

Revisiting Recommended BGP Route Flap Damping Configurations

Clemens Mosig^{*}, Randy Bush[†], Cristel Pelsser[‡], Thomas C. Schmidt[§], Matthias Wählisch^{*}

^{*}Freie Universität Berlin, [†]Arrcus / IJJ, [‡]Université de Strasbourg, [§]HAW Hamburg,
{clemens.mosig, m.waehlich}@fu-berlin.de, randy@psg.com, pelsser@unistra.fr, t.schmidt@haw-hamburg.de

Abstract—BGP Route Flap Damping (RFD) is recommended to suppress BGP churn. Current configuration recommendations for RFD, however, are based on a study from 2010. Since then, BGP churn increased by one order of magnitude, which may lead to outdated RFD parameters and introduce more loss of reachability of stable networks. In this paper, we revisit current recommendations to configure RFD. First, we develop an accurate and scalable emulation of Cisco and Juniper RFD implementations and make it publicly available. Second, we successfully reproduce the 2010 measurement study that justified the current RFD recommendations using current data. Third, we consider the RFD implementation of an additional major router vendor (Juniper), which penalizes BGP churn differently compared to the previously studied Cisco implementation. Fourth, we include IPv6 data from 2020. Our results show that the recommended RFD configuration parameters from 2010, though seemingly rarely used, still hold today in IPv4 and IPv6 and across vendors, even though BGP churn increased significantly. Our study revises metrics to assess the impact of incorrectly configured RFD, discusses collateral damage, and gives insights into sweet spots when damping routes.

Index Terms—Internet, BGP, RFD, Measurements

I. INTRODUCTION

BGP Route Flap Damping (RFD) [1] is considered an effective mechanism to prevent oscillating routes from Internet-wide propagation [2], [3]. An RFD-enabled BGP router maintains a *penalty* per IP prefix per BGP session that additively increases with each prefix announcement or withdrawal, and decreases exponentially in between (see Figure 1). The penalty, which increases for each update type, is predefined, while the decay speed (*half-life*) can be adjusted in current RFD implementations. A prefix is suppressed when the penalty exceeds a configurable *suppress-threshold* and released when the penalty decays below a configurable *reuse-threshold*. Suppression of a prefix that was learned via one adjacent AS can lead to an explicit withdrawal of the prefix if no alternative routes learned via other adjacent ASes are available. Therefore, operators should carefully configure RFD because incorrect thresholds *etc.* can lead to unwanted suppression of prefixes and thus unreachability of IP networks.

Seven years after the first RFD implementation was available in a major router product, it was shown that the vendor default parameters cause collateral damage [4] because the common BGP convergence process can cause enough updates

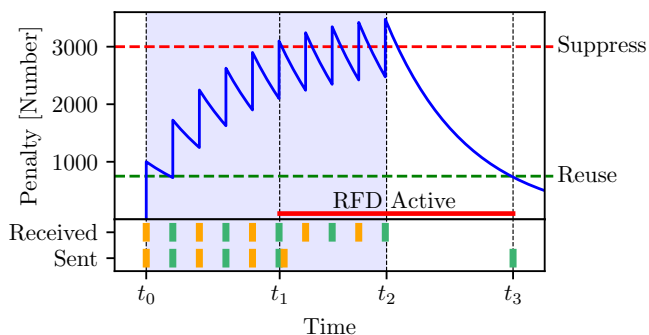


Fig. 1. RFD router perspective: The penalty for a prefix that oscillates between announcement (green) and withdrawal (orange). The dashed, horizontal lines represent suppress and reuse thresholds.

to trigger RFD. The current recommendation for the suppress-threshold [2], [3], the major knob to fine-tune RFD behavior in routers, is based on BGP data from 2010 [5], even though the Internet changed significantly during the last ten years. Those changes include more networks and higher churn in IPv4 and different BGP dynamics in IPv6.

A recent study [6] confirms RFD deployment in different types of networks but the research and operator communities lack an up-to-date view on proper RFD configuration.

In this paper, we revisit parameter recommendations for RFD by making the following contributions:

- 1) We give an up-to-date view on BGP Churn from almost 2000 BGP vantage points.
- 2) We develop an accurate and scalable emulation of Cisco and Juniper RFD implementations and make it publicly available.
- 3) We reproduce the original study by Pelsser *et al.* [5] that justifies current RFD recommendations for IPv4 using a selected subset of vantage points.
- 4) We analyze IPv6, which shows a different churn signature compared to IPv4, and compare to recent data in 2020.
- 5) We incorporate an additional router vendor implementation (Juniper), which implements a different RFD behavior compared to the vendor studied in the past (Cisco).
- 6) We extend and revise metrics used in the original

study by extensively analyzing the duration over which prefixes are suppressed.

Pelsser *et al.* [5] *estimated* the damping duration and churn reduction of RFD for different suppress-thresholds. We present an accurate, scalable emulation of the Cisco and Juniper RFD implementation that enables us to assess the collateral damage of RFD on the global routing system. Additionally, we analyze the cumulative suppress duration of a prefix relative to its impact on global routing churn. We identify a suppress-threshold which ensures that heavily flapping prefixes are suppressed, but also minimizes collateral damage on the rest of the global routing system. Table I summarizes our contributions compared to prior work that is most related to our analysis.

Our results are still on par with previous RFD recommendations. Current recommendations are sufficiently robust to cope with different vendor implementations, both IP versions, as well as enhanced BGP churn today.

The remainder of this paper is structured as follows. In Section II, we investigate BGP churn from multiple perspectives including different vantage points, IP versions, and time. In Section III, we reproduce and extend the study by Pelsser *et al.* [5]. Finally, in Section IV, we introduce a new metric for quantifying the impact of RFD and present the rationales why current RFD recommendations are still valid. We discuss related work in Section V and conclude in Section VI.

Our toolset is publicly available on <https://git.rg.net/bgp-rfd/tma-2021-auxiliary-data>.

II. BGP UPDATE CHURN OVERVIEW

BGP is noisy [7]–[16]. To assess the level of churn, we first analyze the distribution of BGP updates (*i.e.*, announcements and withdrawals) across address space, time, and location. We consider BGP updates for IPv4 and IPv6 prefixes from *all* peers of three route collector projects, Isolario [17], RIPE RIS [18], and RouteViews [19], in June 1-7, 2020. Additionally, where necessary, we compare our results with data from the same period in 2010.

