Revisiting Recommended BGP Route Flap Damping Configurations

Clemens Mosig, Randy Bush, Cristel Pelsser, Thomas C. Schmidt, Matthias Wählisch

BGP Route Flap Damping

Route

10.20.30.0/24

Flap

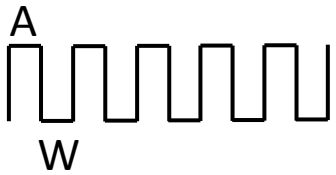

Damping

Last year at RIPE ...

RIPE 80
Virtual | 12 - 14 May

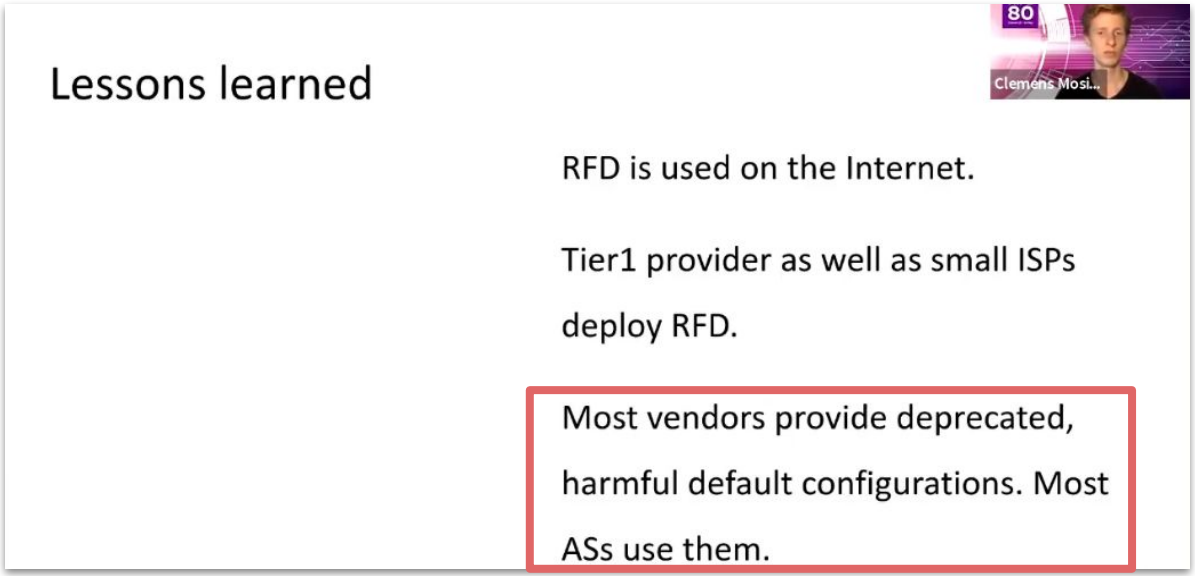

Last year at RIPE ...

RIPE 80
Virtual | 12 - 14 May

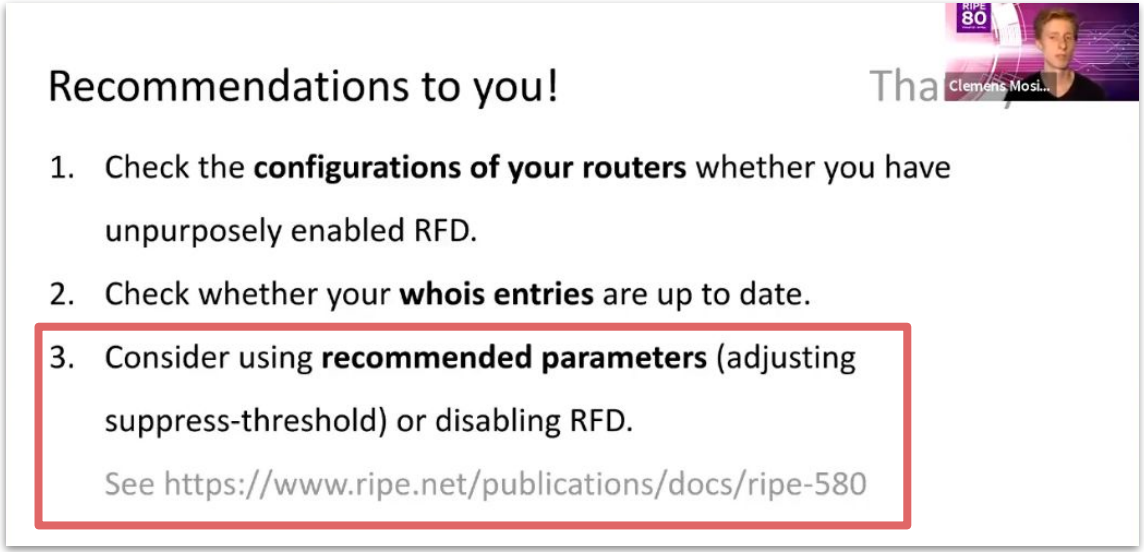

IETF/RIPE Recommendations are based on ...

Measurement	Pelsser *et al.* [5]
Year	2010
IP version	IPv4
RFD implementation	Cisco
Vantage point ASes	NTT, Equinix
Damping duration	estimated
RFD impact on BGP churn	✓
Collateral damage	✘
Sweet spot analysis	✘

IETF/RIPE Recommendations are based on ...

Measurement	Pelsser *et al.* [5]
Year	2010
IP version	IPv4
RFD implementation	Cisco
Vantage point ASes	NTT, Equinix
Damping duration	estimated
RFD impact on BGP churn	✓
Collateral damage	✘
Sweet spot analysis	✘

IETF/RIPE Recommendations are based on ...

Measurement	Pelsser *et al.* [5]
Year	2010
IP version	IPv4
RFD implementation	Cisco
Vantage point ASes	NTT, Equinix
Damping duration	estimated
RFD impact on BGP churn	✓
Collateral damage	✘
Sweet spot analysis	✘

IETF/RIPE Recommendations are based on ...

Measurement	Pelsser *et al.* [5]
Year	2010
IP version	IPv4
RFD implementation	Cisco
Vantage point ASes	NTT, Equinix
Damping duration	estimated
RFD impact on BGP churn	✓
Collateral damage	✘
Sweet spot analysis	✘

IETF/RIPE Recommendations are based on ...

Measurement	Pelsser *et al.* [5]
Year	2010
IP version	IPv4
RFD implementation	Cisco
Vantage point ASes	NTT, Equinix
Damping duration	estimated
RFD impact on BGP churn	✓
Collateral damage	✘
Sweet spot analysis	✘

We took a fresh look ...

Measurement	Pelsser *et al.* [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	✘	✓

We took a fresh look ...

Measurement	Pelsser *et al.* [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	✘	✓

We took a fresh look ...

Measurement	Pelsser *et al.* [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	✘	✓

We took a fresh look ...

