The Impact of Router Outages on the AS-Level Internet

Matthew Luckie* - University of Waikato
Robert Beverly - Naval Postgraduate School

*work started while at CAIDA, UC San Diego

SIGCOMM 2017, August 24th 2017

Internet Resilience

Where are the Single Points of Failure?

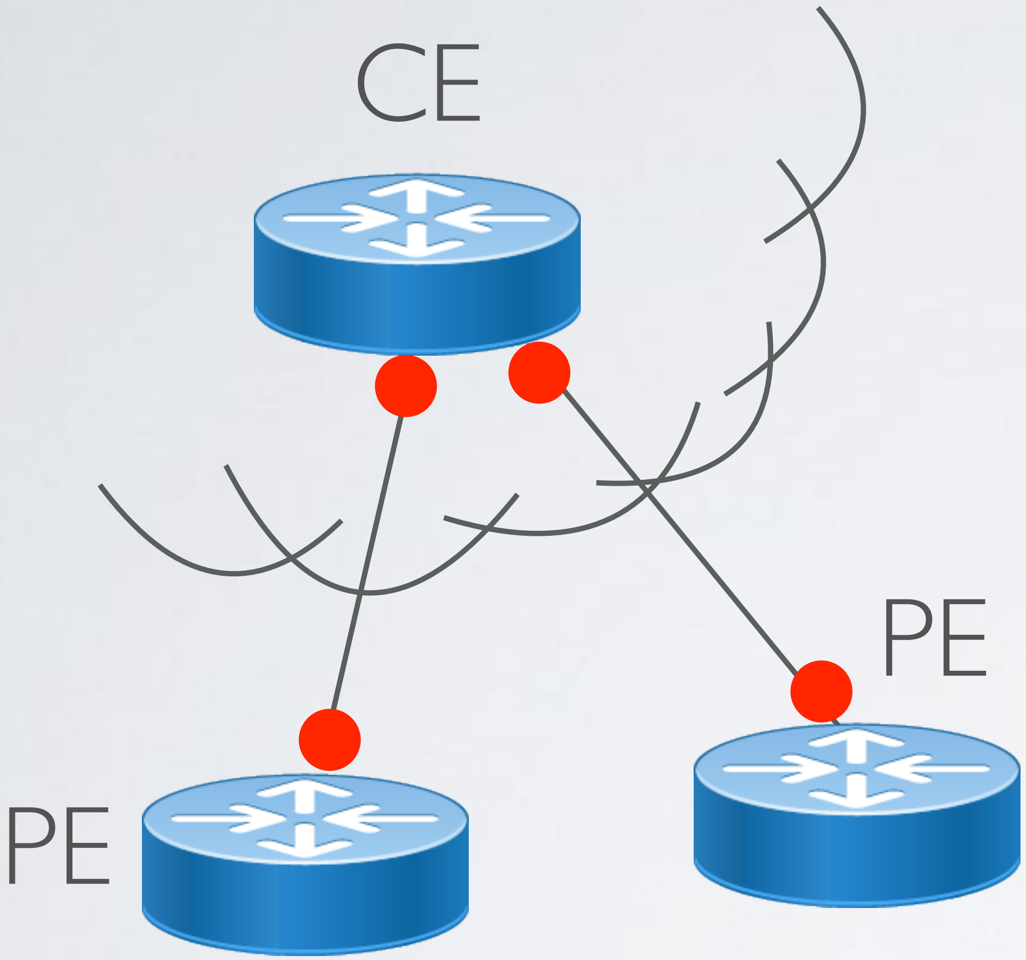

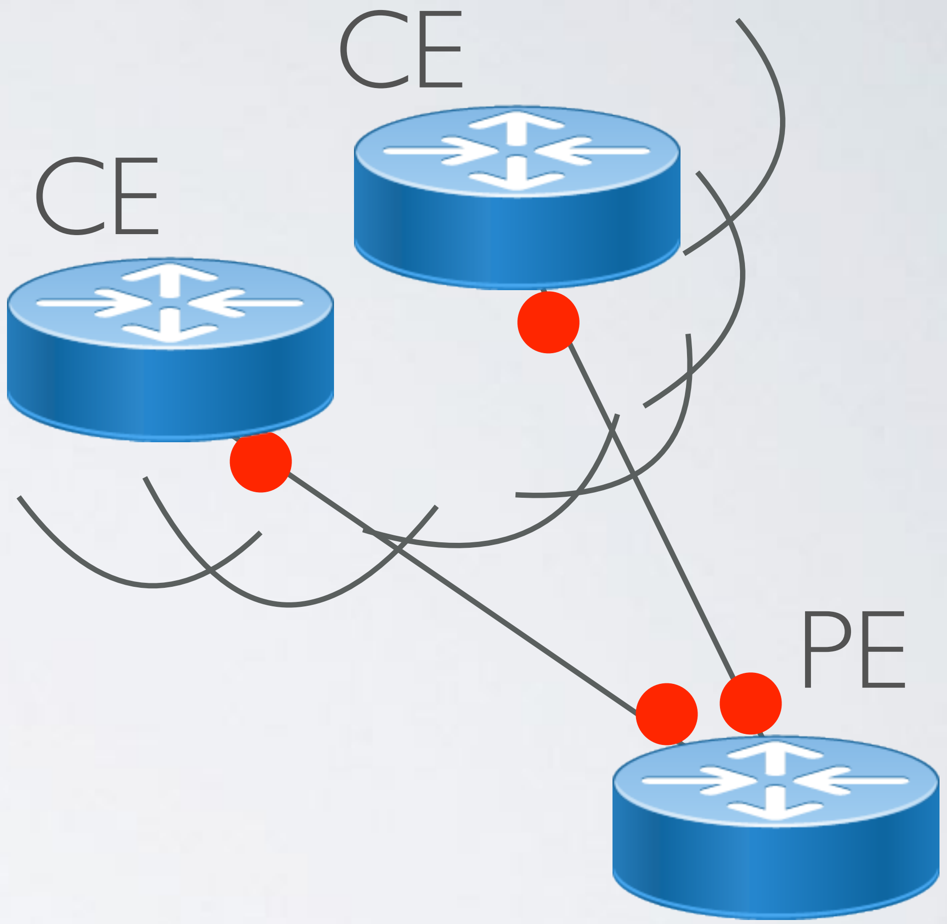

Example #A

Example #B

CE: Customer Edge
PE: Provider Edge

Internet Resilience

Where are the Single Points of Failure?

