Network Working Group James Ng
Internet Draft (editor)
Expiration Date: March 2003 Cisco Systems
File Name: draft-ng-sobgp-bgp-extensions-00.txt October 2002

Extensions to BGP to Support Secure Origin BGP (soBGP)
draft-ng-sobgp-bgp-extensions-00.txt

1. Contributors

A large number of people contributed to this document; we've tried to include all of them here (in no particular order), but might have missed a few. From Cisco, Russ White, Alvaro Retana, Dave Cook, John Scudder, Martin Djernaes, Chris Lonvick, Brian Weis, Tim Gage, Scott Fanning, Barry Friedman, Jim Duncan, and Robert Adams. From Genuity, Tony Tauber.

2. Abstract

There is a great deal of concern over the security of routing systems within the Internet, particularly in relation to the Border Gateway Protocol [BGP], which is used to provide routing information between autonomous systems. This document proposes a system where the origin of any advertisement within BGP can be verified and authenticated, preventing the advertisement of prefix blocks by unauthorized networks, and verifying that the final destination in the path is actually within the autonomous system to which the packets are being routed.

This proposal is based on the threat models currently being worked on in the RPSEC working group; it does not:

- Attempt to provide information on how such a security system could or should be deployed; readers are referenced to [SOBGP-DEPLOY] for this discussion.

- Attempt to determine what sorts of keys should be used within such a system, nor how any sort of trust relationship can or should be built between the entities cooperating within the routing system.

This document primarily focuses on extensions to the BGP protocol itself to support such a security system. This document is to solve several vulnerabilities within a BGP routing system given the following constraints:

- Any signaling which must take place to provide security should be specified in as flexible a way as possible. For instance, any certificates exchanged may be solely contained with the BGP protocol, or through some other transport/distribution mechanism. [SOBGP-DEPLOY] contains more discussion on this issue.

- The proposed solution should be incrementally deployable, and require minimal changes in a network to support it (software changes only along the path where support is required, for instance).

- The proposed solution should not rely on the routing system it is securing to reach information necessary to secure the routing system.

- The operator may configure various levels of trust, and prefer faster convergence with slower security verification, or

stronger security verification over reachability and faster con-
vergence.

- All information received from peering relationships within an autonomous system are implicitly trusted.

3. Overview

This document deals with security within a BGP routing system in several pieces; each piece is discussed in a separate section:

- Carrying security information within BGP
- Certificates used in the security system
- Authenticating the identities of entities within the routing system
- Authorizing entities to advertise given blocks of address space
- Aggregation of prefixes within the routing system
- Verifying the path of any given advertisement.

It's important to note that any given security system is a tradeoff between several factors, including the actual security provided and the difficulty entailed in deploying the system; different solutions may provide different levels of protection between these, leaving some parts of the problem unsolved.

This proposal does not protect against:

- Attacks through all forms of autonomous system spoofing. This proposal can be subverted if an attacker is able to masquerade as an autonomous system under some network conditions. It is assumed that management of inter-AS connections and policy implementation can resolve these possible attacks.
- Path Authentication. While this system allows the path any given update takes to be checked against possible paths, it does not ensure this particular update provides the correct path. It is observed that there are many possible valid paths within a BGP routing system, and BGP itself is designed to find the correct path among the possible paths.
- Authentication of path attributes, such as the community and MED attributes passed with a prefix.

4. Carrying Security Information in BGP

5. The Security Message

This document proposes a new message type, the SECURITY message, which is to be used for carrying security information within the BGP protocol. The SECURITY message is type 6. The security message is used to transport three types of certificates and a request format for requesting security certificates. These certificates types are:

- The Entity Certificate (EC)
- The Policy Certificate (PC)
- The Authorization Certificate (AC)

  Each of these certificates are described in the section "Certificates Used in the Security System."

5.1. Negotiating Security Capability

The ability to exchange SECURITY messages shall be negotiated at session startup, as described in [CAPABILITY]. The capability code is <to be assigned by IANA>. A pair of BGP speakers MAY negotiate to send messages in the SECURITY message type only (exchange no NLRI or forwarding information).

If two BGP speakers have negotiated the exchange of security messages, they should exchange the Authorization Certificates contained in their local databases. Entity Certificates and Policy Certificates may also be exchanged, as determined by local configuration of the BGP speakers. Certificates may be filtered by either speaker in a peering relationship.

5.2. The Security Message Format

The SECURITY message is formatted as described in [BGP], with a type code of 6. Within each message is a series of TLVs, or security message blocks, formatted as:

- Type: A two octet unsigned integer describing the type of

information contained within the data field.

