



Enhancing BGP

IETF IDR Current Drafts

Rob Shakir (RJS-RIPE)

UKNOF16 - April 2010

Aim

- BGP is a key protocol for all operators.
- But it's not perfect!
- Protocols must evolve to support network development.
- IETF Inter-Domain Routing WG “looks after” BGP.
- A look at the drafts that are in IDR currently.
- Encouragement to get involved!

IETF IDR?

- Working group at the IETF.
 - Operators, vendors, implementors...
 - Open to contributions.
- Ideas put into Internet-Drafts.
 - Drafts considered, refined, and adopted by WG.
 - Once consensus reached, proposed as RFC.

Capabilities (RFC5492)

- BGP originally carried only IPv4 routing information.
- Requirement to extend the protocol to carry additional information.
- Widely used - IPv6, VPNv4, VPNv6, VPLS etc.
- Capabilities are used to indicate supported information.
- MP-BGP is a capability - SAFI/AFI codes.
- Other features - AS4, G-R etc.

draft-ietf-idr-dynamic-cap

- Capabilities currently are established in OPEN.
- At the start of the session only!
- Aim to remove the requirement to bounce a session to add functionality.
- Advertised via a capability at session OPEN.
- Updates done by CAPABILITY message.
- Some consideration for error handling.

Operational Example



Current:



With Draft:



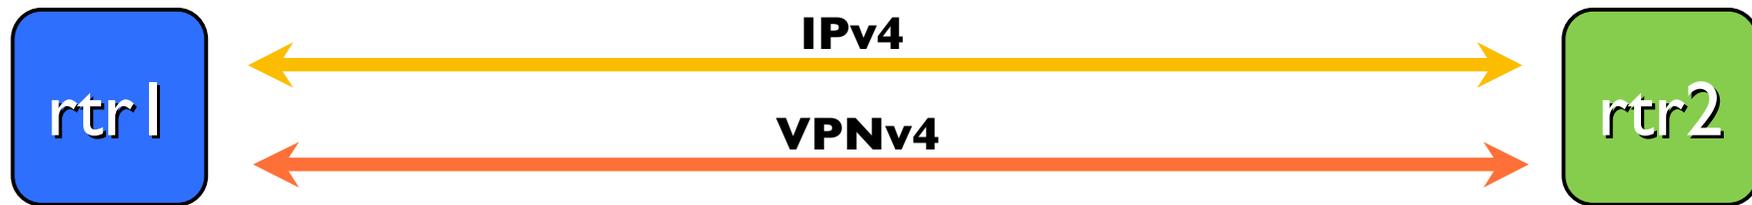
Current WG Position on Dyn-Cap

- Adopted as an IDR Draft - but still in discussion.
- Pros:
 - Allows BGP to grow, and functionality to increase without operational impact.
- Cons:
 - Complex implementation?
 - Solving a non-existent problem?

Multi-Session BGP

- Currently BGP uses a single TCP session to carry routing information.
- Introduce support for multiple transport sessions.
- Brings session robustness.
- NOTIFICATION does not affect all AFI/SAFI.
- Process distribution.
- Single 'bgpd' process not required.

Benefit of Multi-Session (I)



