

Revised Error Handling for BGP Updates from External Neighbors

draft-chen-ebgp-error-handling-00.txt

Enke Chen, Pradosh Mohapatra, Keyur Patel

IETF 79, November 2010, Beijing, China

Problem

- RFC4271 specifies that a malformed attribute resets the session
- Relaxed for optional transitive attributes
- There are other cases not covered such as
 - AS-PATH
 - Length of an optional transitive attribute
- Session rest is disruptive
 - No distinction between good / bad prefixes
 - Multiple networks / ASes

Goals

- Minimize the scope of networks being impacted
- Limit the impact to only the routes involved
- Maintain the protocol correctness

Proposal

- Partially revise error handling for updates from EBGP
- Do not reset an EBGP session if all possible
- Treat as withdraw
 - As long as the prefixes can be parsed
 - Safe, like applying a filter
 - Debugging facility must be provided
- About session integrity
 - NLRI field must be parsed completely
 - Validation of the BGP message header (synchronization header, message length and type)

Parsing of the NLRI Fields

- It is fundamental to the proposed approach
- Format the prefixes in the very front of the message
 - MP-REACH / MP-UNREACH as the first attribute
 - Use MP-REACH / MP-UNREACH for IPv4 Unicast

What about IBGP

- In general “treat-as-withdraw” is unsafe for IBGP
- Error handling (i.e., session reset) remains unchanged
- Once prefixes are identified, apply filters at the ingress
- Not pursuing automation of the procedures in protocol
 - Issues from EBGP typically dominate
 - Require significant changes to the protocol