- Length: A two octet unsigned integer describing the length of the data field, in octets.

- Data: The data, as described within this and other documents which describe information to be carried within the SECURITY message type.

Four TLVs are currently defined within the SECURITY message.

5.2.1. The Entity Certificate TLV

Entity Certificates MAY be encapsulated within a TLV type 1 and transmitted within the SECURITY message type.

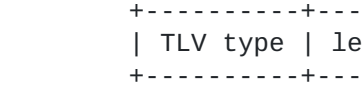

- TLV type: (2 octets), 1 (0x0001)

- length: (2 octets), set to the length of the Entity Certificate in octets.

- EC: The Entity Certificate

5.2.2. The Policy Certificate TLV

Policy Certificates should be encapsulated within a TLV type 2 and transmitted within the SECURITY message type.

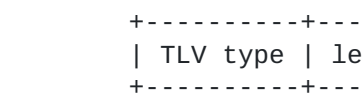

- TLV type: (2 octets), 2 (0x0002)

- length: (2 octets), set to the length of the Policy Certificate.

- PC: The Policy Certificate

5.2.3. The Authorization Certificate TLV

Authorization Certificates SHOULD be distributed using the SECURITY message TLV type 3.

```
+----------+--------+----------
| TLV type | length | AC
+----------+--------+------------
```

- TLV type: (2 octets), 3 (0x0003)
- Length: (2 octets), set to the total length of the Authorization Certificate in octets.
- AC: The Authorization Certificate.

5.2.4. The Request TLV

A BGP speaker may request a certificate from a peer using the SECURITY message TLV type 4.

```
+----------+--------------+--------+------------------------
| TLV type | request type | length | request indicator SubTV
+----------+--------------+--------+--------------------------
```

- TLV type: (2 octets), 4 (0x0004)
- Request Type: (1 octet), treated as an unsigned integer indicating the type of the type of information requested. The enumerated options are describe below.
- Length: (2 octets), set to the total length of the request in octets.
- Request Indicator: The information indicated by the request type bit field.

The request type is a two octet bit field denoting the type of information requested.

- Bit 1: If set, indicates the sender is requesting a set of Entity Certificates be returned.
- Bit 2: If set, indicates the sender is requesting a set of Policy Certificates be returned.
- Bit 3: If set, indicates the sender is requesting a set of

Authorization Certificates be returned.

Request indicator SubTVs restrict the set of certificates returned; there may be one or more request indicator SubTVs included in a request. Each SubTV consists of a 2 octet type field, treated as an unsigned integer, and a fixed length field containing the request indicator.

- Type 1: A four octet origin/authorized AS Number; two octet AS numbers shall be right justified within this field (placed in the two least significant octets).
- Type 2: A four octet signer/authorizer AS Number; two octet AS numbers shall be right justified within this field (placed in the two least significant octets).
- Type 3: A four octet IPv4 address is included in the request indicator.
- Type 4: A sixteen octet IPv6 address is included in the request indicator.
- Type 5: A four octet starting serial number is included in the request indicator.
- Type 6: A four octet ending serial number is included in the request indicator.

6. Certificates Used in the Security System

Three different certificates are used in soBGP, each of which carries a different piece of information. The Operation section describes how these certificates are used.

6.1. The Entity Certificate

Entity Certificates are used to verify, through a trust model, the existence of an entity within the routing system, and the value of that entity's public key for use in the routing system. Each entity within the routing system MUST generate a public/private key pair. The public key portion of this pair is then signed, verifying that anyone using this public key is actually the entity in question. This signature may be provided by various other trusted parties within the routing system, including (but not limited to):

- The authority which issued the autonomous system number.
- An external commercial authority which provides authentication certificates for other commercial transactions.
- Any other trusted party within the domain of Internet routing, such as a well known Service Provider.
- Self-signed if the entity is well known within the routing system.

This public key is used to verify the validity of other messages transmitted by this entity within the routing system.

The public key, along with other verifying information, are formatted into an Entity Certificate.

Any device receiving this certificate can verify it by checking the signature(s) on the certificate, along with the verifying information, and can use the public key contained in this certificate to verify messages purportedly sent by this entity. Self-signed certificates may be used but are not recommended.

The format of each item in the Entity Certificate is outside the scope of this document; they will be defined in a separate document, not yet published, discussing the exact semantics of the certificates and the certificate exchange mechanisms.