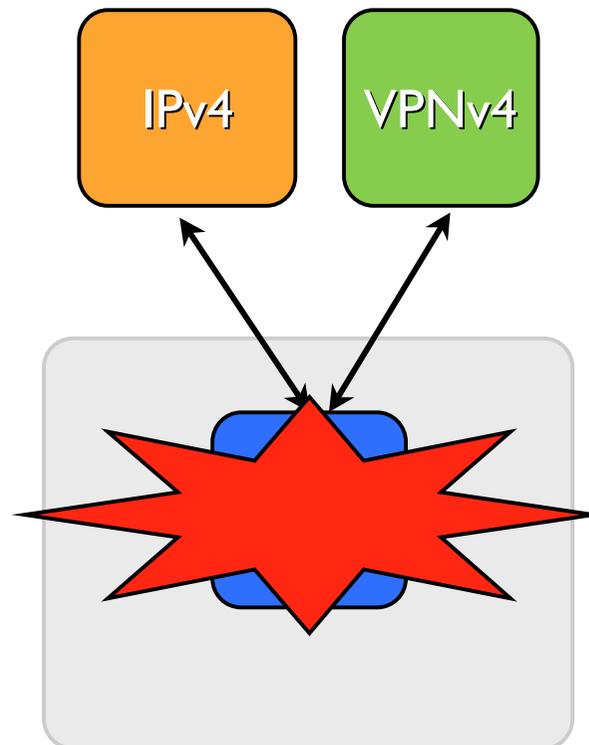
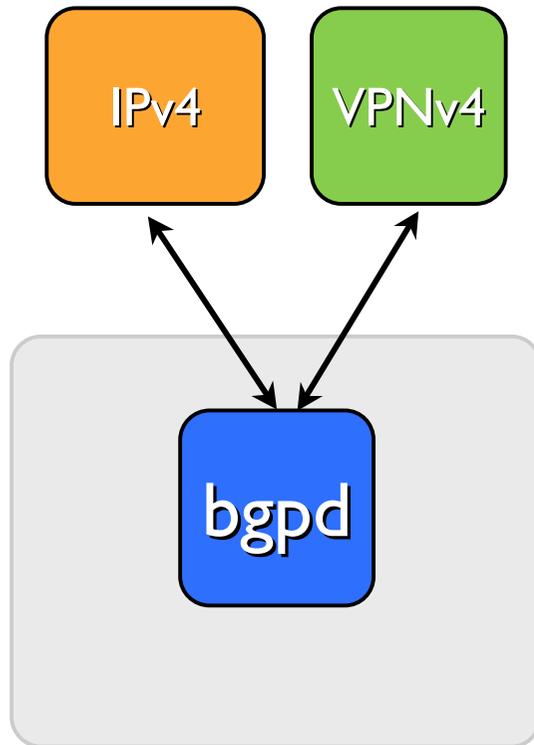
Error occurs in IPv4 UPDATE requiring NOTIFICATION.



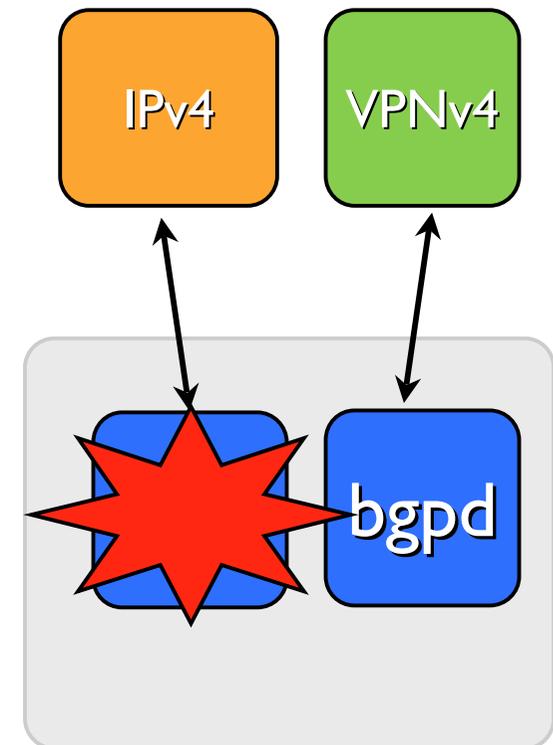
Only IPv4 services affected.

Where (e)iBGP - this can mean commercial implications of DFZ anomaly is much reduced.

Benefits of Multi-Session (2)



High utilisation
cannot be
managed.

