




6	Traffic Class	Flow Label	
14 Minutes	58	64	
 Florian::Holzbauer ,  Markus::Maier ,  Johanna::Ullrich			
ff05::IMC24			
128	0	Checksum	
Identifier		Sequence Number	

| Destination Reachable:
 | What ICMPv6 Error Messages Reveal
 | About Their Sources

++-+-+--

Our Goal:

Analyze ICMPv6 error messages
beyond their source addresses

Our Contribution:

1. **Verification** of error message
type & code usage and **network
activity scans** of **11 Billion IPv6
prefixes**
2. **Rate-limiting** measurements
of **1.4M routers** to
**enumerate vendor defaults
& kernel versions**

```
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|  6   | Traffic Class |                               Flow Label                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Payload Length                               |                               58                               |                               56                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               2001:db8:200:1000::ab                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               2001:db8:5ba::1                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               1,3                               |                               Code                               |                               Checksum                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Identifier                               |                               Sequence Number                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

| ICMPv6 Error Message

```
| ... Original Destination IP: 2001:db8:200::random IID
+-----+-----
```

What Do ICMPv6 Error Messages Reveal About Their Sources?

1. Networks

2. Routers

Methods Overview

1. Controlled Environment

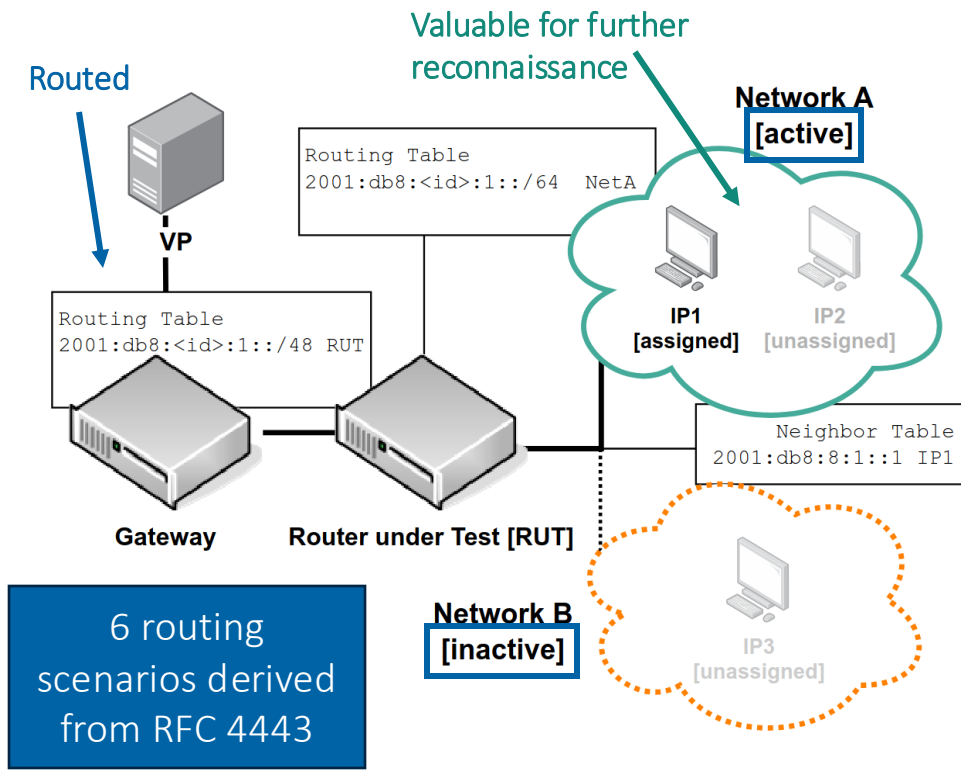
2. Verification in the IPv6 Internet

3. IPv6-wide Measurements

Router Lab

Typ	Router / Router OS
1	Cisco XRv, IOS-XE, IOS
2	Juniper Junos
3	Huawei NetEngine
4	HPE VSR
5	Arista
6	VyOs
7	Mikrotik (2 Versions)
8	OpenWRT (2 Versions)
9	Aruba

Typ	Firewall
1	PfSense
2	Fortigate



Result I: Router Lab

Mikrotik

Forward chain in iptables

Scenario	1	2	3	4	5	6	
Description	Active Network	Inactive Network	Active Netw. ACL	Inactive Netw. ACL	Null Route	Routing Loop	Classification
No Route (S2)	○ 0	● 14	● 1	● 2	● 2	○ 0	Ambiguous