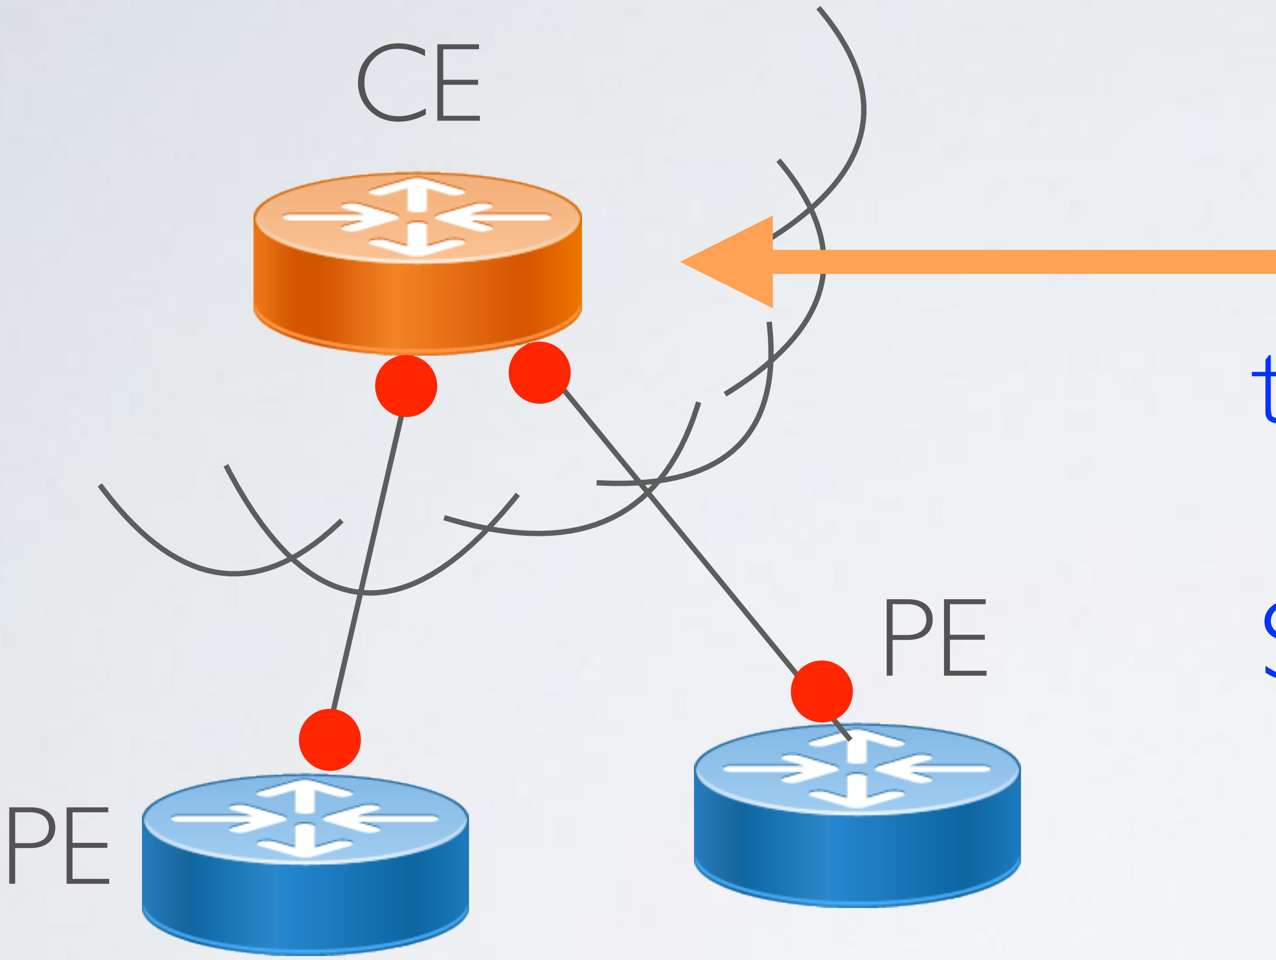

If the CE router fails,
the network is disconnected,
so the CE router is a
Single Point of Failure (SPoF)

Example #A

CE: Customer Edge
PE: Provider Edge

Internet Resilience

Where are the Single Points of Failure?

