

BGP Advisory Message Draft

draft-scholl-idr-advisory-00

Authors:

T. Scholl
AT&T

J. Scudder
Juniper

With valuable input from:

Richard Steenbergen
David Freedman

...and others

History

Proposed at IETF SF, read later slide on the response.

Has similarities to soft-notify and inform, both of which were never accepted. Probably should have been the first clue that this would be difficult...

Problem Statement

In the settlement-free peering world, communication between peers via traditional methods do not always work so well.

Peers (hopefully) use email to notify other networks of scheduled maintenance and other administrative notices such as a request to increase maximum-prefix thresholds.

However, sometimes email doesn't work so well.

Problem Statement (cont'd)

Basically, there needs to be a way to communicate advisory notices to BGP neighbors via the protocol that would appear in the neighbors router log.

Router logs traditionally use SYSLOG and are exported to network management and some automated ticketing systems.

This is not only restricted to the settlement-free peering world but also applies to the provider -> customer space as well.

Original Proposal

Give an operator the ability to transmit a text message to a BGP neighbor. This message would appear in the log of the neighbor.

The capability to send such a message will require a new BGP capability type.

This message would not be transitive.

Do you want other networks sending BGP messages to your customers offering them lower prices?

Do you want to see v1agr@ advertisements in your syslog?

BGP Twitter?

twitter



BB1_CHCGIL

```
Jun 2 15:41:25.486 UTC: %BGP-3-  
NOTIFICATION: sent to neighbor  
146.20.5.1 3/1 (update malformed) 0  
bytes FFFF FFFF FFFF FFFF FFFF  
FFFF FFFF
```

half a minute ago from web

Applications of this draft (With fake CLI examples)

- Notify BGP peers of an upcoming maintenance, specifying length of downtime, NOC ticket number and phone number.

```
Mar 11 09:21:35.118 UTC: %BGP-3-ADVISORY: received from neighbor 69.31.111.242  
(AS4436) We are up upgrading our router, should return at 0940Z. Ticket  
#1462016, 312-698-4800.
```

- Notify BGP peers if their advertisement count is approaching the configured maximum-prefix count

```
Feb 12 20:13:13 EX1.EQCHIL rpd[4618]: %DAEMON-4-RPD_BGP_NEIGHBOR_ADVISORY:  
BGP peer 206.223.119.18 (External AS 19151) You are approaching the maximum-  
prefix setting of 1,500. We are receiving 1,208 from you.
```

Applications of this draft

- Notify BGP peers of any problems noticed with a session

```
Mar  4 15:35:28.264 UTC: %BGP-3-ADVISORY: received from neighbor
206.223.119.43 (AS40009) - UPDATE message for prefix 146.20.0.0/16
contained invalid attribute $FOO value $BAR - Ignoring
```

- Notify BGP peers of administrative action

```
Feb 22 11:37:41  EX1.EQCHIL rpd[4618]: %DAEMON-4-RPD_BGP_NEIGHBOR_ADVISORY:
BGP peer 154.54.11.237 (External AS 174) The check was not in the mail
as you had indicated. Goodbye.
```

```
Feb 22 11:37:44  EX1.EQCHIL rpd[4618]: %DAEMON-4-RPD_BGP_NEIGHBOR_STATE_CHANGED:
BGP peer 154.54.11.237 (External AS 174) changed state from Established to Idle
(event RecvNotify)
```

```
Feb 22 11:37:44  EX1.EQCHIL rpd[4618]: %DAEMON-4: bgp_read_v4_message:8469:
NOTIFICATION received from 154.54.11.237 (External AS 174): code 6 (Cease)
subcode 3 (Peer Unconfigured)
```


Applications of this draft

- **Notify BGP peers of stated peering policy**

```
Mar  4 12:34:33.381 UTC: %BGP-3-ADVISORY: received from neighbor
206.223.119.97 (AS4436) - Welcome to peering with nLayer. Please adhere to our
peering policy at http://www.nlayer.net/network/peering Thanks.
```

- **Notify BGP peers of imminent reload (not operator initiated)**

```
Feb 12 14:22:34  EX1.EQCHIL rpd[4618]: %DAEMON-4-RPD_BGP_NEIGHBOR_ADVISORY:
BGP peer 206.223.119.107 (External AS 29791) Warning - disabling all BGP sessions
and executing reload due to software fault
```

Initiating messages to neighbors

- Operators should have the capability to generate messages to one, multiple or all BGP neighbors:

```
EX1.EQCHIL# bgp message ?
```

```
*           Send to all peers
```

```
A.B.C.D     BGP neighbor address to message
```

```
peer-group  Send to all neighbors in a peer-group
```

```
EX1.EQCHIL# bgp message peer-group PEERS We are going down for software upgrade.  
We should be back in 15 minutes.
```

```
Update sent!
```

```
EX1.EQCHIL#
```

- Reload commands should permit the insertion of BGP messages when shutting down:

```
tom@EX3.EQABVA> request system reboot bgp-message We are going down for a software  
upgrade.
```

Good ideas that should be added

- Provide several types of messages, which can be useful when parsing log output to classify:
 - Administrative
 - Operational
 - The above two would be structured message types.
 - Freeform
 - This is where you can specify a custom text message
- Allow the operator to pick which types of messages they want to accept and which ones to dump.
 - This way, an operator can ignore what they don't want to see.

Good ideas that should be added

- Concept of sticky messages
 - An idea where you can transmit an advisory message to “stick” in the BGP neighbor state. This way, upon viewing “show bgp nei x.x.x.x”, you can see the latest sticky message transmitted by the peer.
 - Would not be sent to the log.
 - Useful when you want to notify peers of some action and want this message associated with the session (until the next one comes and erases the old one).
- Message identifiers (multi-part messages).

The IETF SF response:

- Laughter, comparisons made to BGP SMS
- Why don't you use SYSLOG to send messages between routers
 - Who wants to open up UDP/514 to the world?
- If you don't trust open SYSLOG, use an IPSEC tunnel for SYSLOG
 - Getting people to deploy MD5 was hard enough, you think IPSEC SYSLOG will go over better?
 - If IPSEC is such a secure method, why aren't we using it for BGP already?
- Have your routers (the router itself) send emails to the peer
- You should know your peers better & custom messages are hard to parse and automate.
 - Sure, if you are a "tier 1" with < 20 peers, okay.
 - If you have thousands of peers, maybe this isn't for you.
 - But there is a huge amount of networks that fall between a traditional "tier 1" and a massive content provider with thousands of sessions.