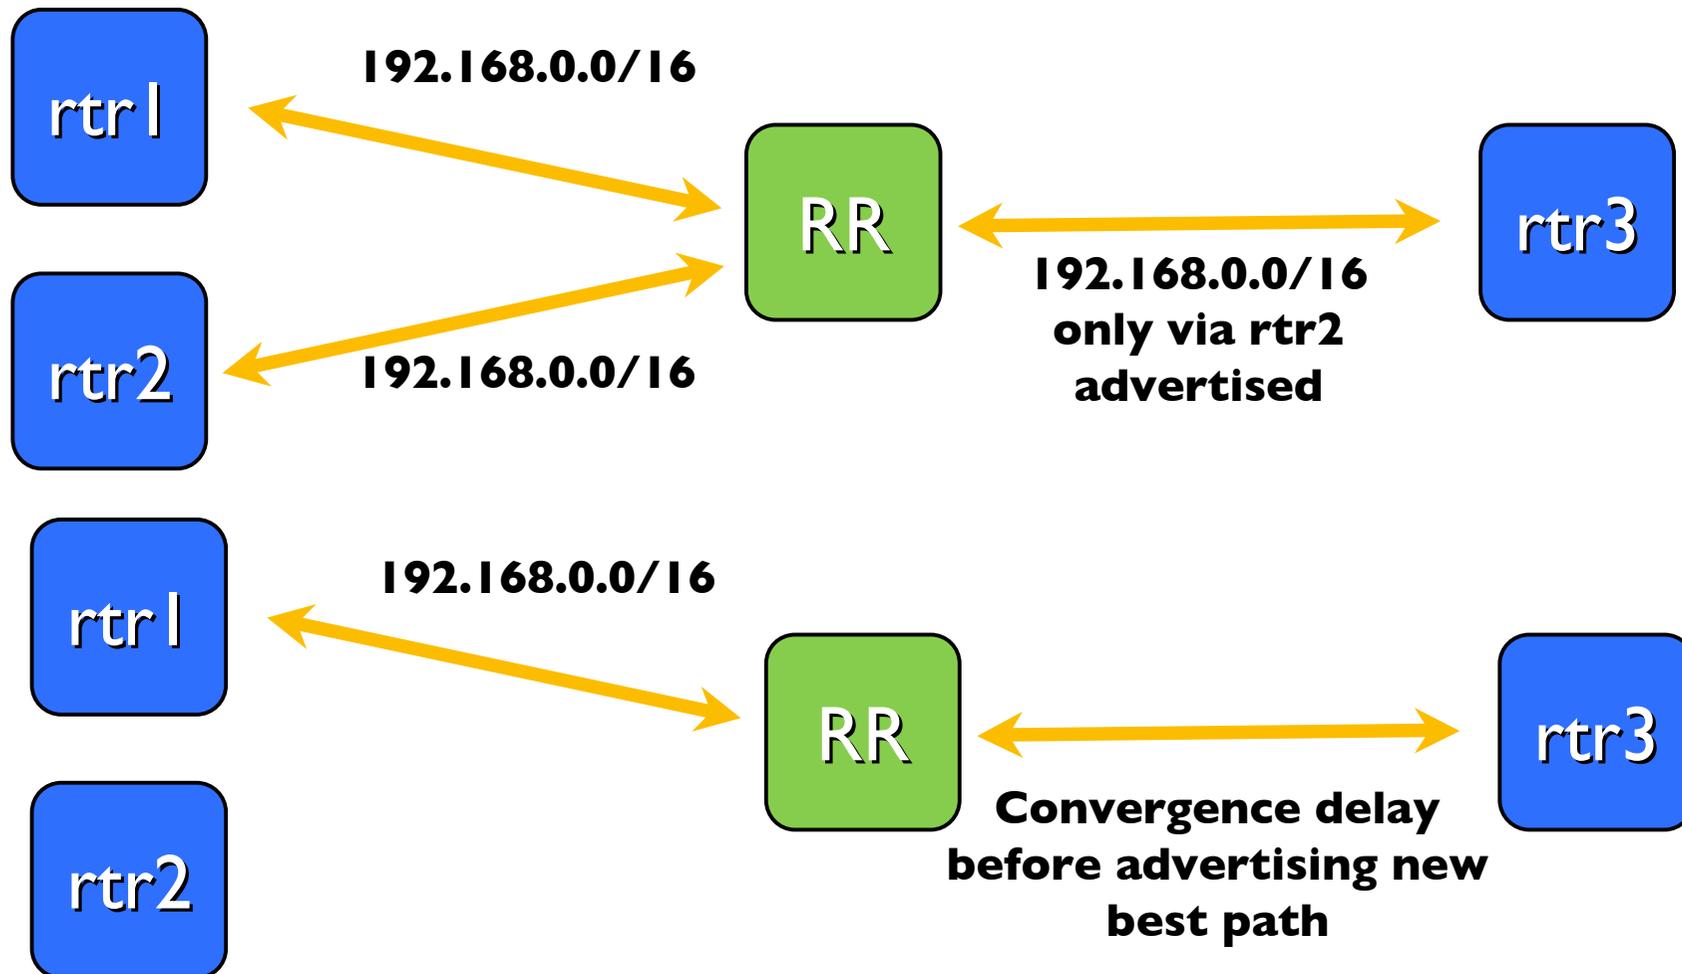


Process scheduling
handles run-away
processes.

draft-ietf-idr-bgp-multisession

- John Scudder's original draft.
 - Relatively mature, supported in some software.
- Not the only approach, introduces new BGP ports.
 - “New local TCP port for each new connection”
- Affect on CoPP - operational complexity.
 - draft-varlashkin-idr-multisession-same-port
 - Ilya Varlashkin's draft - multiplex onto port 179.

