

Secure Inter-Domain Routing
Working Group
Internet Draft
Intended status: Informational
Expires: May 3, 2009

M. Lepinski
S. Kent
BBN Technologies
November 3, 2008

An Infrastructure to Support Secure Internet Routing
draft-ietf-sidr-arch-04.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on August 3, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2008).

Abstract

This document describes an architecture for an infrastructure to support improved security of Internet routing. The foundation of this architecture is a public key infrastructure (PKI) that represents the allocation hierarchy of IP address space and Autonomous System Numbers; and a distributed repository system for storing and disseminating the data objects that comprise the PKI, as well as

other signed objects necessary for improved routing security. As an initial application of this architecture, the document describes how a holder of IP address space can explicitly and verifiably authorize one or more ASes to originate routes to that address space. Such verifiable authorizations could be used, for example, to more securely construct BGP route filters.

Table of Contents

1.	Introduction.....	3
2.	PKI for Internet Number Resources.....	4
2.1.	Role in the overall architecture.....	5
2.2.	CA Certificates.....	5
2.3.	End-Entity (EE) Certificates.....	7
2.4.	Trust Anchors.....	8
2.5.	Default Trust Anchor Considerations.....	8
2.6.	Representing Early-Registration Transfers (ERX).....	10
3.	Route Origination Authorizations.....	10
3.1.	Role in the overall architecture.....	11
3.2.	Syntax and semantics.....	11
4.	Repositories.....	13
4.1.	Role in the overall architecture.....	14
4.2.	Contents and structure.....	14
4.3.	Access protocols.....	16
4.4.	Access control.....	16
5.	Manifests.....	17
5.1.	Syntax and semantics.....	17
6.	Local Cache Maintenance.....	18
7.	Common Operations.....	18
7.1.	Certificate issuance.....	19
7.2.	ROA management.....	20
7.2.1.	Single-homed subscribers (without portable allocations)	20
7.2.2.	Multi-homed subscribers.....	21
7.2.3.	Portable allocations.....	21
7.3.	Route filter construction.....	22
8.	Security Considerations.....	23
9.	IANA Considerations.....	24
10.	Acknowledgments.....	24
11.	References.....	25
11.1.	Normative References.....	25
11.2.	Informative References.....	25
	Authors' Addresses.....	26
	Intellectual Property Statement.....	26
	Disclaimer of Validity.....	27

1. Introduction

This document describes an architecture for an infrastructure to support improved security for BGP routing [2] for the Internet. The architecture encompasses three principle elements:

- . a public key infrastructure (PKI)
- . digitally-signed routing objects to support routing security
- . a distributed repository system to hold the PKI objects and the signed routing objects

The architecture described by this document enables an entity to verifiably assert that it is the legitimate holder of a set of IP addresses or a set of Autonomous System (AS) numbers. As an initial application of this architecture, the document describes how a holder of IP address space can explicitly and verifiably authorize one or more ASes to originate routes to that address space. Such verifiable authorizations could be used, for example, to more securely construct BGP route filters. In addition to this initial application, the infrastructure defined by this architecture also is intended to provide future support for security protocols such as S-BGP [10] or soBGP [11]. This architecture is applicable to the routing of both IPv4 and IPv6 datagrams. IPv4 and IPv6 are currently the only address families supported by this architecture. Thus, for example, use of this architecture with MPLS labels is beyond the scope of this document.

In order to facilitate deployment, the architecture takes advantage of existing technologies and practices. The structure of the PKI element of the architecture corresponds to the existing resource allocation structure. Thus management of this PKI is a natural extension of the resource-management functions of the organizations that are already responsible for IP address and AS number resource allocation. Likewise, existing resource allocation and revocation practices have well-defined correspondents in this architecture. To ease implementation, existing IETF standards are used wherever possible; for example, extensive use is made of the X.509 certificate profile defined by PKIX [3] and the extensions for IP Addresses and AS numbers representation defined in RFC 3779 [5]. Also CMS [4] is used as the syntax for the newly-defined signed objects required by this infrastructure.

As noted above, the infrastructure is comprised of three main components: an X.509 PKI in which certificates attest to holdings of IP address space and AS numbers; non-certificate/CRL signed objects

(including route origination authorizations and manifests) used by the infrastructure; and a distributed repository system that makes all of these signed objects available for use by ISPs in making routing decisions. These three basic components enable several security functions; this document describes how they can be used to improve route filter generation, and to perform several other common operations in such a way as to make them cryptographically verifiable.

1.1. Terminology

It is assumed that the reader is familiar with the terms and concepts described in "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile" [3], and "X.509 Extensions for IP Addresses and AS Identifiers" [5].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [1].

2. PKI for Internet Number Resources

Because the holder of a block IP address space is entitled to define the topological destination of IP datagrams whose destinations fall within that block, decisions about inter-domain routing are inherently based on knowledge the allocation of the IP address space. Thus, a basic function of this architecture is to provide cryptographically verifiable attestations as to these allocations. In current practice, the allocation of IP address is hierarchic. The root of the hierarchy is IANA. Below IANA are five Regional Internet Registries (RIRs), each of which manages address and AS number allocation within a defined geopolitical region. In some regions the third tier of the hierarchy includes National Internet Registries and (NIRs) as well as Local Internet Registries (LIRs) and subscribers with so-called 'portable' (provider-independent) allocations. (The term LIR is used in some regions to refer to what other regions define as an ISP. Throughout the rest of this document we will use the term LIR/ISP to simplify references to these entities.) In other regions the third tier consists only of LIRs/ISPs and subscribers with portable allocations.