Measurement	Pelsser *et al.* [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	✘	✓

3% of all IPv4 prefixes cause **53.9%** of BGP updates

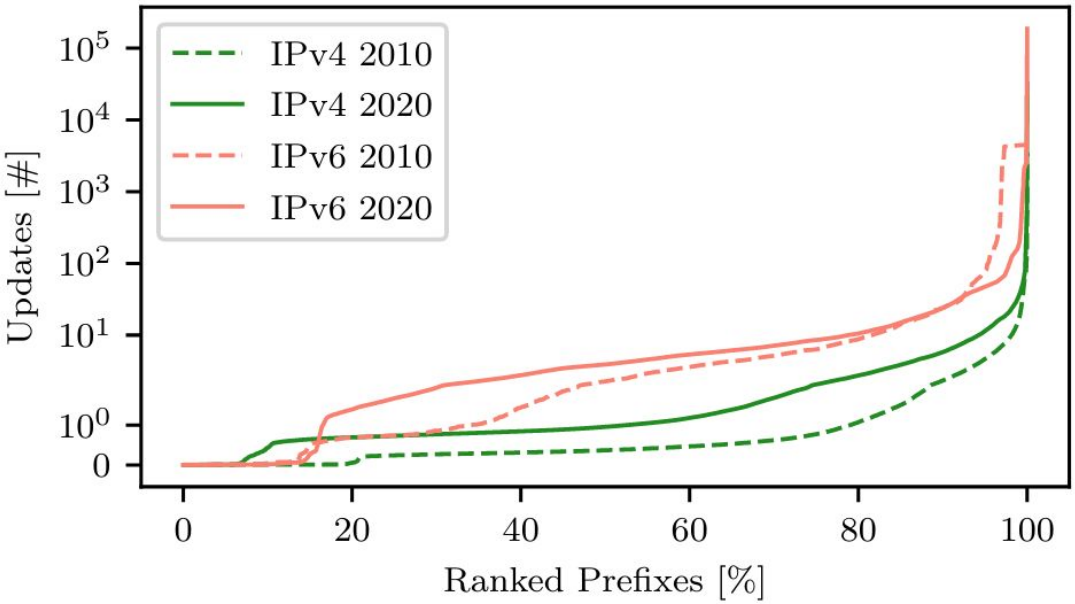

3% of all IPv4 prefixes cause **53.9%** of BGP updates

