Graceful/Hitless restart of Routing Protocols

Prasad Babu

pbabu@juniper.net

Juniper™ NETWORKS

Why does control plane restart ?

- Due to control plane software upgrades
  - aka "planned restart"
- Due to control plane software bugs
  - aka "unplanned restart"

Both exist in the real life -> need a solution that handles both !!!

Impact of control plane restart - current situation

- Disrupt services by disrupting data path used by the services on the restarting node
- Control plane restart on PE is especially disruptive
  - As all the VPN sites that have connectivity just to that PE loose connectivity to other VPN sites
  - Disruption lasts as long as it takes for the PE to restart and to reacquire all the routing information
    - Both from other routers within the service provider network (both PEs and Ps) as well as from the directly connected CEs

Impact of control plane restart - current situation (cont.)

- Also disrupt services due to transient forwarding loops that could happen during routing convergence in response to control plane restart
  - Disruption happens twice: once when the control plane goes down, and once when the control plane comes back
  - Disruption involves multiple nodes, not just the node whose control plane restarts
    - The scope of the affected nodes is hard to predict (other than providing the worst case scenario)
  - Disruption lasts for as long as it takes routing to converge
    - The time it takes routing to converge is hard to predict
      - Because the scope of the nodes that have to converge is hard to predict (other than providing the worst case scenario)

⇨ The scope and the duration of the disruption is hard to predict (other than providing the worst case scenario)

Impact of control plane restart - current situation (cont.)

- Increases the load on the control plane
  - Involves multiple nodes, not just the node whose control plane restarts
    - The scope of the affected nodes is hard to predict (other than providing the worst case scenario)
  - Adversely impacts the scalability of the control plane
  - Adversely impacts the adaptability and convergence of the control plane

Bottom line: control plane restart adversely impacts service availability !!!

Graceful Restart - objectives:

- Improve service availability by minimize disruption of services (e.g., 2547 VPNs, L2 VPNs, VPLS, Internet, etc…) due to the control plane restarts
- Restart could happen anywhere in the network that delivers the service
  - Either at the edge (PE), or in the middle (P)
- Handle either planned (e.g., control plane software upgrade), or unplanned (control plane crash) restart

Graceful Restart - How ?