Admin. Prohib. (S3,S4)	○ 0	○ 0	● 4	● 5	● 3	○ 0	Ambiguous
Addr. Unreach. (S1)	● 14	○ 0	○ 0	○ 0	● 1	○ 0	Ambiguous
Port Unreach. ()	○ 0	○ 0	● 3	● 2	○ 0	○ 0	Ambiguous
Failed Policy (S3,S4)	○ 0	● 1	● 1	● 2	○ 0	○ 0	Ambiguous
Reject Route (S5)	○ 0	○ 0	○ 0	○ 0	● 2	○ 0	Inactive
Time Exceeded (S6)	○ 0	○ 0	○ 0	○ 0	○ 0	● 15	Inactive
∅	● 1	○ 0	● 4	● 3	● 9	○ 0	

() ... Expected Scenario based on RFC 4443

RUTs may return multiple types in a single scenario

Error messages in S are classified as

Active
 Inactive

Result I: Router Lab

Juniper

Scenario	1	2	3	4	5	6	
Description	Active Network	Inactive Network	Active Netw. ACL	Inactive Netw. ACL	Null Route	Routing Loop	Classification
No Route (S2)	○ 0	● 14	● 1	● 2	● 2	○ 0	Ambiguous
Admin. Prohib. (S3,S4)	○ 0	○ 0	● 4	● 5	● 3	○ 0	Ambiguous
Addr. Unreach. (S1)	● 14	○ 0	○ 0	○ 0	● 1	○ 0	Ambiguous
Port Unreach. ()	○ 0	○ 0	● 3	● 2	○ 0	○ 0	Ambiguous
Failed Policy (S3,S4)	○ 0	● 1	● 1	● 2	○ 0	○ 0	Ambiguous
Reject Route (S5)	○ 0	○ 0	○ 0	○ 0	● 2	○ 0	Inactive
Time Exceeded (S6)	○ 0	○ 0	○ 0	○ 0	○ 0	● 15	Inactive
∅	● 1	○ 0	● 4	● 3	● 9	○ 0	

Solution: RTTs! AU in S1 shows delays of 2, 3 and 18 seconds, in S5 it is returned immediately

Error messages in S are classified as

Active
 Inactive

Result I: Router Lab

Scenario	1	2	3	4	5	6	
Description	Active Network	Inactive Network	Active Netw. ACL	Inactive Netw. ACL	Null Route	Routing Loop	Classification
No Route (S2)	○ 0	● 14	● 1	● 2	● 2	○ 0	Ambiguous
Admin. Prohib. (S3,S4)	○ 0	○ 0	● 4	● 5	● 3	○ 0	Ambiguous
Addr. Unreach. _{RTT≥1sec}	● 14	○ 0	○ 0	○ 0	○ 0	○ 0	Active
Addr. Unreach. _{RTT<1sec}	○ 0	○ 0	○ 0	○ 0	● 1	○ 0	Inactive
Port Unreach. ()	○ 0	○ 0	● 3	● 2	○ 0	○ 0	Ambiguous
Failed Policy (S3,S4)	○ 0	● 1	● 1	● 2	○ 0	○ 0	Ambiguous
Reject Route (S5)	○ 0	○ 0	○ 0	○ 0	● 2	○ 0	Inactive
Time Exceeded (S6)	○ 0	○ 0	○ 0	○ 0	○ 0	● 15	Inactive
∅	● 1	○ 0	● 4	● 3	● 9	○ 0	

Methods Overview


1. Controlled Environment

2. Verification in the IPv6 Internet

3. IPv6-wide Measurements

BValue Steps

- Goal: Collect error messages for IPv6 **networks** known to be **active and inactive**
- Starting point: hitlist addresses, for which we know at least **one** responsive **assigned** IP => BValue Steps



Original hitlist address:

2001:db8:1234:abcd:1234:abcd:1234:0101

Generated addresses:

<original bits> **<random bits>**

B127 2001:db8:1234:abcd:1234:abcd:1234:010**0**