In general, the holder of a set of IP addresses may sub-allocate portions of that set, either to itself (e.g., to a particular unit of the same organization), or to another organization, subject to contractual constraints established by the registries. Because of this structure, IP address allocations can be described naturally by a hierarchic public-key infrastructure, in which each certificate

attests to an allocation of IP addresses, and issuance of subordinate certificates corresponds to sub-allocation of IP addresses. The above reasoning holds true for AS number resources as well, with the difference that, by convention, AS numbers may not be sub-allocated except by regional or national registries. Thus allocations of both IP addresses and AS numbers can be expressed by the same PKI. Such a PKI is a central component of this architecture.

2.1. Role in the overall architecture

Certificates in this PKI are called Resource Certificates, and conform to the certificate profile for such certificates [6]. Resource certificates attest to the allocation by the (certificate) issuer of IP addresses or AS numbers to the subject. They do this by binding the public key contained in the Resource Certificate to the IP addresses or AS numbers included in the certificate's IP Address Delegation or AS Identifier Delegation Extensions, respectively, as defined in [RFC 3779](#) [5].

An important property of this PKI is that certificates do not attest to the identity of the subject. Therefore, the subject names used in certificates are not intended to be 'descriptive.' That is, this PKI is intended to provide authorization, but not authentication. This is in contrast to most PKIs where the issuer ensures that the descriptive subject name in a certificate is properly associated with the entity that holds the private key corresponding to the public key in the certificate. Because issuers need not verify the right of an entity to use a subject name in a certificate, they avoid the costs and liabilities of such verification. This makes it easier for these entities to take on the additional role of CA.

Most of the certificates in the PKI assert the basic facts on which the rest of the infrastructure operates. CA certificates within the PKI attest to IP address space and AS number holdings. End-entity (EE) certificates are issued by resource holder CAs to delegate the authority attested by their allocation certificates. The primary use for EE certificates is the validation of Route Origination Authorizations (ROAs). Additionally, signed objects called manifests will be used to help ensure the integrity of the repository system, and the signature on each manifest will be verified via an EE certificate.

2.2. CA Certificates

Any holder of Internet resources who is authorized to sub-allocate them must be able to issue Resource Certificates to correspond to these sub-allocations. Thus, for example, CA certificates will be

associated with IANA and each of the RIRs, NIRs, and LIRs/ISPs. A CA certificate also is required to enable a resource holder to issue ROAs, because it must issue the corresponding end-entity certificate used to validate each ROA. Thus some subscribers also will need to have CA certificates for their allocations, e.g., subscribers with portable allocations, to enable them to issue ROAs. (A subscriber who is not multi-homed, whose allocation comes from an LIR/ISP, and who has not moved to a different LIR/ISP, need not be represented in the PKI. Moreover, a multi-homed subscriber with an allocation from an LIR/ISP may or may not need to be explicitly represented, as discussed in [Section 6.2.2](#))

Unlike in most PKIs, the distinguished name of the subject in a CA certificate is chosen by the certificate issuer. If the subject of a certificate is an RIR or IANA, then the distinguished name of the subject will be chosen to convey the identity of the registry and should consist of (a subset of) the following attributes: country, organization, organizational unit, and common name. For example, an appropriate subject name for the APNIC RIR might be:

- . Country: AU
- . Organization: Asia Pacific Network Information Centre
- . Common Name: APNIC Resource Certification Authority

If the subject of a certificate is not an RIR or IANA, (e.g., the subject is a NIR, or LIR/ISP) the distinguished name **MUST** consist only of the common name attribute and must not attempt to convey the identity of the subject in a descriptive fashion. Additionally, the subject's distinguished name must be unique among all certificates issued by a given authority. In this PKI, the certificate issuer, being an internet registry or LIR/ISP, is not in the business of verifying the legal right of the subject to assert a particular identity. Therefore, selecting a distinguished name that does not convey the identity of the subject in a descriptive fashion minimizes the opportunity for the subject to misuse the certificate to assert an identity, and thus minimizes the legal liability of the issuer. Since all CA certificates are issued to subjects with whom the issuer has an existing relationship, it is recommended that the issuer select a subject name that enables the issuer to easily link the certificate to existing database records associated with the subject. For example, an authority may use internal database keys or subscriber IDs as the subject common name in issued certificates.

Each Resource Certificate attests to an allocation of resources to its holder, so entities that have allocations from multiple sources

will have multiple CA certificates. A CA also may issue distinct certificates for each distinct allocation to the same entity, if the CA and the resource holder agree that such an arrangement will facilitate management and use of the certificates. For example, an LIR/ISP may have several certificates issued to it by one registry, each describing a distinct set of address blocks, because the LIR/ISP desires to treat the allocations as separate.

2.3. End-Entity (EE) Certificates

The private key corresponding to public key contained in an EE certificate is not used to sign other certificates in a PKI. The primary function of end-entity certificates in this PKI is the verification of signed objects that relate to the usage of the resources described in the certificate, e.g., ROAs and manifests. For ROAs and manifests there will be a one-to-one correspondence between end-entity certificates and signed objects, i.e., the private key corresponding to each end-entity certificate is used to sign exactly one object, and each object is signed with only one key. This property allows the PKI to be used to revoke these signed objects, rather than creating a new revocation mechanism. When the end-entity certificate used to sign an object has been revoked, the signature on that object (and any corresponding assertions) will be considered invalid, so a signed object can be effectively revoked by revoking the end-entity certificate used to sign it.

A secondary advantage to this one-to-one correspondence is that the private key corresponding to the public key in a certificate is used exactly once in its lifetime, and thus can be destroyed after it has been used to sign its one object. This fact should simplify key management, since there is no requirement to protect these private keys for an extended period of time.

Although this document describes only two uses for end-entity certificates, additional uses will likely be defined in the future. For example, end-entity certificates could be used as a more general authorization for their subjects to act on behalf of the holder of the specified resources. This could facilitate authentication of inter-ISP interactions, or authentication of interactions with the repository system. These additional uses for end-entity certificates may require retention of the corresponding private keys, even though this is not required for the private keys associated with end-entity certificates keys used for verification of ROAs and manifests, as described above.

