

Network Working Group  
Internet Draft  
Expiration Date: October 2004  
File Name: [draft-white-sobgparchitecture-00.txt](#)

Russ White  
(editor)  
Cisco Systems  
April 2004

Architecture and Deployment Considerations for Secure Origin BGP (soBGP)  
[draft-white-sobgparchitecture-00.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a "working draft" or "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>.

Abstract

There is a great deal of concern over the security of the Border Gateway Protocol, which is used to provide routing information to the Internet and other large internetworks. This draft provides an architecture for a secure distributed registry of routing information to address these concerns. The draft begins with an overview of the operation of this system, and then follows with various deployment scenerios, starting with what we believe will be the most common deployment option.

## 1. Background

There are two fundamental pieces of a routing system that need to be secured:

- o Adjacencies between devices running the routing protocol
- o Information carried within the routing protocol.

While security between BGP [[BGP](#)] speakers has been addressed in a number of ways, including cryptographic authentication [BGP-MD5] and limiting the attack radius through TTL mechanisms [GTSH], security for the information carried within BGP is not considered a solved problem.

This draft proposes a possible solution to securing the information within BGP, using the certificates and protocol extensions proposed in [SOBGP-BGPTRANSPORT], [[SOBGP-CERTIFICATE](#)], and [[SOBGP-RADIUS](#)].

A large number of people contributed to this draft; we've tried to include all of them here (but might have missed a few): James Ng, Tim Gage, Alvaro Retana, Dave Cook, Brian Weis, and Iljitsch van Beijnum.

## 2. General Theory

soBGP provides a secure registry mechanism against which a BGP speaker can check:

- o The authorization of the AS listed as the originating AS in any received update to advertise reachability to the prefix listed in the update.
- o The validity of the AS Path contained in the update.

We use the term validity in reference to the AS Path, in this document, to indicate the plausibility of the AS Path listed. As shown in [[PATH-CONSIDER](#)], it isn't possible to communicate authorization through an AS Path; only the existence or nonexistence of the AS Path listed can be proven.

soBGP operates by distributing a set of signed certificates, described in [[SOBGP-CERTIFICATE](#)], containing the information required to validate the two pieces of information given above. These certificates MAY be distributed using the mechanisms described in [SOBGP-BGPTRANSPORT], or some other mechanism. Once these certificates have



been received and processed (signatures validated, etc, as described in [[SOBGP-CERTIFICATE](#)], they form a database containing:

- o A listing of IP address blocks and the AS authorized to originate them.
- o Policies related to specific prefixes and blocks of addresses.
- o A list of autonomous systems connected to each autonomous system within the internetwork. This connection list is used to build a graph of AS interconnectivity within the internetwork, as described in the section Building the AS Connectivity Graph, below.

This effectively forms a secure registry of routing information which can be used to check the validity of routing information received from BGP peers. This database is termed the "authorization database." No assumption about the location of the authorization database is made within this document.

When soBGP is supported, a BGP speaker MUST have access to the authorization database. Possible methods of access include:

- o Have a local copy of this authorization database, and perform the checks described later in this document against that local database.
- o Pass received routing information to a locally maintained server for validation against that server's copy of the authorization database.
- o Accept filters built from a copy of the authorization database contained on a locally maintained server.

As BGP updates are processed, a security preference is assigned to each prefix, as described further in the Security Preference section of this document. BGP update processing is described in the Receiving and Processing Updates section of this document.



### **3. soBGP Operation**

Each section below provides detailed information on some aspect of soBGP operation.

#### **3.1. The Security Preference**

Rather than simply noting a given prefix should be dropped (not trusted) or retained (trusted), soBGP extends the concept of locally generated and maintained policy in BGP by assigning each prefix a Security Preference. This allows the local operator to drop prefixes not meeting certain security criteria, while simply lowering their preference for prefixes meeting some security criteria. This allows operators some flexibility in their implementation of security policies, especially as the security system is being tested, or while the security system isn't fully deployed.

