Secure Inter-Domain Routing
Working Group
Internet Draft
Intended status: Informational
Expires: August 2007

R. Barnes
S. Kent
BBN Technologies

February 23, 2007

An Infrastructure to Support Secure Internet Routing
draft-ietf-sidr-arch-00.txt

Status of this Memo

Copyright Notice

Abstract

This document describes an architecture for an infrastructure to support secure 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;

certificates from this PKI are used to verify signed objects that authorize autonomous systems to originate routes for specified IP address prefixes. The data objects that comprise the PKI, as well as other signed objects necessary for secure routing, are stored and disseminated through a distributed repository system. This document also describes at a high level how this architecture can be used to add security features to common operations such as IP address space allocation and route filter construction.

Conventions used in this document

In examples, "C:" and "S:" indicate lines sent by the client and server respectively.

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].

Table of Contents

1. Introduction

This document describes an architecture for an infrastructure to support improved security for BGP routing [2] for the Internet. The architecture described by this document supports, at a minimum, two types of routing security: It enables an entity to verifiably assert that it is the legitimate holder of a set IP addresses or a set of Autonomous System (AS) numbers, and it allows the holder of IP address space to explicitly and verifiably authorize an AS to originate routes to that address space. In addition to these initial applications, however, this architecture could also support, without extension, more advanced security protocols such as S-BGP [7] or soBGP [8]. This architecture is applicable to routing of both IPv4 and IPv6 datagrams.

In order to facilitate deployment, the architecture takes advantage of existing technologies and practices. The structure of the architecture corresponds to the existing resource allocation structure, so that management of this architecture is a natural extension of the resource-management functions of 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 defined in RFC 3779 [4].

The architecture is comprised of three main components: An X.509 public-key infrastructure (PKI) where certificates attest to holdings

of IP address space and AS numbers; signed objects called Route Origination Authorizations (ROAs) that enable an address space holder to explicitly authorize an AS to originate routes to portions of the IP address space; and a distributed repository system that makes these 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.

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 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. 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 signing 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 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 Certificate, and conform to the certificate profile for such certificates [5]. 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. An important property of this PKI is that the names assigned to certificate issuers and subjects are not intended to be meaningful; this is in

contrast to most PKIs where considerable effort is expended to ensure that the subject name in a certificate is properly associated with the entity that holds the private key corresponding to the public key in the certificate. This PKI is different because it is an authorization PKI, not an authentication PKI. 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. Only the basic PKI requirement, that a CA not associate the same name with two distinct subjects to whom it issues certificates, is imposed.

The certificates in the PKI assert the basic facts on which the rest of the infrastructure operates. Certificates within the CA hierarchy

attest to IP address space and AS number holdings. End-entity certificates are issued by resource holders to delegate the authority attested by their allocation certificates, the primary use for this being the signing of ROAs. These certificates and corresponding certificate revocation lists will comprise a significant portion of the data stored in the repository system.

2.2. CA Certificates

Any holder of Internet Number 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 each of the Regional Internet Registries, National Internet Registries, and Local Internet Registries, as well as with all ISPs. A CA certificate is also necessary for a resource holder to issue ROAs (because it must also issue the corresponding end-entity certificates used to validate ROAs), so many subscribers also will need to have CA certificates for their allocations (in particular, multi-homed subscribers, and subscribers with portable allocations).

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 issuer and the resource holder agree that such an arrangement will facilitate management and use of the certificates.)

2.3. End-Entity Certificates

Although the private key corresponding to public key contained in an end-entity certificate is not used to sign other certificates in the 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. For this purpose, there is 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 one key. This property allows the PKI itself to be used to revoke these signed objects. 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 only the public portions of end-entity certificates will need to be retained for any significant length of time.

Although this document defines only one use 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 indicated resources. This could facilitate authentication of inter-ISP interactions, or authentication of interactions with the repository system. These additional uses for end-entity certificates, however, 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, as described above.

2.4. Trust Anchors

In any PKI, each relying party (RP) is free to choose its own set of trust anchors. In this case, the hierarchy of this PKI is structured according to the IP address space and AS number allocation hierarchy, so since administrative control of the IP address space (the root of

the allocation hierarchy) rests with IANA and the RIRs , these entities form a natural set of default trust anchors for this PKI. Nonetheless, every relying party is free to choose a different set of trust anchors to use for certificate validation operations.

For example, an RP could create one or more self-signed certificates 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 these trust anchors to whatever entities in the PKI it wishes, with the result that the certificate paths terminating at these locally-installed trust anchors 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.

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 block of IP address space is authorized to do so by the holder of that block; the PKI contains no information about these

authorizations. A Route Origination Authorization (ROA) make 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 that address space.