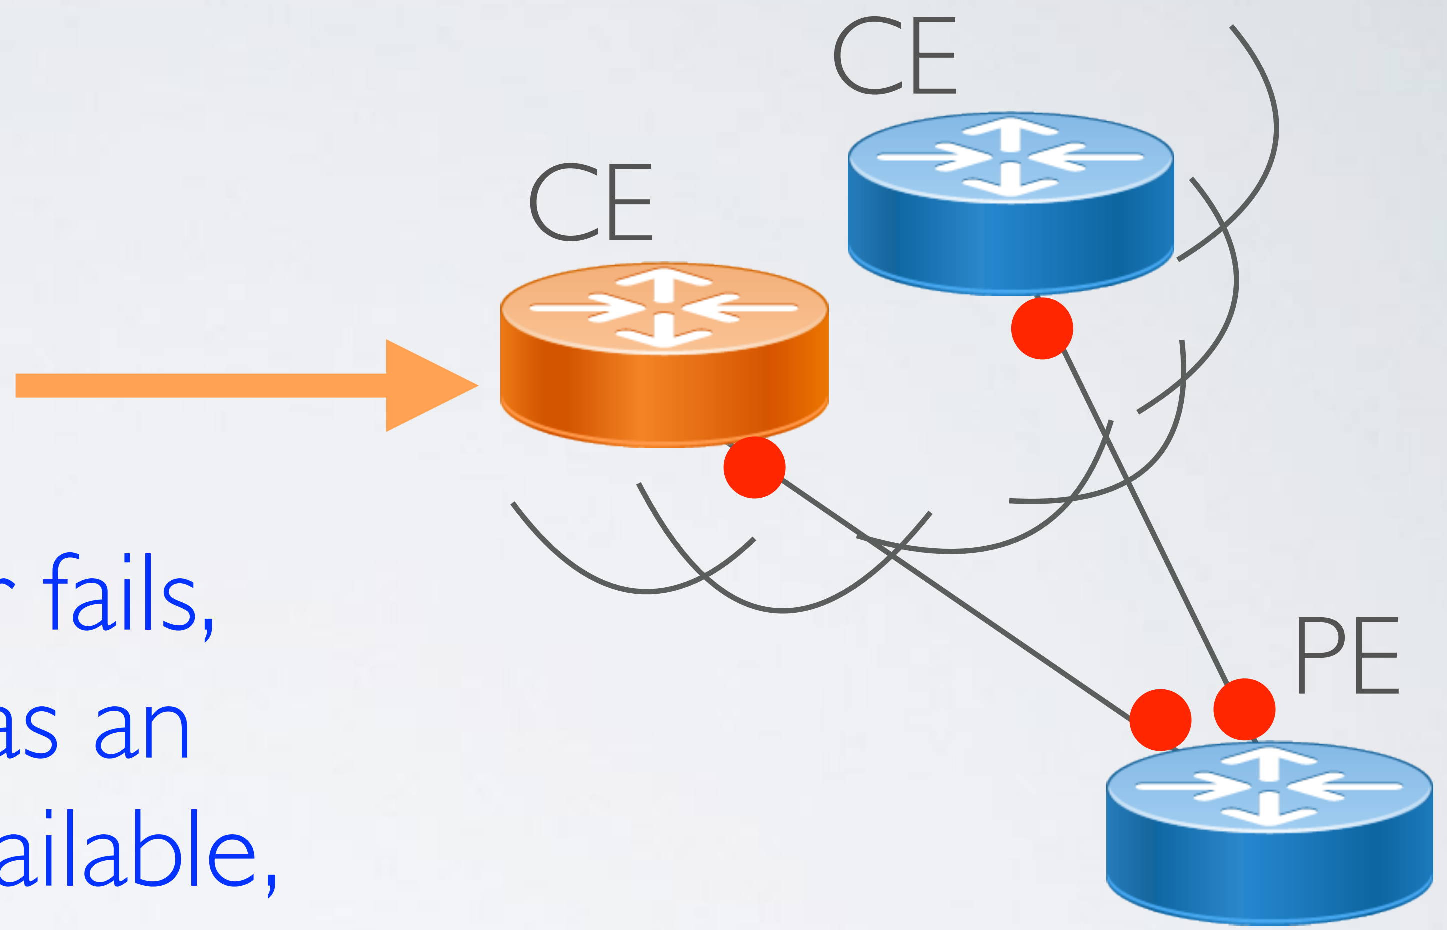

If the CE router fails,
the network has an
alternate path available,
so the CE router is NOT a
Single Point of Failure (SPoF)

Example #B

CE: Customer Edge
PE: Provider Edge

Internet Resilience

Where are the Single Points of Failure?