B120 2001:db8:1234:abcd:1234:abcd:1234:01**e8**

B112 2001:db8:1234:abcd:1234:abcd:1234:**6aa1**

B104 2001:db8:1234:abcd:1234:abcd:12**21:f38d**

...

B48 2001:db8:abcd:**5276:d080:ccd6:7fc3:311c**

B40 2001:db8:ab**3e:3eb7:4c66:7f16:ade5:2b3d**

B32 2001:db8:**7438:221f:b244:476c:66bb:8da5**

- Applied to one hitlist address/subnet per routed BGP prefix
- 47,923 pass ICMPv6 responsiveness check
- Active and inactive in 44% (ICMP)

Result II: BValue Steps

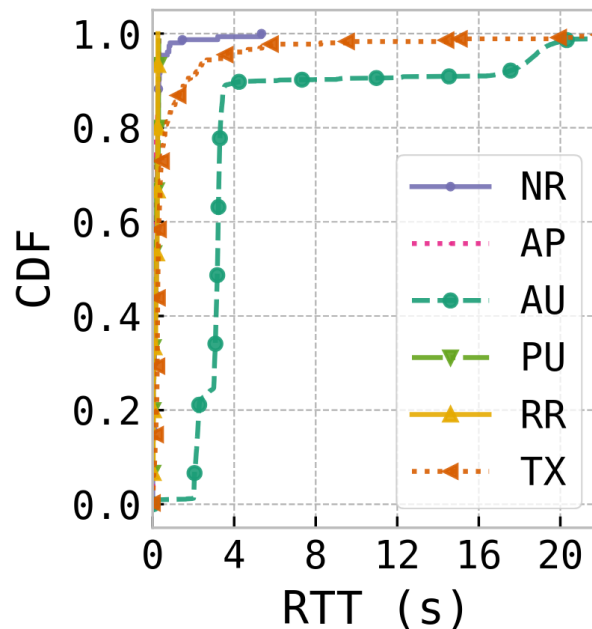
		labeled active			labeled inactive		
		Netw.	σ	%	Netw.	σ	%
active	ICMPv6	17,361	109	95.1%	471	11	4.6%
	TCP	14,522	112	93.7%	620	12	7.4%
	UDP	12,490	82	56.2%	3,687	35	32.0%
ambig.	ICMPv6	352	10	1.9%	1,645	12	15.9%
	TCP	566	10	3.7%	1,552	14	18.6%
	UDP	9,377	91	42.2%	1,455	7	12.6%
inactive	ICMPv6	537	13	2.9%	8,230	34	79.5%
	TCP	405	8	2.6%	6,191	26	74.0%
	UDP	337	12	1.5%	6,396	49	55.4%

NOTE: σ Standard deviation over five days.

True
Positives

False
Positives

True
Negatives



Methods Overview

1. Controlled
Environment

2. Verification in the
IPv6 Internet

3. IPv6-wide
Measurements

Result III: IPv6-wide Measurements


38% of measured networks

Measured /64s	6,085,410,816 (100%)
Responsive	1,368,371,825 (22.5%)
Unresponsive	4,717,038,991 (77.5%)

Responsive	/64s	/48s
Active	355,616,627 = 26%	13.5%
Ambiguous	209,970,342 = 15%	31%
Inactive	802,784,856 = 59%	55.5%

At least 1 /64 in 61% of responsive networks

YARRP % for comparison



What Do ICMPv6 Error Messages Reveal About Their Sources?

1. Networks

2. Routers

+ Active

- $AU_{RTT \geq 1 \text{ sec}}$

+ Inactive

- RR, TX & $AU_{RTT < 1 \text{ sec}}$

What Do ICMPv6 Error Messages Reveal About Their Sources?

1. Networks

2. Routers

+ Active

- $AU_{RTT \geq 1 \text{ sec}}$

+ Inactive

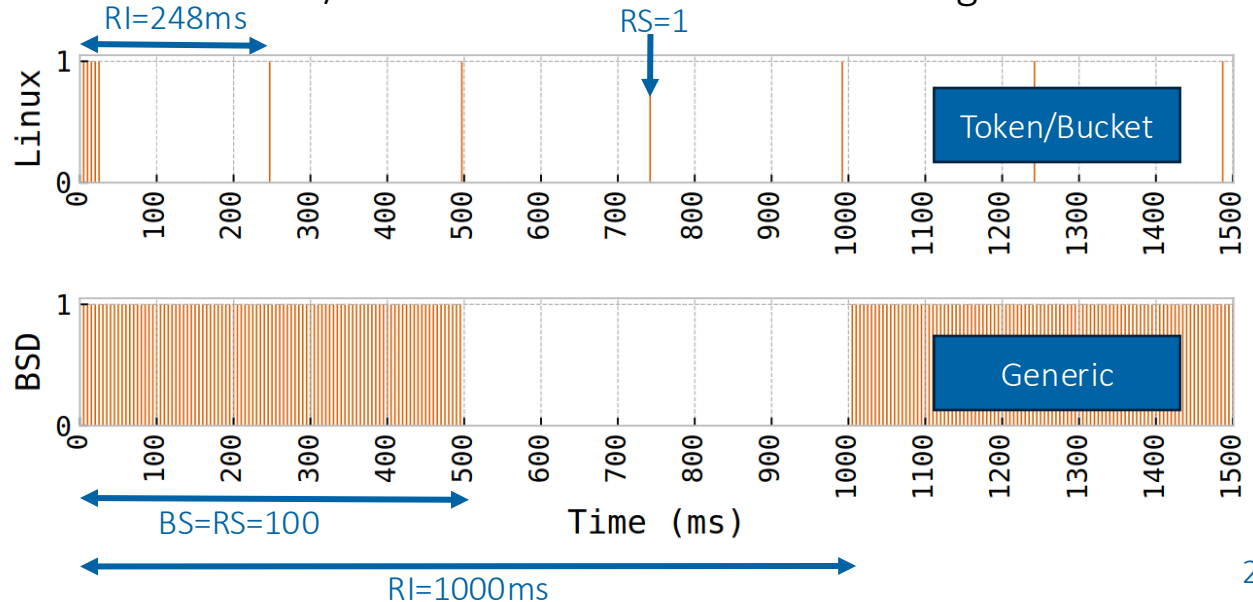
- RR, TX & $AU_{RTT < 1 \text{ sec}}$

Methods Overview

1. Controlled Environment

BS... Bucket Size
RI ... Refill Interval (ms)
RS ... Refill Size
NR(t) ... Total error messages over timespan t

- Regular Probing Interval over 10 seconds
 - ZMap with request number in payload
 - 100, 200, 1000 PPS = 1 Packet every 10, 5 and 1ms
- Extract Token/Bucket and Generic Rate Limiting Parameters



Result I: Router Lab

Vendor Differences

Kernel Defaults

	Router OS	iTTL		Bucket Size			Refill Interval (σ)			Refill Size			# Error Messages			Per Src
		All	AU	TX	NR	AU	TX	NR	AU	TX	NR	AU	TX	NR	AU	
