

Operational Requirements for BGP Error Handling

draft-ietf-grow-ops-reqs-for-bgp-error-handling-03

IETF 83.
Paris, France.
March 2012.

Rob Shakir (BT)

Intention of the Draft.

- Numerous outages and incidents in the Internet DFZ and private networks based on BGP error handling behaviour.
- Aiming to document what the role of BGP is in SP networks and motivate change to how error handling is considered.
- Draft splits the requirements into a number of discrete sets:

Addressing underlying cause.

Reducing impact or recovering from particular failure modes.

Requirements for monitoring BGP.

Applicability/Scaling of Recovery Mechanisms.

What has changed since Québec?

2. Errors within BGP-4 UPDATE Messages	6
2.1. Classifying BGP Errors and Expected Error Handling	7
2.1.1. Critical BGP Errors	7
2.1.2. Semantic BGP Errors	8

- Addressing WG and GROW chairs' suggestion that the draft more clearly defines what type of BGP errors there are, and how they be handled.

7. Operational Complexities Introduced by Altering RFC4271	19
7.1. Reducing the Network Impact of Session Teardown	21

- Discuss applicability of BGP Graceful Shutdown if and when a session is moved to a "hold down" state.

Author's Address

Rob Shakir
 Cable&Wireless Worldwide

Author's Address

Rob Shakir
 BT

- Administrative changes in -03.

Categories of BGP Error.

- Previous comment that draft was not clear enough about what type of errors are “recoverable”.
- Detailed analysis with Alton Lo to work through error conditions identified in parsing engine, to identify where ‘targeted’ error handling means are applicable.

“Critical” Errors.

Error whereby NLRI attribute cannot be reliably extracted from the UPDATE message.

“Semantic” Errors.

Case where NLRI can be extracted – error located in particular path attribute.

- Nomenclature aimed to provide some form of classification of error types.

IDENTIFYING ERROR TYPES.

Critical Errors.

- Only where we cannot identify NLRI contents => avoid session reset wherever possible.
 - Overall UPDATE message length errors.
 - Errors parsing MP_(REACH|UNREACH)_NLRI or IPv4-Unicast Advertised or Withdrawn.

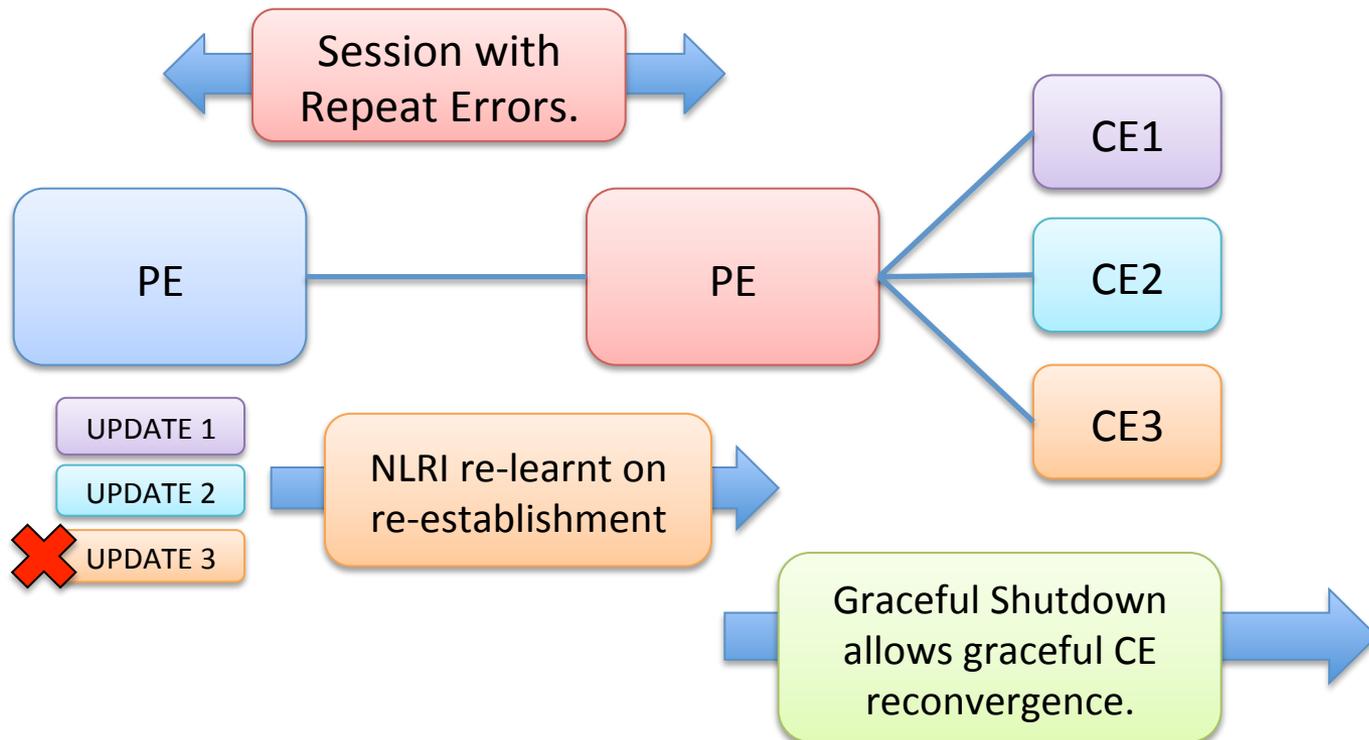
Semantic Errors.