It is worth noting that the Isolario project did not exist in 2010. In the 2020 dataset, we also removed three vantage points connected to Isolario because they exported excessive amounts of updates. All of these vantage points are located in the same AS, and for one vantage point we were able to

TABLE I
COMPARISON OF MOST RELATED AND OUR WORK.

Measurement	Pelsser <i>et al.</i> [5]	This work
Year	2010	2010, 2020
IP version	IPv4	IPv4, IPv6
RFD implementation	Cisco	Cisco, Juniper
Vantage point ASes	NTT, Equinix	5 Tier-1, 20 Random ASes
Damping duration	estimated	emulated
RFD impact on BGP churn	✓	✓
Collateral damage	✗	✓
Sweet spot analysis	✗	✓

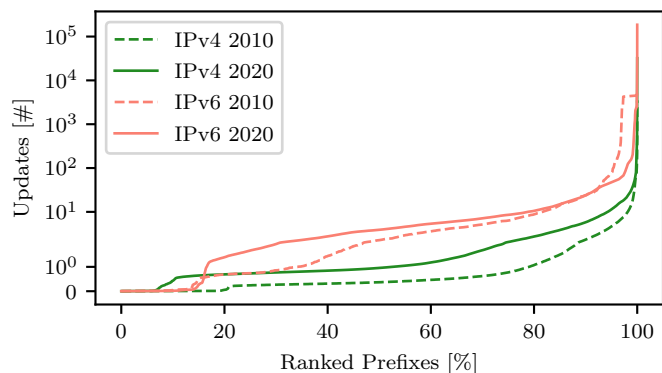


Fig. 2. Average number of announcements and withdrawals per prefix across all vantage points in June 1-7, 2010 and 2020.

find an error in the router configuration. This error has been reported and was then resolved.

In our analysis, we remove BGP duplicates (*i.e.*, two consecutive updates with identical path attributes) because they do not trigger best path selection in routers [20] and do not increment the RFD penalty. Duplicates are likely related to iBGP update activities in the vantage point network, configuration of rate limiting timers, or transitive attribute filtering [9], [21].

A. Churn from Three Perspectives

In total, we observe 4.6 billion IPv4 announcements and 3.6 billion IPv6 announcements. Based on the number of prefixes, this results in a surprising 8.4 times more updates per prefix in IPv6 compared to IPv4.

Prefixes. To better understand how BGP churn distributes across prefixes, Figure 2 shows the number of updates per IPv4 and IPv6 prefix in 2010 and 2020, averaged over all vantage points and ranked by the number of updates. As the number of prefixes differ in IPv4 and IPv6, we assign each prefix an ID relative to the overall number of prefixes per IP version. The large increase in churn over the years for both IP versions has been reported to be directly proportional to the topology size (#ASes) [16], hence we do not further investigate these differences.

In general, we observe three groups of prefixes, which exhibit similar sizes across years and IP versions: (*i*) few prefixes that are very rarely updated, (*ii*) the largest group leading to a medium amount of updates, which is the plateau in the middle, and (*iii*) the smallest group, which are heavy hitters introducing excessive amounts of BGP updates. This observation is consistent in 2010 and 2020 and confirms prior measurements [5], [22].

Heavy hitters contribute significantly to update churn, which is much more pronounced in IPv6. A (possibly changing) group of heavy hitter prefixes has existed in the Internet for at least 20 years [5], [12], [14], [22]. In 2020, 3% of the prefixes were responsible for 53.9% of IPv4 updates (74.8% in IPv6). A similar trend was already visible in 2010, when 3% (109) of

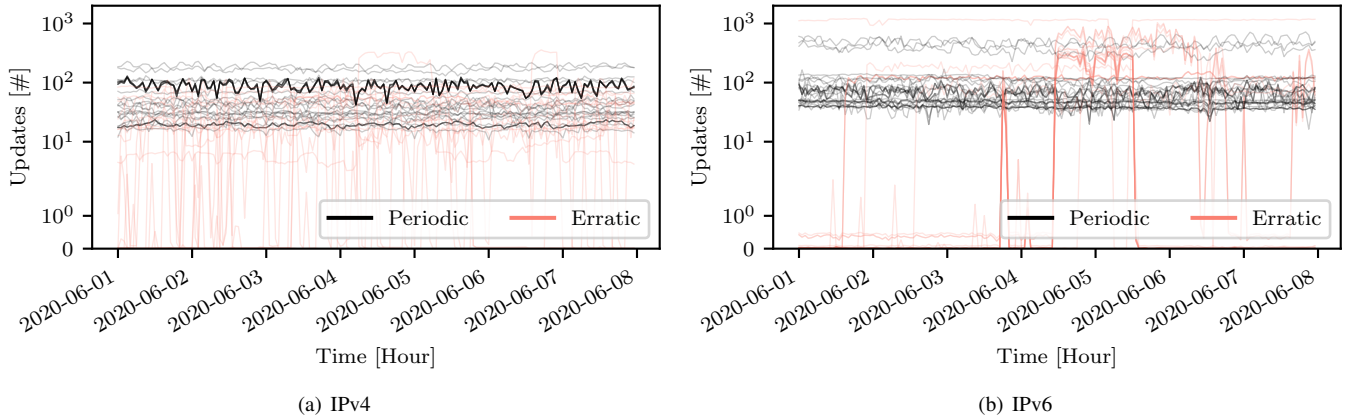


Fig. 3. Number of updates binned by hour for the 50 noisiest prefixes (ranked by cumulative update count) normalized by the vantage point count. A darker color indicates multiple prefixes having the exact same update rate because the lines are drawn with a low alpha value. The churn behavior of a prefix is *periodic* if the update rate is above 10 updates per hour during the entire measurement period, and otherwise *erratic*.

IPv6 prefixes accounted for 93.0% of updates (55.9% in IPv4). It is worth noting that these shares can vary by more than 10% depending on vantage point selection, hence, the above numbers serve as a rough intuition for the distribution.