Diff AU/NR/TX	CiscoXR9000	64	18	10	10	10	1,000	1,000	1,000	1	1	1	19	19	0 [^]	
	CiscoIOS 15.9	64	3	10	10	10	~100	~100	3,800 [*]	1	1	10	~105	~105	22 [*]	
	CiscoCSR1000 17.03	64	3	10	10	10	~100	~100	3,000 [*]	1	1	10	~105	~105	22 [*]	
	Juniper 17.1	64	2	52	12	12	~1,000	10,000	10,000	52	12	12	~520 ^o	12	12	
	HPE VSR1000	64	3	∞	∞	[*]	∞	∞	[*]	∞	∞	[*]	∞	∞	[*]	
	Huawei NE40	64	3	100-200	8	/	1,000	1,000	/	100	8	/	1,000-1,100	88	/	
Diff for AU/NR/TX	Arista 4.28	64	3		∞			∞			∞			∞		
	VyOS 1.3	64	3		6			250 [*]			1			45 [*]		✓
	Mikrotik 6.48	64,255	3		6			1000			1			15		✓
	Mikrotik 7.7	64	3		6			250 [*]			1			45 [*]		✓
	OpenWRT 19.07	64	3		6			250 [*]			1			45 [*]		✓
	OpenWRT 21.02	64	3		6			250 [*]			1			45 [*]		✓
	ArubaOS 10.09	64	3		6			250 [*]			1			45 [*]		✓
	Fortigate 7.2.0	255	3		6			10			1			1000		✓
PIsense 2.6.0	64	3		100			1000			100			1000			

~ ... Refill interval is less stable / ... The response type is not returned by the RUT. ^{*} ... Affected by the Neighbor Discovery Process. ^{*} ... /48 destination prefix; for other prefix sizes see Table 7 ∞ ... RUT is either not rate-limited or > *scanrate* (tested up to 10K pps). ^o ... Juniper's Neighbor Discovery for hop limit 0 packets causes a 2-second delay also for TX.

Kernels: Linux, Wind River Linux and FreeBSD.

Result I: Router Lab

Token/Bucket (Default)

Generic Rate

Global Rate

Peer-based

	Router OS	iTTL	Delay	Bucket Size			Refill Interval (σ)			Refill Size			# Error Messages			Per Src
		All	AU	TX	NR	AU	TX	NR	AU	TX	NR	AU	TX	NR	AU	
Diff AU/NR/TX	CiscoXRV9000	64	18	10	10	10	1,000	1,000	1,000	1	1	1	19	19	0*	
	CiscoIOS 15.9	64	3	10	10	10	~100	~100	3,800*	1	1	10	~105	~105	22*	
	CiscoCSR1000 17.03	64	3	10	10	10	~100	~100	3,000*	1	1	10	~105	~105	22*	
	Juniper 17.1	64	2	52	12	12	~1,000	10,000	10,000	52	12	12	~520 $^\diamond$	12	12	
	HPE VSR1000	64	3	∞	∞	*	∞	∞	*	∞	∞	*	∞	∞	*	
	Huawei NE40	64	3	100-200	8	/	1,000	1,000	/	100	8	/	1,000-1,100	88	/	
No diff for AU/NR/TX	Arista 4.28	64	3		∞			∞			∞			∞		<input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input checked="" type="checkbox"/> <input type="checkbox"/>
	VyOS 1.3	64	3		6			250*			1			45*		
	Mikrotik 6.48	64,255	3		6			1000			1			15		
	Mikrotik 7.7	64	3		6			250*			1			45*		
	OpenWRT 19.07	64	3		6			250*			1			45*		
	OpenWRT 21.02	64	3		6			250*			1			45*		
	ArubaOS 10.09	64	3		6			250*			1			45*		