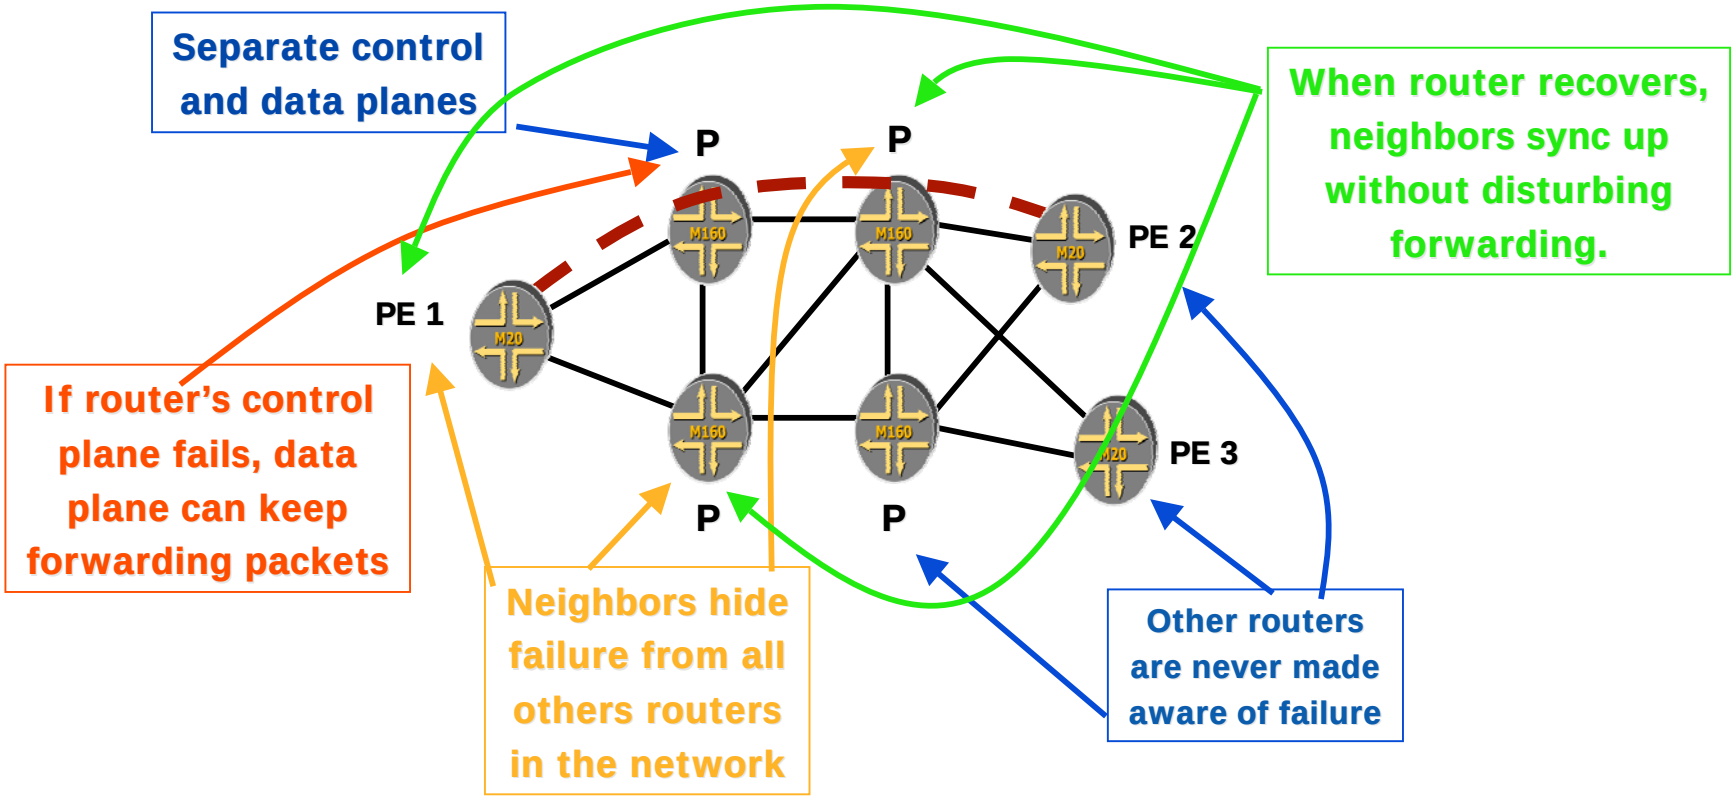

Graceful Restart - How ? (cont.)

- On the restarting node separate control component from forwarding component:
  - e.g., RE – control component (control plane)
  - e.g., PFE – forwarding component (forwarding state)
- On the restarting node preserve the forwarding state (forwarding component) across the restart of the control plane (control component)
- Localize the knowledge that the node's control plane restarts to only the routing peers of the restarting node
- On the routing peers of the restarting node preserve routing information associated/received from the restarting node across the restart of the control plane of the restarting node
- Restarting node (re)learns its routing information from its routing peers

Graceful Restart - How ? (cont.)

- Graceful restart mechanisms are protocol specific:
  - BGP – see draft-ietf-idr-restart-10.txt
  - ISIS – RFC 3847
  - OSPF – RFC 3623
  - LDP – RFC 3478
  - BGP/MPLS – see draft-ietf-mpls-bgp-mpls-restart-05.txt
  - RSVP - draft-ietf-mpls-generalized-rsvp-te-09.txt
  - RIP – already build in !!!
- No preservation of any of the protocol-related state across the restart on the restarting node
  - For all of the above protocols !!!

Graceful Restart – results:

- No disruption in the data path on the restarting node
  - Due to restarting node preserving its forwarding component
- No disruption in the data path on the routing peers of the restarting node
  - Due to the routing peers of the restarting node preserving routing information associated/received from the restarting node across the restart of the control plane of the restarting node
    - Implies that the routing peers don't modify their forwarding state in response to the restart of the control plane of the restarting node
- No disruption of the data path elsewhere
  - Due to nodes other than the routing peers of the restarting node being unaware of the restart of the control plane of the restarting node

Graceful restart – results (cont.):

- No change in the traffic pattern
  - preserves steady state traffic pattern
  - no impact on jitter, latency, packet ordering, route optimality
- Improved control plane scalability
  - By limiting the scope of the nodes that are aware of the restart to only the routing peers of the restarting node
- Improved control plane convergence/adaptability
  - By limiting the scope of the nodes that are aware of the restart to only the routing peers of the restarting node

Thank You

www.Juniper.net