Network Working Group
Internet Draft
Expiration Date: October 2004
File Name: draft-ng-sobgp-bgpextensions-00.txt

James Ng
(editor)
Cisco Systems
April 2004

Extensions to BGP Transport soBGP Certificates
draft-ng-sobgp-bgpextensions-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.

1. Contributors

A large number of people contributed to or provided valuable feedback on this document; we've tried to include all of them here (in no particular order), but might have missed a few: Russ White, Alvaro Retana, Dave Cook, John Scudder, David Ward, Martin Djernaes, Chris Lonvick, Brian Weis, Tim Gage, Scott Fanning, Barry Friedman, Jim Duncan, Yi Yang, Robert Adams, Tony Tauber, Iljitsch van Beijnum, and Jonathan Natale.

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, verifying that the final destination in the path is actually within the autonomous system to which the packets are being routed, and proving the validity of the AS Path contained in the update.

This document does not:

- Attempt to provide information on how such a security system could or should be deployed; readers are referenced to [SOBGP-ARCH] 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. These are considered in [SOBGP-CERTIFICATE].

- Attempt to analyse the performance, memory utilization, or other impacts on devices running this protocol; these are addressed in [SOBGP-ARCH].

- Attempt to analyze the security protection provided by the proposed security system. This may be address in a future draft.

This document primarily focuses on extensions to the BGP protocol itself to support such a security system through the transport of the certificates described in [SOBGP-CERTIFICATE].

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

3. Definitions

- Entity: A participant in the internetwork routing system.

4. 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 [TBD]. The SECURITY message is used to transport the certificates described in [SOBGP-CERTIFICATE].

4.1. Negotiating Security Capability

The ability to exchange SECURITY messages MAY be negotiated at session startup, as described in [CAPABILITY]. The capability code is <to be assigned by IANA>.

- Speakers MAY negotiate the exchange of SECURITY information only or SECURITY and NLRIs.
- If the exchange of SECURITY messages is negotiated, the SECURITY option message MUST be exchanged before any other SECURITY messages are exchanged. The option bits in this message determine if SECURITY messages or NLRIs will be exchanged first.
- If two BGP speakers have negotiated to exchange SECURITY messages, they SHOULD exchange the soBGP certificates contained in their local databases.

4.2. The Security Message Format

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

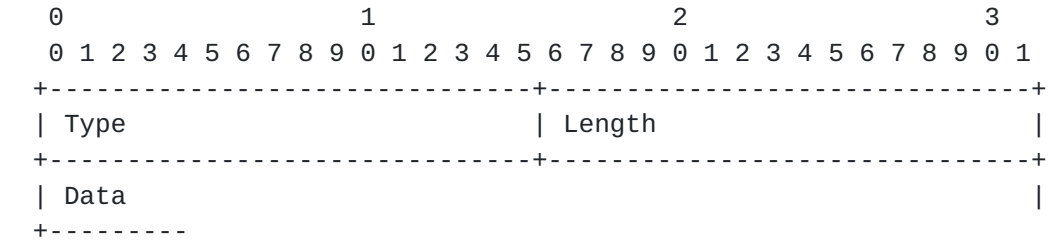

- Type: A two octet unsigned integer describing the type of information contained within the data field.

- 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.

Two TLVs are currently defined within the SECURITY message. Further TLVs are defined for carrying certificates in [SOBGP-CERTIFICATE].

4.2.1. The SECURITY Option TLV

The SECURITY Option TLV provides a way for exchanging speakers to inform their peers about local configurations which may pertain to the peering session. SECURITY Option TLVs are encapsulated within a TLV Type 1, and transmitted within the SECURITY message type.

If SECURITY Option TLVs are transmitted, they MUST be transmitted before the transmission of any other SECURITY data.