For this proposal to succeed, however, the Entity Certificate must contain:

- AS: A four octet unsigned integer containing the autonomous system number of the originating entity. Two octet AS numbers shall be right aligned in this four octet field.
- Key: The public key of this entity.
- Signer AS: A four octet unsigned integer indicating the autonomous system of the trusted signer. Two octet AS numbers shall be right aligned in this four octet field. Each entity which signs Entity Certificates MUST be assigned an AS number, even if they do not originate routes into the internetwork.
- Serial Number: A serial number which identifies this certificate; this should be taken from an circular number space of at least 32 bits in length.

- Lowest Valid Entity Certificate Serial Number: The lowest Entity

Certificate serial number which can be considered valid; this field should be four octets.

- Begin Authentication Certificate Serial: A four octet unsigned integer containing the first serial number considered valid for an Authorization Certificate for this key and autonomous system.

- End Authentication Certificate Serial: A four octet unsigned integer containing the last serial number considered valid for an Authorization Certificate for this key and autonomous system.

- Trusted Signature: A trusted signature validating the contents of the Entity Certificate.

Entity Certificates MUST be able to be saved in a text file format, ASCII encoded, and copied to a device's local configuration memory for bootstrapping. Any entity certificates manually configured and copied to a device's local configuration are implicitly trusted as being previously verified and authenticated.

6.2. The Policy Certificate

A second certificate, called the Policy Certificate (PC), which provides information about AS' which originate prefixes, is also required; this certificate is formatted as a series of TLVs.

Each TLV will includes a type, which is treated as a 16 bit (two octet) unsigned integer, a length, which is also two octets, and a variable length data field. TLVs MUST be placed in the Policy Certificate in type order.

6.2.1. The Autonomous System

- TLV type: 1 (0x0001)

- Length: Set to 4.

- Autonomous System: (four octets), the autonomous system which originated this certificate. Two octet AS numbers MUST be placed in the least significant position within the four octet field (the right most octets).

6.2.2. The Serial Number

```
+----------+--------+---------------+
| TLV type | length | serial number |
+----------+--------+---------------+
```

- TLV type: 2 (0x0002)
- Length: Set to 4.
- Serial Number: (four octets), A serial number which identifies this Policy Certificate, taken from a 32 bit circular number space.

6.2.3. Attached Autonomous Systems

```
+----------+--------+-------------------+
| TLV type | length | autonomous system |
+----------+--------+-------------------+
```

- TLV type: 3 (0x0003)
- Length: Set to 4.
- Autonomous System: (four octets), autonomous systems which are connted to the originating autonomous system through some form of peering arrangement. Two octet AS numbers MUST be placed in the least significant position within the four octet field (the right most octets).

One or more Attached AS TLVs may be included in the Policy Certificate. Each type 3 TLV indicates an AS which is connected to the AS which originates this Policy Certificate through a BGP peering relationship.

6.2.4. Revoked Authorization Certificate Serial

```
+----------+--------+-------------------+------------------
| TLV type | length | revoked AC serial | revoked AC serial
+----------+--------+-------------------+--------------------
```

- TLV type: 4 (0x0004)
- Length: length of TLV data (the list of revoked Authorization Certificates) in octets

- Revoked Authorization Certificate Serial: A list of serial numbers of Authorization Certificates which have been revoked, each serial number being four octets in length.

A single Revoked Authorization Certificate Serial TLV may be included in a Policy Certificate.

6.2.5. Authorization Certificate Policies

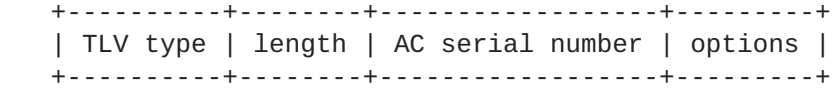

- TLV type: 5 (0x0005)

- Length: Set to 6.

- Authorization Certificate Serial Number: (four octets), the serial number of the Authorization Certificate which these policies apply to.

- Options: (two octets), a bit field describing various policies which should be applied to the prefixes indicated.

The options bit field describes policies which should be applied to the address block described in the TLV. These options are:

- Bit 0: Path Check. If this bit is set, the receiver should not accept any prefix for which the path cannot be verified as described in the section Verifying the Path, below.

- Bit 1: Second Hop Check. If this bit is set, the receiver should not accept any prefix for which the second entry in the AS PATH cannot be verified as described in the section Verifying the Second Hop, below.

- Bits 2-15: Reserved for future use.

A Policy Certificate may contain more than one Authorization Certificate Policy TLV.

6.2.6. Signature

```
+----------+----------------+--------+-------
| TLV type | signature type | length | signature
+----------+----------------+--------+------------
```

- TLV type: 65535 (0x00FE)

- Signature Type: A two byte unsigned integer denoting the type of signature (the algorithm used to build this signature). Each possible signing algorithm is assigned an integer from this field.

