Internet Engineering Task Force (IETF)
Request for Comments: 7196
Category: Standards Track
ISSN: 2070-1721

C. Pelsser
R. Bush
Internet Initiative Japan
K. Patel
Cisco Systems
P. Mohapatra
Sproute Networks
O. Maennel
Loughborough University

May 2014

Making Route Flap Damping Usable

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. Based on experimental measurement, this document recommends adjusting a few RFD algorithmic constants and limits in order to reduce the high risks with RFD. The result is damping a non-trivial amount of long-term churn without penalizing well-behaved prefixes' normal convergence process.

Status of This Memo

Table of Contents

1. 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]. Based on the measurements of [PELSSER2011], [RIPE580] now recommends that RFD is usable with some changes to the parameters. Based on the same measurements, this document recommends adjusting a few RFD algorithmic constants and limits. The result is 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, the measurements in [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 provider 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.

1.1. 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]. A survey of Japanese operators' use of RFD and their desires is reported in [RFD-SURVEY].

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in RFC 2119 [RFC2119] only when they appear in all upper case. They may also appear in lower or mixed case as English words, without normative meaning.

3. RFD Parameters

The following RFD parameters are common to all implementations. Some may be tuned by the operator, some not. There is currently no consensus on a single set of default values.

Parameter	Tunable?	Cisco	Juniper
Withdrawal	No	1,000	1,000
Re-Advertisement	No	0	1,000
Attribute Change	No	500	500
Suppress Threshold	Yes	2,000	3,000
Half-Life (min.)	Yes	15	15
Reuse Threshold	Yes	750	750
Max Suppress Time (min.)	Yes	60	60

Note: Values without units specified are dimensionless constants.

Table 1: Default RFD Parameters of Juniper and Cisco

4. 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. [PELSSER2011] studies the number of prefixes damped over a week between September 29, 2010 and October 6, 2010. With this 6,000 suppress threshold, 90% fewer prefixes are damped compared to use of a 2,000 threshold. That is, far fewer well-behaved prefixes are damped.

Setting the suppress threshold to 12,000 leads to very few damped prefixes (0.22% of the prefixes were damped with a threshold of 12,000 in the experiments in [PELSSER2011], yielding an average hourly update reduction of 11% compared to not using RFD).

Suppress Threshold	Damped Prefixes	% of Table Damped	Update Rate (one-hour bins)
2,000	43,342	13.16%	53.11%
4,000	11,253	3.42%	74.16%
6,000	4,352	1.32%	81.03%
8,000	2,104	0.64%	84.85%
10,000	1,286	0.39%	87.12%
12,000	720	0.22%	88.74%
14,000	504	0.15%	89.97%
16,000	353	0.11%	91.01%
18,000	311	0.09%	91.88%
20,000	261	0.08%	92.69%

Note: the current default Suppress Threshold (2,000) is overly agressive.

Table 2: Damped Prefixes vs. Churn, from [PELSSER2011]

5. Maximum Penalty

It is important to understand that the parameters shown in Table 1 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', and the current computed penalty may not exceed it.

6. Recommendations

Use of the following values is 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, existing BGP implementations, including the examples in Table 1, SHOULD NOT change their default values.

Minimum Suppress Threshold: Operators that want damping that is much less destructive than the current damping, 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 can 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. Acknowledgments

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

9. References

9.1. Normative References

[MAO2002] Mao, Z., Govidan, R., Varghese, G., and R. Katz, "Route Flap Damping Exacerbates Internet Routing Convergence", In Proceedings of SIGCOMM, August 2002, <http://conferences.sigcomm.org/sigcomm/2002/papers/routedampening.pdf>.

[PELSSER2011] Pelsser, C., Maennel, O., Mohapatra, P., Bush, R., and K. Patel, "Route Flap Damping Made Usable", PAM 2011: Passive and Active Measurement Conference, March 2011, <http://pam2011.gatech.edu/papers/pam2011--Pelsser.pdf>.

[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.

[RIPE378] Smith, P. and P. Panigl, "RIPE Routing Working Group Recommendations On Route-flap Damping", RIPE 378, May 2006, <http://www.ripe.net/ripe/docs/ripe-378>.

9.2. Informative References

[HUSTON2006] Huston, G., "2005 - A BGP Year in Review", RIPE 52, 2006, <http://meetings.ripe.net/ripe-52/presentations/ripe52-plenary-bgp-review.pdf>.

[RFD-SURVEY] Tsuchiya, S., Kawamura, S., Bush, R., and C. Pelsser, "Route Flap Damping Deployment Status Survey", Work in Progress, June 2012.

[RIPE178] Barber, T., Doran, S., Karrenberg, D., Panigl, C., and J. Schmitz, "RIPE Routing-WG Recommendation for Coordinated Route-flap Damping Parameters", RIPE 178, February 1998, <http://www.ripe.net/ripe/docs/ripe-178>.

[RIPE580] Bush, R., Pelsser, C., Kuhne, M., Maennel, O., Mohapatra, P., Patel, K., and R. Evans, "RIPE Routing Working Group Recommendation for Route Flap Damping", RIPE 580, January 2013, <http://www.ripe.net/ripe/docs/ripe-580>.

Authors' Addresses

Cristel Pelsser
Internet Initiative Japan
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
5147 Crystal Springs
Bainbridge Island, Washington 98110
US

EMail: randy@psg.com

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

EMail: keyupate@cisco.com

Pradosh Mohapatra
Sproute Networks
41529 Higgins Way
Fremont, CA 94539
US

EMail: mpradosh@yahoo.com

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

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