While the amount by which the Security Preference is increased or decreased for any operation described in this draft is locally significant to the autonomous system. All devices processing routes against soBGP information **MUST** use the same mechanisms and values of the Security Preference to ensure consistent routing within the autonomous system.

If the Security Preference is set to a value precluding a route from further consideration in the decision process, the route should be discarded at that point, rather than continuing with the decision process.

The 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.



### **3.2. Building the AS Connectivity Graph**

Each ASPolicyCert advertised by a member of the internetwork contains a list of the autonomous systems the advertising AS is connected to, along with possible policy information about that connection. From this information, a graph of AS connectivity within the internetwork is built.

Any AS can be used as the starting point for building this graph, thus multiple disconnected graphs (representing section of the internetwork running soBGP and providing interconnection information) are possible. If every AS within the internetwork is providing interconnection information, one graph can be built containing all the internetwork's interconnections.

The process of creating this graph is:

- o Examine the list of connected autonomous systems advertised by the current AS.
- o Examine the ASPolicyCert of each AS the current AS is advertising as connected, and determine if that AS is advertising a connection back to the current AS. This is termed the two way connectivity check.
- o If the two way connectivity check passes, the connection SHOULD be added to the interconnection graph, and marked as trustable.
- o If the two way connectivity check fails, the connection MAY be added to the interconnection graph, but marked so a lower security preference will be assigned to AS\_PATHs traversing this link.
- o Repeat this process for each ASPolicyCert in the authorization database.

The resulting graph is called the internetwork graph.

### **3.3. Validating Routing Information**

For each prefix within a given BGP UPDATE message:

- o The local authorization database is examined, and the AuthCert with the longest prefix length encompassing the range of addresses described by the prefix is chosen.



- o If there is no entry in the local authorization database 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). The rest of this process is ignored in these cases.
- o The second hop in the AS\_PATH attribute is examined.
  - o If the second hop in the AS\_PATH is advertised as connected by the originating AS, the Security Preference for this prefix SHOULD be increased.
  - o If the second hop in the AS\_PATH is not advertised as connected by the originating AS, the Security Preference for this prefix SHOULD be decreased.
  - o If the second hop in the AS\_PATH is not advertised as connected by the originating AS and the originator's policy indicates the second hop MUST be validated, the prefix should be removed from further consideration.
- o The AS\_PATH attribute is compared to the internetwork graph.
  - o If the AS\_PATH described is contained within the internetwork graph, the Security Preference SHOULD be increased.
  - o If the AS\_PATH described is not contained within the internetwork graph, the Security Preference SHOULD be decreased.
  - o If the AS\_PATH traverses a connection which is only described by one of the two autonomous systems, this is a one way connection. Local policy may be used to determine if the security preference should be increased in this case.
  - o If the AS\_PATH described is not contained within the internetwork graph, and the originator indicated the AS\_PATH MUST be checked, the prefix should be removed from further consideration.
- o The AuthCert chosen at the first step is examined.
  - o If the authorized AS in the AuthCert matches the originating AS in the AS\_PATH, the Security Preference SHOULD be increased.
  - o If the authorized AS in the AuthCert does not match the originating AS in the AS\_PATH, the Security Preference SHOULD be



set low enough to cause the route to be discarded.

- o Other policies contained in the local authorization database should be applied as directed by the policy.

#### **3.4. Validating Received BGP UPDATES**

As BGP UPDATES are received, they may be processed in one of several ways:

- o Each prefix may be validated according to the process outlined in Validating Routing Information before they are installed in the Adj-RIB-IN.
- o Each prefix may be validated according to the process outlined in Validating Routing Information after they are installed in the Adj-RIB-In, but before they are considered in the BGP Best Path calculation.
- o Each prefix may be validated according to the process outlined in Validating Routing Information after they are run through the Best Path algorithm, but before they are installed in the local RIB.
- o Routes may be installed in the local RIB, and then validated using the process outlined in Validating Routing Information. Once validation is accomplished, adjustments to the local RIB and routes advertised to BGP peers may need to be adjusted.