- Length: A two byte unsigned integer denoting the length of the signature which follows.

- Signature: The signature itself.

  The signature is calculated using the private key of the authorizing entity across all the TLVs within the Policy Certificate. It MUST be built using the same signature algorithm which is specified in the Entity Certificate.

Policy Certificates must be savable as ASCII encoded strings, and loadable through configuration at bootstrap. Any Policy Certificates manually configured MUST be implicitly trusted as accurate by the device on which they are configured.

6.3. The Authorization Certificate

The Authorization Certificate will be formatted as a series of TLVs. Each Authorization Certificate TLV includes a type, which is treated as a 16 bit (two octet) unsigned integer. The TLVs described must be placed within the Authorization Certificate in type order; every Authorization Certificate should begin with a TLV type 1 (Autonomous System and Options).

6.3.1. The Authorizing AS

```
+----------+--------+----+
| TLV type | length | AS |
+----------+--------+----+
```

- TLV type: 1 (0x0001)

- Length: Set to 4.

- AS: Four octets, the autonomous system authorizing other entities to advertise prefixes within this block. Two octet AS numbers MUST be placed in the least significant position within the four octet field (the right most octets).

Each authorizing entity MUST have an autonomous system number, used as a unique identifier, even though they may not advertise prefixes into the routing system.

6.3.2. Authorized Originator

```
+----------+--------+----+
| TLV type | length | AS |
+----------+--------+----+
```

- TLV type: 2 (0x0002)

- Length: Set to 4.

- AS: Four octets, the autonomous system of an entity authorized to advertise prefixes within this block. Two octet AS numbers should be placed in the least significant octets of this four octet field (the two rightmost octets).

Multiple authorized originator TLVs may be included in the Authorization Certificate.

6.3.3. The Serial Number TLV

```
+----------+--------+---------------+
| TLV type | length | serial number |
+----------+--------+---------------+
```

- TLV type: 3 (0x0003)

- Length: Set to 4.

- Serial Number: A four octet unsigned integer indicating the serial number of this Authorization certificate.

6.3.4. Uniform Resource Locator

```
+----------+--------+---
| TLV type | length | URL
+----------+--------+-------
```

- TLV type: 4 (0x0004)
- Length: A two byte unsigned integer denoting the length of the following URL in octets.
- URL: A uniform resource locator indicating a location where the public key of the entity which signed this certificate can be found, encoded in ASCII format.

An Authorization Certificate may contain zero or more URL TLVs (the URL TLV is optional).

6.3.5. The Address Block TLV

The address block TLV shall define blocks of address within which the authorized AS' are allowed to advertise prefixes (or routes).

```
+----------+---------------
| TLV type | NLRI data
+----------+----------------
```

- TLV Type: 14 (0x000D)
- NLRI Data: An address block as described in [MULTIPROTOCOL-BGP].

6.3.6. Signature

```
+----------+----------------+--------+-------
| TLV type | signature type | length | signature
+----------+----------------+--------+------------
```

- TLV type: 65535 (0x00FE)
- Signature Type: A two byte unsigned integer denoting the type of signature (the algorithm used to build this signature). Each possible signing algorithm is assigned an integer from this field.
- Length: A two byte unsigned integer denoting the length of the signature which follows.

o Signature: The signature itself.

The signature is calculated using the private key of the authorizing entity across all the TLVs within the Authorization Certificate.

7. Operation of the Security System

There are four steps in the operation of this security system; each is described in detail in the sections below. Each Entity Certificate must be validated, each Authorization Certificate must be validated, the information contained in the each Policy Certificate must be correlated with the information in the Authorization Certificate database, and each prefix must be validated against the Authorization Certificate database.

Note that any and all information which is manually configured on a device, including certificates and routes, MUST be implicitly trusted as accurate and valid.

Assume we begin with some small set of Entity Certificates received through manual configuration, and a device running soBGP has connected to other devices running soBGP. As the device receives Entity Certificates from other devices (as described in [SOBGP-DEPLOY] and other documents), each Entity Certificate is validated against known Entity Certificates already in a local Entity Certificate database. As Policy Certificates are received, they are validated using the data in this Entity Certificate database, and placed in a local database as well.

Each entity which provides a block of addresses to another entity will create a signed Authorization Certificate, and may provide it to the receiving AS or distribute this certificate to other peers on behalf of the authorized AS.

As each Authorization Certificate is received, the public key of the authorizing entity, obtained from the Entity Certificate database, is used to verify the origin AS and the prefix block, and this information is placed in an Authorization Certificate database stored locally on the router. Note that individual prefixes are not advertised using Authorization Certificates, but blocks of addresses. Thus, if an entity is authorized to advertise all the addresses within the 10.0.0.0/8 address space, but advertises multiple blocks

within this address space, only one certificate that covers the entire block is required.