2.4. Trust Anchors

In any PKI, each relying party (RP) is free to choose its own set of trust anchors. This general property of PKIs applies here as well. There is an extant IP address space and AS number allocation hierarchy. IANA is the obvious candidate to be the TA, but operational considerations may argue for a multi-TA PKI, e.g., one in which both IANA and the RIRs form a default set of trust anchors. Nonetheless, every relying party is free to choose a different set of trust anchors to use for certificate validation operations.

For example, an RP (e.g., an LIR/ISP) could create a trust anchor to which all address space and/or all AS numbers are assigned, and for which the RP knows the corresponding private key. The RP could then issue certificates under this trust anchor to whatever entities in the PKI it wishes, with the result that the certificate paths terminating at this locally-installed trust anchor will satisfy the [RFC 3779](#) validation requirements.

An RP who elects to create and manage its own set of trust anchors may fail to detect allocation errors that arise under such circumstances, but the resulting vulnerability is local to the RP.

2.5. Default Trust Anchors

The profile for resource certificates [6] specifies a format for a putative trust anchor to distribute to relying parties trust anchor material consisting of both a self-signed certificate (which would form the root of certification paths in the PKI) along with an additional 'trust anchor' certificate used to validate the self-signed certificate. Any entity claiming authoritative information regarding the allocation of a portion of IP address space may offer itself up in the role of a putative trust anchor by distributing such material (in an out-of-band fashion). Given the extant IP address space and AS number allocation hierarchy, it is envisioned that IANA and the five RIRs will provide trust anchor information to relying parties and that relying parties will generally accept trust anchors from this set.

IANA forms the root of the extant IP address space and AS number allocation hierarchy. Therefore, it is expected that IANA will provide to relying parties trust anchor material whose self-signed certificate has [RFC 3779](#) extensions corresponding either to the entirety of IP address space, or alternatively that portion of IP address space that has not been sub-allocated to any of the five RIRs.

As an example of the use of IANA as a trust anchor, consider the use of private IP address space (i.e., 10/8, 172.16/12, and 192.168/16 in IPv4 and FC00::/7 in IPv6). IANA could issue a CA certificate for these blocks of private address space and then destroy the private key corresponding to the public key in the certificate. In this way, any relying party who configured IANA as their sole trust anchor would automatically reject any ROA containing private addresses, appropriate behavior with regard to routing in the public Internet. On the other hand, such an approach would not interfere with an organization that wishes to use private address space in conjunction with BGP and this PKI technology. Such an organization could configure its relying parties with an additional, local trust anchor that issues certificates for private addresses used within the organization. In this manner, BGP advertisements for these private addresses would be accepted within the organization but would be rejected if mistakenly sent outside the private address space context in question.

In the DNSSEC context, IANA (as the root of the DNS) is already experimenting with the operational procedures needed to digitally sign the root zone. This is very much analogous to the role it would play if it were to act as the default trust anchor for the RPKI, even though DNSSEC does not make use of X.509 certificates. Nonetheless, it is appropriate consider alternative default trust anchor models, if IANA does not act in this capacity. This motivates the consideration of alternative default trust anchor options for RPKI relying parties.

Essentially all allocated IP address and AS number resources are sub-allocated by IANA to one of the five RIRs. Therefore, it is expected that each of the five RIRs will provide trust anchor material provide to relying parties trust anchor material whose self-signed certificate has [RFC 3779](#) extensions corresponding to the IP address and AS number resources that they manage.

One issue that the RIRs will need to consider when providing trust anchor material is how to handle the approximately 49 /8 prefixes containing legacy IPv4 allocations that are not each allocated to a single RIR. Currently, for the purpose of administering reverse DNS zones, each of these prefixes is administered by a single RIR who delegates authority for allocations within the prefix as appropriate. This existing arrangement could be used as the template for the assignment of administrative responsibility for the certification of these address blocks in the RPKI. Such an arrangement would in no way alter the administrative arrangements and the associated policies that apply to the individual legacy allocations that have been made from these address blocks.

2.6. Representing Early-Registration Transfers (ERX)

Currently, IANA allocates IPv4 address space to the RIRs at the level of /8 prefixes. However, there exist allocations that cross these RIR boundaries. For example, A LACNIC customer may have an allocation that falls within a /8 prefix administered by ARIN. Therefore, the resource PKI must be able to represent such transfers from one RIR to another in a manner that permits the validation of certificates with [RFC 3779](#) extensions.

