

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: April 20, 2010

T. Scholl
ATT
J. Scudder
Juniper Networks
R. Steenbergen
Server Central / nLayer
D. Freedman
Claranet Limited
October 17, 2009

BGP Advisory Message
draft-ietf-idr-advisory-00.txt

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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 and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 20, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet-Draft

BGP Advisory Message

October 2009

Abstract

The BGP routing protocol is used with external as well as internal neighbors to propagate route advertisements. In the case of external BGP sessions, there is typically a demarcation of administrative responsibility between the two entities. Provisioning, maintenance and administrative actions are communicated via off-line methods such as email or telephone calls. While these methods have been used for many years, it can be troublesome for an operator to correlate a BGP-related event in the network with a notice that was transmitted in email.

This document proposes a new BGP message type, the Advisory message, which can be used to convey advisory information to a BGP speaker's peer. A capability is used to ensure that the recipient of the Advisory message is capable of supporting it.

1. Introduction

The BGP routing protocol is used with external as well as internal neighbors to propagate route advertisements. In the case of external BGP sessions, there is typically a demarcation of administrative responsibility between the two entities. While initial configuration and troubleshooting of these sessions is handled via offline means such as email or telephone calls, there is gap when it comes to advising a BGP neighbor of a behavior that is occurring or will occur momentarily. There is a need for operators to transmit a message to a BGP neighbor to notify them of a variety of types of messages. These messages typically would include those related to a planned or unplanned maintenance action. These advisory messages could then be interpreted by the remote party and either parsed via logging mechanisms or viewed by a human on the remote end via the CLI. This capability will improve operator NOC-to-NOC communication by providing a communications medium on an established and trusted BGP session between two autonomous systems.

The reason that this method is preferred for NOC-to-NOC communications is that other offline methods do fail for a variety of reasons. Emails to NOC aliases ahead of a planned maintenance may have ignored the mail or may have not recorded it properly within an internal tracking system. Even if the message was recorded properly, the staff that is on-duty at the time of the maintenance

event typically was not the same staff who received the maintenance notice several days prior. In addition, the staff on duty at the time of the event may not even be able to find the recorded event in their internal tracking systems. The end result is that during a planned event, some subset of eBGP peers will respond to a session/

peer down event with additional communications to the operator who is initiating the maintenance action. This can be via telephone or via email, but either way, it may result in a sizable amount of replies inquiring as to why the session is down. The result of this is that the NOC responsible for initiating the maintenance can be inundated with calls/emails from a variety of parties inquiring as to the status of the BGP session. The NOC initiating the maintenance may have to further inquire with engineering staff (if they are not already aware) to find out the extent of the maintenance and communicate this back to all of the NOCs calling for additional information. The above scenario outlines what is typical in a planned maintenance event. In an unplanned maintenance event (the need for and immediate router upgrade/reload), the number of calls and emails will dramatically increase as more parties are unaware of the event.

With the BGP advisory capability, an operator can transmit an advisory message just prior to initiating the maintenance specifying what event will happen, what ticket number this event is associated with and the expected duration of the event. This message would be received by BGP peers and stored in their router syslog as well as any monitoring system if they have this capability. Now, all of the BGP peers have immediate access to the information about this session, why it went down, what ticket number this is being tracked under and how long they should wait before assuming there is an actual problem. Even smaller networks without the network management capabilities to correlate BGP events and advisory messages would typically have an operator login to a router and examine the logs via the CLI.

There are several problems with e-mail only notifications:

- Up-to-date contact information is fairly difficult to maintain. Some networks who have very open peering policies may peer with up to 1,000 unique ASNs.

A NOC e-mail address does not always reach its way to the proper individuals at the NOC. A large amount of e-mail received at NOC aliases are typically spam or issues not appropriate for a typical NOC queue.

E-mail is not real time. In some environments, e-mail processing can be delayed and when looking at unplanned maintenances, some operators do not have the time to draft an e-mail as well as the distribution list.

There are several advantages to the advisory capability to operators:

Scholl, et al.

Expires April 20, 2010

[Page 3]

Internet-Draft

BGP Advisory Message

October 2009

There is no requirement for an external contact database. Contact databases are important, but this capability provides a way for an operator to transmit a message about a specific BGP session with no external contact information being required.