Time. We are now interested in understanding whether heavy hitters are responsible for a constant amount of updates over time or for short-term bursts. Figure 3 shows the update count per hour for the top 50 prefixes, ranked using the cumulative update count. The values (y-axis) are normalized by the number of vantage points. Darker lines indicate multiple prefixes having the same update pattern because each line is drawn with a low alpha value. The color black indicates prefixes with an update rate consistently above 10 updates per hour during the entire measurement period (*periodic*), while red represents prefixes that did not exhibit this feature (*erratic*).

For 64% and 60% of these prefixes in IPv4 and IPv6, we observe a constant rate of more than 10 updates per hour. One would expect some kind of connection between these prefixes. But these prefixes are announced by 32 unique ASes in IPv4 (35 in IPv6) from 12 different countries. BGP Beacons [23] also exhibit this behavior and the most known examples are the RIPE Beacons [24] and the RFD Beacons [6]. One IPv4 prefix (84.205.66.0/24) in Figure 3(a) is part of the RIPE NCC address space used for BGP Beacons but it is currently not listed as active [24]. All of the remaining prefixes do not belong to publicly known Beacon projects.

For IPv4, Figure 3(a) shows a thicker line which consists of multiple prefixes having the exact same update rate. The majority of those prefixes are assigned in the same geographical region and their origin ASes use the same upstream provider. It is unlikely multiple distinct ASes send updates at the same rate, hence, we suspect their upstream provider to be causing the churn.

In a previous study, Livadariu *et al.* [22] focussed on a much larger set (top 1%) of prefixes instead of the top 50 prefixes. In contrast to our setup, they used a much smaller set of 5 vantage

points, of which the majority was located in the Internet core. They found that only 5% of IPv4 prefixes (20% in IPv6), in the top 1%, ranked by update count, are active for more than one week in a one month measurement period. To confirm this, we computed the share of prefixes in the top 1% for which daily updates are available throughout our 7-day measurement period at half of the vantage points where one of the RIPE Beacons was also visible. For IPv4 this share is 10% and 20% for IPv6. Using a very different metric, Oliveira *et al.* [12] have observed a similar erratic behavior for the largest share of highly active prefixes in 2005. Ariemma *et al.* [25] hunt for long BGP update sequences during the entire year of 2019 using the method of Discrete Wavelet Transforms. In contrast, our work analyses churn behavior in detail with focus on BGP-specific features.

Vantage points. To validate that the measurement infrastructure is not biased by a single vantage point, we analyze the distribution of churn across vantage points (see Figure 4). In contrast to the update distribution across prefixes, the plateau in this plot is much closer to the maximum. This indicates that the churn we see is the result of common Internet behavior

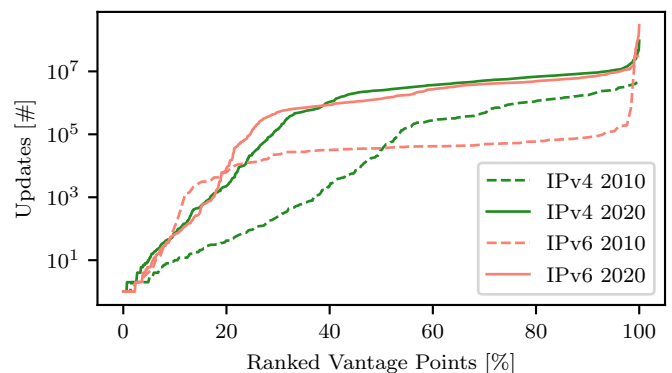


Fig. 4. Number of announcements and withdrawals per vantage point.

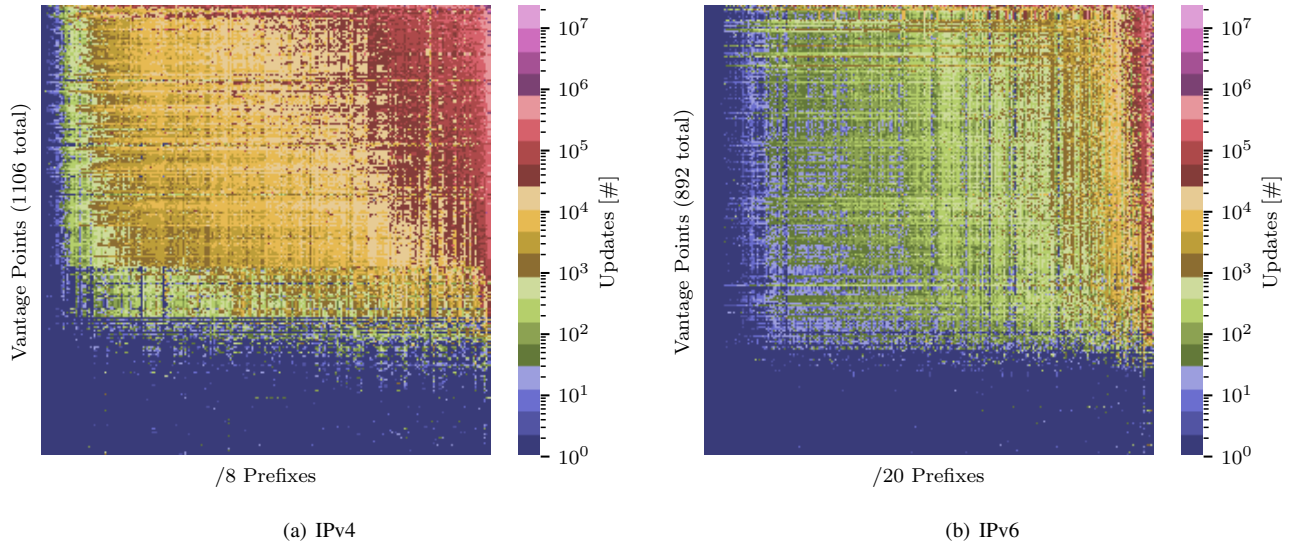


Fig. 5. Churn distribution in IPv4 and IPv6 address spaces across all vantage points sorted by 95th percentile. All values in IPv6 that are larger than the maximum value in IPv4 are assigned to the same color (light pink).

instead of malfunctioning route collectors or vantage points.

B. Churn Map

A small percentage of prefixes is responsible for most of the BGP updates we observe at all vantage points of common route collector projects. It is unclear, though, whether different vantage points see churn from different address spaces.

