



BGP Monitoring Protocol (BMP)

**John Scudder <jgs@juniper.net>
GROW WG, IETF-73, November 19, 2008**

Draft Information

- **Draft-scudder-bmp-01.txt**
- **Authors:**
 - John Scudder <jgs@juniper.net>
 - Rex Fernando <rex@juniper.net>
 - Stephen Stuart <sstuart@google.com>

What Is It?

- **A way for a monitoring station to get a complete dump of the routes received from a peer or peers (including all peers)**
- **And to get ongoing updates about all routes received from that peer or peers**
- **Currently, this information is only available by some variation on “show route” and parsing ASCII (“screen scraping”)**
- **Useful for looking glasses, network analysis, etc.**

History

- **Draft-scudder-bmp-00 submitted in 2005**
 - Turned out to be too hard to implement (despite best intentions!)
 - Set aside
- **Renewed interest recently**
 - New approach, provides similar benefits in an implementable fashion
 - Implementation has been done (lab only at present)

Why Not Use Plain BGP?

- **BGP only provides *best* paths**
- **BMP provides *all* paths**

Overview of Operation

- **Router configured with management station identity, list of peers of interest (could be all peers)**
- **Connects to management station, sends initial dump of all routes for those peers**
 - Formatted as BGP UPDATE messages wrapped in a BMP header
- **As peers advertise/withdraw routes, sends additional updates to management station**

Information Provided

- **In addition to usual BGP UPDATE information, BMP header has**
 - Peer identity (address, BGP Identifier, RD if applicable, AS number)
 - Time stamp (when route or route withdrawal was received, microsecond granularity)

Statistics Reporting

- **Pro-actively report stats of interest**
 - Threshold or timer driven
 - Optional
- **Defined counters:**
 - Prefixes rejected by inbound policy
 - Duplicate prefix advertisements
 - Duplicate withdraws
 - Updates invalidated due to CLUSTER_LIST loop
 - Updates invalidated due to AS_PATH loop
- **Stats message is extensible (TLV) to add new counters**

Peer Notification

- **Notification message sent when peering session goes down**
- **Includes BGP NOTIFICATION data, if any**

Characteristics

- **BMP messages are not bit-for-bit clones of received UPDATES**
 - Messages are regenerated according to usual BGP UPDATE generation logic
 - However, data is taken from Adj-RIB-In, not Loc-RIB
- **Implications**
 - Not every received UPDATE will necessarily result in a BMP message being sent
 - During busy times, some UPDATES might be suppressed if obsoleted by newer UPDATES
 - However, BMP messages will generally be the same as or very close to received UPDATES
- **BMP *will* converge to the correct set of routes**

Summary

- **Allows a management station to track routes received from one or more peers**
 - Even routes which are not “best”
- **Updates not “cloned” but regenerated**
- **Also provides some convenience counters**
- **Not suitable for use as a routing protocol**
- **Implementation works in lab**

Next Steps

- **Re-adopt as GROW WG Item?**

Juniper *your* Net™