

Distance Vector and Path Vector Notvia

RTGWG IETF69 July 2007

Stewart Bryant & Mike Shand
{stbryant, mshand}@cisco.com


The IPFRR story so far...

- Work on IPFRR has focused on link state routing protocols.
- The working assumption has been that notvia needed complete topology information to calculate the repair topologies.
- The aim here is to clarify a small addendum to the notvia draft that explains how NV can operate in a distance vector or path vector environment.


What about non-link state protocols?

- We make the observation that non-link state protocols (e.g. RIP, BGP) perform a distributed computation of the optimum path and that -
- By selectively filtering routing information for specific “not-via” addresses, such that a not-via address is never advertised across the failure to which it relates a distributed computation of the optimum not-via path can be achieved.

Normal DV




Not-via DV



Filtering

- For link protection the not-via advertisement can be filtered either by the source B, or the recipient A.
- For node-protection ALL not-via advertisements referring to the protected node must be prevented from forming a path through it.
- Simplest to filter at the protected node.

Node protecting Not-via DV




Additional information for node protection

- For link protection the next hop, and hence required not-via address, is obvious
- For node protection we need next-next-hop information.
 - i.e. which of the protected node's neighbors should we repair the traffic to?
- Various possible solutions, e.g....
 - draft shen-nnhop
 - BGP AS path list

Step-wise deployment

- Not-via advertisement is just another address
- Only protecting routers need to understand special semantics
- Correct not-via routes will automatically be computed by legacy routers.

BGP example



BGP

- Various possibilities for
 - Inter-AS Link protection
 - Border Node protection
 - Etc.
- All using DV not-via propagation techniques

Why is BGP FRR interesting?

- BGP convergence is much slower than the link state IGPs that are used in the interior of the AS's that form the Internet.
- BGP FRR may therefore provide a real benefit to users accessing material outside their own AS.
- By getting clusters of AS's to provide mutual cover against local failure to provide BGP connectivity we may be able to reduce the churn and hence improve the stability of the Internet.

Questions?