

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 8, 2011

C. Pelsser
R. Bush
IIJ
K. Patel
P. Mohapatra
Cisco Systems
O. Maenel
Loughborough University
March 7, 2011

Making Route Flap Damping Usable
draft-ymbk-rfd-usable-00

Abstract

Route Flap Damping (RFD) was first proposed to reduce BGP churn in routers. Unfortunately, RFD was found to severely penalize sites for being well-connected because topological richness amplifies the number of update messages exchanged. Many operators have turned RFD off. This document recommends adjusting a few RFD algorithmic constants and limits, to reduce the high risks with RFD, with the result being damping a non-trivial amount of long term churn without penalizing well-behaved prefixes' normal convergence process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 8, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Suggested Reading	4
2. Introduction	4
3. Suppress Threshold Versus Churn	4
4. RFD Parameters	5
5. Maximum Penalty	6
6. Recommendations	6
7. Security Considerations	6
8. IANA Considerations	7
9. Acknowledgments	7
10. References	7
10.1. Normative References	7
10.2. Informative References	7
Authors' Addresses	8

1. Suggested Reading

It is assumed that the reader understands BGP, [RFC4271] and Route Flap Damping, [RFC2439]. This work is based on the measurements in the paper [pelsser2011].

2. Introduction

Route Flap Damping (RFD) was first proposed (see [ripe178] and [RFC2439]) and subsequently implemented to reduce BGP churn in routers. Unfortunately, RFD was found to severely penalize sites for being well-connected because topological richness amplifies the number of update messages exchanged, see [mao2002]. Subsequently, many operators turned RFD off, see [ripe378]. This document recommends adjusting a few RFD algorithmic constants and limits, with the result being damping of a non-trivial amount of long term churn without penalizing well-behaved prefixes' normal convergence process.

Very few prefixes are responsible for a large amount of the BGP messages received by a router, see [huston2006] and [pelsser2011]. For example, [pelsser2011] showed that only 3% of the prefixes were responsible for 36% percent of the BGP messages at a router with real feeds from a Tier-1 and an Internet Exchange Point during a one week experiment. Only these very frequently flapping prefixes should be damped. The values recommended in Section 6 achieve this. Thus, RFD can be enabled, and some churn reduced.

The goal is to, with absolutely minimal change, ameliorate the danger of current RFD implementations and use. It is not a panacea, nor is it a deep and thorough approach to flap reduction.

3. Suppress Threshold Versus Churn

By turning RFD back on with the values recommended in Section 6 churn is reduced. Moreover, with these values, prefixes going through normal convergence are generally not damped.

[pelsser2011] estimates that, with a suppress threshold of 6,000, the BGP update rate is reduced by 19% compared to a situation without RFD enabled. With this 6K suppress threshold, 90% fewer prefixes are damped compared to use of a 2K threshold. I.e. far fewer well-behaved prefixes are damped.

Setting the suppress threshold to 12K leads to very few damped prefixes (1.7% of the prefixes damped with a threshold of 2K, in the experiments in [pelsser2011] yielding an average hourly update

reduction of 11% compared to not using RFD.

Suppress Threshold	Damped Instances	Update Rate (one hour bins)
2k	43342	53.11%
4k	11253	74.16%
6k	4352	81.03%
8k	2104	84.85%
10k	1286	87.12%
12k	720	88.74%
14k	504	89.97%
16k	353	91.01%
18k	311	91.88%
20k	261	92.69%

Damped Prefixes Versus Churn

Table 1

4. RFD Parameters

The following RFD parameters are common to all implementations. Some may be adjusted by the operator, some not.

Parameter	Tunable?	Cisco	Juniper
Withdrawal	No	1000	1000
Re-Advertisement	No	0	1000
Attribute Change	No	500	500
Suppress Threshold	Yes	2000	3000
Half-Life (min)	Yes	15	15
Reuse Threshold	Yes	750	750
Max Suppress Time (min)	Yes	60	60

RFD Parameters of Juniper and Cisco

Table 2

5. Maximum Penalty

It is important to understand that the parameters shown in Table 2, and the implementation's sampling rate, impose an upper bound on the penalty value, which we can call the 'computed maximum penalty'.

In addition, BGP implementations have an internal constant which we will call the 'maximum penalty' which the current computed penalty may not exceed.

6. Recommendations

The following changes are recommended:

Router Maximum Penalty: The internal constant for the maximum penalty value **MUST** be raised to at least 50,000.

Default Configurable Parameters: In order not to break existing operational configurations, BGP implementations **SHOULD NOT** change the default values in Table 2.

Minimum Suppress Threshold: Operators wishing damping which is much less destructive than current, but still somewhat aggressive **SHOULD** configure the Suppress Threshold to no less than 6,000.

Conservative Suppress Threshold: Conservative operators **SHOULD** configure the Suppress Threshold to no less than 12,000.

Calculate But Do Not Damp: Implementations **MAY** have a test mode where the operator could see the results of a particular configuration without actually damping any prefixes. This will allow for fine tuning of parameters without losing reachability.

7. Security Considerations

It is well known that an attacker can generate false flapping to cause a victim's prefix(es) to be damped.

As the recommendations merely change parameters to more conservative values, there should be no increase in risk.

In fact, the parameter change to more conservative values should slightly mitigate the false flap attack.

8. IANA Considerations

This document has no IANA Considerations.

9. Acknowledgments

Nate Kushman initiated this work some years ago. Seiichi Kawamura and Erik Muller contributed useful suggestions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2439] Villamizar, C., Chandra, R., and R. Govindan, "BGP Route Flap Damping", RFC 2439, November 1998.
- [RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [mao2002] Mao, Z. M., Govindan, R., Varghese, G., and Katz, R., "Route Flap Damping Exacerbates Internet Routing Convergence", In Proceedings of SIGCOMM , August 2002, <<http://www.acm.org/sigcomm/sigcomm2002/papers/routedampening.pdf>>.
- [pelsser2011] Pelsser, C., Maennel, O., Mohapatra, P., Bush, R., and Patel, K., "Route Flap Damping Made Usable", Passive and Active Measurement (PAM), March 2011, <<http://archive.psg.com/110103.pam-rfd.pdf>>.
- [ripe378] Panigl, P. and Smith, P., "RIPE Routing Working Group Recommendations On Route-flap Damping", 2006, <<http://www.ripe.net/ripe/docs/ripe-378>>.

10.2. Informative References

- [huston2006] Huston, G., "BGP Extreme Routing Noise", RIPE 52 , 2006, <<http://meetings.ripe.net/ripe-52/presentations/ripe52-plenary-bgp-review.pdf>>.
- [ripe178] Barber, T., Doran, S., Karrenberg, D., Panigl, C., and

Schmitz, J., "RIPE Routing-WG Recommendation for Coordinated Route-flap Damping Parameters", 2001, <<http://www.ripe.net/ripe/docs/ripe-178>>.

Authors' Addresses

Cristel Pelsser
Internet Initiative Japan, Inc.
Jinbocho Mitsui Buiding, 1-105
Kanda-Jinbocho, Chiyoda-ku, Tokyo 101-0051
JP

Phone: +81 3 5205 6464
Email: cristel@iij.ad.jp

Randy Bush
Internet Initiative Japan, Inc.
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

Phone: +1 206 780 0431 x1
Email: randy@psg.com

Keyur Patel
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
US

Email: keyupate@cisco.com

Pradosh Mohapatra
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134
US

Email: pmohapat@cisco.com

Olaf Maennel
Loughborough University
Department of Computer Science - N.2.03
Loughborough
UK

Phone: +44 115 714 0042
Email: o@maennel.net