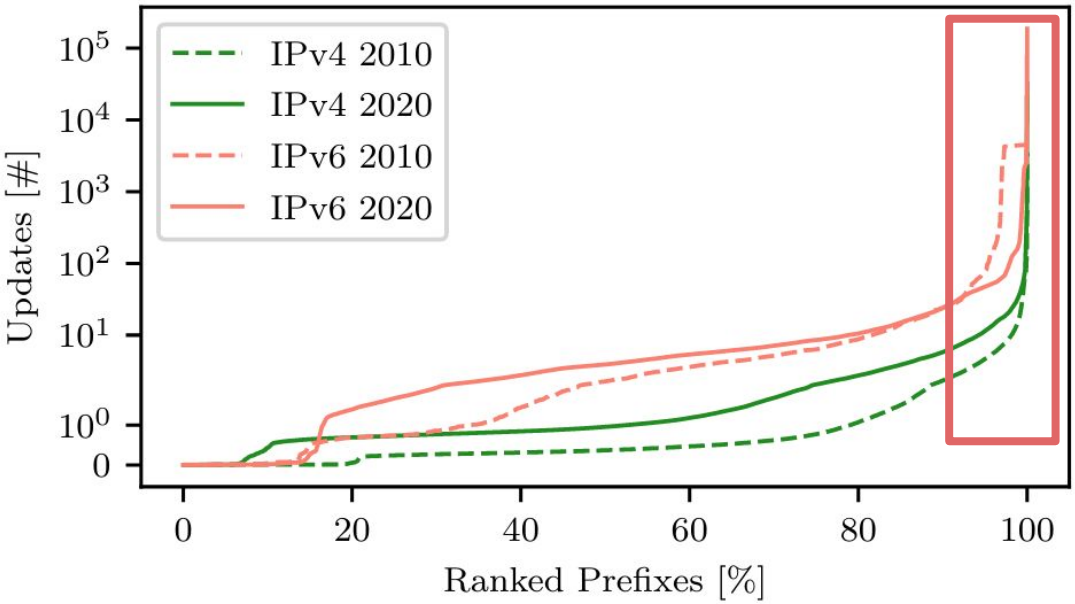

3% of all IPv4 prefixes cause **53.9%** of BGP updates

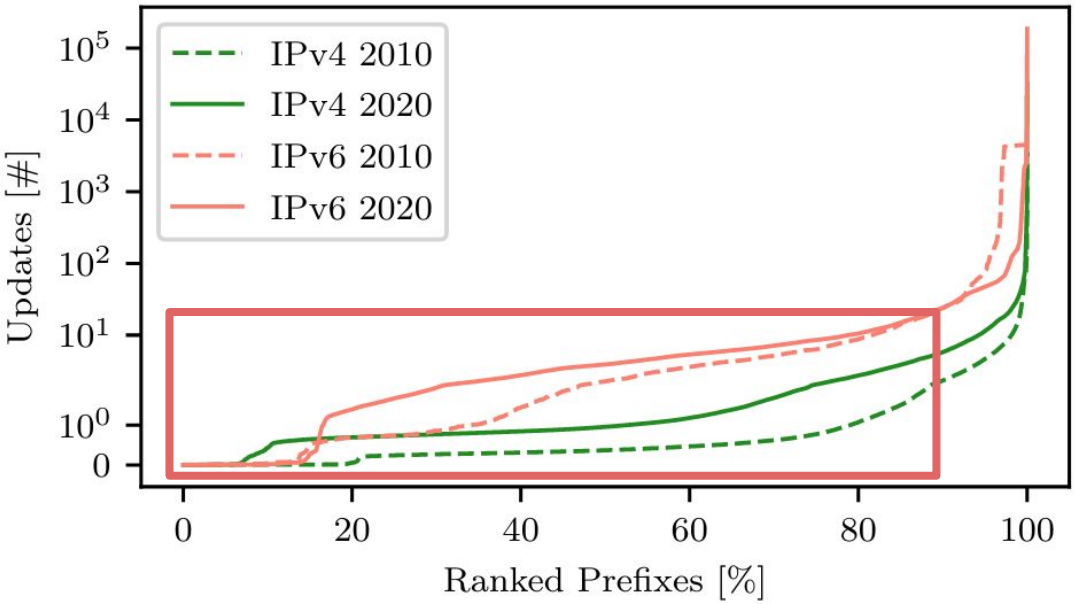

Route Flap Damping Mechanism

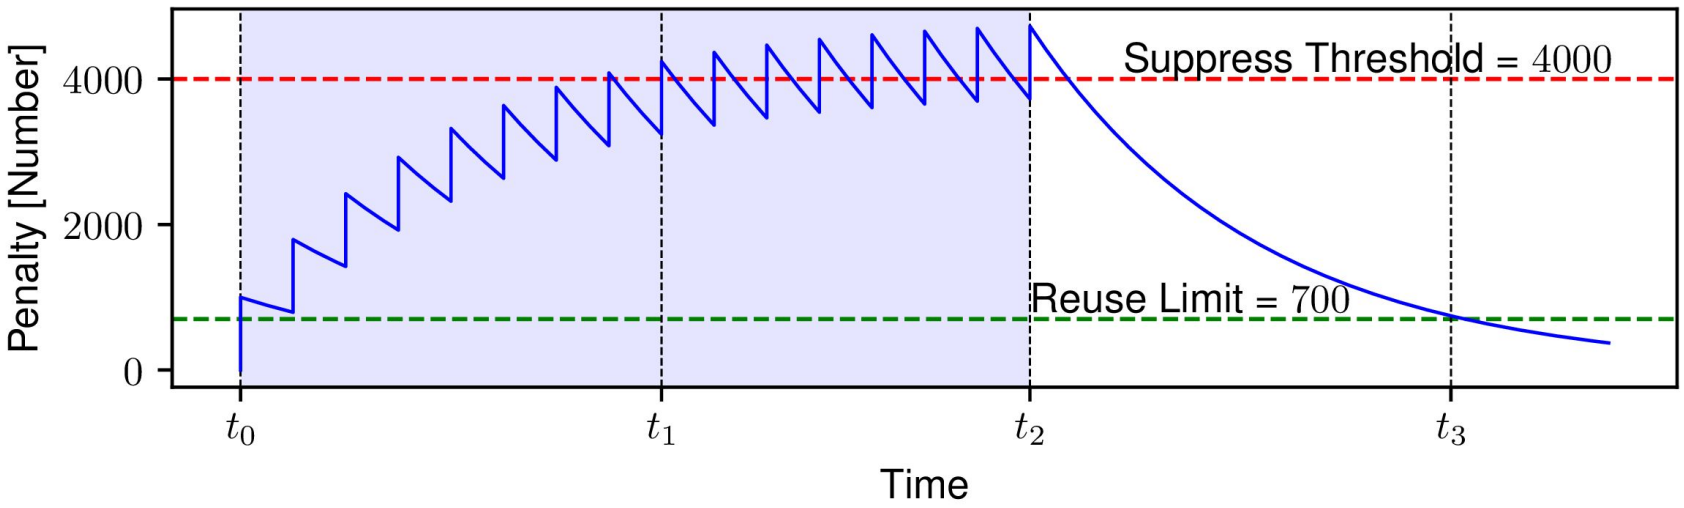