Combining the authorization information in the Authorization

Certificates with the policy information in the Policy Certificates, a device can validate received prefixes and determine the policies which should be applied to them.

7.1. Receiving and Validating Entity Certificates

As each Entity Certificate is received:

- The local validated Entity Certificate database is examined, and any certificates which contain an AS number equal to the Signer AS within the certificate being validated are taken.
- The public key of each of the Entity Certificates matching the criteria above would be used to attempt to validate the signature on the certificate being examined.
- If any of these signatures validates the certificate being examined, the certificate is placed in the local database of validated Entity Certificates.
- If none of these signatures validate the certificate being examined, it is discarded.
- If no Entity Certificate is found with an AS number equal to the Signer AS in the certificate being examined, it should be discarded. The local system may request an Entity Certificate which contains the correct public key to validate this certificate from its peers. An unvalidated Entity Certificate may also be kept in a local database of unvalidated certificates for transmission to other peers.

A note on trust: The user may configure the device to trust only those Entity Certificates signed by a few trusted entities, as represented in the Signer AS field of the Entity Certificates received. Or, the user may configure the device to accept only a certain "depth" of signatures, trusting second party signatures, but not third party signatures. This level of trust associated with Entity Certificates and their signatures is outside the scope of this document, and will be dealt with in a forthcoming document describing these certificates more fully.

Once the Entity Certificate has been validated:

- The serial number of this certificate is examined. If this

serial number matches the serial number of an Entity Certificate originated by the same AS already in the database, and all other fields in the certificates match, this will be treated as a

duplicate, and be discarded. No other action needs to be taken.

- If the serial number of the certificate matches the serial number of an Entity Certificate originated by the same AS which is already in the database, and any other field of the certificates do not match, this will be treated as a protocol error, and both Entity Certificates will be discarded.

- The serial number and lowest valid serial number fields are examined. Of the Entity Certificates contained in the local Entity Certificate database with the same AS, the certificate with the highest serial number is considered the newest.

- The Lowest Valid Entity Certificate Serial Number is taken from the newest certificate in the validated Entity Certificate database, and any certificates with serial numbers lower than this serial number are discarded.

- If an Entity Certificate is discarded, the Authorization Certificate database shall be examined, and any Authorization Certificates which were validated using the discarded certificate should be revalidated, or removed from the database if they no can no longer be validated using the currently available Entity Certificates. See the section "Validating Authorization Certificates" below for information on validating impacted Authorization Certificates. See the section "Discarding Authorization Certificates" below for information on what actions to take when discarding an Authorization Certificate.

- If an Entity Certificate is discarded, the Policy Certificate database shall be examined, and the Policy Certificate for the originating AS shall be revalidated using the currently available valid Entity Certificates. If the Policy Certificate cannot be revalidated, it shall be discarded; see the section "Discarding Policy Certificates" below.

7.2. Receiving and Validating Policy Certificates

As each Policy Certificate is received:

- The database of validated Entity Certificates should be examined for Entity Certificates which contain the same AS as this Policy Certificate.

o For each matching Entity Certificate found, the signature on the certificate should be verified using the public key contained in the Entity Certificate. If any single matching Entity

Certificate validates this certificate, it is considered valid.

- If no matching Entity Certificate is found, the device may request all known Entity Certificates for this AS and attempt to validate the Policy Certificate with any new information received. A Policy Certificate for which there is no matching Entity Certificate to validate it may be kept in a local database of unvalidated certificates for transmission to peers.

- If the signature cannot be verified using any key from a matching Entity Certificate, the Policy Certificate should be discarded.

Once the Policy Certificate has been validated:

- If a Policy Certificate exists in the database which has the same serial number and AS, and all other fields match, it shall be treated as a duplicate, and discarded. No further action is needed.

- If a Policy Certificate exists in the database which has the same serial number and AS, and other fields do not match, this shall be treated as a protocol error, and both certificates shall be discarded.

- Any Policy Certificates with lower serial numbers with the same AS shall be discarded.

- The Policy Certificate should be linked to each Entity Certificate with a matching AS number.

- The list of revoked Authorization Certificates shall be examined, and with a serial number listed should be removed from the local validated Authorization Certificate database. Actions taken when discarding an Authorization Certificate are described in the section "Discarding Authorization Certificates" below.

- Once any Discarded Policy Certificates are processed, each Authorization Certificate Policy TLV contained in the Policy Certificate shall be associated with the correct certificate in the valid Authorization Certificate database.