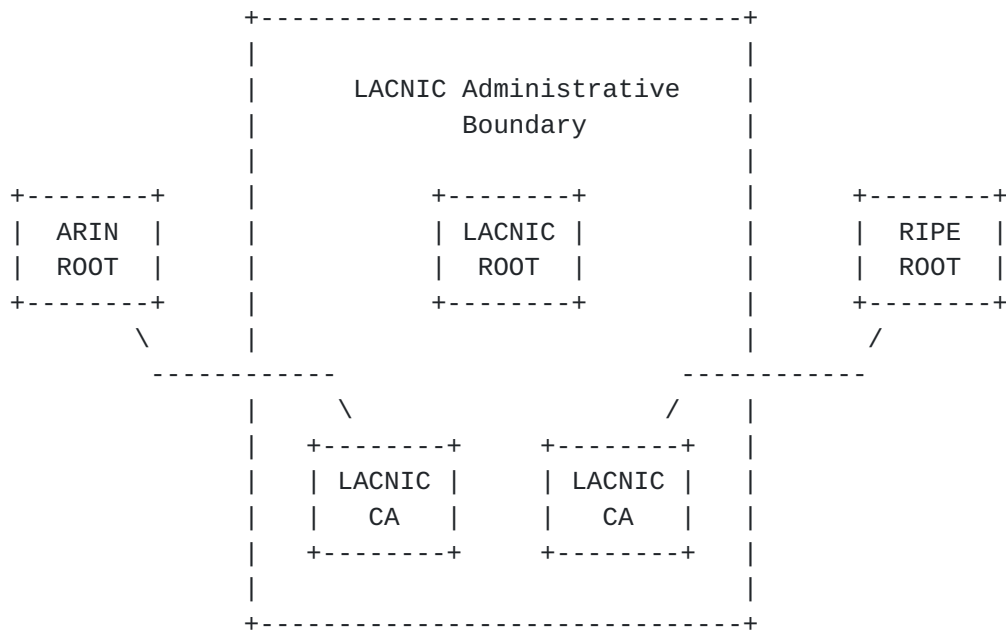


FIGURE 1: Representing ERX

To represent such transfers, RIRs will need to manage multiple CA certificates, each with distinct public (and corresponding private) keys. Each RIR will have a single 'root' certificate (e.g., a self-signed certificate or a certificate signed by IANA, see [Section 2.5](#)), plus one additional CA certificate for each RIR from which it receives a transfer. Each of these additional CA certificates will be issued under the 'root' certificate of the RIR from which the transfer is received. This means that although the certificate is bound to the RIR that receives the transfer, for the purposes of certificate path construction and validation, it does not appear under that RIR's 'root' certificate (see Figure 1).

3. Route Origination Authorizations

The information on IP address allocation provided by the PKI is not, in itself, sufficient to guide routing decisions. In particular, BGP

is based on the assumption that the AS that originates routes for a particular prefix is authorized to do so by the holder of that prefix (or an address block encompassing the prefix); the PKI contains no information about these authorizations. A Route Origination Authorization (ROA) makes such authorization explicit, allowing a holder of address space to create an object that explicitly and verifiably asserts that an AS is authorized originate routes to prefixes.

3.1. Role in the overall architecture

A ROA is an attestation that the holder of a set of prefixes has authorized an autonomous system to originate routes for those prefixes. A ROA is structured according to the format described in [7]. The validity of this authorization depends on the signer of the ROA being the holder of the prefix(es) in the ROA; this fact is asserted by an end-entity certificate from the PKI, whose corresponding private key is used to sign the ROA.

ROAs may be used by relying parties to verify that the AS that originates a route for a given IP address prefix is authorized by the holder of that prefix to originate such a route. For example, an ISP might use ROAs as inputs to route filter construction for use by its BGP routers. These filters would prevent importation of any route in which the origin AS of the AS-PATH attribute is not an AS that is authorized (via a valid ROA) to originate the route. (See [Section 6.3](#) for more details.)

Initially, the repository system will be the primary mechanism for disseminating ROAs, since these repositories will hold the certificates and CRLs needed to verify ROAs. In addition, ROAs also could be distributed in BGP UPDATE messages or via other communication paths, if needed to meet timeliness requirements.

3.2. Syntax and semantics

A ROA constitutes an explicit authorization for a single AS to originate routes to one or more prefixes, and is signed by the holder of those prefixes. A detailed specification of the ROA syntax can be found in [7] but, at a high level, a ROA consists of (1) an AS number; (2) a list of IP address prefixes; and, optionally, (3) for each prefix, the maximum length of more specific (longer) prefixes that the AS is also authorized to advertise. (This last element facilitates a compact authorization to advertise, for example, any prefixes of length 20 to 24 contained within a given length 20 prefix.)

Note that a ROA contains only a single AS number. Thus, if an ISP has multiple AS numbers that will be authorized to originate routes to the prefix(es) in the ROA, an address space holder will need to issue multiple ROAs to authorize the ISP to originate routes from any of these ASes.

A ROA is signed using the private key corresponding to the public key in an end-entity certificate in the PKI. In order for a ROA to be valid, its corresponding end-entity (EE) certificate must be valid and the IP address prefixes of the ROA must exactly match the IP address prefix(es) specified in the EE certificate's [RFC 3779](#) extension. Therefore, the validity interval of the ROA is implicitly the validity interval of its corresponding certificate. A ROA is revoked by revoking the corresponding EE certificate. There is no independent method of invoking a ROA. One might worry that this revocation model could lead to long CRLs for the CA certification that is signing the EE certificates. However, routing announcements on the public internet are generally quite long lived. Therefore, as long as the EE certificates used to verify a ROA are given a validity interval of several months, the likelihood that many ROAs would need to be revoked within time that is quite low.

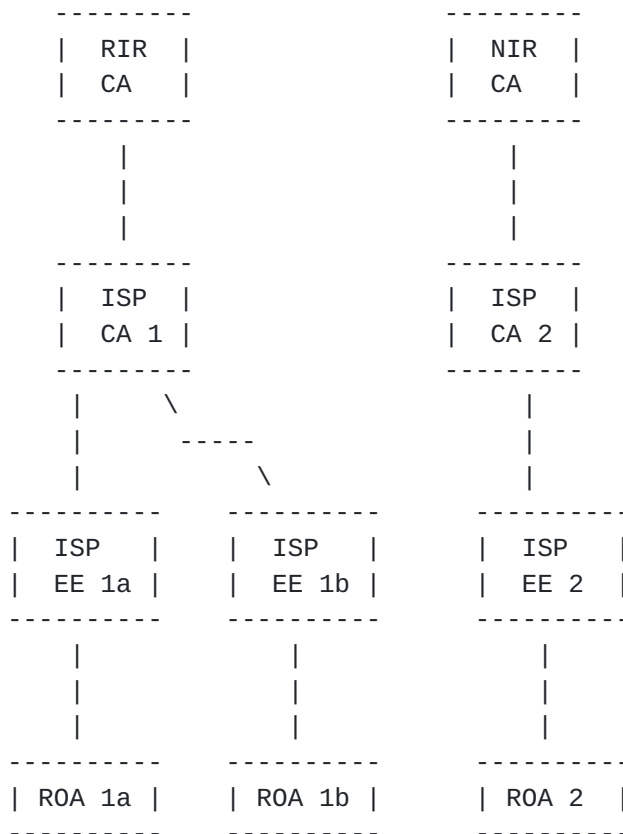


FIGURE 2: This figure illustrates an ISP with allocations from two sources (and RIR and an NIR). It needs two CA certificates due to [RFC 3779](#) rules.

