apex label and a GNS2DNS record is explicitly requested by the application, such records MUST still be returned, even if DNS support is disabled by the GNS resolver configuration.

7.3.3. BOX

When a BOX record is received, a GNS resolver must unbox it if the name to be resolved continues with "_SERVICE._PROTO". Otherwise, the BOX record is to be left untouched. This way, TLSA (and SRV) records do not require a separate network request, and TLSA records become inseparable from the corresponding address records.

7.3.4. Zone Delegation Records

When the resolver encounters a record of a supported zone delegation record type (such as PKEY or EDKEY) and the remainder of the name is not empty, resolution continues recursively with the remainder of the name in the GNS zone specified in the delegation record.

Whenever a resolver encounters a new GNS zone, it MUST check against the local revocation list (see Section 4.2) whether the respective zone key has been revoked. If the zone key was revoked, the resolution MUST fail with an empty result set.

Implementations MUST NOT allow multiple different zone delegations under a single label (except if some are shadow records). Implementations MAY support any subset of ztypes. Implementations MUST NOT process zone delegation records stored under the apex label ("@"). If a zone delegation record is encountered under the apex label, resolution fails and an error MUST be returned. The implementation MAY choose not to return the reason for the failure, merely impacting troubleshooting information for the user.

If the remainder of the name to resolve is empty and a record set was received containing only a single delegation record, the recursion is continued with the record value as authoritative zone and the apex label "@" as remaining name. Except in the case where the desired record type as specified by the application is equal to the ztype, in which case the delegation record is returned.

7.3.5. NICK

NICK records are only relevant to the recursive resolver if the record set in question is the final result which is to be returned to the application. The encountered NICK records can either be supplemental (see Section 5) or non-supplemental. If the NICK record is supplemental, the resolver only returns the record set if one of the non-supplemental records matches the queried record type. It is possible that one record set contains both supplemental and non-supplemental NICK records.

The differentiation between a supplemental and non-supplemental NICK record allows the application to match the record to the authoritative zone. Consider the following example:

```
Query: alice.example.gns.alt (type=A)
Result:
A: 192.0.2.1
NICK: eve (non-supplemental)
```

In this example, the returned NICK record is non-supplemental. For the application, this means that the NICK belongs to the zone "alice.example.gns.alt" and is published under the apex label along with an A record. The NICK record is interpreted as: The zone defined by "alice.example.gns.alt" wants to be referred to as "eve". In contrast, consider the following:

```
Query: alice.example.gns.alt (type=AAAA)
Result:
AAAA: 2001:DB8::1
NICK: john (supplemental)
```

In this case, the NICK record is marked as supplemental. This means that the NICK record belongs to the zone "example.gns.alt" and is published under the label "alice" along with an AAAA record. Here, the NICK record should be interpreted as: The zone defined by "example.gns.alt" wants to be referred to as "john". This distinction is likely useful for other records published as supplemental.

8. Internationalization and Character Encoding

All names in GNS are encoded in UTF-8 [RFC3629]. Labels MUST be canonicalized using Normalization Form C (NFC) [Unicode-UAX15]. This does not include any DNS names found in DNS records, such as CNAME record data, which is internationalized through the IDNA specifications [RFC5890].

9. Security and Privacy Considerations

9.1. Availability

In order to ensure availability of records beyond their absolute expiration times, implementations MAY allow to locally define relative expiration time values of records. Records can then be published recurringly with updated absolute expiration times by the implementation.

Implementations MAY allow users to manage private records in their zones that are not published in the storage. Private records are considered just like regular records when resolving labels in local zones, but their data is completely unavailable to non-local users.

9.2. Agility

The security of cryptographic systems depends on both the strength of the cryptographic algorithms chosen and the strength of the keys used with those algorithms. The security also depends on the engineering of the protocol used by the system to ensure that there are no non-cryptographic ways to bypass the security of the overall system. This is why developers of applications managing GNS zones SHOULD select a default ztype considered secure at the time of releasing the software. For applications targeting end users that are not expected to understand cryptography, the application developer MUST NOT leave the ztype selection of new zones to end users.

This document concerns itself with the selection of cryptographic algorithms used in GNS. The algorithms identified in this document are not known to be broken (in the cryptographic sense) at the current time, and cryptographic research so far leads us to believe that they are likely to remain secure into the foreseeable future. However, this is not necessarily forever, and it is expected that new revisions of this document will be issued from time to time to reflect the current best practices in this area.

In terms of crypto-agility, whenever the need for an updated cryptographic scheme arises to, for example, replace ECDSA over Ed25519 for PKEY records, it can simply be introduced through a new record type. Zone administrators can then replace the delegation record type for future records. The old record type remains and zones can iteratively migrate to the updated zone keys. To ensure that implementations correctly generate an error message when encountering a ztype that they do not support, current and future delegation records must always have the CRITICAL flag set.

9.3. Cryptography

The following considerations provide background on the design choices of the ztypes specified in this document. When specifying new ztypes as per Section 4, the same considerations apply.

GNS PKEY zone keys use ECDSA over Ed25519. This is an unconventional choice, as ECDSA is usually used with other curves. However, standardized ECDSA curves are problematic for a range of reasons described in the Curve25519 and EdDSA papers [ed25519]. Using EdDSA directly is also not possible, as a hash function is used on the private key which destroys the linearity that the key blinding in GNS depends upon. We are not aware of anyone suggesting that using Ed25519 instead of another common curve of similar size would lower the security of ECDSA. GNS uses 256-bit curves because that way the encoded (public) keys fit into a single DNS label, which is good for usability.

In order to ensure ciphertext indistinguishability, care must be taken with respect to the initialization vector in the counter block. In our design, the IV always includes the expiration time of the record block. When applications store records with relative expiration times, monotonicity is implicitly ensured because each time a block is published into the storage, its IV is unique as the expiration time is calculated dynamically and increases monotonically with the system time. Still, an implementation **MUST** ensure that when relative expiration times are decreased, the expiration time of the next record block **MUST** be after the last published block. For records where an absolute expiration time is used, the implementation **MUST** ensure that the expiration time is always increased when the record data changes. For example, the expiration time on the wire could be increased by a single microsecond even if the user did not request a change. In case of deletion of all resource records under a label, the implementation **MUST** keep track of the last absolute expiration time of the last published resource block. Implementations **MAY** define and use a special record type as a tombstone that preserves the last absolute expiration time, but then **MUST** take care to not publish a block with such a tombstone record. When new records are added under this label later, the implementation **MUST** ensure that the expiration times are after the last published block. Finally, in order to ensure monotonically increasing expiration times the implementation **MUST** keep a local record of the last time obtained from the system clock, so as to construct a monotonic clock in case the system clock jumps backwards.

9.4. Abuse Mitigation

GNS names are UTF-8 strings. Consequently, GNS faces similar issues with respect to name spoofing as DNS does for internationalized domain names. In DNS, attackers can register similar sounding or looking names (see above) in order to execute phishing attacks. GNS zone administrators must take into account this attack vector and incorporate rules in order to mitigate it.

Further, DNS can be used to combat illegal content on the Internet by having the respective domains seized by authorities. However, the same mechanisms can also be abused in order to impose state censorship. Avoiding that possibility is one of the motivations behind GNS. In GNS, TLDs are not enumerable. By design, the start zone of the resolver is defined locally and hence such a seizure is difficult and ineffective in GNS.

9.5. Zone Management

In GNS, zone administrators need to manage and protect their zone keys. Once a private zone key is lost, it cannot be recovered and the zone revocation message cannot be computed anymore. Revocation messages can be pre-calculated if revocation is required in case a private zone key is lost. Zone administrators, and for GNS this includes end-users, are required to responsibly and diligently protect their cryptographic keys. GNS supports signing records in advance ("offline") in order to support processes (such as air gaps) which aim to protect private keys.

Similarly, users are required to manage their local start zone configuration. In order to ensure integrity and availability of names, users must ensure that their local start zone information is not compromised or outdated. It can be expected that the processing of zone revocations and an initial start zone is provided with a GNS implementation ("drop shipping"). Shipping an initial start zone configuration effectively establishes a root zone. Extension and customization of the zone is at the full discretion of the user.

While implementations following this specification will be interoperable, if two implementations connect to different remote storages they are mutually unreachable. This can lead to a state where a record exists in the global namespace for a particular name, but the implementation is not communicating with the remote storage that contains the respective block and is hence unable to resolve it. This situation is similar to a split-horizon DNS configuration. Which remote storages are implemented usually depends on the application it is built for. The remote storage used will most likely depend on the specific application context using GNS resolution. For example, one application is the resolution of hidden services within the Tor network, which would suggest using Tor routers for remote storage. Implementations of "aggregated" remote storages are conceivable, but are expected to be the exception.

9.6. DHTs as Remote Storage

This document does not specify the properties of the underlying remote storage which is required by any GNS implementation. It is important to note that the properties of the underlying remote storage are directly inherited by the GNS implementation. This includes both security as well as other non-functional properties such as scalability and performance. Implementers should take great care when selecting or implementing a DHT for use as remote storage in a GNS implementation. DHTs with reasonable security and performance properties exist [R5N]. It should also be taken into consideration that GNS implementations which build upon different DHT overlays are unlikely to be interoperable with each other.

9.7. Revocations

Zone administrators are advised to pre-generate zone revocations and to securely store the revocation information in case the zone key is lost, compromised or replaced in the future. Pre-calculated revocations can cease to be valid due to expirations or protocol changes such as epoch adjustments. Consequently, implementers and users must take precautions in order to manage revocations accordingly.