Route Flap Damping Mechanism

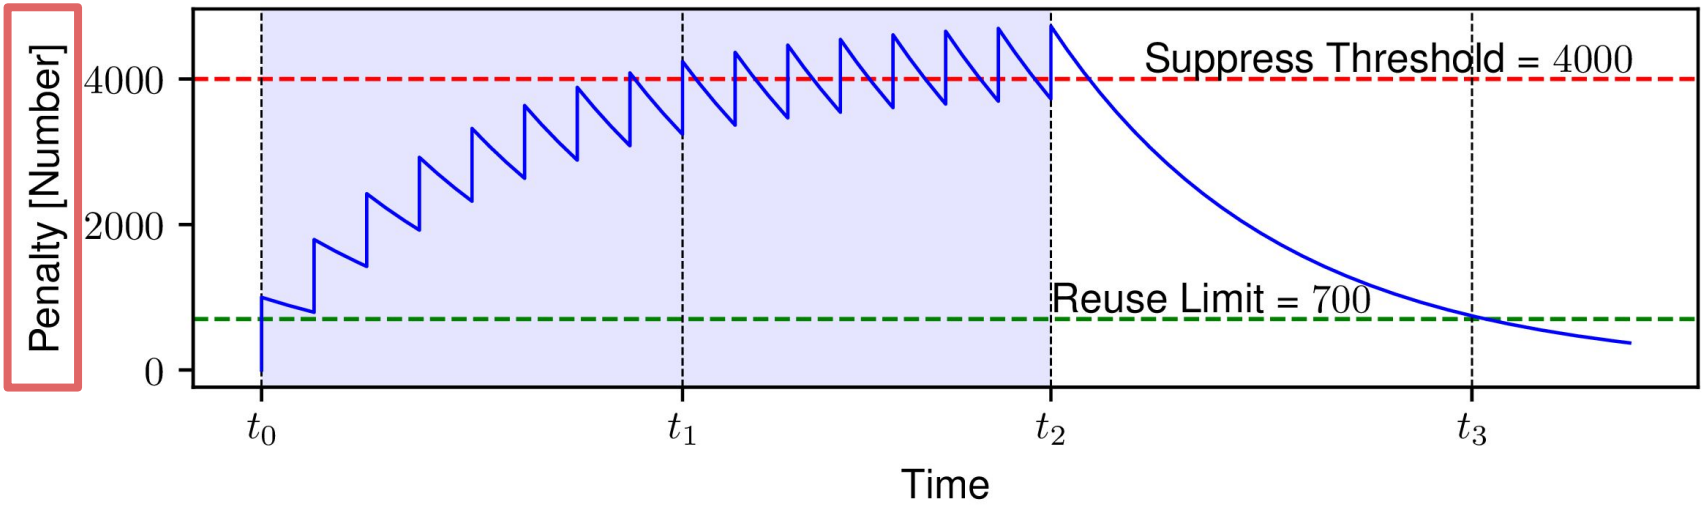

Route Flap Damping Mechanism

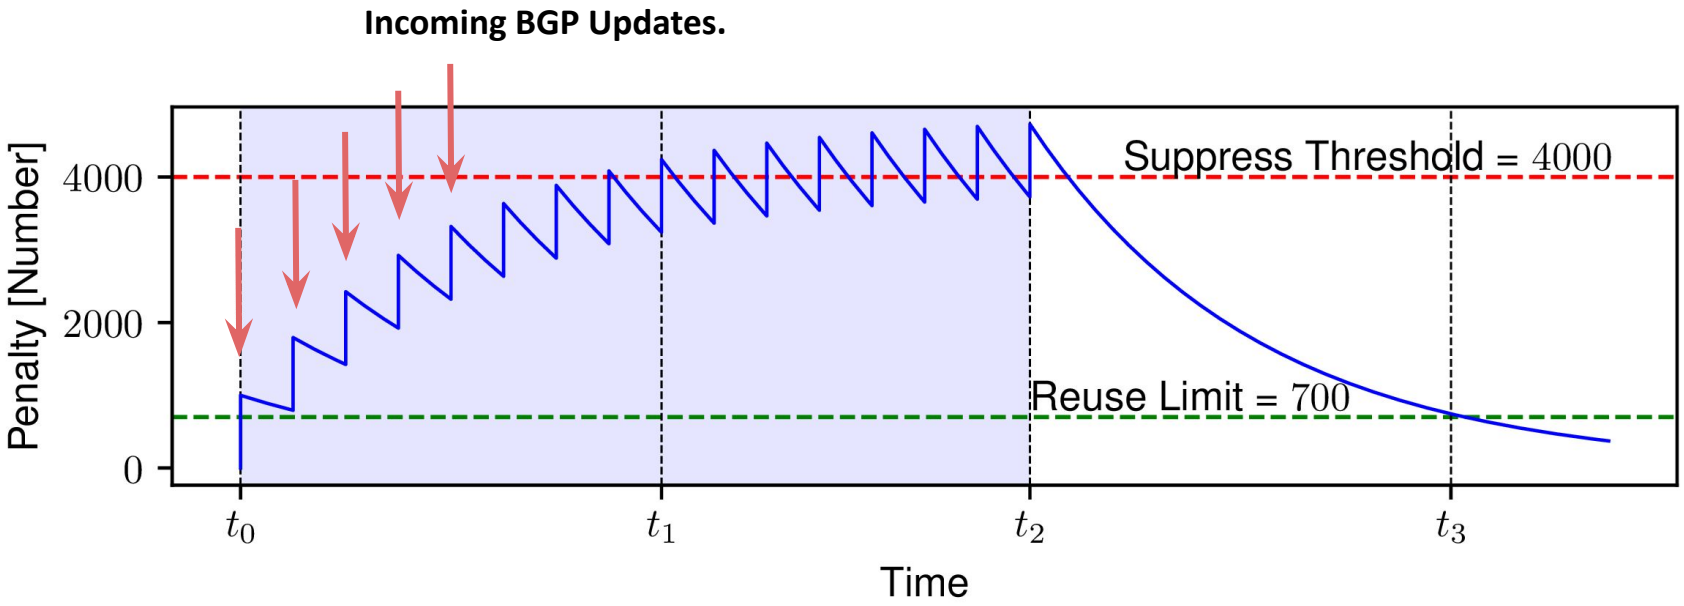

Route Flap Damping Mechanism

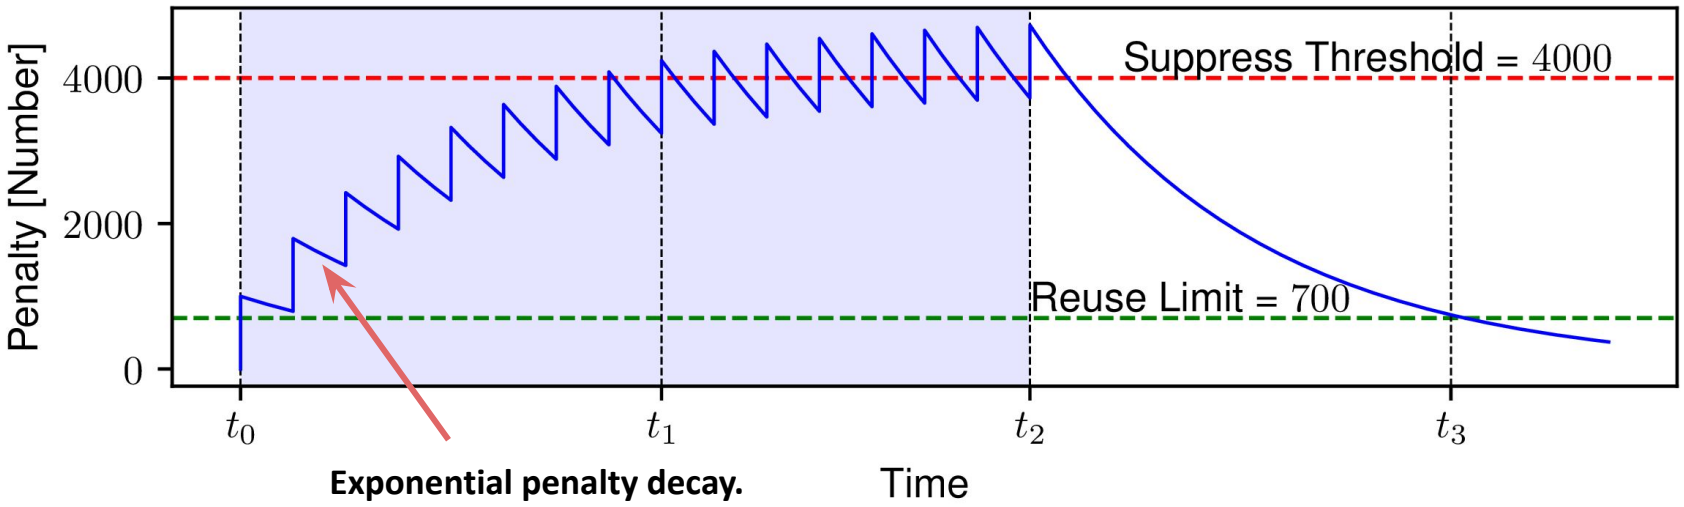