Because each ROA is associated with a single end-entity certificate, the set of IP prefixes contained in a ROA must be drawn from an allocation by a single source, i.e., a ROA cannot combine allocations from multiple sources. Address space holders who have allocations from multiple sources, and who wish to authorize an AS to originate routes for these allocations, must issue multiple ROAs to the AS.

4. Repositories

Initially, an LIR/ISP will make use of the resource PKI by acquiring and validating every ROA, to create a table of the prefixes for which each AS is authorized to originate routes. To validate all ROAs, an LIR/ISP needs to acquire all the certificates and CRLs. The primary function of the distributed repository system described here is to store these signed objects and to make them available for download by LIRs/ISPs. Note that this repository system provides a mechanism by which relying parties can pull fresh data at whatever frequency they deem appropriate. However, it does not provide a mechanism for

pushing fresh data to relying parties (e.g. by including resource PKI objects in BGP or other protocol messages) and such a mechanism is beyond the scope of the current document.

The digital signatures on all objects in the repository ensure that unauthorized modification of valid objects is detectable by relying parties. Additionally, the repository system uses manifests (see [Section 5](#)) to ensure that relying parties can detect the deletion of valid objects and the insertion of out of date, valid signed objects.

The repository system is also a point of enforcement for access controls for the signed objects stored in it, e.g., ensuring that records related to an allocation of resources can be manipulated only by authorized parties. The use access controls prevents denial of service attacks based on deletion of or tampering to repository objects. Indeed, although relying parties can detect tampering with objects in the repository, it is preferable that the repository system prevent such unauthorized modifications to the greatest extent possible.

[4.1.](#) Role in the overall architecture

The repository system is the central clearing-house for all signed objects that must be globally accessible to relying parties. When certificates and CRLs are created, they are uploaded to this repository, and then downloaded for use by relying parties (primarily LIRs/ISPs). ROAs and manifests are additional examples of such objects, but other types of signed objects may be added to this architecture in the future. This document briefly describes the way signed objects (certificates, CRLs, ROAs and manifests) are managed in the repository system. As other types of signed objects are added to the repository system it will be necessary to modify the description, but it is anticipated that most of the design principles will still apply. The repository system is described in detail in [\[9\]](#).

[4.2.](#) Contents and structure

Although there is a single repository system that is accessed by relying parties, it is comprised of multiple databases. These databases will be distributed among registries (RIRs, NIRs, LIRs/ISPs). At a minimum, the database operated by each registry will contain all CA and EE certificates, CRLs, and manifests signed by the CA(s) associated with that registry. Repositories operated by LIRs/ISPs also will contain ROAs. Registries are encouraged maintain copies of repository data from their customers, and their customer's customers (etc.), to facilitate retrieval of the whole repository

contents by relying parties. Ideally, each RIR will hold PKI data from all entities within its geopolitical scope.

For every certificate in the PKI, there will be a corresponding file system directory in the repository that is the authoritative publication point for all objects (certificates, CRLs, ROAs and manifests) verifiable via this certificate. A certificate's Subject Information Authority (SIA) extension provides a URI that references this directory. Additionally, a certificate's Authority Information Authority (AIA) extension contains a URI that references the authoritative location for the CA certificate under which the given certificate was issued. That is, if certificate A is used to verify certificate B, then the AIA extension of certificate B points to certificate A, and the SIA extension of certificate A points to a directory containing certificate B (see Figure 2).

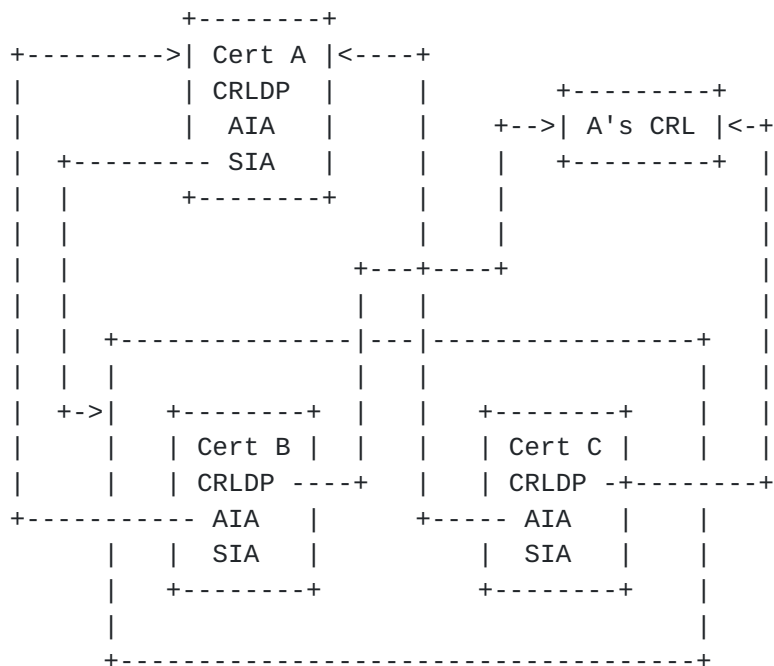


FIGURE 3: In this example, certificates B and C are issued under certificate A. Therefore, the AIA extensions of certificates B and C point to A, and the SIA extension of certificate A points to the directory containing certificates B and C.