- The attached AS list contained in the certificate should be placed in the path directed graph database, and the algorithm which computes the directed graph should be recalculated.

Discarding Policy Certificates

- The attached AS list in any Policy Certificate which is thus discarded should be examined, and those AS' on the list should be removed from the directed graph database.

- Each Authorization Certificate TLV within the Policy Certificate should be examined, and any association with Authorization Certificates removed.

If all the Policy Certificates for a given AS are discarded due to protocol errors or some other reason, each Authorization Certificate which authorizes the advertisement of addresses from the AS are examined, and their policies are set to default values:

- The Path Check bit is cleared.

- The Second Hop Check is set.

- The revocation list is left the same as the last known good Policy Certificate.

7.3. Receiving and Validating Authorization Certificates

As each Authorization Certificate is received:

- The database of validated Entity Certificates should be examined for an Entity Certificate which matches the authorizing AS, the Start Authorization Certificate Serial Number is lower than the certificate's serial number, and the End Authorization Certificate Serial Number is higher than the certificate's serial number.

- For any (or every) Entity Certificate which matches this criteria, the public key is taken.

- If no Entity Certificate matches this criteria, the device may request an appropriate Entity Certificate to validate this certificate (from its peers), or it may use any URL embedded in the certificate to retrieve the necessary Entity (and possibly Policy) Certificate. If no Entity Certificate can be found with which to validate this certificate, it should be discarded, although it may be kept in a local database of unvalidated certificates for transmission to peers.

- The certificate's signature is verified against each of these public keys.

- If the certificate's signature can be verified using one of

these public keys, the Authorization Certificate should be placed in a local database of validated Authorization Certificates.

- If the certificate's signature cannot be verified using one of these public keys, it MUST be discarded.

- Once the certificate's signature has been validated, the serial number of the validated Authorization Certificate should be checked against the list of revoked Authorization Certificates contained in the Policy Certificate associated with the Entity Certificate used to validate the certificate. If this certificate's serial number is contained in the revoked list, it MUST be discarded.

Once an Authorization Certificate is placed in the local database of validated Authorization Certificates:

- Any Authorization Certificate with an identical address space (prefix and prefix length), authorized originating AS', and a lower serial number should be discarded.

- If any certificate's are discarded, the actions described in the section "Discarding Authorization Certificates," below, should be taken.

- The list of Authorization Certificate Policy TLVs contained in the Policy Certificate associated with the Entity Certificate used to validate this certificate shall be examined for a policy TLV which matches this certificate's serial number. This policy TLV shall be associated with the certificate.

- For any Authorization Certificate which has changed in any way other than the signature, the Adj-RIB-In should be examined for those prefixes which fall into the ranges permitted by the certificate. These prefixes should be revalidated through the process described in "Receiving and Validating Prefixes" below.

7.4. Discarding Authorization Certificates

For each Authorization Certificate which is discarded:

- The Adj-RIB-In should be examined for any routes which fall within the address range described by any discarded certificate's and are sourced from one of the authorized AS' listed in the certificate.

- Any prefix found which matches those conditions described above should be revalidated using the process described in "Receiving and Processing Prefixes," below.

- If a prefix is found to be invalid in this process, it should be removed from the Adj-RIB-In, and steps taken as described in [BGP] to withdraw the route, as needed.

7.5. Receiving and Validating Prefixes

Validation of prefixes received by a device may take place as the prefixes are received, periodically, or at some time after the prefix is received.

- The local Authorization Certificate database is examined, and the certificates with the longest prefix length which encompasses the range of addresses described by the prefix is chosen.

- This set of Authorization Certificates is examined to determine if the route is originated from one of the AS' listed in the set of certificates.

- If the route is originated from one of the AS' listed in the set of certificates, the route is said to be validated.

- If the is not originated from one of the AS' listed in the set of certificates, the route is said to be invalid, and should be discarded.

- If the Path Check bit is set in the policy TLV associated with the Authorization Certificate, the AS PATH of the route should be verified, as described in the section "Verifying the AS PATH," below, should be followed.

- If the Second Hop Check bit is set in the policy TLV associated with the Authorization Certificate, the second hop of the AS PATH should be verified as described in the section "Verifying the Second Hop," below.

- If there is no Authorization Certificate which encompasses the range of addresses described by the prefix, then the route is said to be unverified, and should be handled according to local policy (either discarded, or have its security preference lowered).

7.6. Aggregation

Aggregation is a difficult problem with any method which attempts to verify the origin of any given prefix, since aggregation removes the relationship between prefixes originated and originators. Prefixes may only be aggregated by an entity which is otherwise authorized to advertise the aggregated prefix.