Figure 5 visualizes churn across the entire address space and vantage points. We aggregate all prefixes into $/8$ IPv4 and $/20$ IPv6 blocks to prevent ambiguity. We discard less specific prefixes ($< /8$ and $< /20$) as they are either deprecated [26] or specific to very few providers [27]. All prefixes are sorted by the 95th percentile from left to right and bottom to top in ascending order.

For both IPv4 and IPv6, about one quarter of vantage points exports very few updates to the route collectors. In IPv4, more than 95% (80% in IPv6) of the bottom quarter are partial feeders, *i.e.*, only exporting parts of their routing table, and all other vantage points are full feeders.

In Figure 5(a), almost all colors are visible for large chunks for vantage points, which indicates that all levels of churn occur in IPv4. Most prefixes have very different churn levels at two different vantage points, *i.e.*, 100k updates at few vantage points (red), but also 100 updates at some other vantage points (green). The largest share of vantage points exports 1k to 10k (yellow) updates for most prefixes. Overall, for most of the IPv4 prefixes, we conclude that the level of churn for a given prefix heavily depends on the vantage point or location in the Internet.

Figure 5(b), in contrast to Figure 5(a), does not show diagonal patterns but vertical lines suggesting a more uniform churn distribution across vantage points for the same IPv6 prefix. This confirms the findings by Jia *et al.* [16] who show, using Kendall’s correlation coefficient [28], that correlation

between vantage points, regarding the number of updates, is higher in IPv6 compared to IPv4. One possible explanation is the smaller size of the IPv6 topology causing a more homogeneous view from multiple vantage points [10], [16]. Also, mainly in IPv6, a small number of vantage points export excessive amounts of updates relative to the other vantage points for the largest part of the address space. This could be caused by malfunctioning routers or BGP optimizers.

III. REPRODUCING AND EXTENDING RFD MEASUREMENTS FROM 2010

The current recommendation of the RFD suppress-threshold is based on a measurement study by Pelsser *et al.* [5] conducted ten years ago. This study used BGP feeds from different networks (*e.g.*, NTT (AS 2914) and the Equinix IXP) to monitor RFD penalty values in a Cisco router for all prefixes during one week in late 2010. As introduced in the study by Pelsser *et al.*, the penalty value at a given point in time for one prefix is called *instance*. Based on this dataset they were recommending a new suppress-threshold to reduce negative side effects of RFD. They also showed that the results are independent of the specific BGP feed. *This study, however, (i) is limited to the RFD implementation of Cisco, which is different compared to Juniper, another major router vendor, and (ii) only considers BGP updates of IPv4 prefixes.* In this section, we reproduce the study by Pelsser *et al.* [5] and extend our view by including the behavior of Juniper routers and IPv6 data from 2010 and 2020.

A. Setup

In detail, we select five Tier-1 networks and randomly choose 20 ASes that provide both IPv4 and IPv6 feeds and exported at least 1k updates to the collector projects Isolario, RIPE RIS, or RouteViews during our measurement period, the

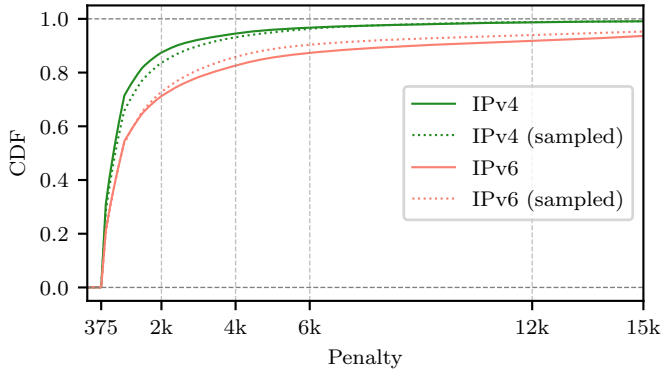


Fig. 6. CDF of penalty values visualizing the impact of using 50 vantage points (dotted) compared to all vantage points (solid).

first week of June 2020. This leads to overall 50 vantage points in 25 ASes. Limiting the set of vantage points does not affect our results but helps to reduce computational complexity. Figure 6 shows the distribution of penalty instances for our chosen subset of vantage points compared to the entire set of 1998 vantage points of all three route collector projects. The results of both data sets diverge by 6% max.

In contrast to Pelsser *et al.* [5], we do not run our measurements on Cisco (or Juniper) hardware but emulate the different RFD implementations. Essentially, we emulate RFD for each of the vantage points, using the updates they export to the route collectors. With RFD deployment at about 10% in the Internet core [6], it is possible that vantage points or other routers on the path could already implement RFD and suppress announcements, hence the update activity we see is a lower bound.

Emulation of the RFD mechanism has two advantages. First, we do not conflict with limited hardware resources on real routers. Gathering penalty values on real routers leads to inconsistent snapshot intervals at peak update rates because of internal prioritization of system processes. The 95-th quantile of the snapshot interval length was under 10 minutes [5] in the previous experiment. In our emulation, we consistently use 1 minute snapshots of penalty values. Second, our software, which is publicly released, will allow revisiting RFD behavior without special hardware in the future.

We implemented the RFD mechanism as standardized in [1] and validated the emulation using a Juniper router by testing a large set of assumptions, *e.g.*, speed of penalty decay or penalty increase with different BGP attribute changes. Unfortunately, we do not have access to the source code of Juniper or Cisco routers, which means there can be small differences between our emulation and the vendor implementation, *e.g.*, different data structures for performance improvements. We argue that these differences are negligible because, during our validation process, there were no measurable differences. Additionally, the penalty is discarded after decreasing below 375, identical to the Cisco router implementation [29].