Route Flap Damping Mechanism

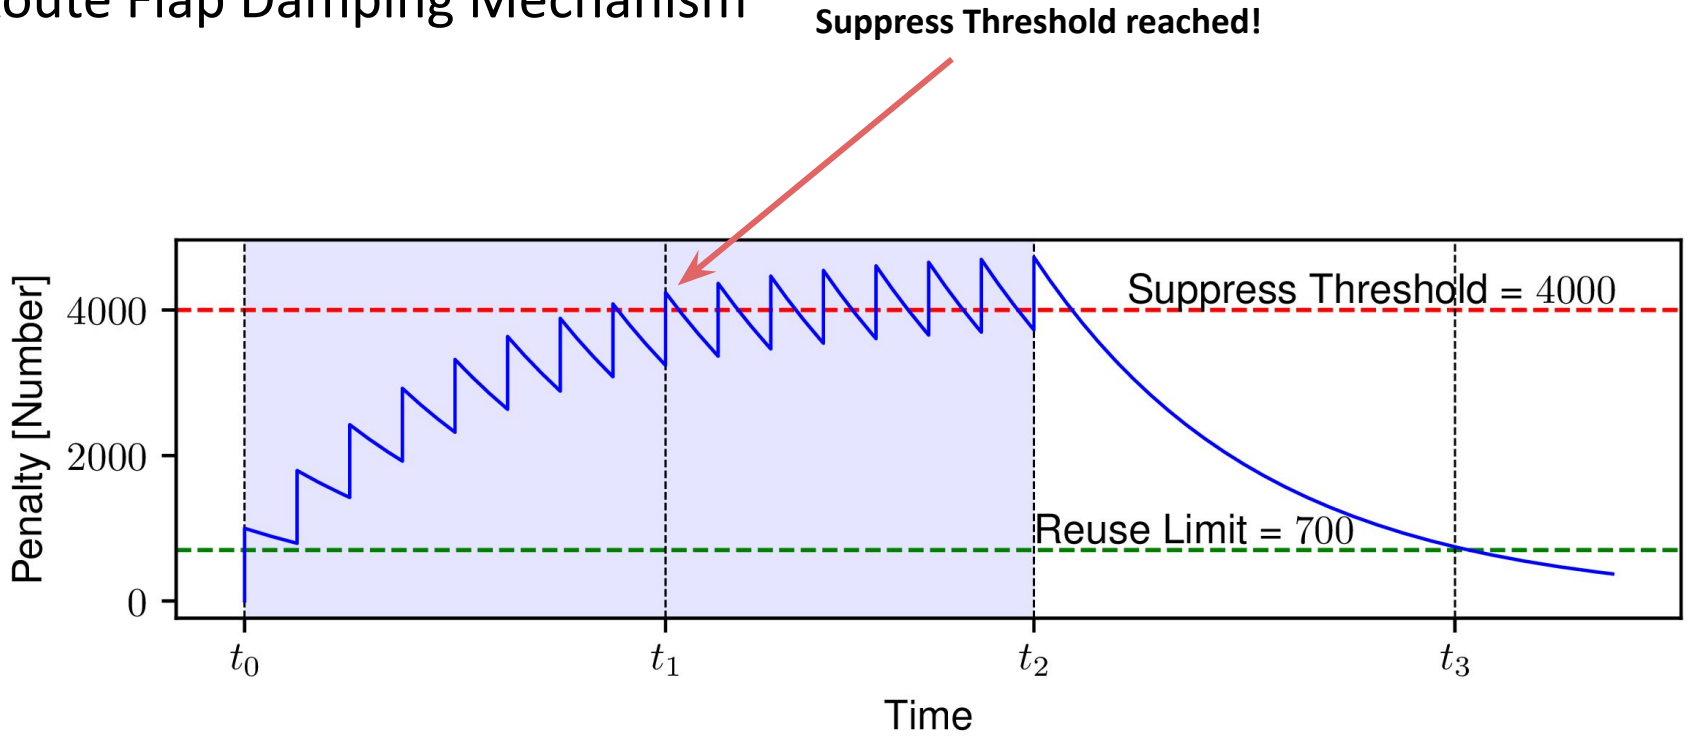

Route Flap Damping Mechanism

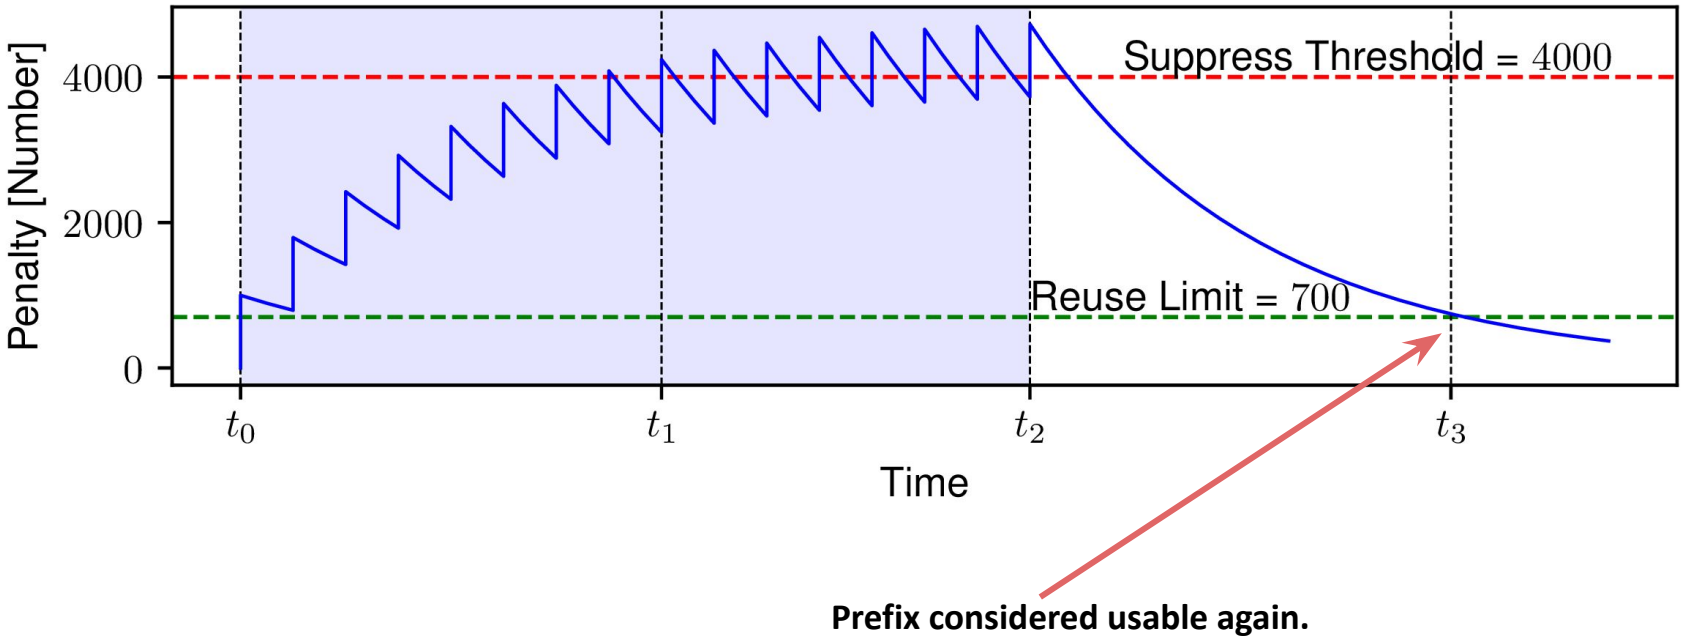

Route Flap Damping Configuration depends on the Suppress Threshold.