3.1. Role in the overall architecture

A ROA is an attestation that the holder of a set of IP addresses has authorized an autonomous system to originate routes for that set of IP addresses. A ROA is structured according to format described in [6]. The validity of this authorization depends on the issuer of the ROA being the owner of the set of IP addresses 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. The repository system will be the primary mechanism for disseminating ROAs, since the operators of repositories already provide other types routing information. In addition, ROAs could also be distributed in BGP UPDATE messages or via other communication paths, since route

filtering is their primary application.

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 ROA thus have three essential components:

1. An AS number

2. One or more IP address prefixes

3. A digital signature

In addition, a ROA has a version number, to accommodate changes in syntax (or semantics) over time. The AS number contained in a ROA is that of an AS authorized to originate routes for the indicated IP address prefixes. Only one AS number is contained in a ROA in order to simplify ROA management, e.g., to avoid the complexity that might arise is AS numbers for multiple ISPs were referenced from a single ROA. If an ISP serving a subscriber has multiple AS numbers, and wants the address space holder to authorize advertisement of the same set of prefixes by any of these ASes, the ISP should request the subscriber to issue multiple ROAs, each specifying a distinct AS number. Similarly, a multi-homed address space holder must generate multiple ROAs, one for each ISP that will originate routes for it.

A ROA is signed using the private key whose public key is contained in an end-entity certificate in the PKI, from which the ROA inherits

two properties. First, The IP prefixes listed in the ROA are the ones that the indicated AS is authorized to originate; in order for this authorization (i.e., the ROA) to be valid, the prefixes contained in a ROA must be exactly the same as the set of IP addresses in the IP Address Delegation Extension of the end-entity certificate used to sign the ROA. Second, the ROA is valid only as long as the certificate used to sign it is valid; a ROA is invalidated by revoking the end-entity certificate corresponding to the public key used to verify it, and the validity interval of the ROA is implicitly that of the validity interval of the corresponding certificate.

Address holders that have allocations from multiple sources must

issue multiple ROAs. If an address holder has allocations from multiple sources, then these allocations will be described by multiple CA certificates in the PKI, each issued by the provider of the respective allocation; the sets of IP addresses contained in end-entity certificates issued by this address holder are required to be subsets of these allocations. Because end-entity certificates are in one-to-one correspondence with ROAs, this means that the set of IP addresses contained in a ROA must be drawn from an allocation by a single source; hence ROAs cannot combine allocations from multiple sources.

3.3. Revocation

If an address holder decides that an AS should no longer originate routes for addresses that it holds (e.g., if the address holder transfers from one ISP to another), then it will be necessary to invalidate the ROAs that attest to any such authorization. Since ROAs are in one-to-one correspondence with end-entity certificates, the standard method for revoking a ROA is to revoke the corresponding end-entity certificate in the PKI. There is no independent revocation mechanism for ROAs.

4. Repository system

In order for ROAs (and other objects to be verified using certificates from the PKI) to be validated, the objects necessary to validate them must be universally accessible. The primary function of the distributed repository system described here is to store these objects and to make them available for download. 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. This requirement exists to prevent denial of service attacks based on deletion of or tampering with repository objects. Although

any unauthorized modification is detectable by relying parties, because all the objects are digitally signed, it is preferable that the repository system prevent unauthorized modifications.

4.1. Role in the overall architecture

The repository system is the central clearing-house for the objects required for validation of signed objects like ROAs. When

certificates and CRLs are created, they are uploaded to this repository, and then downloaded for use by relying parties. In addition, signed objects that require universal distribution can also be made accessible through the repository system; ROAs are the only such objects defined by this document, but other types of signed objects may be added to this architecture in the future. The repository system also must ensure the integrity of the data it contains by enforcing appropriate controls on access to the repository and on modifications to entries in it. This document describes the controls necessary for PKI objects and ROAs, but does not assume that they are applicable to other types of objects; if other types of objects are to be included in the repository system in the future, any necessary controls on them must be defined.

4.2. Contents and structure

The primary function of the repository system is to provide universal distribution of objects necessary for the function of this architecture. First among these are the objects that comprise the PKI, namely Resource Certificates and their corresponding CRLs; these objects require universal distribution so that all relying parties have access to the PKI components required to validate signed objects used by this architecture. In addition, it may be necessary to make other types of signed objects available through the repository system. ROAs are a prime example of such a type, since routes whose origination is authorized by a ROA are distributed through the entire routing infrastructure, any component of which may, by local policy, examine the route origin for consistency with the ROA.