The maximum-suppress-time is a configurable parameter

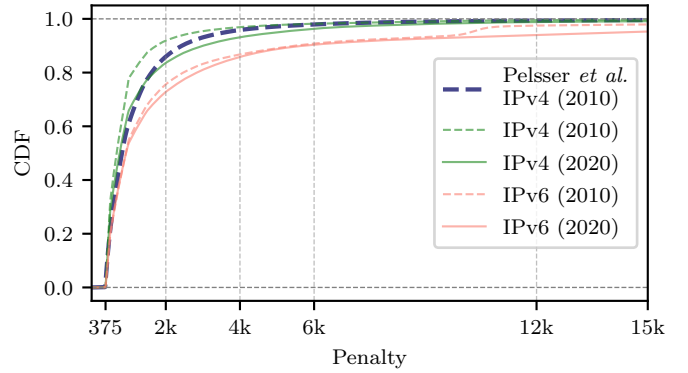


Fig. 7. CDF of reproduced penalty values for all recorded instances for a Cisco router for IPv4 (green) and IPv6 (red) in 2010 (dashed) and 2020 (solid). The original study is shown in bold and dashed blue.

which represents how long a prefix is allowed to be suppressed after the period of instability. Analogous to the setup of Pelsser *et al.* [5], we chose to not implement the maximum-suppress-time because, with vendor default values, it limits the maximum possible penalty to 12000. Changing the maximum-suppress-time dynamically based on the experiment would introduce a new level of unwanted complexity.

When we compare our results to previous measurements, we refer to data from the *BGP live-feed* [5] as provided by the authors.

B. Penalty Distribution

Figure 7 visualizes the key plot of the study we reproduce and shows the distribution of penalty values for the given input data of BGP updates. In this figure, we compare the original results from 2010 (bold, blue dashed-line) [5, Figure 5] with our data from 2010 (dashed) and 2020 (solid) for IPv4 (green) and IPv6 (red). As we consider multiple RFD vantage points in our setup, we present the average over all vantage points.

Only instances with penalty values larger than 375 are shown because Cisco routers remove penalties that are smaller than half of the RFD reuse-threshold [29]. In 2020, 99.5% of all instances in IPv4 are below 375 (98.8% in IPv6).

Fortunately, we are able to reproduce the previous results. Minor differences between both 2010 datasets are caused by the selection of vantage points (see Figure 6). Those differences were also acknowledged by Pelsser *et al.* [5] when they compared different BGP feeds. Overall, the change of BGP churn rates over the last ten years has negligible impact on the resulting RFD penalties. This can be attributed to the exponential penalty decay of the RFD mechanism which implies that maintaining higher penalties for long periods of time also requires exponentially higher churn rates.

On the other hand, it is clearly visible that IPv6 shows a different characteristic compared to IPv4, which may suggest the use of different RFD configurations in IPv6. In Section IV, we show that updating parameter recommendations for IPv6 is not necessary, though, because a few heavy hitter prefixes disproportionately contribute to the results.

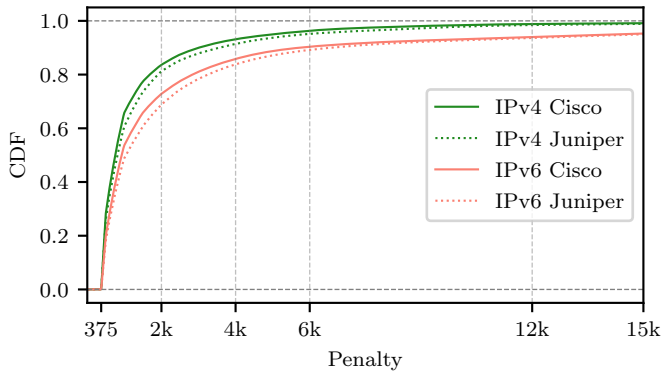


Fig. 8. CDF of penalty values for all recorded instances comparing the RFD implementation of Cisco (solid) and Juniper (dotted) in 2020.

To better understand implementation choice of different router vendors, Figure 8 compares the penalty values of a Cisco RFD implementation with a Juniper implementation. The RFD implementation in Juniper routers differ in two important ways from Cisco routers: First, Juniper routers penalize not only announcements but also withdrawals. Second, Juniper uses a default suppress-threshold of 3000 instead of 2000. Recommending vendor-specific suppress-thresholds, however, is not needed because both implementations result into distributions of penalty values with negligible differences (see Figure 8).

C. Suppressed Address Space

To fully understand the impact of a misconfigured router, Figure 9 shows the share of prefixes damped at least once across our set of vantage points. 100% would mean every routed IP prefix is being suppressed at the respective vantage point at least once in our measurement period. The dashed lines indicate share of prefixes that has been damped by at least one vantage point. With the vendor default suppress-threshold 29% IPv4 prefixes and 37% IPv6 prefixes have been damped, and therefore unreachable, by at least one vantage point!

Since the range across vantage points in Figure 9 is quite large we can conclude that the level of churn differs significantly across the Internet, which could be caused by RFD deployment. Also, the difference between the maximum and the cumulative share is large, which indicates that the vantage points see updates for different prefixes. The vantage points suppressing almost no prefixes have either limited visibility or deploy RFD.

IV. CUMULATIVE SUPPRESS DURATION

Intuitively, RFD should suppress prefixes with significant, long-lasting churn for long periods of time, but leave prefixes that trigger few updates unsuppressed. RFD models this by defining a prefix noisy when the penalty passes the suppress-threshold and remains above the reuse-threshold. The *suppress duration* describes how long a prefix is damped.

Pelsser *et al.* [5] analyzed the cumulative time a prefix penalty exceeds the suppress-threshold. This, however, does

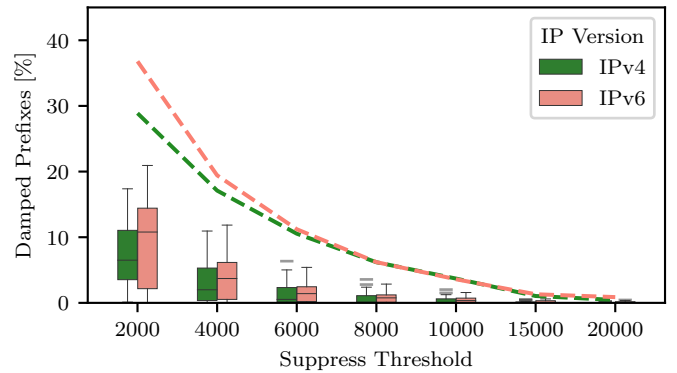


Fig. 9. Boxplot across vantage points, showing the share of the global RIB that has been damped at least once. One half of the data lies within the box, split by the median and whiskers are placed at 1.5 IQR. The dashed lines represent the total share of prefixes that has been damped by *at least one* vantage point.

not accurately reflect the duration a prefix is suppressed. In the RFD mechanism, prefixes are released only after they decrease below the *reuse-threshold* [1]. For example, assuming Cisco default values and five path changes in a short time period, the prefix is suppressed for about 26 minutes but its penalty stays above the suppress-threshold for less than 5 minutes.