	Fortigate 7.2.0	255	3		6			10			1			1000		
	PfSense 2.6.0	64	3		100			1000			100			1000		

~ ... Refill interval is less stable / ... The response type is not returned by the RUT. * ...Affected by the Neighbor Discovery Process. * ... /48 destination prefix; for other prefix sizes see Table 7 ∞ ... RUT is either not rate-limited or > *scanrate* (tested up to 10K pps). \diamond ... Juniper's Neighbor Discovery for hop limit 0 packets causes a 2-second delay also for TX.

Kernels: Linux, Wind River Linux and FreeBSD.

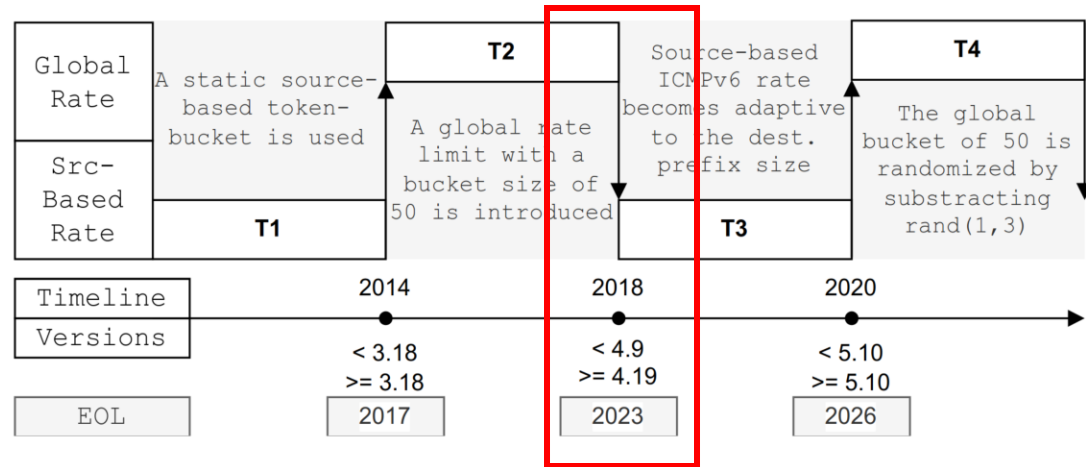
Result I: Changes in the Kernel

	Kernel Version	Release	IPv4	IPv6
Linux	2.6.26-1-2	2008	15	15
	3.16.0-4-6	2014	15	15
	4.9.0-3-13	2016	15	15
	4.19.0-5-21	2018	15	45
	5.10.0-8-22	2020	15	45
	6.1.0-9	2022	15	45
Freebsd	11.0	2016	2000	1000
Netbsd	8.2	2020	1000	1000

- Number of error messages over 10 seconds
- Static and dynamic testing of the Linux kernel shows a change for IPv6 in kernel version 4.

Three rate-limiting changes over time:

1. Introduction of global rate
2. Peer-based becomes adaptive to dest. prefix
3. Global bucket is randomized




Methods Overview

2. Verification in the IPv6 Internet

- Extract SNMPv3 vendor labels for 476K IPv6 routers (Albakour, 2021)
 - 50,952 exist in our tracerouting data, for which we can collect rate limit parameters
 - requires: destination behind router, hop limit
 - 1) Validate lab defaults
 - 2) Extend lab defaults with new vendor-specific rate-limits

Result II: SNMPv3 Label Comparison



- Huawei labeled overlap lab default
- Cisco labeled match XRv and IOS lab defaults

- Majority of Juniper labeled is limited above measurement scanrate of 200pps
- Mikrotik matches Linux kernel defaults


Methods Overview

3. IPv6-wide Measurements

- We collected rate-limits for 1.4M of the IPv6 routers
 - 200 PPS, Time Exceeded, 10 seconds
 - Centrality Score
 - count the number of paths to a /48 destination a router is seen on
 - Distance-based Matching
 - Bin the number of responses for each second (T_1, T_2, \dots, T_{10}) and compute distance
 - Match rate limiting parameters (Bucket Size, Refill Interval, Refill Size) for rates within adaptive threshold (10 to 100 based on $NR(10)$)

Result III: Rate Limit Matching

- Routers on multiple paths (green), more **distinguishable** rate limits
 - o Cisco (18% IOS, 4.2% XRv)
 - Requires AU to separate IOS further