Revocation payloads do not include a 'new' key for key replacement. Inclusion of such a key would have two major disadvantages:

1. If a revocation is published after a private key was compromised, allowing key replacement would be dangerous: if an adversary took over the private key, the adversary could then broadcast a revocation with a key replacement. For the replacement, the compromised owner would have no chance to issue a revocation. Thus, allowing a revocation message to replace a private key makes dealing with key compromise situations worse.
2. Sometimes, key revocations are used with the objective of changing cryptosystems. Migration to another cryptosystem by replacing keys via a revocation message would only be secure as long as both cryptosystems are still secure against forgery. Such a planned, non-emergency migration to another cryptosystem should be done by running zones for both cipher systems in parallel for a while. The migration would conclude by revoking the legacy zone key only once it is deemed no longer secure, and hopefully after most users have migrated to the replacement.

9.8. Zone Privacy

GNS does not support authenticated denial of existence of names within a zone. Record data is published in encrypted form using keys derived from the zone key and record label. Zone administrators should carefully consider if a label and zone key are public, or if one or both of these should be used as a shared secret to restrict access to the corresponding record data. Unlike public zone keys, low-entropy labels can be guessed by an attacker. If an attacker knows the public zone key, the use of well known or guessable labels effectively threatens the disclosure of the corresponding records.

It should be noted that the guessing attack on labels only applies if the zone key is somehow disclosed to the adversary. GNS itself does not disclose it during a lookup or when resource records are published (as only the blinded zone keys are used on the network). However, zone keys do become public during revocation.

It is thus RECOMMENDED to use a label with sufficient entropy to prevent guessing attacks if any data in a resource record set is sensitive.

9.9. Zone Governance

While DNS is distributed, in practice it relies on centralized, trusted registrars to provide globally unique names. As the awareness of the central role DNS plays on the Internet rises, various institutions are using their power (including legal means) to engage in attacks on the DNS, thus threatening the global availability and integrity of information on the Internet. While a wider discussion of this issue is out of scope for this document, analyses and investigations can be found in recent academic research works including [SecureNS].

GNS is designed to provide a secure, privacy-enhancing alternative to the DNS name resolution protocol, especially when censorship or manipulation is encountered. In particular, it directly addresses concerns in DNS with respect to query privacy. However, depending on the governance of the root zone, any deployment will likely suffer from the issues of a "Single Hierarchy with a Centrally Controlled Root" and "Distribution and Management of Root Servers" as raised in [RFC8324]. In DNS, those issues are a direct result from the centralized root zone governance at the Internet Corporation for Assigned Names and Numbers (ICANN) which allows it to provide globally unique names.

In GNS, start zones give users local authority over their preferred root zone governance. It enables users to replace or enhance a trusted root zone configuration provided by a third party (e.g. the implementer or a multi-stakeholder governance body like ICANN) with secure delegation of authority using local petnames while operating under a very strong adversary model. In combination with zTLDs, this provides users of GNS with a global, secure and memorable mapping without a trusted authority.

Any GNS implementation MAY provide a default governance model in the form of an initial start zone mapping.

9.10. Namespace Ambiguity

Technically, the GNS protocol can be used to resolve names in the namespace of the global DNS. However, this would require the respective governance bodies and stakeholders (e.g. IETF and ICANN) to standardize the use of GNS for this particular use case.

However, this capability implies that GNS names may be indistinguishable from DNS names in their respective common display format [RFC8499] or other special-use domain names [RFC6761] if a local start zone configuration maps suffixes from the global DNS to GNS zones. For applications, it is then ambiguous which name system should be used in order to resolve a given name. This poses a risk when trying to resolve a name through DNS when it is actually a GNS name as discussed in [RFC8244]. In such a case, the GNS name is likely to be leaked as part of the DNS resolution.

In order to prevent disclosure of queried GNS names it is RECOMMENDED that GNS-aware applications try to resolve a given name in GNS before any other method taking into account potential suffix-to-zone mappings and zTLDs. Suffix-to-zone mappings are expected to be configured by the user or local administrator and as such the resolution in GNS is in line with user expectations even if the name could also be resolved through DNS. If no suffix-to-zone mapping for the name exists and no zTLD is found, resolution MAY continue with other methods such as DNS. If a suffix-to-zone mapping for the name exists or the name ends with a zTLD, it MUST be resolved using GNS and resolution MUST NOT continue by any other means independent of the GNS resolution result.

Mechanisms such as the Name Service Switch (NSS) of Unix-like operating systems are an example of how such a resolution process can be implemented and used. It allows system administrators to configure host name resolution precedence and is integrated with the system resolver implementation.

For use cases where GNS names may be confused with names of other name resolution mechanisms (in particular DNS), the ".gns.alt" domain SHOULD be used. For use cases like implementing sinkholes to block malware sites or serving DNS domains via GNS to bypass censorship, GNS MAY be deliberately used in ways that interfere with resolution of another name system.

10. GANA Considerations

GANA has assigned signature purposes in its "GNUnet Signature Purpose" registry as listed in Figure 24.

Purpose	Name	References	Comment
3	GNS_REVOCATION	[This.I-D]	GNS zone key revocation
15	GNS_RECORD_SIGN	[This.I-D]	GNS record set signature

Figure 24: Requested Changes in the GANA GNUnet Signature Purpose Registry.

10.1. GNS Record Types Registry

GANA [GANA] manages the "GNS Record Types" registry. Each entry has the following format:

- * Name: The name of the record type (case-insensitive ASCII string, restricted to alphanumeric characters). For zone delegation records, the assigned number represents the ztype value of the zone.
- * Number: 32-bit, above 65535
- * Comment: Optionally, a brief English text describing the purpose of the record type (in UTF-8)
- * Contact: Optionally, the contact information of a person to contact for further information.
- * References: Optionally, references describing the record type (such as an RFC).

The registration policy for this registry is "First Come First Served". This policy is modeled on that described in [RFC8126], and describes the actions taken by GANA:

Adding new entries is possible after review by any authorized GANA contributor, using a first-come-first-served policy for unique name allocation. Reviewers are responsible to ensure that the chosen "Name" is appropriate for the record type. The registry will define a unique number for the entry.

Authorized GANA contributors for review of new entries are reachable at `gns-registry@gnunet.org`.

Any request **MUST** contain a unique name and a point of contact. The contact information **MAY** be added to the registry given the consent of the requester. The request **MAY** optionally also contain relevant references as well as a descriptive comment as defined above.

GANAs has assigned numbers for the record types defined in this specification in the "GNU Name System Record Types" registry as listed in Figure 25.

Number	Name	Contact	References	Comment
65536	PKEY	(*)	[This.I-D]	GNS zone delegation (PKEY)
65537	NICK	(*)	[This.I-D]	GNS zone nickname
65538	LEHO	(*)	[This.I-D]	GNS legacy hostname
65540	GNS2DNS	(*)	[This.I-D]	Delegation to DNS
65541	BOX	(*)	[This.I-D]	Boxed records
65551	REDIRECT	(*)	[This.I-D]	Redirection record.
65556	EDKEY	(*)	[This.I-D]	GNS zone delegation (EDKEY)

(*): `gns-registry@gnunet.org`

Figure 25: The GANA Resource Record Registry.

10.2. .alt Subdomains Registry

GANAs [GANAs] manages the ".alt Subdomains" registry. Each entry has the following format:

- * **Label:** The label of the subdomain (in DNS LDH format as defined in Section 2.3.1 of [RFC5890]).
- * **Comment:** Optionally, a brief English text describing the purpose of the subdomain (in UTF-8)
- * **Contact:** Optionally, the contact information of a person to contact for further information.
- * **References:** Optionally, references describing the record type (such as an RFC).

The registration policy for this registry is "First Come First Served". This policy is modeled on that described in [RFC8126], and describes the actions taken by GANA:

Adding new entries is possible after review by any authorized GANA contributor, using a first-come-first-served policy for unique subdomain allocation. Reviewers are responsible to ensure that the chosen "Subdomain" is appropriate for the purpose.

Authorized GANA contributors for review of new entries are reachable at alt-registry@gnunet.org.

Any request MUST contain a unique subdomain and a point of contact. The contact information MAY be added to the registry given the consent of the requester. The request MAY optionally also contain relevant references as well as a descriptive comment as defined above.

GANA has assigned the subdomain defined in this specification in the ".alt subdomains" registry as listed in Figure 26.

Subdomain	Contact	References	Comment
gns	(*)	[This.I-D]	The .alt subdomain for GNS.

(*): alt-registry@gnunet.org

Figure 26: The GANA .alt Subdomains Registry.

11. IANA Considerations

This document makes no requests for IANA action. This section may be removed on publication as an RFC.

12. Implementation and Deployment Status

There are two implementations conforming to this specification written in C and Go, respectively. The C implementation as part of GNUnet [GNUnetGNS] represents the original and reference implementation. The Go implementation [GoGNS] demonstrates how two implementations of GNS are interoperable if they are built on top of the same underlying DHT storage.

Currently, the GUNet peer-to-peer network [GUNet] is an active deployment of GNS on top of its [R5N] DHT. The [GoGNS] implementation uses this deployment by building on top of the GUNet DHT services available on any GUNet peer. It shows how GNS implementations can attach to this existing deployment and participate in name resolution as well as zone publication.

The self-sovereign identity system re:claimID [reclaim] is using GNS in order to selectively share identity attributes and attestations with third parties.

The Ascension tool [Ascension] facilitates the migration of DNS zones to GNS zones by translating information retrieved from a DNS zone transfer into a GNS zone.