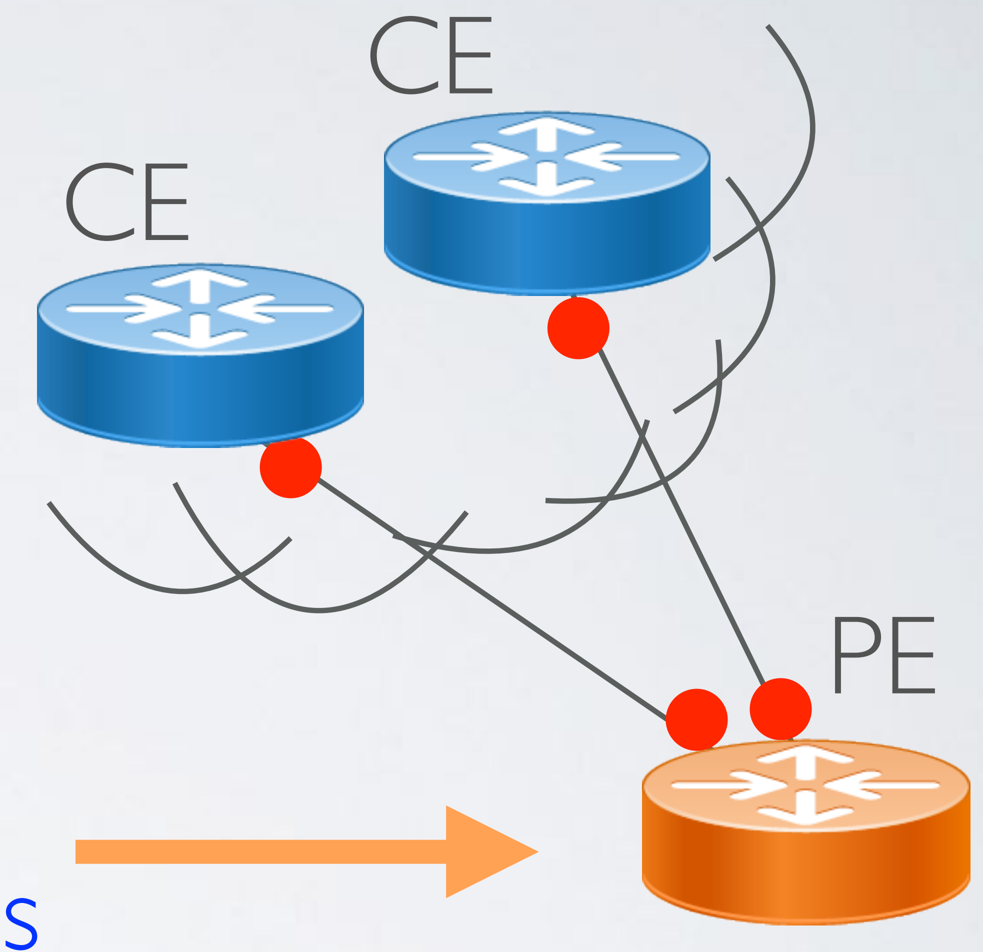

If the PE router fails,
the customer network is
disconnected, so the PE router is
a Single Point of Failure (SPoF)

Example #B

CE: Customer Edge
PE: Provider Edge

Challenges in topology analysis

- Prior approaches analyzed static AS-level and router-level topology graphs,
  - e.g.: Nature 2000
- Important AS-level and router-level **topology might be invisible to measurement**, such as backup paths,
  - e.g: INFOCOM 2002
- A router that appears to be central to a network's connectivity might not be
  - e.g.: AMS 2009

What we did

Large-scale (**Internet-wide**) longitudinal (**2.5 years**) measurement study to characterize prevalence of Single Points of Failure (**SPoF**):

1. Efficiently inferred **IPv6 router outage time windows**
2. **Associated** routers with **IPv6 BGP prefixes**
3. **Correlated** router outages with **BGP control plane**
4. **Correlated** router outages with **data plane**
5. **Validated inferences** of SPoF with network operators

