

Network Working Group
Request for Comments: 3392
Obsoletes: 2842
Category: Standards Track

R. Chandra
Redback Networks
J. Scudder
Cisco Systems
November 2002

Capabilities Advertisement with BGP-4

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

This document defines a new Optional Parameter, called Capabilities, that is expected to facilitate the introduction of new capabilities in the Border Gateway Protocol (BGP) by providing graceful capability advertisement without requiring that BGP peering be terminated.

This document obsoletes RFC 2842.

1. Introduction

Currently BGP-4 requires that when a BGP speaker receives an OPEN message with one or more unrecognized Optional Parameters, the speaker must terminate BGP peering. This complicates introduction of new capabilities in BGP.

2. Specification of Requirements

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 [RFC2119].

3. Overview of Operations

When a BGP speaker [BGP-4] that supports capabilities advertisement sends an OPEN message to its BGP peer, the message MAY include an Optional Parameter, called Capabilities. The parameter lists the capabilities supported by the speaker.

A BGP speaker determines the capabilities supported by its peer by examining the list of capabilities present in the Capabilities Optional Parameter carried by the OPEN message that the speaker receives from the peer.

A BGP speaker that supports a particular capability may use this capability with its peer after the speaker determines (as described above) that the peer supports this capability.

A BGP speaker determines that its peer doesn't support capabilities advertisement, if in response to an OPEN message that carries the Capabilities Optional Parameter, the speaker receives a NOTIFICATION message with the Error Subcode set to Unsupported Optional Parameter. In this case the speaker SHOULD attempt to re-establish a BGP connection with the peer without sending to the peer the Capabilities Optional Parameter.

If a BGP speaker that supports a certain capability determines that its peer doesn't support this capability, the speaker MAY send a NOTIFICATION message to the peer, and terminate peering (see Section "Extensions to Error Handling" for more details). The Error Subcode in the message is set to Unsupported Capability. The message SHOULD contain the capability (capabilities) that causes the speaker to send the message. The decision to send the message and terminate peering is local to the speaker. If terminated, such peering SHOULD NOT be re-established automatically.

4. Capabilities Optional Parameter (Parameter Type 2):

This is an Optional Parameter that is used by a BGP speaker to convey to its BGP peer the list of capabilities supported by the speaker.

The parameter contains one or more triples <Capability Code, Capability Length, Capability Value>, where each triple is encoded as shown below:

```
+-----+
| Capability Code (1 octet) |
+-----+
| Capability Length (1 octet) |
+-----+
| Capability Value (variable) |
+-----+
```

The use and meaning of these fields are as follows:

Capability Code:

Capability Code is a one octet field that unambiguously identifies individual capabilities.

Capability Length:

Capability Length is a one octet field that contains the length of the Capability Value field in octets.

Capability Value:

Capability Value is a variable length field that is interpreted according to the value of the Capability Code field.

BGP speakers SHOULD NOT include more than one instance of a capability with the same Capability Code, Capability Length, and Capability Value. Note however, that processing of multiple instances of such capability does not require special handling, as additional instances do not change the meaning of announced capability.

BGP speakers MAY include more than one instance of a capability (as identified by the Capability Code) with non-zero Capability Length field, but with different Capability Value, and either the same or different Capability Length. Processing of these capability instances is specific to the Capability Code and MUST be described in the document introducing the new capability.

5. Extensions to Error Handling

This document defines new Error Subcode - Unsupported Capability. The value of this Subcode is 7. The Data field in the NOTIFICATION message SHOULD list the set of capabilities that cause the speaker to send the message. Each such capability is encoded the same way as it would be encoded in the OPEN message.

6. IANA Considerations

This document defines a Capability Optional Parameter along with an Capability Code field. IANA maintains the registry for Capability Code values. Capability Code value 0 is reserved. Capability Code values 1 through 63 are to be assigned by IANA using the "IETF Consensus" policy defined in RFC 2434. Capability Code values 64 through 127 are to be assigned by IANA, using the "First Come First Served" policy defined in RFC 2434. Capability Code values 128 through 255 are for "Private Use" as defined in RFC 2434.

7. Security Considerations

This extension to BGP does not change the underlying security issues inherent in the existing BGP [Heffernan].

8. Acknowledgements

The authors would like to thank members of the IDR Working Group for their review and comments.

9. Comparison with RFC 2842

In addition to several minor editorial changes, this document also clarifies how to handle multiple instances of the same capability.

10. References

- [BGP-4] Rekhter, Y. and T. Li, "A Border Gateway Protocol 4 (BGP-4)", RFC 1771, March 1995.
- [Heffernan] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

11. Authors' Addresses

Ravi Chandra
Redback Networks Inc.
350, Holger Way
San Jose, CA 95134

EMail: rchandra@redback.com

John G. Scudder
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134

EMail: jgs@cisco.com

12. Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