#### **3.5. 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.

#### **3.6. Requirements for Systems Running soBGP**

This section describes requirements for autonomous systems running soBGP, requirements for BGP speakers forming external adjacencies from within such autonomous systems, and devices exchanging soBGP certificates.



- o Any peering session along the border of an autonomous system running soBGP SHOULD be authenticated through some means such as [BGP-MD5], IPsec ([ESP], [AH]), or through some other current, effective means of protecting BGP sessions from being hijacked, or otherwise abused.
- o Any peering session along which soBGP certificates are exchanged SHOULD be authenticated through some means such as [BGP-MD5], IPsec ([ESP], [AH]), or through some other current, effective means of protecting BGP sessions from being hijacked, or otherwise abused.
- o The AS\_PATH of any routing information received from any BGP peer outside the autonomous system MUST be checked to validate the next hop AS is the AS the update was received from. If the next hop AS in any received update does not match the configured AS the route is learned from, the update MUST be discarded.

#### **4. soBGP Deployment**

This section begins by describing what we believe to be the most practical deployment of this secure registry of routing information. Following sections describe some other deployment options that may prove useful in some situations, or may prove to be more practical than the deployment outlined in this section.

##### **4.1. Deploying soBGP on Distributed Registry Servers**

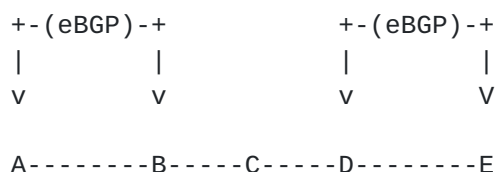
This deployment scenerio works within three constraints:

- o It may not be not desirable to combine routing and cryptographic processing of soBGP certificates on the same device.
- o The system should be distributed, using as few centralized resources as possible.
- o Trust relationships should be based on existing business and working relationships, rather than building new relationships specifically for securing the routing system.

Assume we have a small internetwork, as shown below:

```
S1 - - - - - S2 - - - S3
10.1.1.0/24---A---B-----C---D-----E---F
```







```

      ^           ^
      |           |
      +---(iBGP)---+

```

In this network, A is sending certificates it has learned from other sources to B using the mechanisms described in [[SOBGP-BGPEXT](#)]. It is passing these certificates to D via iBGP, and D is passing these certificates to E via eBGP. Each edge router, B and D, process these certificates locally, building the databases required to validate received routing information from them.

### 4.3. Multihoming Deployment

Multihoming presents a special challenge to the deployment of soBGP within a large scale internetwork.

```

      (-----)           (-----)
      ( AS65401 )         ( AS65402 )
      (           )       (           )
      (           )       (           )
      (---A---)          (---B---)
      |                  |
      \                  /
      \-----+          +-----/
              |          |
              (---C-----D---)
              (           )
              (  No-AS    )
              (-----)

```

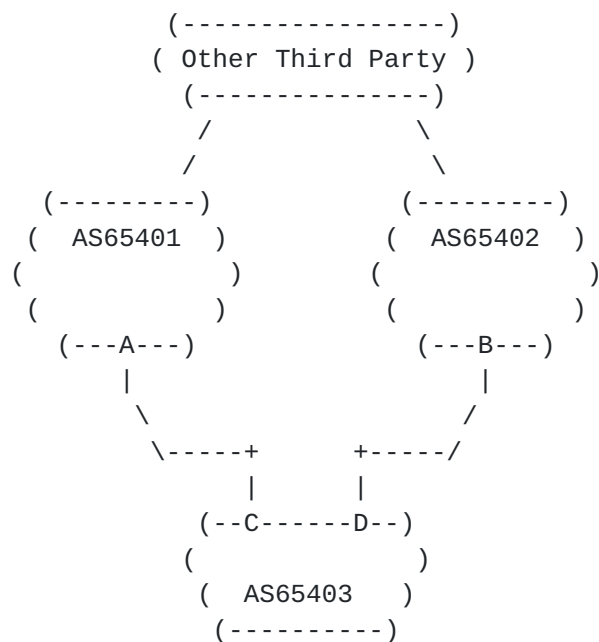
Assume No-AS has obtained a block of addresses, 10.1.1.0/24, from AS65401, and would like to advertise that same block of addresses through AS65402. Since No-AS has no AS number, it cannot generate any soBGP certificates, and must rely on its upstream providers to work out the security impact in some way. The simplest solution would be, of course, for NOAS to obtain an AS number, and fully participate in soBGP, but barring that, what other solutions are there?