BGP Advertisement/Convergence



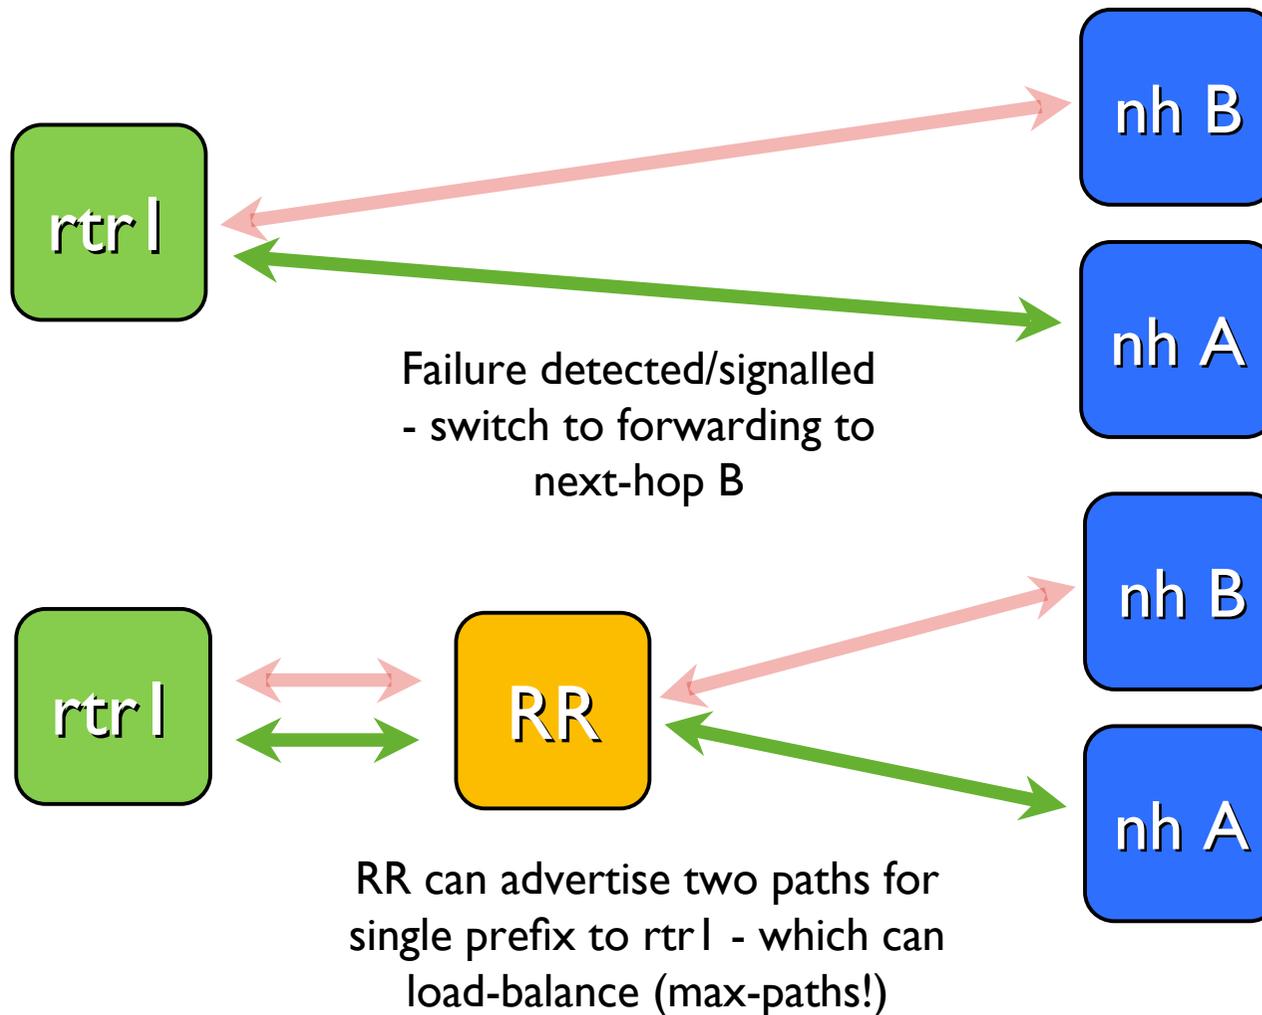
Some Implications of Single Path

- Forwarding implications:
 - One path per prefix, UPDATE results in implicit withdraw.
 - Means that we must wait for UPDATE to re-converge to new path.
 - In the interim, forwarding outage.
- Load-Balancing:
 - Only one path learnt from upstream.

draft-ietf-idr-add-paths

- Add a new “Path Identifier” to each prefix.
 - Rather like RD in RFC4364 VPNs (2547bis).
- Adjust implicit withdraw behaviour using Path ID.
- Some resource complexities.
 - Longer best-path calculations (more candidates?)
 - Control plane resources.

Add-Paths Usage



Error Handling

- Presented about this before!
- draft-ietf-optional-transitive (Errors in Opt Transitive attributes).
 - Address 'tunnelled' errors being sent NOTIFICATION.
- draft-ietf-idr-advisory (BGP Advisory capability).
 - Provide in-band signalling mechanism for ops.

Other Drafts

- BGP MIB v2
 - draft-ietf-idr-bgp4-mibv2-10 (10th version!)
 - IPv6 support, operationally useful counters
- Extended Communities for 32-bit ASNs
 - draft-ietf-idr-as4octet-extcomm-generic-subtype
 - Support <as4>:<community>
 - Increasing ops concern during ASN32 growth?

Some Encouragement

- Lots of vendors in IETF IDR.
- Operator's view is always interesting.
- After all, we are the ones using the features.
- It's really not hard to get involved!
 - Relatively low traffic mailing list(s).
 - Interesting discussions - good way to put pressure on vendors to implement features.

Questions, Comments, Corrections?

Many thanks to:

Steve Colam (GX Networks)

Tom Hodgson (Datahop)

Ben White (Timico)

Sven Huster (Huster Networks)

Questions, or comments later?

rjs@eng.gxn.net or rjs@rob.sh

RJS-RIPE

Public Comments?

IETF IDR - idr@ietf.org

(To Subscribe: idr-request@ietf.org, In Body: subscribe idr-post)

List Archive - <http://www.ietf.org/pipermail/idr/>