In this section, we analyze the cumulative suppress duration per-prefix with different suppress thresholds. This helps us to understand whether a specific suppress threshold mainly damps prefixes that contribute to most of the BGP updates.

Our results are based on our data set from June 1-7, 2020 with views from 25 IPv4 and 25 IPv6 vantage points (details see Section III). We emulate the penalty decay with vendor defaults, *i.e.*, the half-life is set to 15 minutes and the reuse-threshold is set to 750. We cumulate suppress durations for each prefix separately.

For each prefix (*x*-axis) and suppress threshold (*y*-axis), Figure 10 depicts the cumulative damp duration (color), averaged over all vantage points. Each prefix is ranked by the number of updates it triggers, *i.e.*, prefixes with low churn are on the left and prefixes with high churn are on the right side. Prefixes which have not been damped by any vantage point are not drawn.

Figure 10 clearly shows that a remarkable number of prefixes in the lower ranks (left) are being suppressed by at least one router. Considering current suppress thresholds of Cisco (2000) and Juniper (3000), up to 29% of IPv4 (37% IPv6) prefixes would be suppressed by at least one of the 25 vantage points. This is a significant share of the global address space and suggests that the current vendor suppress-thresholds are not suitable.

The current best practices [2], [3] suggest a suppress-threshold between 6000 and 12000. Considering that 3% of the prefixes are responsible for 59.5% of the IPv4 (and 86.5% of IPv6) updates, we believe a suppress-threshold of at least 10000 is more suitable. To give a sense of scale, it is worth noting that the penalty takes one hour to decay from 12000 to

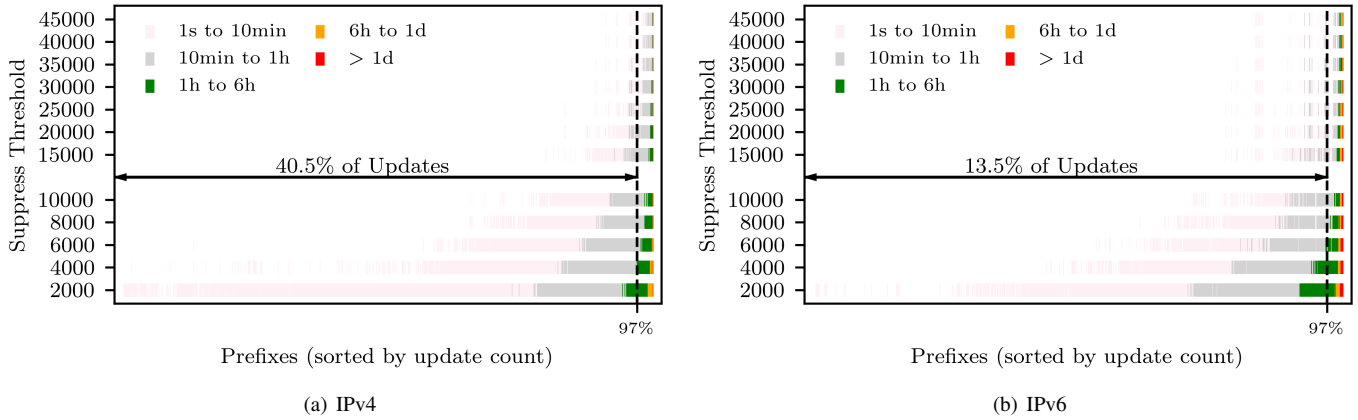


Fig. 10. Mean cumulative damp duration (color) for each prefix at different suppress-thresholds. Prefixes are sorted by the total number updates.

the reuse-threshold.

The update distribution in IPv6 has a much larger share of the mass in the top 3% compared to IPv4. For the remaining 97% of prefixes the color distribution in Figure 10 looks rather similar for both IP versions. This is because the median update count across prefixes is identical in both IP versions while the mean is $\sim 4\times$ higher in IPv6. Therefore, we argue distinction between IP versions is not necessary when configuring the RFD suppress-threshold despite different BGP churn patterns.

Only the top 3% of prefixes (ranked by update counts) need to be damped to significantly reduce churn on routers without compromising connectivity. We determined these top 3% prefixes for each vantage point and computed their average cumulative damp duration (see Figure 11). The majority of prefixes is being damped for a suppress-threshold of ≥ 2000 . Surprisingly, even a suppress-threshold of 12000 damps less than 70% of the top 3% prefixes, which suggests that 12000 is a reasonable suppress-threshold to massively reduce collateral damage of the quiet prefixes. Prefixes in the top 3% that are not being damped are oscillating with a low frequency, thus not reaching high penalty values, but are in the top 3% due to the steady flow of updates adding up over time.

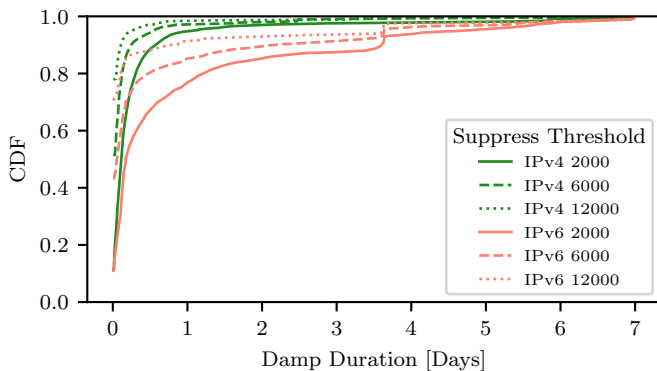


Fig. 11. CDF of damp duration across the respective top 3% of prefixes for each vantage point.

V. RELATED WORK

BGP Churn. BGP churn has been analyzed thoroughly for more than 20 years [7]–[11], [15], [16], [25]. In the most recent study, Jia *et al.* [16] discuss the longitudinal increase in churn between 1998 and 2016 for IPv4 and IPv6. They observe that BGP churn remains relatively constant over time, relative to the Internet topology size in terms of AS numbers for both IP versions.