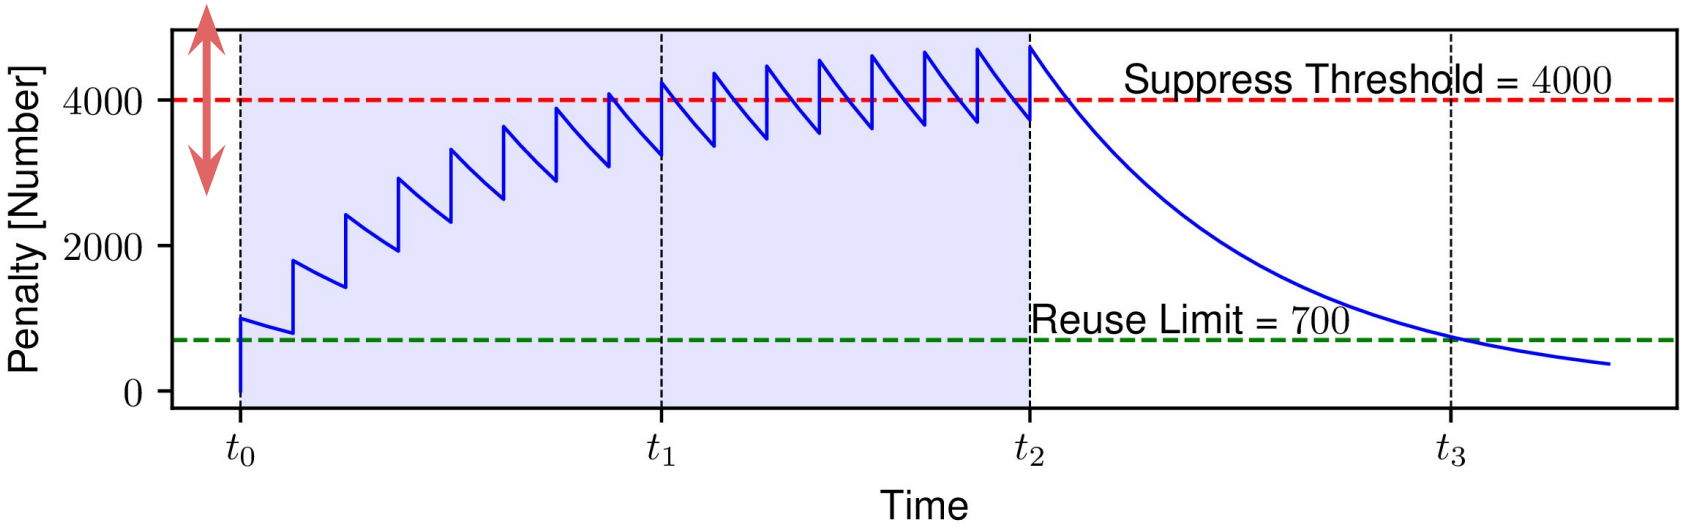

Suppress Duration for different thresholds

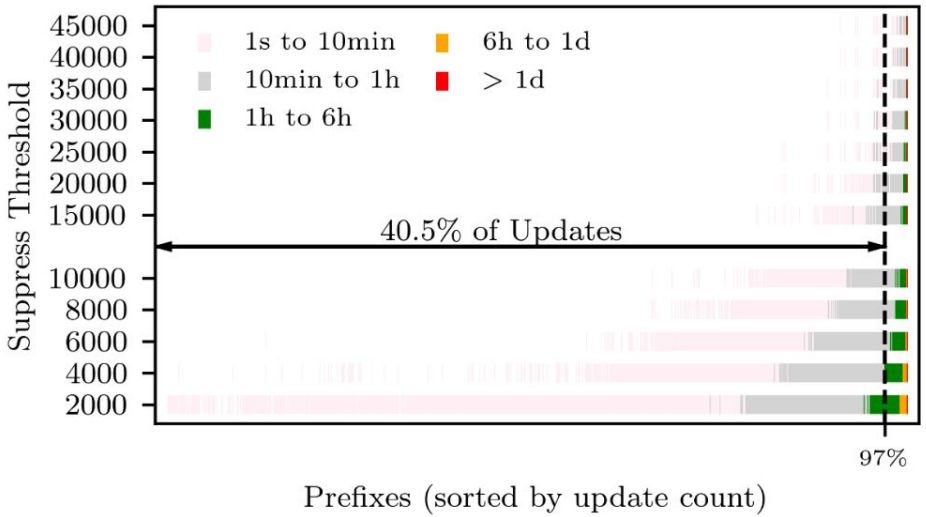

IPv4

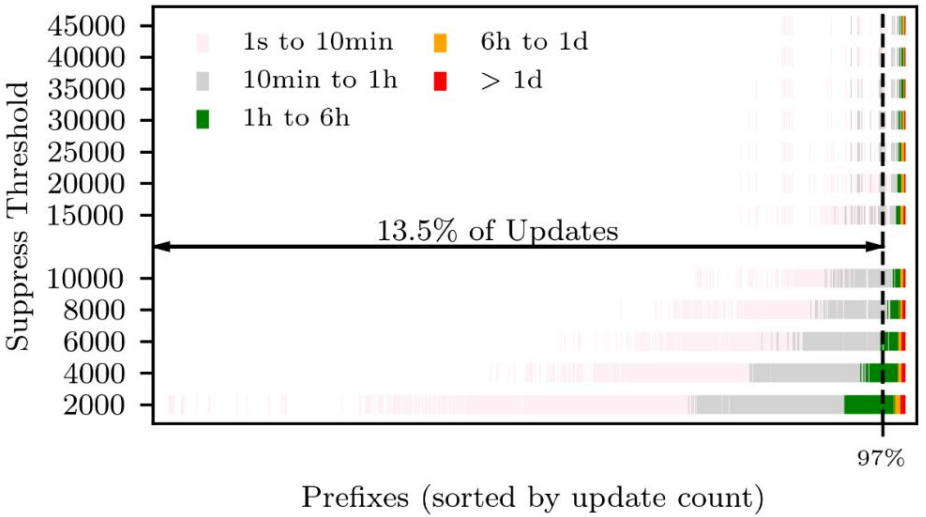

IPv6

Suppress Duration for different thresholds

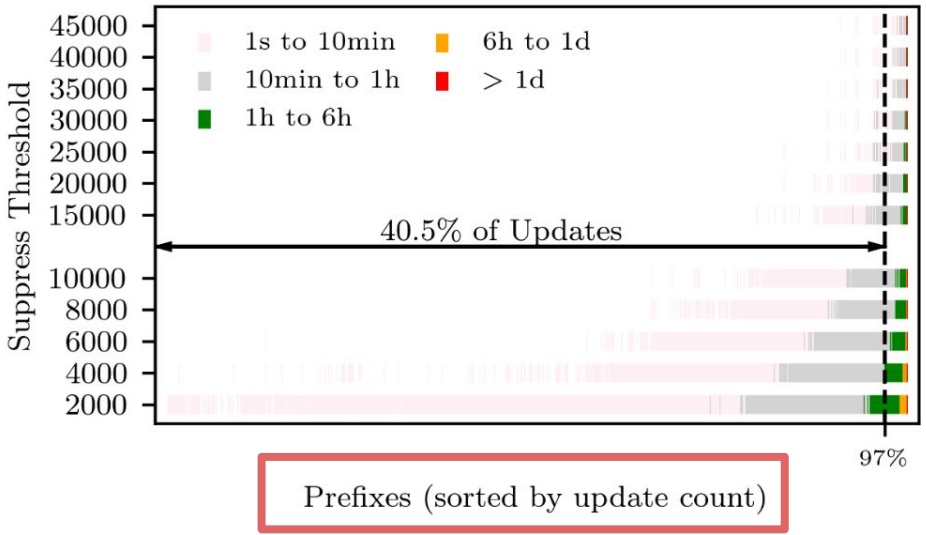

IPv4

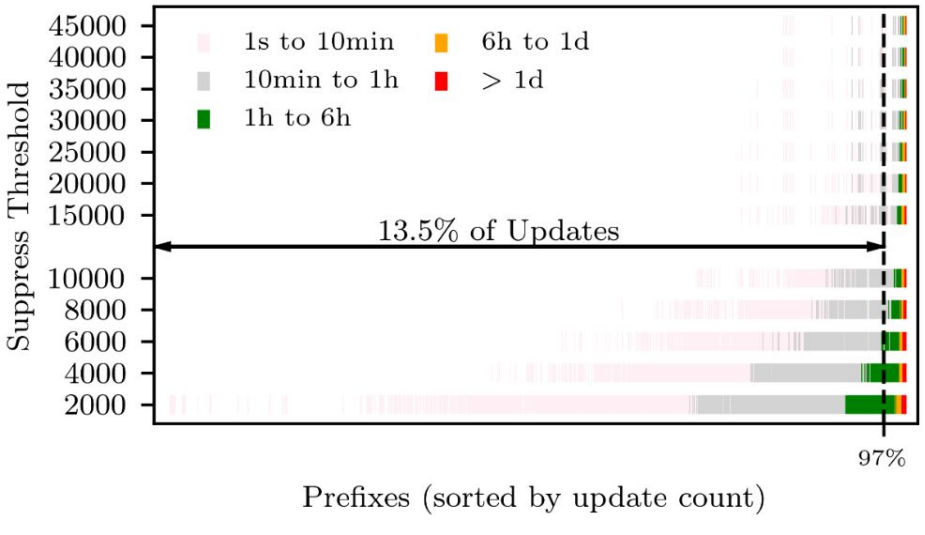

IPv6

Suppress Duration for different thresholds

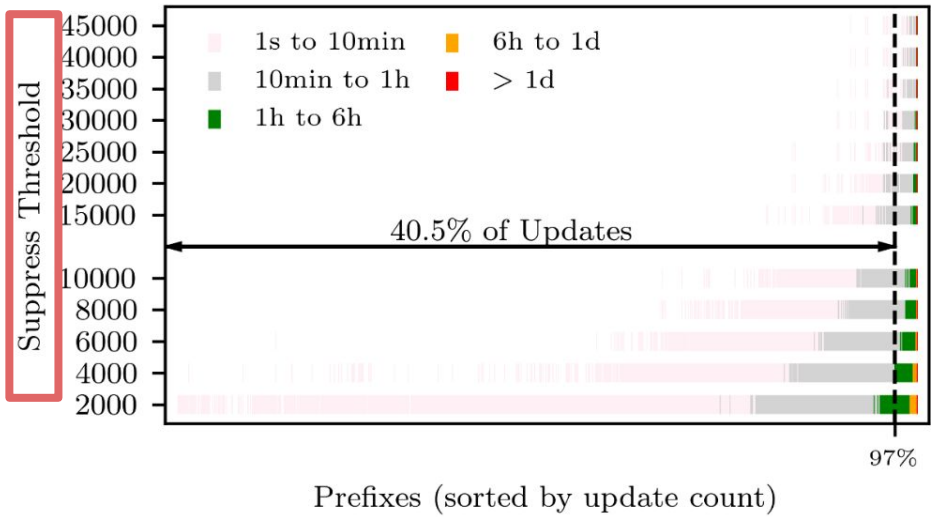

IPv4