If a CA certificate is reissued with the same public key, it should not be necessary to reissue (with an updated AIA URI) all certificates signed by the certificate being reissued. Therefore, a certification authority SHOULD use a persistent URI naming scheme for issued certificates. That is, reissued certificates should use the

same publication point as previously issued certificates having the same subject and public key, and should overwrite such certificates.

[4.3. Access protocols](#)

Repository operators will choose one or more access protocols that relying parties can use to access the repository system. These protocols will be used by numerous participants in the infrastructure (e.g., all registries, ISPs, and multi-homed subscribers) to maintain their respective portions of it. In order to support these activities, certain basic functionality is required of the suite of access protocols, as described below. No single access protocol need implement all of these functions (although this may be the case), but each function must be implemented by at least one access protocol.

Download: Access protocols MUST support the bulk download of repository contents and subsequent download of changes to the downloaded contents, since this will be the most common way in which relying parties interact with the repository system. Other types of download interactions (e.g., download of a single object) MAY also be supported.

Upload/change/delete: Access protocols MUST also support mechanisms for the issuers of certificates, CRLs, and other signed objects to add them to the repository, and to remove them. Mechanisms for modifying objects in the repository MAY also be provided. All access protocols that allow modification to the repository (through addition, deletion, or modification of its contents) MUST support verification of the authorization of the entity performing the modification, so that appropriate access controls can be applied (see [Section 4.4](#)).

Current efforts to implement a repository system use RSYNC [[12](#)] as the single access protocol. RSYNC, as used in this implementation, provides all of the above functionality. A document specifying the conventions for use of RSYNC in the PKI will be prepared.

[4.4. Access control](#)

In order to maintain the integrity of information in the repository, controls must be put in place to prevent addition, deletion, or modification of objects in the repository by unauthorized parties. The identities of parties attempting to make such changes can be authenticated through the relevant access protocols. Although specific access control policies are subject to the local control of repository operators, it is recommended that repositories allow only the issuers of signed objects to add, delete, or modify them.

Alternatively, it may be advantageous in the future to define a formal delegation mechanism to allow resource holders to authorize other parties to act on their behalf, as suggested in [Section 2.3](#) above.

5. Manifests

A manifest is a signed object listing of all of the signed objects issued by an authority responsible for a publication in the repository system. For each certificate, CRL, or ROA issued by the authority, the manifest contains both the name of the file containing the object, and a hash of the file content.

As with ROAs, a manifest is signed by a private key, for which the corresponding public key appears in an end-entity certificate. This EE certificate, in turn, is signed by the CA in question. The EE certificate private key may be used to issue one or more manifests. If the private key is used to sign only a single manifest, then the manifest can be revoked by revoking the EE certificate. In such a case, to avoid needless CRL growth, the EE certificate used to validate a manifest SHOULD expire at the same time that the manifest expires. If an EE certificate is used to issue multiple (sequential) manifests for the CA in question, then there is no revocation mechanism for these individual manifests.

Manifests may be used by relying parties when constructing a local cache (see [Section 6](#)) to mitigate the risk of an attacker who deletes files from a repository or replaces current signed objects with stale versions of the same object. Such protection is needed because although all objects in the repository system are signed, the repository system itself is untrusted.

5.1. Syntax and semantics

A manifest constitutes a list of (the hashes of) all the files in a repository point at a particular point in time. A detailed specification of manifest syntax is provided in [\[8\]](#) but, at a high level, a manifest consists of (1) a manifest number; (2) the time the manifest was issued; (3) the time of the next planned update; and (4) a list of filename and hash value pairs.

The manifest number is a sequence number that is incremented each time a manifest is issued by the authority. An authority is required to issue a new manifest any time it alters any of its items in the repository, or when the specified time of the next update is reached. A manifest is thus valid until the specified time of the next update or until a manifest is issued with a greater manifest number,

whichever comes first. (Note that when an EE certificate is used to sign only a single manifest, whenever the authority issues the new manifest, the CA MUST also issue a new CRL which includes the EE certificate corresponding to the old manifest. The revoked EE certificate for the old manifest will be removed from the CRL when it expires, thus this procedure ought not result in significant CRLs growth.)

6. Local Cache Maintenance

In order to utilize signed objects issued under this PKI (e.g. for route filter construction, see [Section 6.3](#)), a relying party must first obtain a local copy of the valid EE certificates for the PKI. To do so, the relying party performs the following steps:

1. Query the registry system to obtain a copy of all certificates, manifests and CRLs issued under the PKI.
2. For each CA certificate in the PKI, verify the signature on the corresponding manifest. Additionally, verify that the current time is earlier than the time indicated in the nextUpdate field of the manifest.
3. For each manifest, verify that certificates and CRLs issued under the corresponding CA certificate match the hash values contained in the manifest. If the hash values do not match, use an out-of-band mechanism to notify the appropriate repository administrator that the repository data has been corrupted.
4. Validate each EE certificate by constructing and verifying a certification path for the certificate (including checking relevant CRLs) to the locally configured set of TAs. (See [\[6\]](#) for more details.)

Note that when a relying party performs these operations regularly, it is more efficient for the relying party to request from the repository system only those objects that have changed since the relying party last updated its local cache. Note also that by checking all issued objects against the appropriate manifest, the relying party can be certain that it is not missing an updated version of any object.

7. Common Operations

Creating and maintaining the infrastructure described above will entail additional operations as 'side effects' of normal resource allocation and routing authorization procedures. For example, a

subscriber with ''portable'' address space who enters a relationship with an ISP will need to issue one or more ROAs identifying that ISP, in addition to conducting any other necessary technical or business procedures. The current primary use of this infrastructure is for route filter construction; using ROAs, route filters can be constructed in an automated fashion with high assurance that the holder of the advertised prefix has authorized the first-hop AS to originate an advertised route.