Oliveira *et al.* [12], Rexford *et al.* [13], and Broido *et al.* [14] investigate prefixes responsible for the most churn, “heavy hitters”. Analyzing BGP data from 2001–2004, Oliveira *et al.* [12] identify a set of highly active prefixes for each day based on a threshold. They observe a different set of highly active prefixes at each of their vantage points, though, and find that 80% were highly active for shorter than a day. In 2002, Rexford *et al.* [13] correlate BGP measurements with data-plane measurements of a tier 1 provider. They found that highly active prefixes contribute to half of the BGP update events but only carry 1.4% of Internet traffic. On the other hand, those prefixes that receive half of the traffic at their vantage points contribute to only 0.1% of the BGP updates. In 2001, Broido *et al.* [14] observe that BGP churn is significantly skewed across origin ASes such that half of prefix re-announcements (flips) originate at 1% of ASes.

RFD Configurations. Mao *et al.* [4] are the first who carefully analyzed the impact of RFD configuration on the reachability of IP prefixes. They discovered that RFD configured with vendor default values does not only increase BGP convergence time but also may lead to a withdraw of the entire prefix during the convergence process in specific topologies (*e.g.*, a clique of 5 nodes) as the best path selection process increases churn. RFD delays the convergence of succeeding announcements until damping routers consider the prefix usable again, which takes up to an hour with RFD default values. Analyzing real-world BGP samples they also observed that suppressions triggered by withdrawals can be greatly reduced by halving the penalty for a path change.

The importance of RFD timer configurations in the context

of route suppression is also acknowledged by Zhang *et al.* [30], [31]. Based on a numerical analysis, they reveal that different RFD reuse timer configurations of Cisco and Juniper routers trigger route suppressions at different times, which leads to different number of updates at other routers and thus impact their damping behavior.

Pelsser *et al.* [5] revisit this observation and conduct a detailed analysis to make precise recommendations for parameters to make RFD usable again. Their goal was to filter out “elephants” without causing collateral damage, *i.e.*, reachability problems for innocent networks. Surprisingly 14% of prefixes reached an accumulated penalty greater than 2000, which is the default Cisco suppress threshold. Subsequently, this router would have suppressed every second update on average. They concluded that 12000, which suppresses only 0.64% of prefixes, shall be the recommended suppress-threshold.

In this paper, we follow previous observations that RFD needs careful configuration to prevent collateral damage, and that this configuration depends on *current* Internet dynamics. To justify the deployment of recommended RFD configurations from 2010 in 2020, we contribute a fresh view considering common deployments in terms of IP versions and router vendors in 2020.

VI. DISCUSSION AND CONCLUSION

During the last ten years, BGP churn increased by one order of magnitude, and differs in particular between IPv4 and IPv6. To our surprise, current recommendations of configuring route flap damping (RFD)—a mechanism designed to prevent the propagation of noisy prefixes—are based on measurements from 2010 considering a single vendor and only IPv4. Motivated by a recent study that observed deployment of RFD [6], we reproduced and extended prior measurements to reflect the status quo in BGP.

Suggesting the deployment (and common parameters) of RFD is intricate. Our results show that BGP churn heavily depends on the vantage point, which challenges general recommendations. On the positive side, we found that 3% of IP prefixes constantly contribute to 50% of the updates. Damping those prefixes may help to reduce load on constrained routers without introducing too much of collateral damage. Following this perspective, current RFD recommendations still hold, even though the Internet changed. This can largely be attributed to the exponential penalty decay mechanism built into RFD. Our results clearly indicate that the more noisy IPv6 is covered as well, and that different RFD implementations do not require updating RFD recommendations.

Currently, the RFD default parameters of two major router vendors do not comply with the recommended parameters. Default values are especially relevant because, in practice, most network operators who deploy RFD use the default values provided by vendors [6]. The suppress thresholds they deliver lead to many short damps of prefixes which contribute less than 50% of IPv4 BGP updates. This is even more crucial in IPv6 because 3% of IPv6 prefixes introduce 86.5% of BGP churn.

The current RFD parameter recommendations hold and cover both IP versions, but this helps little if they are not used in practice. We argue that the most important step towards making RFD more useful and less harmful is changing the default values in routers. We understand that a sudden change of default parameter configurations for any router mechanism can introduce confusion in the user base. A warning when deploying deprecated parameters may help to raise awareness, and hopefully move network operators towards using proper configurations.

Available Artifacts

We make our tools to emulate RFD implementations of Cisco and Juniper publicly available on <https://git.rg.net/bgp-rfd/tma-2021-auxiliary-data>. All input data is based on public measurement projects, RIPE RIS, RouteViews, and Isolario.

APPENDIX A RFD DEFAULT PARAMETERS

RFD parameter	Cisco	Juniper	RFC 7454
Withdrawal penalty	1000	1000	1000
Readvertisement penalty	0	1000	0/1000
Attributes change penalty	500	500	500
Suppress-threshold	2000	3000	6000
Half-life (min)	15	15	15
Reuse-threshold	750	750	750
Max suppress time (min)	60	60	60

APPENDIX B VALIDATION USING JUNIPER ROUTER

We tested eight assumptions regarding the operational behavior of the Juniper RFD implementation on a router (r_1) running JUNOS 14.2R7.5. Router r_1 was connected to router r_2 running ExaBGP version 4.1.0-2074ac17. We sent different sequences of BGP updates from r_2 to r_1 to verify:

- 1) Penalty decays as defined in RFC 2439 [1] based on the configured half-life.
- 2) A route is considered usable again when the penalty decreases below the configured reuse-threshold.
- 3) BGP duplicates, which are updates were all path attributes match, do not increase the penalty.
- 4) The penalty increases by 1000 when an announcement is received after a withdrawal.
- 5) The penalty increases by 1000 when a withdrawal is received after an announcement.
- 6) The penalty increases by 500 when a path attribute is different from the previous advertisement.
- 7) Flapping with path attributes 8 times in a row increases the penalty to 4000.
- 8) The penalty is maintained below 750.

ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers and our shepherd Balakrishnan Chandrasekaran for their valuable feedback. This work was partly supported by the German Federal Ministry of Education and Research (BMBF) within the project *PRIME*net.