7.7. Verifying the Path

In order to verify the path of any given advertisement, we propose to build a directed graph of all possible transit paths within the Internet, and verifying the path of each advertisement against this graph. Note that for any given prefix, this graph will contain a superset of all valid paths for the prefix. The BGP protocol itself is designed to choose the correct path out of the possible paths.

In order to build the directed graph, each entity within the routing system may advertise the autonomous systems it is connected to using the attached AS' field in the Policy Certificate, described above. This graph is called the PATH database.

7.7.1. Building and Using the Directed Graph

Each device verifying routing information it is receiving may build a directed graph from the information contained in the Policy Certificate for each AS (as described above), which provides a list of all known valid paths through the internetwork. Each link between a pair of AS' may be verified from both ends of the link (using a two way connectivity check), thus ensuring that no single AS may advertise itself as connected to an entity it is not connected to.

As routes are received from BGP peers, the AS PATH contained in the route may be checked against the PATH database for congruence with a known good path. If the AS PATH traverses a link which cannot be verified to exist in both directions (the link fails the two way connectivity check), it's security preference may be lowered or the route may be discarded, as local device configuration directs.

7.8. Verifying the Second Hop

As a prefix is processed by a receiving the device, the originator and second hop in the AS PATH may be checked against the authorized originator and the list of attached AS' advertised by that originator. If the second hop in the AS path (the second entry in the AS

PATH) does not match one of the attached AS' advertised by the originator, the prefix should be treated as invalid, and discarded.

7.9. Receiving and Processing Requests

If a device receives a Request TLV, as described in the section "The Security Message," above, it should:

- Examine the request to ensure it is logically consistent. For instance, requesting an Entity Certificate based on an IPv4 address range is not logically consistent, since these certificates only contain an AS and a Signer AS.
- If the request is not logically consistent, discard it.
- If the request is logically consistent, examine its local databases, and transmit the certificates requested which fulfill the conditions supplied in the request indicator SubTVs.

Logically consistent requests and the returned results include:

- If Entity Certificates are requested, and the request indicator includes an AS number, any Entity Certificates which contain an AS number matching the request indicator should be transmitted to the requester.
- If Entity Certificates are requested, and the request indicator includes a Signer AS, any Entity Certificates which contain an AS number matching the request indicator should be transmitted to the requester.
- If Policy Certificates are requested, and the request indicator includes an origin AS, any Policy Certificates which contain an AS number matching the request indicator should be transmitted to the requester.
- If Authorization Certificates are requested, and the request indicator includes an origin AS, any Authorization Certificates which contain an authorized AS number matching the request indicator should be transmitted to the requester.
- If Authorization Certificates are requested, and the request indicator includes an IPv4 or IPv6 address and prefix length, any Authorization Certificates whose address blocks are encompassed by or match the request indicator should be returned to the requester.

- If Authorization Certificates are requested, and the request indicator includes both an IPv4 or IPv6 address and prefix length, and an origin AS, the set of Authorization Certificates whose address blocks encompassed by or match the address included, and whose authorized AS matches the AS number included in the request indicator, should be transmitted to the requester.

- If Authorization Certificates are requested, and the request indicator includes a Signer AS, any Authorization Certificates which contain an authorizer AS matching the request indicator should be transmitted to the requester.

- If Authorization Certificates are requested, and the request indicator includes both an IPv4 or IPv6 address and prefix length, and a Signer AS, the set of Authorization Certificates whose address blocks encompassed by or match the address included, and whose authorizing AS matches the AS number included in the request indicator, should be transmitted to the requester.

- If start and end serial numbers are included in the request, the set of returned certificates should be restrained by the range of serial numbers indicated.

If more than one of the same request indicator is included in a request message, they shall be treated as an OR condition; if any of the conditions match, the certificate shall match the set.

8. The Security Preference

As indicated in other sections of this document, the user has a wide range of flexibility when determining the level of security to be applied to each prefix. The level of security may be translated into a Security Preference and used to influence the route selection process [BGP].

To determine the Secuirty Preference of a prefix, the following algorithm may be used:

1 **A Security Preference is assigned to each prefix indicating neutral trust. The neutral value is locally significant to the autonomous** system, and all routers in it MUST use the same value.

2 **The Security Preference SHOULD be lowered for any of the following cases:**

- The AS_PATH verification failed at a link other than the second hop.
- The Second Hop verification failed.
- The prefix is unverified.
- The prefix is invalid.

Each of these cases represents a different degree of failure in the validation and verification process. The ammount by which the Security Preference is lowered is locally significant to the autonomous system, and all routers in it MUST use the same value. If the local policy determines that a prefix may be used (even if it may have failed the validation and verification process), it is recommended that the Security Preference be lowered such that a lower value is assigned to invalid paths.