The very existence of the BGP session itself has inherent authentication and message routing properties. An operator immediately knows for every advisory message that it is coming from someone you are directly connected to (and thus have a relationship with) and which particular BGP session this is regarding. This is all completed without any additional human parsing required.

Because there is a BGP session that exists, an operator already has an authenticated session. There is no requirement for further authentication of the BGP session (key exchange).

The advisory message provides for real-time delivery of a message to a BGP neighbor. This will provide a rapid option in comparison to drafting an email to all BGP peers and waiting for the receipt before commencing with an unplanned maintenance event.

This draft aims to provide operators with the capability to transmit an advisory message to BGP peers to assist with daily network operations.

[1.1.](#) Requirements Language

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](#)].

[2.](#) Capability

A new BGP capability [[RFC5492](#)] called Advisory is introduced, with type code TBD. This capability indicates that the router advertising it is capable of receiving and parsing Advisory messages. The capability is of variable length. The data portion of the capability lists the Advisory message subtypes which are supported. The String subtype MUST be supported, which implies that the length MUST be at least 2 if the capability is advertised.

[3.](#) Advisory Message

The Advisory message is a BGP message of type TBD. It consists of a

BGP fixed header followed by a two-byte subtype and a data portion of variable length, calculated according to the Length field in the fixed header. The format of the data portion is dependent on the subtype. This document defines the following subtypes:

0 - Reserved:

MUST NOT be sent, MUST be ignored (other than optionally logging an error) on receipt.

1 - Advisory String:

A message comprised of a string of ASCII characters. The string's length is given by the length of the message, there is no null termination. Upon reception, the string SHOULD be reported to the router's administrator. The means of reporting the string are implementation-specific but could include methods such as syslog.

2 - Static String:

A message comprised of a string of ASCII characters. The

string's length is given by the length of the message, there is no null termination. Upon reception, the string SHOULD be stored in a BGP neighbor statistics field within the router. This string would then be accessible to the operator by executing CLI commands or any other remote method to obtain BGP neighbor statistics (NETCONF, SNMP). The expectation is that the last static message received from a BGP neighbor will be the message visible to the operator (the most current static message).

While this document mandates no particular events for which advisory messages should be generated, there are a variety of applications where the advisory message may be used. Implementations SHOULD provide its users the ability to transmit a free form text message generated by user input.

Implementations MAY choose to define a standard set of advisory messages that are automatically driven rather than requiring a human to enter specific reasons. These messages may be automatically transmitted based upon specific router functions such as a router reload, administrative action (neighbor shutdown) or reconfiguration (new BGP address-family support).

Implementations SHOULD provide router administrators with the ability to filter out specific BGP Advisory message types on a per neighbor or per peer-group basis. This interface should be provided to the

operator to clearly define if they want advisory, static or both types of messages.

Implementations MUST rate-limit the rate in which they transmit and receive advisory messages. Specifically, an implementation MUST NOT allow the handling of advisory messages to negatively impact any other functions on a router such as regular BGP message handling or other routing protocols.

As its name implies the Advisory message is intended to be used to advise a peer of some condition which may be of interest to that peer (or its administrator). It MUST NOT be used as a replacement for the Notification message in fatal error situations (i.e., situations where the integrity of the BGP peering is violated or suspect), although an Advisory message MAY precede a Notification message.

[4.](#) Error Handling

An Advisory message MUST NOT be sent to any peer which has not advertised the Advisory capability indicating support for the relevant subtype. If a router which has advertised the Advisory capability receives an Advisory message with a subtype for which it has not advertised support, it MUST accept and discard that message. It MAY locally log an error when this occurs.

[5.](#) IANA Considerations

IANA is requested to allocate a type code for the Advisory message from the BGP Message Types registry, to allocate a type code for the Advisory Capability from the Capability Codes registry, and to establish and maintain a registry for BGP Advisory message subtypes, to be allocated according to the First Come First Served policy defined in [[RFC5226](#)].

[6.](#) Security Considerations

No new security issues are introduced to the BGP protocol by this specification.

[7.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", [RFC 5492](#), February 2009.

Authors' Addresses

Tom Scholl
ATT

Email: tom.scholl@att.com

John Scudder
Juniper Networks

Email: jgs@juniper.net

Richard Steenbergen
Server Central / nLayer

Email: ras@e-gerbil.net

David Freedman
Claranet Limited

Email: david.freedman@uk.clara.net