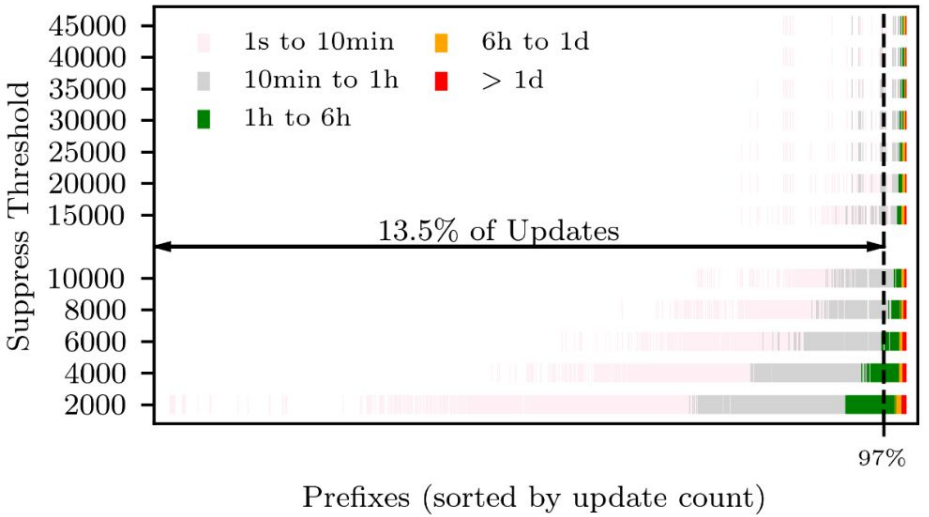

IPv6

Suppress Duration for different thresholds

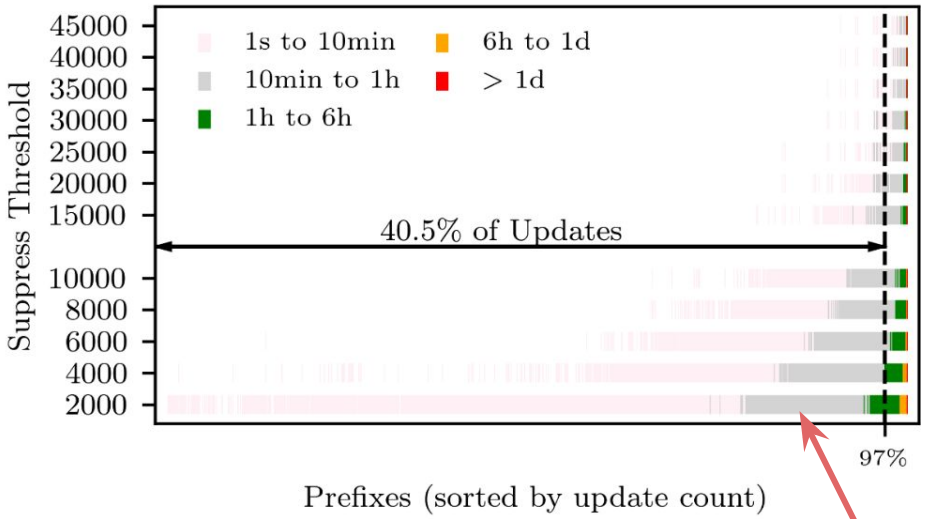

IPv4

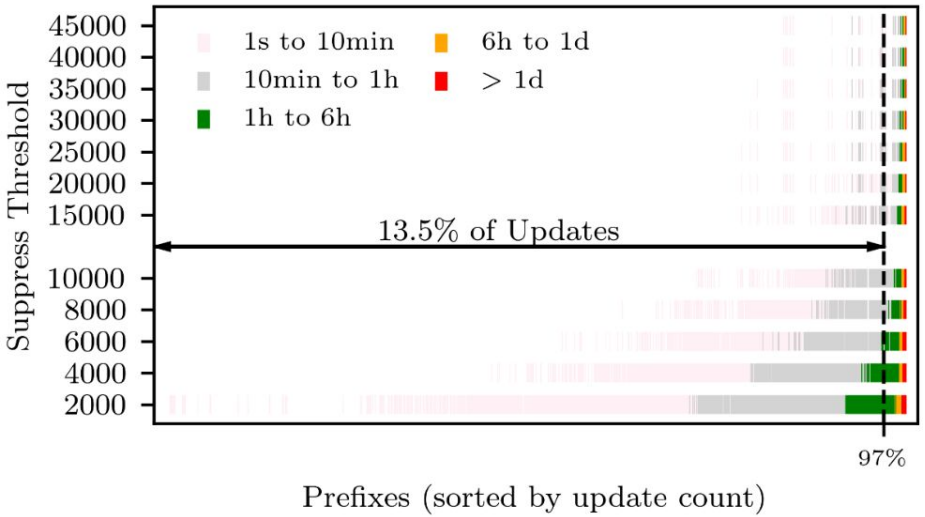

IPv6

Prefix X ranks at 80% and is suppressed between 10min to 1h during the 7 day measurement period.

Suppress Duration for different thresholds

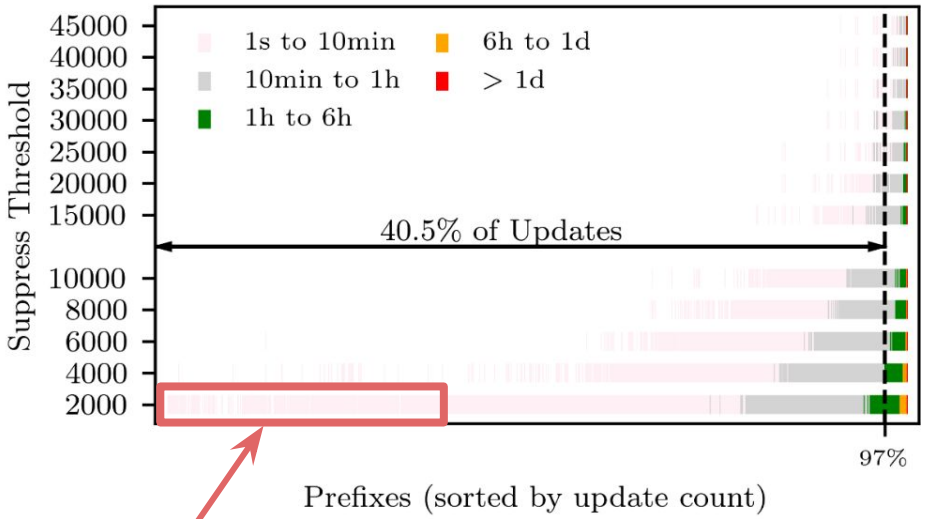

IPv4

Even very quiet prefix get suppressed.

