



BGP Flap Damping

Where to now?

Philip Smith
RIPE50, Stockholm



History

- Early Internet was susceptible to “routing storms”
 - Repeated withdrawal and re-announcement of /24 address blocks
 - Consumed significant CPU on early routers
 - Caused instability in the Internet
- “Flap damping” proposed to mitigate the effects of this instability



History

- Route flap damping was introduced to BGP4
 - RFC2439 describes the algorithm and conditions flap damping is applied under
- Requirements:
 - Fast convergence for normal route changes
 - Suppress oscillating routes
 - Announce stable routes
 - History predicts the future



Issues

- Implementations are highly configurable
- No prior operational experience of the optimum configuration
- Operational experience showed that vendor defaults seemed too aggressive for the operational Internet
 - A couple of prefix flaps resulted in disconnectivity in the order of tens of minutes
 - BGP reset or router restart had severe implications for ISPs in the emerging commercial Internet



Solutions

- RIPE 178 documented the problems and proposed acceptable route flap damping configuration parameters
- Updated by RIPE 210 to include “Golden Networks”
 - The address blocks of the 13 Root Servers
- Further updated by RIPE 229
 - Added website and more configuration examples



New Problems

- We all thought RIPE 229 would solve the problems
- It has not



Research work examples:

- "Route Flap Damping Exacerbates Internet Routing Convergence"
 - Zhuoqing Morley Mao, Ramesh Govindan, George Varghese & Randy H. Katz, August 2002
- "What is the sound of one route flapping?"
 - Tim Griffin, June 2002
- Various work on routing convergence by Craig Labovitz and Abha Ahuja a few years ago
- "Happy Packets"
 - Closely related work by Randy Bush *et al*



Morley Mao *et al*

- Route changes caused by path exploration increments the flap penalty
 - e.g. implementations penalise attribute changes
 - Best path lost → next best path chosen → neighbouring AS sees this as AS_PATH attribute change → penalty incremented
- Natural reaction is to not penalise non-decreasing path changes
 - But this is not immune to local provider policies
- Proposed *selective route flap damping*



Morley Mao *et al*

Selective Route Flap Damping

- Requires sender of route to include (relative) preference of route compared with previous announcement
 - Encoded as a BGP community?
- BGP keeps two bits to store comparative value of last two announcements received
 - 00 – fewer than two routes received
 - 01 – preference values of the route routes the same
 - 10 – latest route is higher preference than previous
 - 11 – latest route is lower preference than previous
- Comparison bits recomputed on fresh announcement
 - New value compared with old value
 - Change in value \Rightarrow route flap
- Simulation results highly successful



What next?

- Should RIPE 229 be declared obsolete? Or modified?
- Is flap damping bad for your network?
 - Do we need flap damping any more?
- Needed at Internet edge?
 - i.e. ISPs who are not providing transit to any other ASNs
- Needed in the Internet core?
 - Transit providers



What next?

- Proposal to reopen Route Flap Damping recommendations as a Routing WG work item
- Aim: New route flap damping recommendations for ISPs