AS65401 could issue a certificate allowing AS65402 to originate just the prefix in question, 10.1.1.0/24, or AS65401 could simply list AS65402 in the certificate covering 10.1.1.0/24 as an authorized originator for this address space (as multiple authorized originators are allowed).



#### 4.4. Proxy Advertisement of Certificates

Note there is no requirement for a given entity which originates routes into the routing system to actually originate the corresponding certificates required for the correct origination of the route to be validated, and the AS Path attached to the route to be verified.



In this case, AS65401, AS65402, or some other third part may actually advertise the certificates necessary for AS65403 to originate validated routes.

## 5. Other Deployment Considerations

In this section, we move from specific deployment scenerios to other deployment considerations, such as key generation and protection, and memory utilization/impact.

### 5.1. Certificate Generation and Private Key Protection

There is only one private/public key pair per autonomous system; certificates are generated as determined by local policy and as required to account for changes in the network. Since the entity's private key is not used in any part of the operations verifying received information, or in generating information to transmit to other devices, these certificates could be generated on some secure central system in the AS, and the results, containing only public keys, can be transmitted throughout the network.



Securing the private key of each entity should be relatively easy in this environment, since the location of the private key can be carefully constrained; no device other than the system which generates the required certificates needs use of the private key.

## **5.2. Impact on Performance and Memory Utilization**

Detailed performance and memory utilization characteristics of soBGP will be the subject of future investigation. However, as this is an important area of consideration, we present some suggested analysis below. (In other words, this is a guess).

In terms of memory, each device running sobGP will need to store:

- o Each of the Entitycerts Received. The maximum number of Entitycerts within the routing system would be the number participating autonomous systems multiplied by the number of outstanding Entitycerts from each autonomous system.
- o Each of the ASPolicycerts Received. The number of ASPolicycerts within the system will probably be similar to the number of Entitycerts within the system.
- o Each of the PrefixPolicycerts Received. The number of PrefixPolicyCerts within the system will depend on the number of address blocks each participant in the routing system advertises, and could double during key rollover.

Performance will depend on the cryptographic processing requirements imposed by the certificate signature methods, as described in [[SOBGP-CERTIFICATE](#)]. However, all of this additional memory and processing would most likely be required on a distributed soBGP server, rather than on routers themselves.

The primary impact on routers and routing protocol convergence will be the memory and processing requirements added from the additional route filters or processing as required by the deployment technique used.



## 6. Normative References

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

[SOBGP-BGPEXT]

Ng J (editor), "Extensions to BGP to Support Secure Origin BGP (soBGP)", [draft-ng-sobgp-bgp-extensions-01.txt](#), April 2004

[SOBGP-CERTIFICATE]

Weis, Brian (editor), "Secure Origin BGP (soBGP) Certificates", [draft-weis-sobgp-certificates-01.txt](#), October 2003

## 7. Informative References

[SOBGP-RADIUS]

Lovnick, C, "RADIUS Attributes for soBGP Support", [draft-lonvick-sobgp-radius-04.txt](#), February 2004

[PATH-CONSIDER]

White, Russ, "Considerations in Validating the Path in Routing Protocols", [draft-white-pathconsiderations-02.txt](#), April 2004

[COST]

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

## 8. Editor's Address

Russ White  
Cisco Systems  
7025 Kit Creek Road  
Research Triangle Park, NC 27709  
riw@cisco.com