13. Acknowledgements

The authors thank all reviewers for their comments. In particular, we thank D. J. Bernstein, S. Bortzmeyer, A. Farrel, E. Lear and R. Salz for their insightful and detailed technical reviews. We thank J. Yao and J. Klensin for the internationalization reviews. We thank NLnet and NGI DISCOVERY for funding work on the GNU Name System.

14. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <<https://www.rfc-editor.org/info/rfc1035>>.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, DOI 10.17487/RFC2782, February 2000, <<https://www.rfc-editor.org/info/rfc2782>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <<https://www.rfc-editor.org/info/rfc3629>>.

- [RFC3686] Housley, R., "Using Advanced Encryption Standard (AES) Counter Mode With IPsec Encapsulating Security Payload (ESP)", RFC 3686, DOI 10.17487/RFC3686, January 2004, <<https://www.rfc-editor.org/info/rfc3686>>.
- [RFC3826] Blumenthal, U., Maino, F., and K. McCloghrie, "The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model", RFC 3826, DOI 10.17487/RFC3826, June 2004, <<https://www.rfc-editor.org/info/rfc3826>>.
- [RFC5237] Arkko, J. and S. Bradner, "IANA Allocation Guidelines for the Protocol Field", BCP 37, RFC 5237, DOI 10.17487/RFC5237, February 2008, <<https://www.rfc-editor.org/info/rfc5237>>.
- [RFC5869] Krawczyk, H. and P. Eronen, "HMAC-based Extract-and-Expand Key Derivation Function (HKDF)", RFC 5869, DOI 10.17487/RFC5869, May 2010, <<https://www.rfc-editor.org/info/rfc5869>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<https://www.rfc-editor.org/info/rfc5890>>.
- [RFC5895] Resnick, P. and P. Hoffman, "Mapping Characters for Internationalized Domain Names in Applications (IDNA) 2008", RFC 5895, DOI 10.17487/RFC5895, September 2010, <<https://www.rfc-editor.org/info/rfc5895>>.
- [RFC6234] Eastlake 3rd, D. and T. Hansen, "US Secure Hash Algorithms (SHA and SHA-based HMAC and HKDF)", RFC 6234, DOI 10.17487/RFC6234, May 2011, <<https://www.rfc-editor.org/info/rfc6234>>.
- [RFC6895] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", BCP 42, RFC 6895, DOI 10.17487/RFC6895, April 2013, <<https://www.rfc-editor.org/info/rfc6895>>.
- [RFC6979] Pornin, T., "Deterministic Usage of the Digital Signature Algorithm (DSA) and Elliptic Curve Digital Signature Algorithm (ECDSA)", RFC 6979, DOI 10.17487/RFC6979, August 2013, <<https://www.rfc-editor.org/info/rfc6979>>.
- [RFC7748] Langley, A., Hamburg, M., and S. Turner, "Elliptic Curves for Security", RFC 7748, DOI 10.17487/RFC7748, January 2016, <<https://www.rfc-editor.org/info/rfc7748>>.

- [RFC8032] Josefsson, S. and I. Liusvaara, "Edwards-Curve Digital Signature Algorithm (EdDSA)", RFC 8032, DOI 10.17487/RFC8032, January 2017, <<https://www.rfc-editor.org/info/rfc8032>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [RFC8499] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", BCP 219, RFC 8499, DOI 10.17487/RFC8499, January 2019, <<https://www.rfc-editor.org/info/rfc8499>>.
- [RFC9106] Biryukov, A., Dinu, D., Khovratovich, D., and S. Josefsson, "Argon2 Memory-Hard Function for Password Hashing and Proof-of-Work Applications", RFC 9106, DOI 10.17487/RFC9106, September 2021, <<https://www.rfc-editor.org/info/rfc9106>>.
- [GANA] GNUnet e.V., "GNUnet Assigned Numbers Authority (GANA)", November 2022, <<https://gana.gnunet.org/>>.
- [MODES] Dworkin, M., "Recommendation for Block Cipher Modes of Operation: Methods and Techniques", December 2001, <<https://doi.org/10.6028/NIST.SP.800-38A>>.
- [CrockfordB32] Douglas, D., "Base32", March 2019, <<https://www.crockford.com/base32.html>>.
- [XSalsa20] Bernstein, D., "Extending the Salsa20 nonce", 2011, <<https://cr.yp.to/snuffle/xsalsa-20110204.pdf>>.
- [Unicode-UAX15] The Unicode Consortium, "Unicode Standard Annex #15: Unicode Normalization Forms, Revision 31", September 2009, <<http://www.unicode.org/reports/tr15/tr15-31.html>>.
- [Unicode-UTS46] The Unicode Consortium, "Unicode Technical Standard #46: Unicode IDNA Compatibility Processing, Revision 27", August 2021, <<https://www.unicode.org/reports/tr46>>.

15. Informative References

- [RFC1928] Leech, M., Ganis, M., Lee, Y., Kuris, R., Koblas, D., and L. Jones, "SOCKS Protocol Version 5", RFC 1928, DOI 10.17487/RFC1928, March 1996, <<https://www.rfc-editor.org/info/rfc1928>>.
- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <<https://www.rfc-editor.org/info/rfc4033>>.
- [RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI 10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/info/rfc6066>>.
- [RFC7363] Maenpaa, J. and G. Camarillo, "Self-Tuning Distributed Hash Table (DHT) for REsource LOcation And Discovery (RELOAD)", RFC 7363, DOI 10.17487/RFC7363, September 2014, <<https://www.rfc-editor.org/info/rfc7363>>.
- [RFC8324] Klensin, J., "DNS Privacy, Authorization, Special Uses, Encoding, Characters, Matching, and Root Structure: Time for Another Look?", RFC 8324, DOI 10.17487/RFC8324, February 2018, <<https://www.rfc-editor.org/info/rfc8324>>.
- [RFC8806] Kumari, W. and P. Hoffman, "Running a Root Server Local to a Resolver", RFC 8806, DOI 10.17487/RFC8806, June 2020, <<https://www.rfc-editor.org/info/rfc8806>>.
- [RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <<https://www.rfc-editor.org/info/rfc6761>>.
- [RFC8244] Lemon, T., Droms, R., and W. Kumari, "Special-Use Domain Names Problem Statement", RFC 8244, DOI 10.17487/RFC8244, October 2017, <<https://www.rfc-editor.org/info/rfc8244>>.
- [I-D.ietf-dnsop-alt-tld]
Kumari, W. A. and P. E. Hoffman, "The ALT Special Use Top Level Domain", Work in Progress, Internet-Draft, draft-ietf-dnsop-alt-tld-25, 4 May 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-alt-tld-25>>.

- [Tor224] Goulet, D., Kadianakis, G., and N. Mathewson, "Next-Generation Hidden Services in Tor", November 2013, <<https://gitweb.torproject.org/torspec.git/tree/proposals/224-rend-spec-ng.txt#n2135>>.
- [SDSI] Rivest, R. and B. Lampson, "SDSI - A Simple Distributed Security Infrastructure", April 1996, <<http://people.csail.mit.edu/rivest/Sdsi10.ps>>.
- [Kademlia] Maymounkov, P. and D. Mazieres, "Kademlia: A peer-to-peer information system based on the xor metric.", 2002, <<http://css.csail.mit.edu/6.824/2014/papers/kademlia.pdf>>.
- [ed25519] Bernstein, D., Duif, N., Lange, T., Schwabe, P., and B. Yang, "High-Speed High-Security Signatures", 2011, <<https://ed25519.cr.yip.to/ed25519-20110926.pdf>>.
- [GNS] Wachs, M., Schanzenbach, M., and C. Grothoff, "A Censorship-Resistant, Privacy-Enhancing and Fully Decentralized Name System", 2014, <https://sci-hub.st/10.1007/978-3-319-12280-9_9>.
- [R5N] Evans, N. S. and C. Grothoff, "R5N: Randomized recursive routing for restricted-route networks", 2011, <<https://sci-hub.st/10.1109/ICNSS.2011.6060022>>.
- [SecureNS] Grothoff, C., Wachs, M., Ermert, M., and J. Appelbaum, "Towards secure name resolution on the Internet", 2018, <<https://sci-hub.st/https://doi.org/10.1016/j.cose.2018.01.018>>.
- [GNUnetGNS] GNUnet e.V., "The GNUnet GNS Implementation", <<https://git.gnunet.org/gnunet.git/tree/src/gns>>.
- [Ascension] GNUnet e.V., "The Ascension Implementation", <<https://git.gnunet.org/ascension.git>>.
- [GNUnet] GNUnet e.V., "The GNUnet Project", <<https://gnunet.org>>.
- [reclaim] GNUnet e.V., "re:claimID", <<https://reclaim.gnunet.org>>.
- [GoGNS] Fix, B., "The Go GNS Implementation", <<https://github.com/bfix/gnunet-go/tree/master/src/gnunet/service/gns>>.

[nsswitch] GNU Project, "System Databases and Name Service Switch",
<https://www.gnu.org/software/libc/manual/html_node/Name-Service-Switch.html>.

Appendix A. Usage and Migration

This section outlines a number of specific use cases which may help readers of the technical specification to understand the protocol better. The considerations below are not meant to be normative for the GNS protocol in any way. Instead, they are provided in order to give context and to provide some background on what the intended use of the protocol is by its designers. Further, this section contains pointers to migration paths.

A.1. Zone Dissemination

In order to become a zone owner, it is sufficient to generate a zone key and a corresponding secret key using a GNS implementation. At this point, the zone owner can manage GNS resource records in a local zone database. The resource records can then be published by a GNS implementation as defined in Section 6. For other users to resolve the resource records, respective zone information must be disseminated first. The zone owner may decide to make the zone key and labels known to a selected set of users only or to make this information available to the general public.