7.1. Certificate issuance

There are several operational scenarios that require certificates to be issued. Any allocation that may be sub-allocated requires a CA certificate, e.g., so that certificates can be issued as necessary for the sub-allocations. Holders of ''portable'' address allocations also must have certificates, so that a ROA can be issued to each ISP that is authorized to originate a route to the allocation (since the allocation does not come from any ISP). Additionally, multi-homed subscribers may require certificates for their allocations if they intend to issue the ROAs for their allocations (see [Section 6.2.2](#)). Other holders of resources need not be issued CA certificates within the PKI.

In the long run, a resource holder will not request resource certificates, but rather receive a certificate as a side effect of the allocation process for the resource. However, initial deployment of the RPKI will entail issuance of certificates to existing resource holders as an explicit event. Note that in all cases, the authority issuing a CA certificate will be the entity who allocates resources to the subject. This differs from most PKIs in which a subject can request a certificate from any certification authority.

If a resource holder receives multiple allocations over time, it may accrue a collection of resource certificates to attest to them. If a resource holder receives multiple allocations from the same source, the set of resource certificates may be combined into a single resource certificate, if both the issuer and the resource holder agree. This is effected by consolidating the IP Address Delegation and AS Identifier Delegation Extensions into a single extension (of each type) in a new certificate. However, if the certificates for these allocations contain different validity intervals, creating a certificate that combines them might create problems, and thus is NOT RECOMMENDED.

If a resource holder's allocations come from different sources, they will be signed by different CAs, and cannot be combined. When a set of resources is no longer allocated to a resource holder, any

certificates attesting to such an allocation **MUST** be revoked. A resource holder **SHOULD NOT** to use the same public key in multiple CA certificates that are issued by the same or differing authorities, as reuse of a key pair complicates path construction. Note that since the subject's distinguished name is chosen by the issuer, a subject who receives allocations from two sources generally will receive certificates with different subject names.

7.2. ROA management

Whenever a holder of IP address space wants to authorize an AS to originate routes for a prefix within his holdings, he **MUST** issue an end-entity certificate containing that prefix in an IP Address Delegation extension. He then uses the corresponding private key to sign a ROA containing the designated prefix and the AS number for the AS. The resource holder **MAY** include more than one prefix in the EE certificate and corresponding ROA if desired. As a prerequisite, then, any address holder that issues ROAs for a prefix must have a resource certificate for an allocation containing that prefix. The standard procedure for issuing a ROA is as follows:

1. Create an end-entity certificate containing the prefix(es) to be authorized in the ROA.
2. Construct the payload of the ROA, including the prefixes in the end-entity certificate and the AS number to be authorized.
3. Sign the ROA using the private key corresponding to the end-entity certificate (the ROA is comprised of the payload encapsulated in a CMS signed message [7]).
4. Upload the end-entity certificate and the ROA to the repository system.

The standard procedure for revoking a ROA is to revoke the corresponding end-entity certificate by creating an appropriate CRL and uploading it to the repository system. The revoked ROA and end-entity certificate **SHOULD BE** removed from the repository system.

7.2.1. Single-homed subscribers (without portable allocations)

In BGP, a single-homed subscriber with a non-portable allocation does not need to explicitly authorize routes to be originated for the prefix(es) it is using, since its ISP will already advertise a more general prefix and route traffic for the subscriber's prefix as an internal function. Since no routes are originated specifically for prefixes held by these subscribers, no ROAs need to be issued under

their allocations; rather, the subscriber's ISP will issue any necessary ROAs for its more general prefixes under resource certificates its own allocation. Thus, a single-homed subscriber with a non-portable allocation is not included in the RPKI, i.e., it does not receive a CA certificate, nor issue EE certificates or ROAs.

7.2.2. Multi-homed subscribers

In order for multiple ASes to originate routes for prefixes held by a multi-homed subscriber, each AS must have a ROA that explicitly authorizes such route origination. There are two ways that this can be accomplished.

One option is for the multi-homed subscriber to obtain a CA certificate from the ISP who allocated the prefixes to the subscriber. The multi-homed subscriber can then create a ROA (and associated end-entity certificate) that authorizes a second ISP to originate routes to the subscriber prefix(es). The ROA for the second ISP generally SHOULD be set to require an exact match, if the intent is to enable backup paths for the prefix. Note that the first ISP, who allocated the prefixes, will want to advertise the more specific prefix for this subscriber (vs. the encompassing prefix). Either the subscriber or the first ISP will need to issue an EE certificate and ROA for the (more specific) prefix, authorizing this ISP to advertise this more specific prefix.

A second option is that the multi-homed subscriber can request that the ISP that allocated the prefixes create a ROA that authorizes the second ISP to originate routes to the subscriber's prefixes. (The ISP also creates an EE certificate and ROA for its own advertisement of the subscriber prefix, as above.) This option does not require that the subscriber be issued a certificate or participate in ROA management. Therefore, this option is simpler for the subscriber, and is preferred if the option is supported by the ISP performing the allocation.

7.2.3. Portable allocations

A resource holder is said to have a portable (provider independent) allocation if the resource holder received its allocation from a regional or national registry. Because the prefixes represented in such allocations are not taken from an allocation held by an ISP, there is no ISP that holds and advertises a more general prefix. A holder of a portable allocation MUST authorize one or more ASes to originate routes to these prefixes. Thus the resource holder MUST generate one or more EE certificates and associated ROAs to enable the AS(es) to originate routes for the prefix(es) in question. This

ROA is required because none of the ISP's existing ROAs authorize it to originate routes to that portable allocation.

7.3. Route filter construction

The goal of this architecture is to support improved routing security. One way to do this is to use ROAs to construct route filters that reject routes that conflict with the origination authorizations asserted by current ROAs. The following is intended to provide a high-level description of how the architecture might be used to construct a route filter. Additional guidance on the use of ROAs to derive inferences about the validity of BGP UPDATE messages is provided in [14]. The guidance in [14] is particularly important during a transition period where not all ISPs implement this architecture (and thus the filter described below which naively rejects all UPDATES without a corresponding ROA would incorrectly reject valid routes originated by ISPs that do not yet implement this architecture).