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 RIRs and ISPs that participate in the architecture. At a minimum, the database operated by each RIR will contain certificates and CRLs issued by that RIR; it may also contain repository objects submitted by holders of addresses that fall within that RIR's scope or copies of data from other RIRs, according to local policy.

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 [9] as the single access protocol. RSYNC, as used in this implementation, provides all of the above functionality.

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 is suggested in Section 2.3 above.

5. Common Operations

Creating and maintaining the infrastructure described above will entail the addition of security operations to normal resource allocation and routing authorization procedures. For example, a multi-homed subscriber entering a relationship with a new ISP will need to issue one or more ROAs to 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.

5.1. Certificate issuance

In order to participate in this infrastructure, resource holders will require certificates in the PKI that attest to their allocations. Each such certificate will show the issuer of the allocation as the certificate issuer, the recipient of the allocation as subject, and a description of the allocated resources in the appropriate RFC 3779 extensions. The two operations defined in this architecture that require a resource holder to have resource certificates for his allocations are (1) issuance of certificates for sub-allocations and (2) management of ROAs (and corresponding end-entity certificates).

In particular, there are several operational scenarios that require certificates to be issued. Any allocation that may be sub-allocated requires a CA certificate so that certificates can be issued as necessary for sub-allocations. Multi-homed subscribers require certificates for their allocations so that they can issue ROAs to their ISPs. Holders of "portable" address allocations 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.

As a resource holder receives multiple allocations over time, it will accrue a collection of resource certificates to attest to them. It may be the case that multiple of a resource holder's allocations are from the same source. A set of resource certificates that are all issued by the same CA could be combined into a single resource certificate by consolidating their IP Address Delegation and AS Identifier Delegation Extensions into a single extension of each type. 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 MAY choose to use the same public key in multiple CA certificates that issued by the same or differing authorities, although such reuse of a key pair does complicate path construction.

5.2. 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 and use the corresponding private key to sign a ROA containing the designated prefix and an AS number identifying the designated AS. 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 prefixes 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 [ROA format I-D])

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.

5.2.1. 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 (or prefixes) 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,

single-homed subscribers with non-portable allocations do not need to issue (or otherwise manage) ROAs.

5.2.2. Multi-homing

If a multi-homed subscriber wants multiple ASes to originate routes for prefixes that it holds, then it must explicitly authorize each of them to do so by issuing a ROA for each AS in question.

5.2.3. Portable allocations

A resource holder is said to have a portable allocation if the resource holder received its allocation from a registry. Because these allocations are not taken from any larger allocations held by an ISP, there is no ISP that holds and advertises a more general prefix. If the holder of a portable allocation wants to authorize an ISP to originate routes to its allocation, then it must issue a ROA to this ISP; none of the ISP's existing ROAs authorize it to originate routes to that portable allocation.

5.3. 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, which can be accomplished with the following procedure:

1. Obtain current certificates, CRLs, and ROAs from the repository system (e.g., update a previous download)

2. Verify each end-entity certificate by constructing and verifying a certification path for the certificate (including checking relevant CRLs).

3. Verify each ROA by verifying that it is signed by a valid end-entity certificate that matches the address allocation in the ROA.

4. Based on validated ROAs, construct a table of prefixes and corresponding authorized origin ASes (or vice versa).

In addition to this basic route-filtering technique, the infrastructure can be used to support more advanced routing-security systems, such as S-BGP [7] and soBGP [8].

The first three steps in the above procedure would incur a prohibitive amount of 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 objects (certificates, CRLs, and ROAs), perform necessary validations, then perform incremental downloads and validations on a regular basis. A typical ISP using the infrastructure might have a daily schedule to download updates from the repository, upload any modifications it has made, and construct route filters.

6. 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

7. IANA Considerations

This document makes no request of IANA.

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

8. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

9. References

9.1. 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] Housley, R., et al., "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 3280, April 2002.

[4] Lynn, C., Kent, S., and K. Seo, "X.509 Extensions for IP Addresses and AS Identifiers", RFC 3779, June 2004.

[5] Huston, G., Michaelson, G., and Loomans, R., "A Profile for X.509 PKIX Resource Certificates", draft-ietf-sidr-res-certs-03, February 2007.

[6] Kong, D., and Kent, S., " A Profile for Route Origin Authorizations (ROA)", draft-ietf-sidr-roa-format-00, February 2007.

9.2. Informative References

[7] [S-BGP]

[8] [soBGP]

[9] [rsync]

Author's Addresses

Richard Barnes
BBN Technologies

Email: rbarnes@bbn.com

Stephen Kent
BBN Technologies

Email: kent@bbn.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.