Sharing zone information directly with specific users not only allows to potentially preserve zone and record privacy, but also allows the zone owner and the user to establish strong trust relationships. For example, a bank may send a customer letter with a QR code which contains the GNS zone of the bank. This allows the user to scan the QR code and establish a strong link to the zone of the bank and with it, for example, the IP address of the online banking web site.

Most Internet services likely want to make their zones available to the general public as efficiently as possible. First, it is reasonable to assume that zones which are commanding high levels of reputation and trust are likely included in the default suffix-to-zone mappings of implementations. Hence dissemination of a zone through delegation under such zones can be a viable path in order to disseminate a zone publicly. For example, it is conceivable that organizations such as ICANN or country-code top-level domain registrars also manage GNS zones and offer registration or delegation services.

Following best practices in particularly those relating to security and abuse mitigation are methods which allow zone owners and aspiring registrars to gain a good reputation and eventually trust. This

includes, of course, diligent protection of private zone key material. Formalizing such best practices is out of scope of this specification and should be addressed in a separate document and take Section 9 into account.

A.2. Start Zone Configuration

A user is expected to install a GNS implementation if it is not already provided through other means such as the operating system or the browser. It is likely that the implementation ships with a default start zone configuration. This means that the user is able to resolve GNS names ending on a zTLD or ending on any suffix-to-name mapping that is part of the default start zone configuration. At this point the user may delete or otherwise modify the implementation's default configuration:

Deletion of suffix-to-zone mappings may become necessary if the zone owner referenced by the mapping has lost the trust of the user. For example, this could be due to lax registration policies resulting in phishing activities. Modification and addition of new mappings are means to heal the namespace perforation which would occur in the case of a deletion or to simply establish a strong direct trust relationship. However, this requires the user's knowledge of the respective zone keys. This information must be retrieved out of band, as illustrated in Appendix A.1: A bank may send the user a letter with a QR code which contains the GNS zone of the bank. The user scans the QR code and adds a new suffix-to-name mapping using a chosen local name for his bank. Other examples include scanning zone information off the device of a friend, from a storefront, or an advertisement. The level of trust in the respective zone is contextual and likely varies from user to user. Trust in a zone provided through a letter from a bank which may also include a credit card is certainly different from a zone found on a random advertisement in the streets. However, this trust is immediately tangible to the user and can be reflected in the local naming as well.

User clients should facilitate the modification of the start zone configuration, for example by providing a QR code reader or other import mechanisms. Implementations are ideally implemented according to best practices and addressing applicable points from Section 9. Formalizing such best practices is out of scope of this specification.

A.3. Globally Unique Names and the Web

HTTP virtual hosting and TLS Server Name Indication are common use cases on the Web. HTTP clients supply a DNS name in the HTTP "Host"-header or as part of the TLS handshake, respectively. This allows the HTTP server to serve the indicated virtual host with a matching TLS certificate. The global uniqueness of DNS names are a prerequisite of those use cases.

Not all GNS names are globally unique. But, any resource record in GNS can be represented as a concatenation of a GNS label and the zTLD of the zone. While not human-readable, this globally unique GNS name can be leveraged in order to facilitate the same use cases. Consider the GNS name "www.example.gns" entered in a GNS-aware HTTP client. At first, "www.example.gns" is resolved using GNS yielding a record set. Then, the HTTP client determines the virtual host as follows:

If there is a LEHO record (Section 5.3.1) containing "www.example.com" in the record set, then the HTTP client uses this as the value of the "Host"-header field of the HTTP request:

```
GET / HTTP/1.1
Host: www.example.com
```

If there is no LEHO record in the record set, then the HTTP client tries to find the zone of the record and translates the GNS name into a globally unique zTLD-representation before using it in the "Host"-header field of the HTTP request:

```
GET / HTTP/1.1
Host: www.000G0037FH3QTBCK15Y8BCCNRVWPV17ZC7TSGB1C9ZG2TPGHZVVF1GMG3W
```

In order to determine the canonical representation of the record with a zTLD, at most two queries are required: First, it must be checked whether "www.example.gns.alt" itself points to a zone delegation record which would imply that the record set which was originally resolved is published under the apex label. If it does, the unique GNS name is simply the zTLD representation of the delegated zone:

```
GET / HTTP/1.1
Host: 000G0037FH3QTBCK15Y8BCCNRVWPV17ZC7TSGB1C9ZG2TPGHZVVF1GMG3W
```

If it does not, the unique GNS name is the concatenation of the label "www" and the zTLD representation of the zone as given in the example above. In any case, this representation is globally unique. As such, it can be configured by the HTTP server administrator as a virtual host name and respective certificates may be issued.

If the HTTP client is a browser, the use of a unique GNS name for virtual hosting or TLS SNI does not necessarily have to be shown to the user. For example, the name in the URL bar may remain as "www.example.gns.alt" even if the used unique name differs.

A.4. Migration Paths

DNS resolution is built into a variety of existing software components. Most significantly operating systems and HTTP clients. This section illustrates possible migration paths for both in order to enable "legacy" applications to resolve GNS names.

One way to efficiently facilitate the resolution of GNS names are GNS-enabled DNS server implementations. Local DNS queries are thereby either rerouted or explicitly configured to be resolved by a "DNS-to-GNS" server that runs locally. This DNS server tries to interpret any incoming query for a name as a GNS resolution request. If no start zone can be found for the name and it does not end in a zTLD, the server tries to resolve the name in DNS. Otherwise, the name is resolved in GNS. In the latter case, the resulting record set is converted to a DNS answer packet and is returned accordingly. An implementation of a DNS-to-GNS server can be found in [GNUnet].

A similar approach is to use operating systems extensions such as the name service switch [nsswitch]. It allows the system administrator to configure plugins which are used for hostname resolution. A GNS name service switch plugin can be used in a similar fashion as the "DNS-to-GNS" server. An implementation of a glibc-compatible nsswitch plugin for GNS can be found in [GNUnet].

The methods above are usually also effective for HTTP client software. However, HTTP clients are commonly used in combination with TLS. TLS certificate validation and in particular server name indication (SNI) requires additional logic in HTTP clients when GNS names are in play (Appendix A.3). In order to transparently enable this functionality for migration purposes, a local GNS-aware SOCKS5 proxy [RFC1928] can be configured to resolve domain names. The SOCKS5 proxy, similar to the DNS-to-GNS server, is capable of resolving both GNS and DNS names. In the event of a TLS connection request with a GNS name, the SOCKS5 proxy can act as a man-in-the-middle, terminating the TLS connection and establishing a secure connection against the requested host. In order to establish a secure connection, the proxy may use LEHO and TLSA records stored in the record set under the GNS name. The proxy must provide a locally trusted certificate for the GNS name to the HTTP client which usually requires the generation and configuration of a local trust anchor in the browser. An implementation of this SOCKS5 proxy can be found in [GNUnet].

Appendix B. Example flows

B.1. AAAA Example Resolution

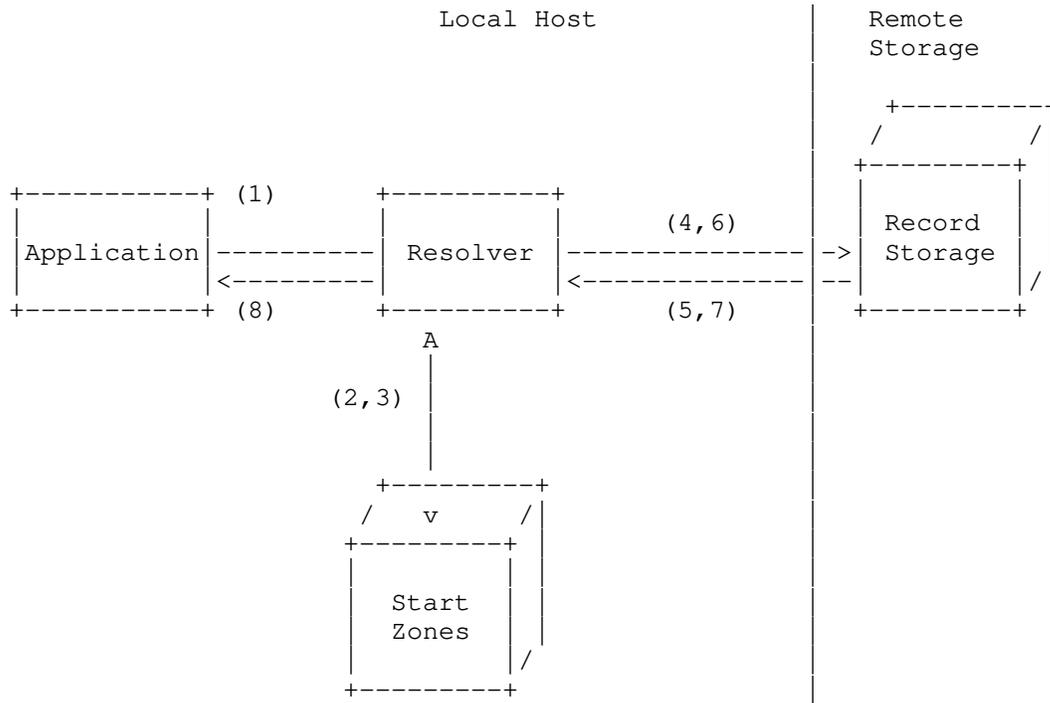


Figure 27: Example resolution of an IPv6 address.

