

Upcoming BGP Changes in response to Operator requests

IDR-chairs

IETF 82

Types of Errors

- Do you desire to change basic Error processing?
 - Question: Poor service or No service
 - RFC 4271: No service on error: “Drop” if error”
 - New: Poor service rather than “No Service”
 - Keep going with Error in Ribs
 - Add Monitoring, Route Refresh of peers , Add 1-Time RR filters
- IETF Operational Docs
 - Requirements for BGP error handling
Grow-ops-reqs-bgp-error-handling-01

Problems with Poor instead of “No Service”


- BGP
 - Junk can hang on
 - Clear reset and flush avoids long-term errors
- Restoring health
 - Limit to affected NLRIs
 - Limit to RIB
 - If nothing else, session reset

Errors with “Treat as withdraw”

- e-bgp – two cases
 - Prefix in Adj-Rib-In of neighbor is use in Global RIB that forwards
 - Not in use
 - Potential issues with e-bgp treat as withdrawn are: Inconsistent BGP tables, L3 VPNs (BGP), L2 VPN endpoints, BGP MAC VPNS
- i-BGP case
 - Use of “treat-as-withdraw” may cause routing loops
 - Operators state: Better than AFI shutdown

How to Recovering RIBs

- Ways
 - Refresh RIB in IBGP/EBGP
 - Partial one-time-filters
 - Zeng-one-time-prefix-orf
 - Enhanced RR
 - Timer based GR with no DROP
 - EOR flags
 - Thread RIB by “this was withdraw/
AS” but I kept it



Existing
and ready

New &
complex

Call to Operators

- Final Call for operators for input on requirements for Error/Monitoring
- IDR will open work on Error / Monitoring specification
 - Grow-ops-reqs-bgp-error-handling-01
 - Draft-ietf-grow-bmp-01
- Catch your window for change now

Future

- IDR will take on Graceful Restart / Graceful Shutdown next
- Requirements
 - Requirements for Graceful Shutdown [RFC 6198]
 - MED considerations [RFC 4451]

Q&A

Operational Toolset

- Message to replace notification
 - I got broken prefixes [prefix, number, what?]
 - I over my broken prefix limit – hit me with a RR
 - It's so broken I'm going to drop AFI
- Diagnostic
 - Razuk-bgp-diagnostic
 - Idr-advisory
 - Grow-BMP
- Comments on Queueing
 - Are we requiring separate queue for messages
 - Do we augment keepalive
 - In-band/snmp/other

grow-ops-reqs-for-bgp-error-handling

- Error handling
 - Syntax check: error length check
 - Valid parsing NLRIs
 - NLRI, MP-REACH, MP-UNREACH OK
 - Validate Path Attributes
 - Bad Flags
 - Opt/Transitive Broken
 - Use Neighbor Complete Flags
 - Well Known broken

Need picture here