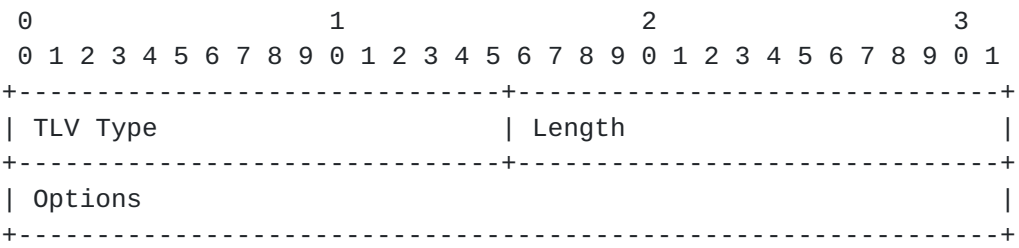

- TLV type: (2 octets), 1 (0x0001)

- Length: (2 octets), set to 2

- Options: (4 octets), a bitfield, described below

The options field is a 32 bit bitfield, allowing up to 32 different options to be specified.

- Bit 0: If set, indicates that SECURITY information should be sent before NLRI information on this session; if cleared, indicates that NLRI information should be sent before SECURITY information.

- Bit 1: If set, indicates that this peer will only transmit

validated certificates of any type along this session. This bit MUST NOT be used for eBGP sessions.

o Bit 2: If set, indicates that this peer will only accept validated certificates of any type along this session (valid only on iBGP sessions).

Bit 0 in the option field allows the operator to configure the local device so it receives all prefixes first, decreasing convergence to the minimum time, or receives all SECURITY information first, allowing all prefixes to be validated before they are installed.

Bits 1 and 2 allow peers along an iBGP session to trust the certifications they receive without validating them. If bit 1 is set on the transmitting peer, bit 2 is set on the receiving peer, and the BGP peering session is an authenticated or encrypted iBGP session, the receiving peer may accept all received certificates from the transmitting peer as already validated. This is called a trusted peering relationship.

4.2.2. The Request TLV

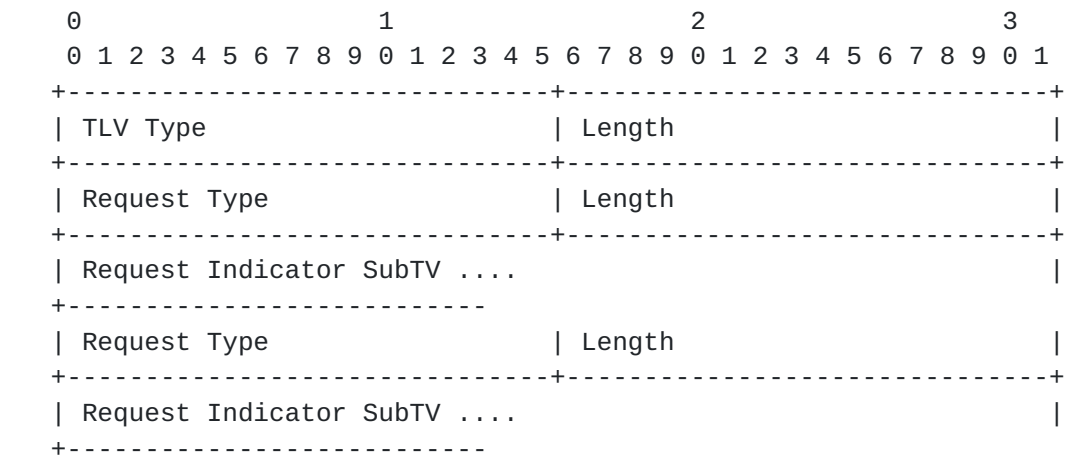

o TLV type: (2 octets), 2

o Length: (2 octets), set to the total length of the request in octets.

o Request Type: (2 octets), treated as an unsigned integer indicating the type of information requested.

o Length: (2 octets), set to the number of requests of the request type included in this request.

- Reserved: (2 octets), set to 0x0000.
- Request Indicator: The information indicated by the request type bit field.

The Request Type field indicates the type of certificates requested. Four request types are defined in this document.

1 Any certificate matching the Request Indicator are requested.

2 EntityCerts matching the Request Indicator are requested.

3 ASPolicyCerts matching the Request Indicator are requested.

4 PrefixPolicyCerts matching the Request Indicator are requested.

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 two 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: An eight octet starting serial number is included in the request indicator.
- Type 6: An eight octet ending serial number is included in the request indicator.