1. Lookup AAAA record for name: `www.example.gnu.gns.alt`.
2. Determine start zone for `www.example.gnu.gns.alt`.
3. Start zone: `zk0` - Remainder: `www.example`.
4. Calculate `q0=SHA512(ZKDF(zk0, "example"))` and initiate `GET(q0)`.
5. Retrieve and decrypt RRBLOCK consisting of a single PKEY record containing `zk1`.
6. Calculate `q1=SHA512(ZKDF(zk1, "www"))` and initiate `GET(q1)`.
7. Retrieve RRBLOCK consisting of a single AAAA record containing the IPv6 address `2001:db8::1`.
8. Return record set to application

B.2. REDIRECT Example Resolution

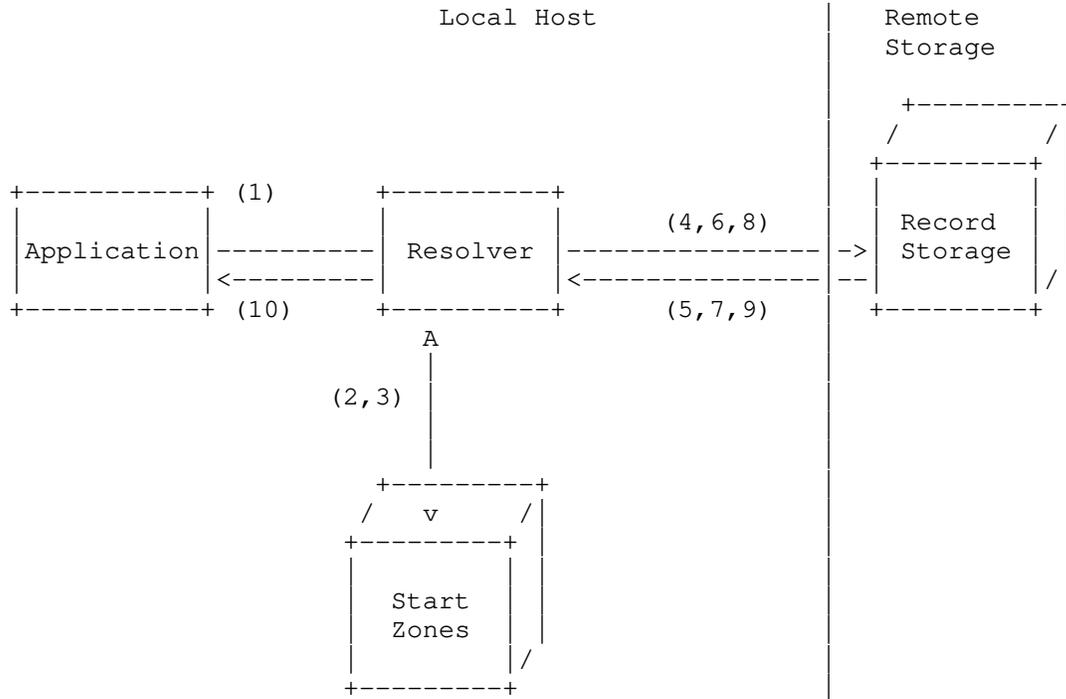


Figure 28: Example resolution of an IPv6 address with redirect.

1. Lookup AAAA record for name: www.example.tld.gns.alt.
2. Determine start zone for www.example.tld.gns.alt.
3. Start zone: zk0 - Remainder: www.example.
4. Calculate $q_0 = \text{SHA512}(\text{ZKDF}(\text{zk}_0, \text{"example"}))$ and initiate GET(q_0).
5. Retrieve and decrypt RRBLOCK consisting of a single PKEY record containing zk1.
6. Calculate $q_1 = \text{SHA512}(\text{ZKDF}(\text{zk}_1, \text{"www"}))$ and initiate GET(q_1).
7. Retrieve and decrypt RRBLOCK consisting of a single REDIRECT record containing www2.+.
8. Calculate $q_2 = \text{SHA512}(\text{ZKDF}(\text{zk}_1, \text{"www2"}))$ and initiate GET(q_2).

9. Retrieve and decrypt RRBLOCK consisting of a single AAAA record containing the IPv6 address 2001:db8::1.
10. Return record set to application.

B.3. GNS2DNS Example Resolution

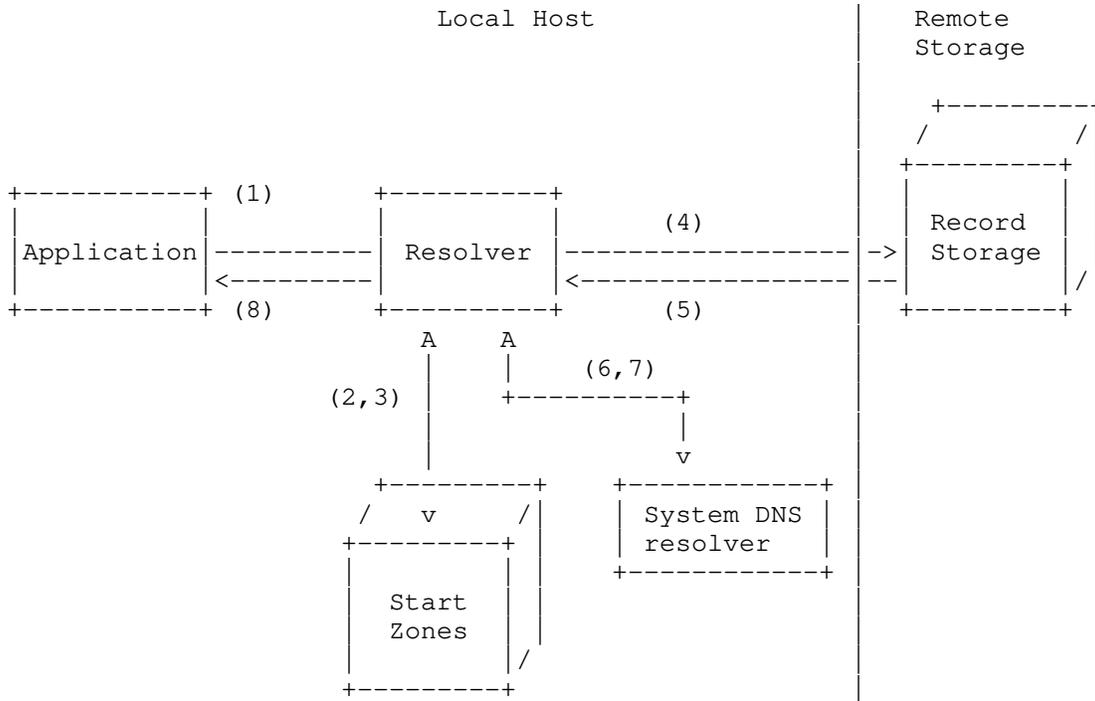


Figure 29: Example resolution of an IPv6 address with DNS handover.

1. Lookup AAAA record for name: www.example.gnu.gns.alt
2. Determine start zone for www.example.gnu.gns.alt.
3. Start zone: zk0 - Remainder: www.example.
4. Calculate $q_0 = \text{SHA512}(\text{ZKDF}(\text{zk}_0, \text{"example"}))$ and initiate GET(q_0).
5. Retrieve and decrypt RRBLOCK consisting of a single GNS2DNS record containing the name example.com and the DNS server IPv4 address 192.0.2.1.
6. Use system resolver to lookup an AAAA record for the DNS name www.example.com.

7. Retrieve a DNS reply consisting of a single AAAA record containing the IPv6 address 2001:db8::1.
8. Return record set to application.

Appendix C. Base32GNS

Encoding converts a byte array into a string of symbols. Decoding converts a string of symbols into a byte array. Decoding fails if the input string has symbols outside the defined set.

This table defines the encode and decode symbols for a given symbol value. Each symbol value encodes 5 bits. It can be used to implement the encoding by reading it as: A symbol "A" or "a" is decoded to a 5 bit value 10 when decoding. A 5 bit block with a value of 18 is encoded to the character "J" when encoding. If the bit length of the byte string to encode is not a multiple of 5 it is padded to the next multiple with zeroes. In order to further increase tolerance for failures in character recognition, the letter "U" MUST be decoded to the same value as the letter "V" in Base32GNS.

Symbol Value	Decode Symbol	Encode Symbol
0	0 O o	0
1	1 I i L l	1
2	2	2
3	3	3
4	4	4
5	5	5
6	6	6
7	7	7
8	8	8
9	9	9
10	A a	A
11	B b	B
12	C c	C
13	D d	D
14	E e	E
15	F f	F
16	G g	G
17	H h	H
18	J j	J
19	K k	K
20	M m	M
21	N n	N
22	P p	P
23	Q q	Q
24	R r	R
25	S s	S
26	T t	T
27	V v U u	V
28	W w	W
29	X x	X
30	Y y	Y
31	Z z	Z

Figure 30: The Base32GNS Alphabet Including the Additional U Encode Symbol.

Appendix D. Test Vectors

The following test vectors can be used by implementations to test for conformance with this specification. Unless indicated otherwise, the test vectors are provided as hexadecimal byte arrays.

D.1. Base32GNS en-/decoding

The following are test vectors for the Base32GNS encoding used for zTLDs. The input strings are encoded without the zero terminator.

Base32GNS-Encode:

Input string: "Hello World"

Output string: "91JPRV3F41BPYWKCCG"

Input bytes: 474e55204e616d652053797374656d

Output string: "8X75A82EC5PPA82KF5SQ8SBD"

Base32GNS-Decode:

Input string: "91JPRV3F41BPYWKCCG"

Output string: "Hello World"

Input string: "91JPRU3F41BPYWKCCG"

Output string: "Hello World"

D.2. Record sets

The test vectors include record sets with a variety of record types and flags for both PKEY and EDKEY zones. This includes labels with UTF-8 characters to demonstrate internationalized labels.