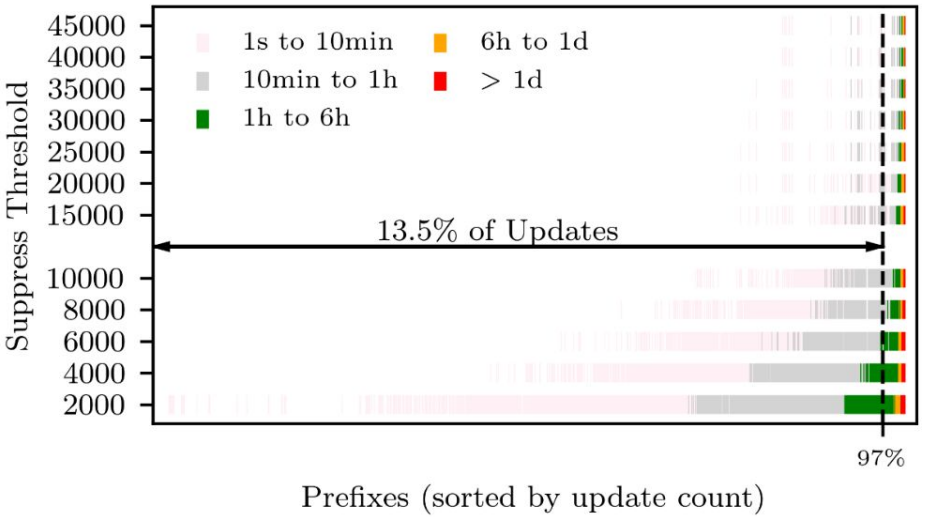

IPv6

Suppress Duration for different thresholds

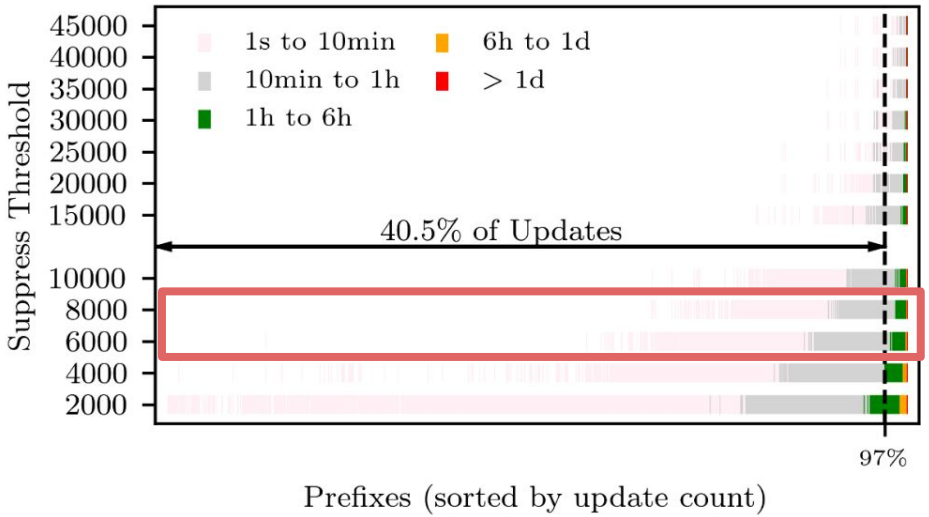

IPv4

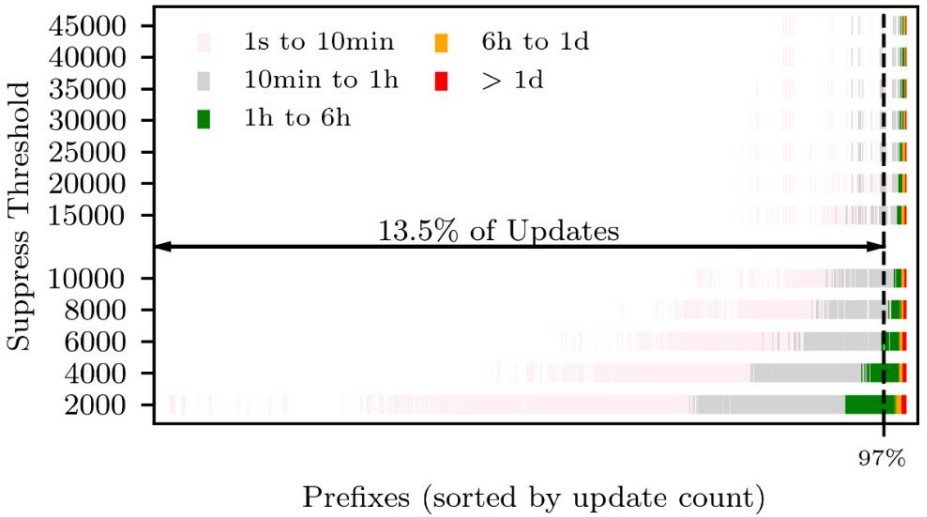

IPv6

Do we need new RFD recommendations?

Do we need new RFD recommendations?

NO.

Current recommendations are still valid

Vendor Defaults = Harmful

Recommendations = Usable

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

Website (rfd.rg.net)

Paper

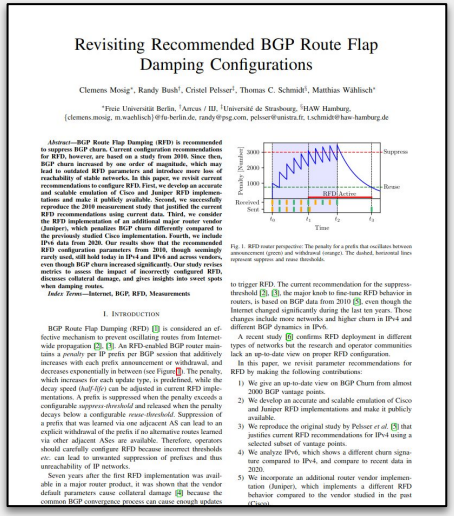

RIPE Labs Article

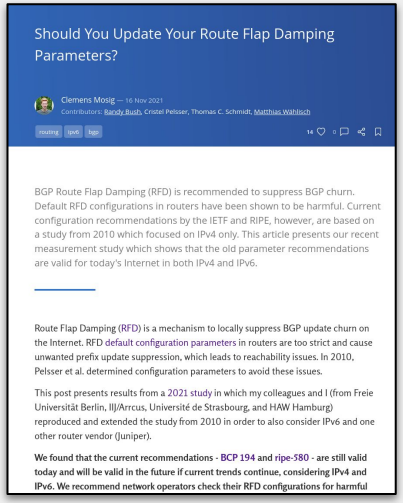

Questions? Contact clemens.mosig@fu-berlin.de!