4.2.3. The Cluster List TLV

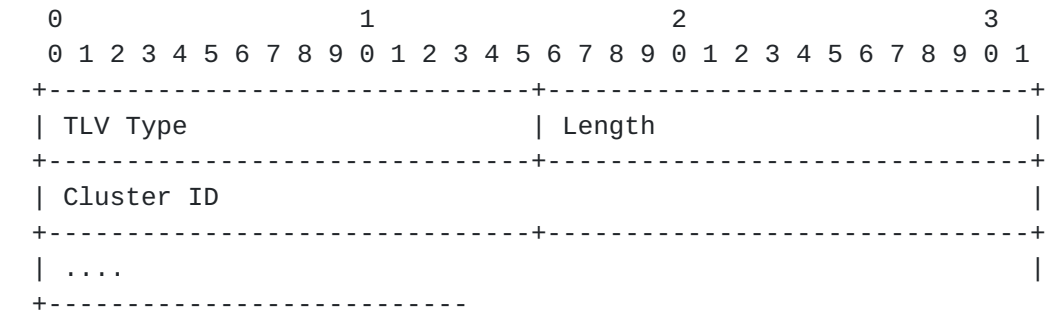

o TLV type: (2 octets), 3

o Length: (2 octets), set to the number of cluster IDs in the TLV

The use of the Cluster List TLV is described in the Reflecting SECURITY messages section below.

5. Receiving and Processing SECURITY messages

Each section below describes the receipt and processing of SECURITY messages.

5.1. Processing SECURITY Messages Containing a Certificate

For each certificate received, the BGP speaker MUST:

o Examine the certificate to determine if a copy of this certificate already exists in the local database. Any certificate which is found to already be held locally MUST be discarded.

o If the certificate is received through an untrusted peering relationship, place the certificate in a local certificate database and process according to [SOBGP-CERTIFICATE].

o If the certificate is received through a trusted peering relationship, place certificate in a local certificate database, treating it as if it is already validated according to [SOBGP-CERTIFICATE].

o If a received certificate is sucessfully validated using the process described in [SOBGP-CERTIFICATE], it should be readvertised to all peers outside the local autonomous system (eBGP peers). If the peering relationship is trusted, the certificate

should be advertised as validated by marking it as indicated in [SOBGP-CERTIFICATE].

5.2. Reflecting SECURITY Messages

A BGP speaker MAY be configured to reflect received SECURITY messages, with or without processing them, in a way similar to the way BGP routing information is reflected among iBGP speakers, described in [BGP-REFLECTION]. When reflecting SECURITY messages, a BGP speaker MUST:

- Examine the SECURITY message for the presence of a Cluster List TLV.
  - If a Cluster List TLV exists, and the local router ID is contained in the list of Cluster IDs, discard the SECURITY message.
  - If a Cluster List TLV exists, and the local router ID is not contained in the list of Cluster IDs, add the local router ID to the list and retransmit the SECURITY message to all BGP peers which have negotiated receipt of SECURITY messages.
  - If a Cluster List TLV does not exist, add a new Cluster List TLV to the SECURITY message, including the local router ID in the new TLV.

5.3. Filtering of Certificates

A BGP speaker may, for reasons of policy, filter soBGP certificates received from a peer.

- If a BGP speaker is part of a transit AS, it SHOULD NOT filter soBGP certificates.
- A BGP speaker MAY discard soBGP certificates which describe the authorization of address space which is being filtered out of the local routing information.

5.4. 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 Entitycert 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.
- 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.

6. Security Considerations

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

7. 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 SECURITY message Type field :

- Type value 0 is reserved.
- Type values 1 through 3 are assigned in this document.
- Type values 4 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 TLV Request Type Field:

- Types 1 through 3 are assigned in this document.
- Types 4 thru 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].

8. Normative 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-ARCH]
White, R. (editor), "Architecture and Deployment Considerations for Secure Origin BGP (soBGP)", draft-white-sobgp-deployment-03, April 2004

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

9. Informative References

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

[BGP-MD5]
Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC2385, August 1998

[ESP] Kent, S., and R. Atkinson, "IP Encapsulating Security Payload", RFC 2406, November 1998.

[AH] Kent, S., and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.

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

[BGP-REFLECTION]
Bates, T, et al, "BGP Route Reflection - An Alternative to Full Mesh IBGP", draft-ietf-idr-rfc2796bis-00.txt, March 2004

10. Editor's Address

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