* (1) PKEY zone with ASCII label and one delegation record*

Zone private key (d, big-endian):

50 d7 b6 52 a4 ef ea df

f3 73 96 90 97 85 e5 95

21 71 a0 21 78 c8 e7 d4

50 fa 90 79 25 fa fd 98

Zone identifier (ztype|zkey):

00 01 00 00 67 7c 47 7d

2d 93 09 7c 85 b1 95 c6

f9 6d 84 ff 61 f5 98 2c

2c 4f e0 2d 5a 11 fe df

b0 c2 90 1f

zTLD:

000G0037FH3QTBCK15Y8BCCNRVWPV17ZC7TSGB1C9ZG2TPGHZVVFV1GMG3W

Label:

74 65 73 74 64 65 6c 65

67 61 74 69 6f 6e

Number of records (integer): 1

Record #0 := (

EXPIRATION: 8143584694000000 us

```
00 1c ee 8c 10 e2 59 80
```

```
DATA_SIZE:
```

```
00 20
```

```
TYPE:
```

```
00 01 00 00
```

```
FLAGS: 00 01
```

```
DATA:
```

```
21 e3 b3 0f f9 3b c6 d3
```

```
5a c8 c6 e0 e1 3a fd ff
```

```
79 4c b7 b4 4b bb c7 48
```

```
d2 59 d0 a0 28 4d be 84
```

```
)
```

```
RDATA:
```

```
00 1c ee 8c 10 e2 59 80
```

```
00 20 00 01 00 01 00 00
```

```
21 e3 b3 0f f9 3b c6 d3
```

```
5a c8 c6 e0 e1 3a fd ff
```

```
79 4c b7 b4 4b bb c7 48
```

```
d2 59 d0 a0 28 4d be 84
```

```
Encryption NONCE|EXPIRATION|BLOCK COUNTER:
```

```
e9 0a 00 61 00 1c ee 8c
```

```
10 e2 59 80 00 00 00 01
```

```
Encryption key (K):
```

```
86 4e 71 38 ea e7 fd 91
```

```
a3 01 36 89 9c 13 2b 23
```

```
ac eb db 2c ef 43 cb 19
```

```
f6 bf 55 b6 7d b9 b3 b3
```

```
Storage key (q):
```

```
4a dc 67 c5 ec ee 9f 76
```

```
98 6a bd 71 c2 22 4a 3d
```

```
ce 2e 91 70 26 c9 a0 9d
```

```
fd 44 ce f3 d2 0f 55 a2
```

```
73 32 72 5a 6c 8a fb bb
```

```
b0 f7 ec 9a f1 cc 42 64
```

```
12 99 40 6b 04 fd 9b 5b
```

```
57 91 f8 6c 4b 08 d5 f4
```

```
ZKDF (zkey):
```

```
18 2b b6 36 ed a7 9f 79
```

```
57 11 bc 27 08 ad bb 24
2a 60 44 6a d3 c3 08 03
12 1d 03 d3 48 b7 ce b6
```

Derived private key (d', big-endian):

```
0a 4c 5e 0f 00 63 df ce
db c8 c7 f2 b2 2c 03 0c
86 28 b2 c2 cb ac 9f a7
29 aa e6 1f 89 db 3e 9c
```

BDATA:

```
0c 1e da 5c c0 94 a1 c7
a8 88 64 9d 25 fa ee bd
60 da e6 07 3d 57 d8 ae
8d 45 5f 4f 13 92 c0 74
e2 6a c6 69 bd ee c2 34
62 b9 62 95 2c c6 e9 eb
```

RRBLOCK:

```
00 00 00 a0 00 01 00 00
18 2b b6 36 ed a7 9f 79
57 11 bc 27 08 ad bb 24
2a 60 44 6a d3 c3 08 03
12 1d 03 d3 48 b7 ce b6
0a d1 0b c1 3b 40 3b 5b
25 61 26 b2 14 5a 6f 60
c5 14 f9 51 ff a7 66 f7
a3 fd 4b ac 4a 4e 19 90
05 5c b8 7e 8d 1b fd 19
aa 09 a4 29 f7 29 e9 f5
c6 ee c2 47 0a ce e2 22
07 59 e9 e3 6c 88 6f 35
00 1c ee 8c 10 e2 59 80
0c 1e da 5c c0 94 a1 c7
a8 88 64 9d 25 fa ee bd
60 da e6 07 3d 57 d8 ae
8d 45 5f 4f 13 92 c0 74
e2 6a c6 69 bd ee c2 34
62 b9 62 95 2c c6 e9 eb
```

(2) PKEY zone with UTF-8 label and three records

Zone private key (d, big-endian):

```
50 d7 b6 52 a4 ef ea df
f3 73 96 90 97 85 e5 95
21 71 a0 21 78 c8 e7 d4
50 fa 90 79 25 fa fd 98
```

Zone identifier (ztype|zkey):

```
00 01 00 00 67 7c 47 7d
2d 93 09 7c 85 b1 95 c6
f9 6d 84 ff 61 f5 98 2c
2c 4f e0 2d 5a 11 fe df
b0 c2 90 1f
```

zTLD:

```
000G0037FH3QTBCK15Y8BCCNRVWPV17ZC7TSGB1C9ZG2TPGHZVFV1GMG3W
```

Label:

```
e5 a4 a9 e4 b8 8b e7 84
a1 e6 95 b5
```

Number of records (integer): 3

Record #0 := (

```
EXPIRATION: 8143584694000000 us
00 1c ee 8c 10 e2 59 80
```

DATA_SIZE:

```
00 10
```

TYPE:

```
00 00 00 1c
```

FLAGS: 00 00

DATA:

```
00 00 00 00 00 00 00 00
00 00 00 00 de ad be ef
```

)

Record #1 := (

```
EXPIRATION: 17999736901000000 us
00 3f f2 aa 54 08 db 40
```

DATA_SIZE:

```
00 06
```

TYPE:

```

00 01 00 01

FLAGS:  00 00

DATA:
e6 84 9b e7 a7 b0

)

Record #2 := (
  EXPIRATION: 11464693629000000 us
  00 28 bb 13 ff 37 19 40

  DATA_SIZE:
  00 0b

  TYPE:
  00 00 00 10

  FLAGS:  00 04

  DATA:
  48 65 6c 6c 6f 20 57 6f
  72 6c 64

)

```

```

RDATA:
00 1c ee 8c 10 e2 59 80
00 10 00 00 00 00 00 1c
00 00 00 00 00 00 00 00
00 00 00 00 de ad be ef
00 3f f2 aa 54 08 db 40
00 06 00 00 00 01 00 01
e6 84 9b e7 a7 b0 00 28
bb 13 ff 37 19 40 00 0b
00 04 00 00 00 10 48 65
6c 6c 6f 20 57 6f 72 6c
64 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00

```

```

Encryption NONCE|EXPIRATION|BLOCK COUNTER:
ee 96 33 c1 00 1c ee 8c
10 e2 59 80 00 00 00 01

```

Encryption key (K):

```
fb 3a b5 de 23 bd da e1
99 7a af 7b 92 c2 d2 71
51 40 8b 77 af 7a 41 ac
79 05 7c 4d f5 38 3d 01
```

Storage key (q):

```
af f0 ad 6a 44 09 73 68
42 9a c4 76 df a1 f3 4b
ee 4c 36 e7 47 6d 07 aa
64 63 ff 20 91 5b 10 05
c0 99 1d ef 91 fc 3e 10
90 9f 87 02 c0 be 40 43
67 78 c7 11 f2 ca 47 d5
5c f0 b5 4d 23 5d a9 77
```

ZKDF (zkey):

```
a5 12 96 df 75 7e e2 75
ca 11 8d 4f 07 fa 7a ae
55 08 bc f5 12 aa 41 12
14 29 d4 a0 de 9d 05 7e
```

Derived private key (d', big-endian):

```
0a be 56 d6 80 68 ab 40
e1 44 79 0c de 9a cf 4d
78 7f 2d 3c 63 b8 53 05
74 6e 68 03 32 15 f2 ab
```

BDATA:

```
d8 c2 8d 2f d6 96 7d 1a
b7 22 53 f2 10 98 b8 14
a4 10 be 1f 59 98 de 03
f5 8f 7e 7c db 7f 08 a6
16 51 be 4d 0b 6f 8a 61
df 15 30 44 0b d7 47 dc
f0 d7 10 4f 6b 8d 24 c2
ac 9b c1 3d 9c 6f e8 29
05 25 d2 a6 d0 f8 84 42
67 a1 57 0e 8e 29 4d c9
3a 31 9f cf c0 3e a2 70
17 d6 fd a3 47 b4 a7 94
97 d7 f6 b1 42 2d 4e dd
82 1c 19 93 4e 96 c1 aa
87 76 57 25 d4 94 c7 64
b1 55 dc 6d 13 26 91 74
```