- Errors within particular attributes – some cannot be defined (based on logic of contents).
 - Length errors localised to non-NLRI attributes.
 - Invalid data and/or flags for non-NLRI attributes.
 - Missing mandatory, unrecognised non-optional, duplicate or invalid attributes.
 - Next-Hop errors (wrong SAFI, len(o), missing...)

GRACEFUL SHUTDOWN APPLICABILITY.

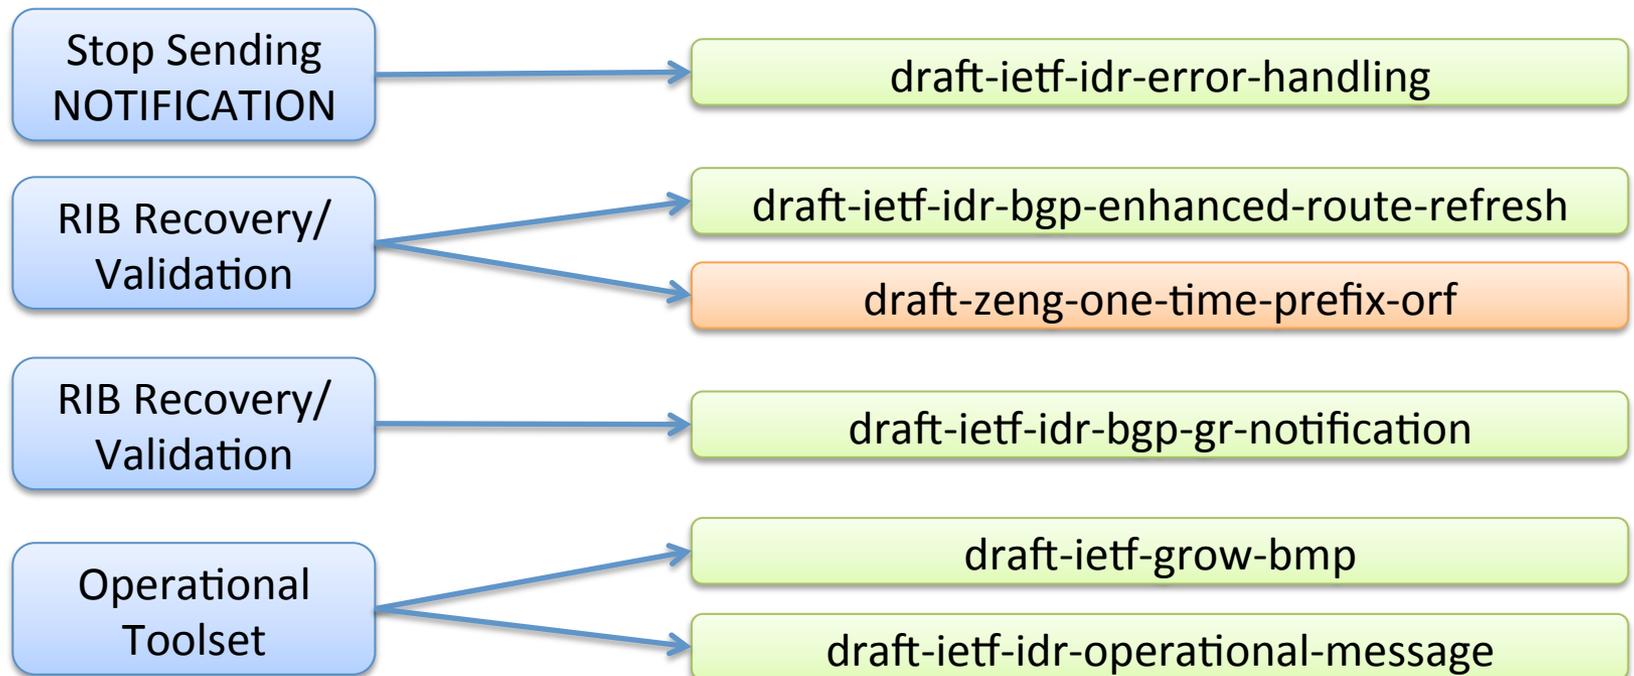
Repeated Error Conditions.

- Recommendation made that where sessions continually reset – enter a “hold down” state.
- Bruno Decraene suggested that in this final shutdown – graceful-shutdown is applied.



Progress.

- Number of work items – and expansion of error handling scope in IDR to address these requirements.
- IDR chairs called for further input to this draft in Taipei.



GOING FORWARD.

WGLC.

- Requirements have stabilised, and limited change since IETF82.
- Feels like consensus that the requirements here document the original problem space.
- Would therefore ask for chairs to kick-off WGLC.

Further Feedback/Comments?

- Please let me know – aim to address these along with those coming from WGLC output.