 - o Huawei (12.6% SNMPv3, 10.3% NE) rate-limits most prominent
 - o 11.8% higher rate limits, Juniper
- Periphery (purple)
 - o 83.4% Linux <4.9 or newer version with small dest. prefix sizes (less likely)
 - o 12.6% ≥ 4.19 with default route



What Do ICMPv6 Error Messages Reveal About Their Sources?

1. Networks

+ Active

- $AU_{RTT \geq 1 \text{ sec}}$

+ Inactive

- RR, TX & $AU_{RTT < 1 \text{ sec}}$

2. Routers

+ Centrality > 1:

- Vendors

+ Centrality = 1:

- Kernel Version

Conclusion

- High share of **ICMPv6 forwarding** IPv6 networks that perform **Neighbor Discovery** and return $AU_{RTT>1sec}$
- Diagnose value of destination unreachable subcodes undermined through use for filtering - e.g. No Route through Mikrotik, Address Unreachable by Juniper
 - Diagnose value **can be restored** if router vendor is known or by monitoring RTTs
- Error Messages are no solution for all networks - **38% remain silent**
- Rate-Limiting based vendor detection fills a **gap**
 - SNMPv3 (Albakour, 2021) labels provide **limited coverage**
 - LFP (Albakour, 2023) is only applicable to **IPv4** (harmonization of Hop limits in IPv6)
 - Time Exceeded unlikely to be disabled

Dr. Johanna Ullrich

University of Vienna/SBA Research

johanna.ullrich@univie.ac.at

Artifacts available:



[sbaresearch/icmpv6-destination-reachable](https://github.com/sbaresearch/icmpv6-destination-reachable)



[sbaresearch/router-lab](https://github.com/sbaresearch/router-lab)

```
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|  6  | Traffic Class |                               Flow Label                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Payload Length                               |          58          |          56          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               2001:db8:200:1000::ab                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               2001:db8:5ba::1                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 129,1,3 |          Code          |                               Checksum                               |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               Identifier                               |          Sequence Number          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

Questions?

| ... Please set Typ and Code accordingly :)
+---+---+---+