REFERENCES

- [1] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," IETF, RFC 2439, November 1998.
- [2] J. Durand, I. Pepelnjak, and G. Doering, "BGP Operations and Security," IETF, RFC 7454, February 2015.
- [3] R. Bush, C. Pelsler, M. Kuhne, O. Maennel, P. Mohapatra, K. Patel, and R. Evans, "RIPE Routing Working Group Recommendations on Route Flap Damping," RIPE, RIPE Document ripe-580, January 2013.
- [4] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz, "Route Flap Damping Exacerbates Internet Routing Convergence," in *Proc. of ACM SIGCOMM*. New York, NY, USA: ACM, 2002, pp. 221–233.
- [5] C. Pelsler, O. Maennel, P. Mohapatra, R. Bush, and K. Patel, "Route Flap Damping Made Usable," in *Proc. of PAM Conf.*, ser. LNCS, vol. 6579. Berlin Heidelberg: Springer, 2011, pp. 143–152.
- [6] C. Gray, C. Mosig, R. Bush, C. Pelsler, M. Roughan, T. C. Schmidt, and M. Wählisch, "BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping," in *Proc. of ACM Internet Measurement Conference (IMC)*. New York: ACM, 2020, pp. 492–505. [Online]. Available: <https://doi.org/10.1145/3419394.3423622>
- [7] A. Elmokashfi, A. Kvalbein, and C. Dovrolis, "BGP Churn Evolution: A Perspective from the Core," *IEEE/ACM Trans. Netw.*, vol. 20, no. 2, pp. 571–584, 2012.
- [8] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkranz, "BGP Routing Dynamics Revisited," *ACM SIGCOMM CCR*, vol. 37, no. 2, pp. 5–16, 2012.
- [9] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability," *IEEE/ACM Trans. Netw.*, vol. 6, no. 5, pp. 515–528, Oct. 1998.
- [10] A. Dhamdhere, M. Luckie, B. Huffaker, K. Claffy, A. Elmokashfi, and E. Aben, "Measuring the deployment of IPv6: topology, routing and performance," in *Proceedings of the 2012 Internet Measurement Conference*, 2012, pp. 537–550.
- [11] C. Labovitz, G. R. Malan, and F. Jahanian, "Origins of Internet routing instability," in *Proc. of IEEE INFOCOM*, vol. 1. IEEE, 1999, pp. 218–226.
- [12] R. V. Oliveira, R. Izhak-Ratzin, Beichuan Zhang, and Lixia Zhang, "Measurement of highly active prefixes in BGP," in *GLOBECOM '05. IEEE Global Telecommunications Conference, 2005.*, vol. 2, 2005, pp. 5 pp.–.
- [13] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP Routing Stability of Popular Destinations," in *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement*, ser. IMW '02. New York, NY, USA: Association for Computing Machinery, 2002, p. 197–202.
- [14] A. Broido, E. Nemeth, and K. Claffy, "Internet expansion, refinement and churn," *European Transactions on Telecommunications*, vol. 13, no. 1, pp. 33–51, 2002.
- [15] A. Elmokashfi and A. Dhamdhere, "Revisiting BGP Churn Growth," *SIGCOMM Comput. Commun. Rev.*, vol. 44, p. 5–12, 2014.
- [16] S. Jia, M. Luckie, B. Huffaker, A. Elmokashfi, E. Aben, K. Claffy, and A. Dhamdhere, "Tracking the deployment of IPv6: Topology, routing and performance," vol. 165, 2019, p. 106947.
- [17] IIT-CNR, "Isolario Project," <https://www.isolario.it/>, 2020.
- [18] "Routing Information Service (RIS)," RIPE NCC, 2020. [Online]. Available: <http://www.ripe.net/projects/ris/rawdata.html>
- [19] University of Oregon, "Route Views Project," <http://www.routeviews.org/>, 2017.
- [20] T. Li, "IRTF Email thread on BGP duplicates," <https://www.mail-archive.com/rrg@irtf.org/msg02714.html>, 2010.
- [21] J. H. Park, D. Jen, M. Lad, S. Amante, D. McPherson, and L. Zhang, "Investigating occurrence of duplicate updates in BGP announcements," in *Proc. of PAM Conf.*, Springer. Springer, 2010, pp. 11–20.
- [22] I. Livadariu, A. Elmokashfi, and A. Dhamdhere, "Characterizing IPv6 control and data plane stability," in *Proc. of IEEE INFOCOM*. IEEE, 2016, pp. 1–9.
- [23] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan, "BGP Beacons," in *Proc. of the 3rd ACM SIGCOMM Conf. on Internet Measurement (IMC'03)*. New York, NY, USA: ACM, 2003, pp. 1–14.
- [24] RIPE NCC, "Current RIS Routing Beacons," <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/current-ris-routing-beacons>, 2020.
- [25] L. Ariemma, S. Liotta, M. Candela, and G. D. Battista, "Long-Lasting Sequences of BGP Updates," in *Proc. of PAM Conf.*, Springer. Springer, 2021, pp. 213–229.
- [26] O. Troan and B. Carpenter, "Deprecating the Anycast Prefix for 6to4 Relay Routers," IETF, RFC 7526, May 2015.
- [27] "Visibility of IPv4 and IPv6 Prefix Lengths in 2019," https://labs.ripe.net/Members/stephen_strowes/visibility-of-prefix-lengths-in-ipv4-and-ipv6, Apr 2019. [Online]. Available: https://labs.ripe.net/Members/stephen_strowes/visibility-of-prefix-lengths-in-ipv4-and-ipv6
- [28] M. Hollander, D. A. Wolfe, and E. Chicken, *Nonparametric statistical methods*. John Wiley & Sons, 2013, vol. 751.
- [29] "Cisco IOS IP Routing: BGP Command Reference," https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html, December 2019. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/command/irg-cr-book/bgp-m1.html
- [30] B. Zhang, D. Massey, and L. Zhang, "BGP Dynamics during Route Flap Damping," USC-CSD, Tech. Rep. 03-805, 2003.
- [31] B. Zhang, D. Pei, D. Massey, and L. Zhang, "Timer Interaction in Route Flap Damping," in *Proc. of IEEE ICDCS*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 393–403.