3 The Security Preference SHOULD be increased for any of the following cases:

- The prefix is validated.
- The Second Hop verification passed.
- The AS_PATH verification passed for the whole path.

The ammount by which the Security Preference is increased is locally significant to the autonomous system, and all routers in it MUST use the same value. It is recommended that the Security Preference be increased such that a higher value is assigned to paths that pass all the validation and verification steps.

The resulting Security Preference value may be used to select among different routes for the same prefix; the higher value MUST be preferred. Any of the following methods may be used:

A Consider the Security Preference prior to calculating the degree of preference [BGP] for a prefix.

B Assign the value of the Security Preference to any of the attributes used in the Decision Process [BGP]. Care must be taken with attributes for which the lower value is preferred.

C Use a Cost Community [COST] and its associated methods to consider the Security Preference at any step in the Decision Process [BGP] without overloading other attributes. Care must be taken as the lowest value in a Cost Community is preferred.

The method selected MUST be consistent through the local Autonomous System.

9. Security Considerations

This document defines extensions to BGP that address specific security concerns for the protocol. While it adds functionality, the flexilibty allows it to not introduce any new security concerns.

10. IANA Considerations

This document defines the Security Message for BGP, which contains a series of TLVs. IANA is expected to maintain a registry of all the values defined, as follows:

The Type field:

- Type value 0 is reserved.
- Type values 1 through 4 are assigned in this document.
- Type values 5 through 16575 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].
- Type values 16576 through 32895 SHOULD be assigned using the "Specification Required" policy defined in RFC2434 [RFC2434].
- Type values 32896 through 65535 are for "Private Use" as defined in RFC2434 [RFC2434].

Request Type Field (in the Request TLV):

- Bits 1 through 3 are assigned in this document.
- Bits 4 thru 8 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].
- Bits 9 thru 16 are for "Private Use" as defined in RFC2434 [RFC2434].

Type Field (in the Request TLV SubTVs):

- Type values 1 through 6 are assigned in this document.
- Type values 7 through 16575 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].

- Type values 16576 through 32895 SHOULD be assigned using the "Specification Required" policy defined in RFC2434 [RFC2434].

- Type values 32896 through 65535 are for "Private Use" as defined in RFC2434 [RFC2434].

The Policy Certificate Type Field:

- Type values 1 through 5 and 65535 are assigned in this document.

- Type values 6 through 16575 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].

- Type values 16576 through 32895 SHOULD be assigned using the "Specification Required" policy defined in RFC2434 [RFC2434].

- Type values 32896 through 65534 are for "Private Use" as defined in RFC2434 [RFC2434].

The Options Field in the Authorization Certificate Policies TLV:

- Bits 0 and 1 are assigned in this document.

- Bits 2 thru 7 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].

- Bits 8 thru 15 are for "Private Use" as defined in RFC2434 [RFC2434].

The Authorization Certificate Type Field:

- Type values 1 through 4, 14 and 65535 are assigned in this document.

- Type values 5 through 13 and 15 through 16575 MUST be assigned using the "IETF Consensus" policy defined in RFC2434 [RFC2434].

- Type values 16576 through 32895 SHOULD be assigned using the "Specification Required" policy defined in RFC2434 [RFC2434].

- Type values 32896 through 65534 are for "Private Use" as defined in RFC2434 [RFC2434].

11. References

[BGP]Rekhter, Y., and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.

[MULTIPROTOCOL-BGP]
Bates, T., Chandra, R., Katz, D., and Rekhter, Y., "Multiprotocol Extensions for BGP-4", RFC 2858, June 2000

[CAPABILITY]
Chandra, R., Scudder, J., "Capabilities Advertisement with BGP-4", RFC2842, May 2000

[SOBGP-DEPLOY]
White R (editor), "Considerations for Secure Origin BGP (soBGP) Deployment", Draft-white-sobgp-deployment-00.doc, October 2002

[COST]
Retana, A., White, R., "BGP Custom Decision Process", Work In Progress (draft-retana-bgp-custom-decision-00.txt), October 2002.

[RFC2434]
Narten, T., Alvestrand, H., "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 2434, October 1998.

[SOBGP-RADIUS]
Lovnick, Chris, "RADIUS Attributes for soBGP Support", draft-ietf-rpsec-radius-sobgp-00.txt, October 2002.

12. Editor's Address

James Ng (Editor) Cisco Systems 7025 Kit Creek Road Research Triangle Park, NC 27709 jamng@cisco.com