

# RIPE Routing Working Group

## Recommendations on Route Flap Damping

---

### Authors:

Randy Bush  
Cristel Pelsser  
Internet Initiative Japan

Mirjam Kuhne  
RIPE NCC

Olaf Maennel  
University of  
Loughborough

Pradosh Mohapatra  
Keyur Patel  
Cisco Systems

Rob Evans  
Janet

**Document ID:** ripe-580

**Updates:** ripe-378

**Date:** January 2013

---

[Introduction](#)

[History and Background](#)

[Analysis](#)

[Recommendations](#)

[References and Further Reading](#)

### Introduction

Route Flap Damping (RFD) [1] is a mechanism for BGP speaking routers that penalises prefixes that exhibit a large number of updates ('flapping'), and suppresses a route when the accumulated penalty exceeds a given threshold. The penalty decays over time until it reaches a lower threshold at which point the route is unsuppressed. RFD is intended to improve the overall stability of the Internet routing table and reduce the load on BGP speaking routers. In ripe-378 [2] it was stated that due to the dynamics of BGP, especially a phenomenon called 'path hunting,' the default configurations of flap damping can do more harm than good as it may suppress a prefix after it has only flapped a few times. Consequently RFD was deprecated due to the problem of over damping (see [2] for more details).

A small number of prefixes on the Internet continue to flap rapidly and cause a disproportionate number of updates to BGP and load on BGP speaking routers. This document uses experimental data gathered from an operational environment to suggest changes to the RFD parameters to suppress the prefixes that flap the most, while minimising the suppression of other prefixes.

This document suggests parameters which would make RFD usable and is based around the work of Cristel Pelsser, Olaf Maennel, Pradosh Mohapatra, Randy Bush, and Keyur Patel presented at PAM2011[3].

## History and Background

In the early 1990s the accelerating growth in the number of prefixes being announced to the Internet (often due to inadequate prefix aggregation), the denser meshing through multiple inter-provider paths, and increased instabilities started to cause significant impact on the performance and efficiency of some Internet backbone routers. Every time a routing prefix altered state because of a single line-flap, the withdrawal was advertised to the whole BGP-Speaking Zone (BSZ) and handled by every router that carried the full Internet routing table.

The load this processing placed on the control planes of routers caused further instability as the routers were not able to process other BGP updates or they dropped traffic transiting the device. This could produce cyclic crashing behaviour.

To overcome this situation RFD was developed in 1993 and has since been integrated into most router BGP software implementations. RFD is described in detail in RFC 2439[1].

When RFD was first implemented in commercial routers, vendor implementations had different default values and different characteristics. As this inconsistency would result in different rates of flap damping, and therefore introduce inconsistent path selection and behavior that was hard to diagnose, the operator community introduced a consistent set of recommendations for flap damping parameters, so that ISPs deploying RFD would treat flapping prefixes in the same way.

This call for consistency resulted in the RIPE Routing Working Group producing first ripe-178, then ripe-210, and finally the ripe-229 documents [2a]. The parameters documented in ripe-229 were considered, at time of publication in 2001, the best current practice. In 2006, this was reviewed again and resulted in ripe-378 [2] which recommended to disable RFD because it created more harm than good.

## Analysis

In the work by Pelsser et al [3], it is shown that 3% of all prefixes cause 36% of BGP updates, and just 0.01% of the prefixes cause 10% of the BGP updates. The aim is to only penalise those prefixes with excessive numbers of updates.

The default values used in current implementations of RFD apply a penalty of 1000 each time a route flaps, and suppresses the prefix when the penalty exceeds a figure in the region of 2000 (Cisco IOS) or 3000 (Juniper JunOS).

The table shows the percentage of prefixes above the suppress threshold and the percentage reduction in BGP churn for various values of suppress threshold. The current default suppress value of 2000 reduces BGP churn by 47%, but it suppressed 14% of the prefixes at some point over the lifetime of the experiment. Significantly larger values of suppress threshold such as 12000, 15000 or 18000 still reduced BGP churn, but suppressed far fewer prefixes which it is believed reduces the risk of penalising otherwise well-behaved prefixes.

<b>Suppress Threshold</b>	<b>% prefixes suppressed</b>	<b>% reduction in BGP churn compared with no damping</b>
2000	14	47
4000	4.2	26
6000	2.1	19
12000	0.63	11.26
15000	0.44	9.51
18000	0.32	8.12

### **Recommendations**

In order to punish the biggest offenders - those prefixes that flap the most – yet without punishing others, the RIPE Routing-WG recommends vendors raise the maximum suppress threshold in router implementations to 50,000 and operators configure a suppress threshold value of at least 6,000. The vendors might also change the default suppress threshold to 6,000. But this might surprise operators who use the default.

This has a number of advantages:

- it is easy to implement
- it will reduce the churn compared to the situation we have now where no RFD is applied
- it spares the smaller offenders.

Changing the default suppress threshold could result in an increase in forwarding table size or announcement rate for operators who use RFD with the default settings. This warrants further discussion.

## References and Further Reading

[1] Curtis Villamizar, Ravi Chandra, Ramesh Govindan  
RFC2439: BGP Route-flap Damping (Proposed Standard)  
<ftp://ftp.ietf.org/rfc/rfc2439.txt>

[2] Most recent RIPE Document  
<ftp://ftp.ripe.net/ripe/docs/ripe-378.txt>

[2a] Older RIPE Documents  
<ftp://ftp.ripe.net/ripe/docs/ripe-178.txt>  
<ftp://ftp.ripe.net/ripe/docs/ripe-210.txt>  
<ftp://ftp.ripe.net/ripe/docs/ripe-229.txt>

[3] Cristel Pelsser, Olaf Maennel, Pradosh Mohapatra, Randy Bush and Keyur Patel. "Route Flap Damping Made Usable". PAM 2011, March 2011.  
<http://www.ij-ii.co.jp/en/lab/researchers/cristel/publications/Pelsser-RFD-PAM2011.pdf>

[4] Zhouqing Mao, Ramesh Govindan, George Varghese, Randy Katz  
Route-flap Damping Exacerbates Internet Routing Congerence SIGCOMM 2002  
<http://www.eecs.umich.edu/~zmao/Papers/sig02.pdf>

[5] Randy Bush, Tim Griffin, Zhouqing Mao Route-flap Damping: Harmful?  
NANOG 26  
<http://www.nanog.org/mtg-0210/ppt/flap.pdf>

[6] Craig Labovitz, Abha Ahuja, Abhijit Bose, Farnam Jihanian  
Delayed Internet Routing Convergence SIGCOMM 2000  
<http://conferences.sigcomm.org/sigcomm/2000/conf/paper/sigcomm2000-5-2.pdf>