1. Obtain a local copy of all currently valid EE certificates, as specified in [Section 5](#).
2. Query the repository system to obtain a local copy of all ROAs issued under the PKI.
3. Verify that the each ROA matches the hash value contained in the manifest of the CA certificate used to verify the EE certificate that issued the ROA and that no ROAs are missing.
4. Validate each ROA by verifying that its signature is verifiable by a valid end-entity certificate that matches the address allocation in the ROA. (See [7] for more details.)
5. Based on the validated ROAs, construct a table of prefixes and corresponding authorized origin ASes (or vice versa).

A BGP speaker that applies such a filter is thus guaranteed that for a given IP address prefix, all routes that the BGP speaker accepts for that prefix were originated by an AS that is authorized by the owner of the prefix to authorize routes to that prefix.

The first three steps in the above procedure might incur a substantial overhead if all objects in the repository system were downloaded and validated every time a route filter was constructed. Instead, it will be more efficient for users of the infrastructure to initially download all of the signed objects and perform the validation algorithm described above. Subsequently, a relying party need only perform incremental downloads and validations on a regular basis. A typical ISP using the infrastructure may choose any frequency it desires for downloading updates from the repository, uploading any modifications it has made, and constructing route filters. However, an ISP might reasonably choose to perform these actions on a daily schedule.

It should be noted that the transition to 4-byte AS numbers (see [RFC 4893](#) [13]) weakens the security guarantees achieved by BGP speakers who do not support 4-byte AS numbers (referred to as OLD BGP speakers). [RFC 4893](#) specifies that all 4-byte AS numbers (except those whose first two bytes are entirely zero) be mapped to the reserved value 23456 before being sent to a BGP speaker who does not understand 4-byte AS numbers. Therefore, when an ISP creates a route filter for use by an OLD BGP speaker, it must allow any 4-byte AS number to advertise routes for an IP address prefix if there exists a ROA that authorizes any 4-byte AS number to advertise routes to that prefix. This means that if an OLD BGP speaker accepts a route that was originated by an AS with a 4-byte AS number, there is no guarantee that it was originated by an authorized 4-byte AS number (unless the route was propagated by an intermediate NEW BGP speaker who performed route filtering as described above).

8. Security Considerations

The focus of this document is security; hence security considerations permeate this specification.

The security mechanisms provided by and enabled by this architecture depend on the integrity and availability of the infrastructure it describes. The integrity of objects within the infrastructure is ensured by appropriate controls on the repository system, as described in [Section 4.4](#). Likewise, because the repository system is structured as a distributed database, it should be inherently resistant to denial of service attacks; nonetheless, appropriate precautions should also be taken, both through replication and backup of the constituent databases and through the physical security of database servers

9. IANA Considerations

This document makes no request of IANA.

Note to RFC Editor: this section may be removed on publication as an RFC

10. Acknowledgments

The architecture described in this draft is derived from the collective ideas and work of a large group of individuals. This work would not have been possible without the intellectual contributions of George Michaelson, Robert Loomans, Sanjaya and Geoff Huston of APNIC, Robert Kisteleki and Henk Uijterwaal of the RIPE NCC, Time Christensen and Cathy Murphy of ARIN, Rob Austein of ISC and Randy Bush of IIR.

Although we are indebted to everyone who has contributed to this architecture, we would like to especially thank Rob Austein for the concept of a manifest, Geoff Huston for the concept of managing object validity through single-use EE certificate key pairs, and Richard Barnes for help in preparing an early version of this document.

11. References

11.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), January 2006
- [3] 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](#), May 2008.
- [4] Housley, R., "'Cryptographic Message Syntax'", [RFC 3852](#), July 2004.
- [5] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", [RFC 3779](#), June 2004.
- [6] Huston, G., Michaelson, G., and Loomans, R., "A Profile for X.509 PKIX Resource Certificates", [draft-ietf-sidr-res-certs-14](#), October 2008.
- [7] Lepinski, M., Kent, S., and Kong, D., "A Profile for Route Origin Authorizations (ROA)", [draft-ietf-sidr-roa-format-04](#), November 2008.
- [8] Austein, R., et al., "'Manifests for the Resource Public Key Infrastructure'", [draft-ietf-sidr-rpki-manifests-04](#), October 2008.

11.2. Informative References

- [9] Huston, G., Michaelson, G., and Loomans, R., "'A Profile for Resource Certificate Repository Structure'", [draft-ietf-sidr-repos-struct-01](#), October 2008.
- [10] Kent, S., Lynn, C., and Seo, K., "Secure Border Gateway Protocol (Secure-BGP)", IEEE Journal on Selected Areas in Communications Vol. 18, No. 4, April 2000.
- [11] White, R., "soBGP", May 2005, <[ftp://ftp-eng.cisco.com/sobgp/index.html](http://ftp-eng.cisco.com/sobgp/index.html)>

- [12] Tridgell, A., "rsync", April 2006,
<<http://samba.anu.edu.au/rsync/>>
- [13] Vohra, Q., and Chen, E., "'BGP Support for Four-octet AS Number Space'", [RFC 4893](#), May 2007.
- [14] Huston, G., Michaelson, G., "'Validation of Route Origination in BGP using the Resource Certificate PKI'", [draft-ietf-sidr-roa-validation-01](#), October 2008.

Authors' Addresses

Matt Lepinski
BBN Technologies
10 Moulton St.
Cambridge, MA 02138

Email: mlepinski@bbn.com

Stephen Kent
BBN Technologies
10 Moulton St.
Cambridge, MA 02138

Email: kent@bbn.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this

specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The IETF Trust (2008).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.