RRBLOCK:

```
00 00 00 f0 00 01 00 00
```

```

a5 12 96 df 75 7e e2 75
ca 11 8d 4f 07 fa 7a ae
55 08 bc f5 12 aa 41 12
14 29 d4 a0 de 9d 05 7e
08 5b d6 5f d4 85 10 51
ba ce 2a 45 2a fc 8a 7e
4f 6b 2c 1f 74 f0 20 35
d9 64 1a cd ba a4 66 e0
00 ce d6 f2 d2 3b 63 1c
8e 8a 0b 38 e2 ba e7 9a
22 ca d8 1d 4c 50 d2 25
35 8e bc 17 ac 0f 89 9e
00 1c ee 8c 10 e2 59 80
d8 c2 8d 2f d6 96 7d 1a
b7 22 53 f2 10 98 b8 14
a4 10 be 1f 59 98 de 03
f5 8f 7e 7c db 7f 08 a6
16 51 be 4d 0b 6f 8a 61
df 15 30 44 0b d7 47 dc
f0 d7 10 4f 6b 8d 24 c2
ac 9b c1 3d 9c 6f e8 29
05 25 d2 a6 d0 f8 84 42
67 a1 57 0e 8e 29 4d c9
3a 31 9f cf c0 3e a2 70
17 d6 fd a3 47 b4 a7 94
97 d7 f6 b1 42 2d 4e dd
82 1c 19 93 4e 96 c1 aa
87 76 57 25 d4 94 c7 64
b1 55 dc 6d 13 26 91 74

```

(3) EDKEY zone with ASCII label and delegation record

Zone private key (d):

```

5a f7 02 0e e1 91 60 32
88 32 35 2b bc 6a 68 a8
d7 1a 7c be 1b 92 99 69
a7 c6 6d 41 5a 0d 8f 65

```

Zone identifier (ztype|zkey):

```

00 01 00 14 3c f4 b9 24
03 20 22 f0 dc 50 58 14
53 b8 5d 93 b0 47 b6 3d
44 6c 58 45 cb 48 44 5d
db 96 68 8f

```

zTLD:

```
000G051WYJWJ80S04BRDRM2R2H9VGQCKP13VCFA4DHC4BJT88HEXQ5K8HW
```

Label:

```
74 65 73 74 64 65 6c 65
67 61 74 69 6f 6e
```

```
Number of records (integer): 1
```

```
Record #0 := (
```

```
  EXPIRATION: 8143584694000000 us
  00 1c ee 8c 10 e2 59 80
```

```
  DATA_SIZE:
```

```
  00 20
```

```
  TYPE:
```

```
  00 01 00 00
```

```
  FLAGS: 00 01
```

```
  DATA:
```

```
  21 e3 b3 0f f9 3b c6 d3
  5a c8 c6 e0 e1 3a fd ff
  79 4c b7 b4 4b bb c7 48
  d2 59 d0 a0 28 4d be 84
```

```
)
```

```
RDATA:
```

```
00 1c ee 8c 10 e2 59 80
00 20 00 01 00 01 00 00
21 e3 b3 0f f9 3b c6 d3
5a c8 c6 e0 e1 3a fd ff
79 4c b7 b4 4b bb c7 48
d2 59 d0 a0 28 4d be 84
```

```
Encryption NONCE|EXPIRATION:
```

```
98 13 2e a8 68 59 d3 5c
88 bf d3 17 fa 99 1b cb
00 1c ee 8c 10 e2 59 80
```

```
Encryption key (K):
```

```
85 c4 29 a9 56 7a a6 33
41 1a 96 91 e9 09 4c 45
28 16 72 be 58 60 34 aa
e4 a2 a2 cc 71 61 59 e2
```

```
Storage key (q):
```

```
ab aa ba c0 e1 24 94 59
75 98 83 95 aa c0 24 1e
```

```
55 59 c4 1c 40 74 e2 55
7b 9f e6 d1 54 b6 14 fb
cd d4 7f c7 f5 1d 78 6d
c2 e0 b1 ec e7 60 37 c0
a1 57 8c 38 4e c6 1d 44
56 36 a9 4e 88 03 29 e9
```

ZKDF(zkey):

```
9b f2 33 19 8c 6d 53 bb
db ac 49 5c ab d9 10 49
a6 84 af 3f 40 51 ba ca
b0 dc f2 1c 8c f2 7a 1a
```

nonce := SHA-256 (dh[32..63] || h):

```
14 f2 c0 6b ed c3 aa 2d
f0 71 13 9c 50 39 34 f3
4b fa 63 11 a8 52 f2 11
f7 3a df 2e 07 61 ec 35
```

Derived private key (d', big-endian):

```
3b 1b 29 d4 23 0b 10 a8
ec 4d a3 c8 6e db 88 ea
cd 54 08 5c 1d db 63 f7
a9 d7 3f 7c cb 2f c3 98
```

BDATA:

```
57 7c c6 c9 5a 14 e7 04
09 f2 0b 01 67 e6 36 d0
10 80 7c 4f 00 37 2d 69
8c 82 6b d9 2b c2 2b d6
bb 45 e5 27 7c 01 88 1d
6a 43 60 68 e4 dd f1 c6
b7 d1 41 6f af a6 69 7c
25 ed d9 ea e9 91 67 c3
```

RRBLOCK:

```
00 00 00 b0 00 01 00 14
9b f2 33 19 8c 6d 53 bb
db ac 49 5c ab d9 10 49
a6 84 af 3f 40 51 ba ca
b0 dc f2 1c 8c f2 7a 1a
9f 56 a8 86 ea 73 9d 59
17 50 8f 9b 75 56 39 f3
a9 ac fa ed ed ca 7f bf
a7 94 b1 92 e0 8b f9 ed
4c 7e c8 59 4c 9f 7b 4e
19 77 4f f8 38 ec 38 7a
8f 34 23 da ac 44 9f 59
```

```
db 4e 83 94 3f 90 72 00
00 1c ee 8c 10 e2 59 80
57 7c c6 c9 5a 14 e7 04
09 f2 0b 01 67 e6 36 d0
10 80 7c 4f 00 37 2d 69
8c 82 6b d9 2b c2 2b d6
bb 45 e5 27 7c 01 88 1d
6a 43 60 68 e4 dd f1 c6
b7 d1 41 6f af a6 69 7c
25 ed d9 ea e9 91 67 c3
```

(4) EDKEY zone with UTF-8 label and three records

Zone private key (d):

```
5a f7 02 0e e1 91 60 32
88 32 35 2b bc 6a 68 a8
d7 1a 7c be 1b 92 99 69
a7 c6 6d 41 5a 0d 8f 65
```

Zone identifier (ztype|zkey):

```
00 01 00 14 3c f4 b9 24
03 20 22 f0 dc 50 58 14
53 b8 5d 93 b0 47 b6 3d
44 6c 58 45 cb 48 44 5d
db 96 68 8f
```

zTLD:

```
000G051WYJWJ80S04BRDRM2R2H9VGQCKP13VCFA4DHC4BJT88HEXQ5K8HW
```

Label:

```
e5 a4 a9 e4 b8 8b e7 84
a1 e6 95 b5
```

Number of records (integer): 3

Record #0 := (

```
EXPIRATION: 8143584694000000 us
00 1c ee 8c 10 e2 59 80
```

DATA_SIZE:

```
00 10
```

TYPE:

```
00 00 00 1c
```

FLAGS: 00 00

```
DATA:
00 00 00 00 00 00 00 00
00 00 00 00 de ad be ef

)

Record #1 := (
  EXPIRATION: 17999736901000000 us
  00 3f f2 aa 54 08 db 40

  DATA_SIZE:
  00 06

  TYPE:
  00 01 00 01

  FLAGS: 00 00

  DATA:
  e6 84 9b e7 a7 b0

)

Record #2 := (
  EXPIRATION: 11464693629000000 us
  00 28 bb 13 ff 37 19 40

  DATA_SIZE:
  00 0b

  TYPE:
  00 00 00 10

  FLAGS: 00 04

  DATA:
  48 65 6c 6c 6f 20 57 6f
  72 6c 64

)

RDATA:
00 1c ee 8c 10 e2 59 80
00 10 00 00 00 00 00 1c
00 00 00 00 00 00 00 00
00 00 00 00 de ad be ef
00 3f f2 aa 54 08 db 40
00 06 00 00 00 01 00 01
```

```
e6 84 9b e7 a7 b0 00 28
bb 13 ff 37 19 40 00 0b
00 04 00 00 00 10 48 65
6c 6c 6f 20 57 6f 72 6c
64 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
00 00 00 00 00 00 00 00
```

Encryption NONCE | EXPIRATION:

```
bb 0d 3f 0f bd 22 42 77
50 da 5d 69 12 16 e6 c9
00 1c ee 8c 10 e2 59 80
```

Encryption key (K):

```
3d f8 05 bd 66 87 aa 14
20 96 28 c2 44 b1 11 91
88 c3 92 56 37 a4 1e 5d
76 49 6c 29 45 dc 37 7b
```

Storage key (q):

```
ba f8 21 77 ee c0 81 e0
74 a7 da 47 ff c6 48 77
58 fb 0d f0 1a 6c 7f bb
52 fc 8a 31 be f0 29 af
74 aa 0d c1 5a b8 e2 fa
7a 54 b4 f5 f6 37 f6 15
8f a7 f0 3c 3f ce be 78
d3 f9 d6 40 aa c0 d1 ed
```

ZKDF (zkey):

```
74 f9 00 68 f1 67 69 53
52 a8 a6 c2 eb 98 48 98
c5 3a cc a0 98 04 70 c6
c8 12 64 cb dd 78 ad 11
```

nonce := SHA-256 (dh[32..63] || h):

```
f8 6a b5 33 8a 74 d7 a1
d2 77 ea 11 ff 95 cb e8
3a cf d3 97 3b b4 ab ca
0a 1b 60 62 c3 7a b3 9c
```

Derived private key (d', big-endian):

```
17 c0 68 a6 c3 f7 20 de
0e 1b 69 ff 3f 53 e0 5d
3f e5 c5 b0 51 25 7a 89
```

a6 3c 1a d3 5a c4 35 58

BDATA:

4e b3 5a 50 d4 0f e1 a4
29 c7 f4 b2 67 a0 59 de
4e 2c 8a 89 a5 ed 53 d3
d4 92 58 59 d2 94 9f 7f
30 d8 a2 0c aa 96 f8 81
45 05 2d 1c da 04 12 49
8f f2 5f f2 81 6e f0 ce
61 fe 69 9b fa c7 2c 15
dc 83 0e a9 b0 36 17 1c
cf ca bb dd a8 de 3c 86
ed e2 95 70 d0 17 4b 82
82 09 48 a9 28 b7 f0 0e
fb 40 1c 10 fe 80 bb bb
02 76 33 1b f7 f5 1b 8d
74 57 9c 14 14 f2 2d 50
1a d2 5a e2 49 f5 bb f2
a6 c3 72 59 d1 75 e4 40
b2 94 39 c6 05 19 cb b1

RRBLOCK:

00 00 01 00 00 01 00 14
74 f9 00 68 f1 67 69 53
52 a8 a6 c2 eb 98 48 98
c5 3a cc a0 98 04 70 c6
c8 12 64 cb dd 78 ad 11
75 6d 2c 15 7a d2 ea 4f
c0 b1 b9 1c 08 03 79 44
61 d3 de f2 0d d1 63 6c
fe dc 03 89 c5 49 d1 43
6c c3 5b 4e 1b f8 89 5a
64 6b d9 a6 f4 6b 83 48
1d 9c 0e 91 d4 e1 be bb
6a 83 52 6f b7 25 2a 06
00 1c ee 8c 10 e2 59 80
4e b3 5a 50 d4 0f e1 a4
29 c7 f4 b2 67 a0 59 de
4e 2c 8a 89 a5 ed 53 d3
d4 92 58 59 d2 94 9f 7f
30 d8 a2 0c aa 96 f8 81
45 05 2d 1c da 04 12 49
8f f2 5f f2 81 6e f0 ce
61 fe 69 9b fa c7 2c 15
dc 83 0e a9 b0 36 17 1c
cf ca bb dd a8 de 3c 86
ed e2 95 70 d0 17 4b 82

```

82 09 48 a9 28 b7 f0 0e
fb 40 1c 10 fe 80 bb bb
02 76 33 1b f7 f5 1b 8d
74 57 9c 14 14 f2 2d 50
1a d2 5a e2 49 f5 bb f2
a6 c3 72 59 d1 75 e4 40
b2 94 39 c6 05 19 cb b1

```

D.3. Zone revocation

The following is an example revocation for a PKEY zone:

Zone private key (d, big-endian):

```

6f ea 32 c0 5a f5 8b fa
97 95 53 d1 88 60 5f d5
7d 8b f9 cc 26 3b 78 d5
f7 47 8c 07 b9 98 ed 70

```

Zone identifier (ztype|zkey):

```

00 01 00 00 2c a2 23 e8
79 ec c4 bb de b5 da 17
31 92 81 d6 3b 2e 3b 69
55 f1 c3 77 5c 80 4a 98
d5 f8 dd aa

```

Encoded zone identifier (zkl = zTLD):

```
000G001CM8HYGYFCRJXXXDET2WRS50EP7CQ3PTANY71QEQ409ACDBY6XN8
```

Difficulty (5 base difficulty + 2 epochs): 7

Signed message:

```

00 00 00 34 00 00 00 03
00 05 ff 1c 56 e4 b2 68
00 01 00 00 2c a2 23 e8
79 ec c4 bb de b5 da 17
31 92 81 d6 3b 2e 3b 69
55 f1 c3 77 5c 80 4a 98
d5 f8 dd aa

```

Proof:

```

00 05 ff 1c 56 e4 b2 68
00 00 39 5d 18 27 c0 00
38 0b 54 aa 70 16 ac a2
38 0b 54 aa 70 16 ad 62
38 0b 54 aa 70 16 af 3e
38 0b 54 aa 70 16 af 93
38 0b 54 aa 70 16 b0 bf

```

```
38 0b 54 aa 70 16 b0 ee
38 0b 54 aa 70 16 b1 c9
38 0b 54 aa 70 16 b1 e5
38 0b 54 aa 70 16 b2 78
38 0b 54 aa 70 16 b2 b2
38 0b 54 aa 70 16 b2 d6
38 0b 54 aa 70 16 b2 e4
38 0b 54 aa 70 16 b3 2c
38 0b 54 aa 70 16 b3 5a
38 0b 54 aa 70 16 b3 9d
38 0b 54 aa 70 16 b3 c0
38 0b 54 aa 70 16 b3 dd
38 0b 54 aa 70 16 b3 f4
38 0b 54 aa 70 16 b4 42
38 0b 54 aa 70 16 b4 76
38 0b 54 aa 70 16 b4 8c
38 0b 54 aa 70 16 b4 a4
38 0b 54 aa 70 16 b4 c9
38 0b 54 aa 70 16 b4 f0
38 0b 54 aa 70 16 b4 f7
38 0b 54 aa 70 16 b5 79
38 0b 54 aa 70 16 b6 34
38 0b 54 aa 70 16 b6 8e
38 0b 54 aa 70 16 b7 b4
38 0b 54 aa 70 16 b8 7e
38 0b 54 aa 70 16 b8 f8
38 0b 54 aa 70 16 b9 2a
00 01 00 00 2c a2 23 e8
79 ec c4 bb de b5 da 17
31 92 81 d6 3b 2e 3b 69
55 f1 c3 77 5c 80 4a 98
d5 f8 dd aa 08 ca ff de
3c 6d f1 45 f7 e0 79 81
15 37 b2 b0 42 2d 5e 1f
b2 01 97 81 ec a2 61 d1
f9 d8 ea 81 0a bc 2f 33
47 7f 04 e3 64 81 11 be
71 c2 48 82 1a d6 04 f4
94 e7 4d 0b f5 11 d2 c1
62 77 2e 81
```

The following is an example revocation for an EDKEY zone:

Zone private key (d):

```
5a f7 02 0e e1 91 60 32
88 32 35 2b bc 6a 68 a8
d7 1a 7c be 1b 92 99 69
a7 c6 6d 41 5a 0d 8f 65
```

Zone identifier (ztype|zkey):

```
00 01 00 14 3c f4 b9 24
03 20 22 f0 dc 50 58 14
53 b8 5d 93 b0 47 b6 3d
44 6c 58 45 cb 48 44 5d
db 96 68 8f
```

Encoded zone identifier (zkl = zTLD):

```
000G051WYJWJ80S04BRDRM2R2H9VGQCKP13VCFA4DHC4BJT88HEXQ5K8HW
```

Difficulty (5 base difficulty + 2 epochs): 7

Signed message:

```
00 00 00 34 00 00 00 03
00 05 ff 1c 57 35 42 bd
00 01 00 14 3c f4 b9 24
03 20 22 f0 dc 50 58 14
53 b8 5d 93 b0 47 b6 3d
44 6c 58 45 cb 48 44 5d
db 96 68 8f
```

Proof:

```
00 05 ff 1c 57 35 42 bd
00 00 39 5d 18 27 c0 00
58 4c 93 3c b0 99 2a 08
58 4c 93 3c b0 99 2d f7
58 4c 93 3c b0 99 2e 21
58 4c 93 3c b0 99 2e 2a
58 4c 93 3c b0 99 2e 53
58 4c 93 3c b0 99 2e 8e
58 4c 93 3c b0 99 2f 13
58 4c 93 3c b0 99 2f 2d
58 4c 93 3c b0 99 2f 3c
58 4c 93 3c b0 99 2f 41
58 4c 93 3c b0 99 2f fd
58 4c 93 3c b0 99 30 33
58 4c 93 3c b0 99 30 82
58 4c 93 3c b0 99 30 a2
58 4c 93 3c b0 99 30 e1
58 4c 93 3c b0 99 31 ce
58 4c 93 3c b0 99 31 de
58 4c 93 3c b0 99 32 12
```

58 4c 93 3c b0 99 32 4e
58 4c 93 3c b0 99 32 9f
58 4c 93 3c b0 99 33 31
58 4c 93 3c b0 99 33 87
58 4c 93 3c b0 99 33 8c
58 4c 93 3c b0 99 33 e5
58 4c 93 3c b0 99 33 f3
58 4c 93 3c b0 99 34 26
58 4c 93 3c b0 99 34 30
58 4c 93 3c b0 99 34 68
58 4c 93 3c b0 99 34 88
58 4c 93 3c b0 99 34 8a
58 4c 93 3c b0 99 35 4c
58 4c 93 3c b0 99 35 bd
00 01 00 14 3c f4 b9 24
03 20 22 f0 dc 50 58 14
53 b8 5d 93 b0 47 b6 3d
44 6c 58 45 cb 48 44 5d
db 96 68 8f 04 ae 26 f7
63 56 5a b7 aa ab 01 71
72 4f 3c a8 bc c5 1a 98
b7 d4 c9 2e a3 3c d9 34
4c a8 b6 3e 04 53 3a bf
1a 3c 05 49 16 b3 68 2c
5c a8 cb 4d d0 f8 4c 3b
77 48 7a ac 6e ce 38 48
0b a9 d5 00

Authors' Addresses

Martin Schanzenbach
Fraunhofer AISEC
Lichtenbergstrasse 11
85748 Garching
Germany
Email: martin.schanzenbach@aisec.fraunhofer.de

Christian Grothoff
Berner Fachhochschule
Hoeheweg 80
CH-2501 Biel/Bienne
Switzerland
Email: christian.grothoff@bfh.ch

Bernd Fix
GNUnet e.V.
Boltzmannstrasse 3
85748 Garching
Germany
Email: fix@gnunet.org