What we did

Identified IPv6 router interfaces from traceroute

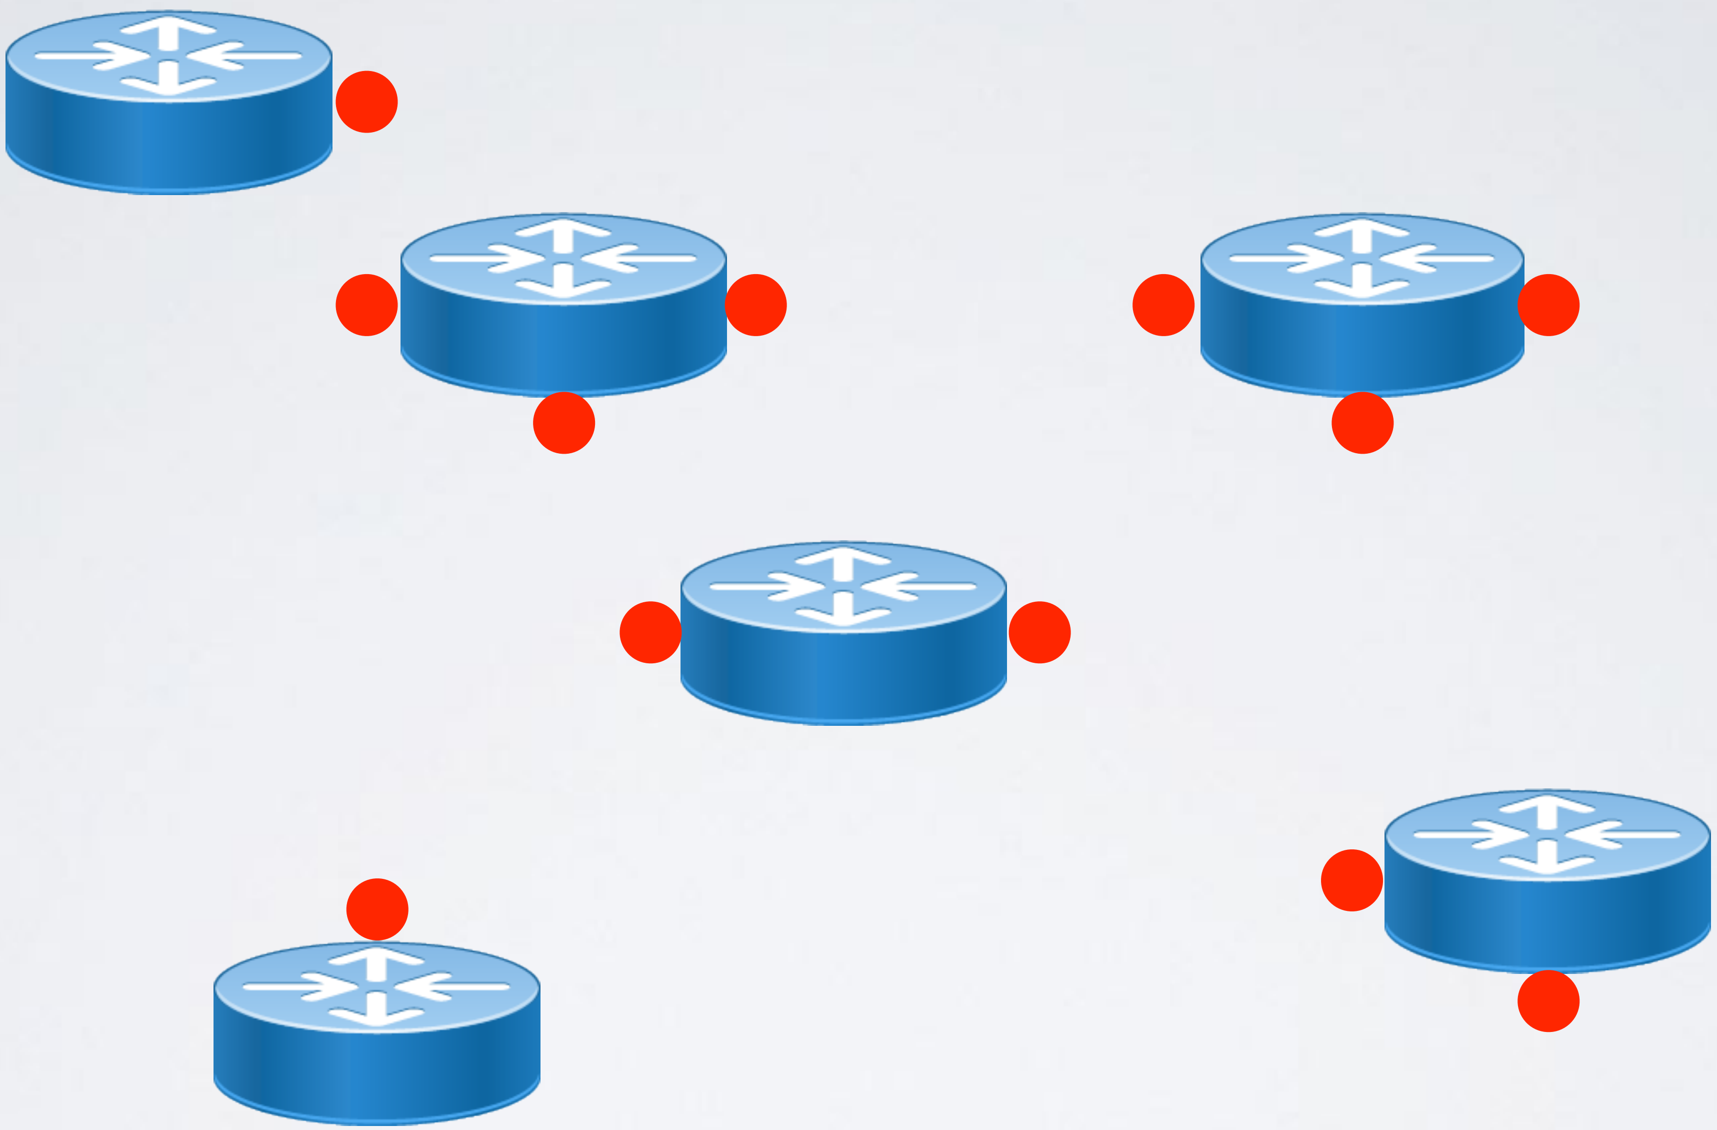

83K to 2.4M interfaces from CAIDA's Archipelago traceroute measurements

What we did

probed router interfaces to infer outage windows

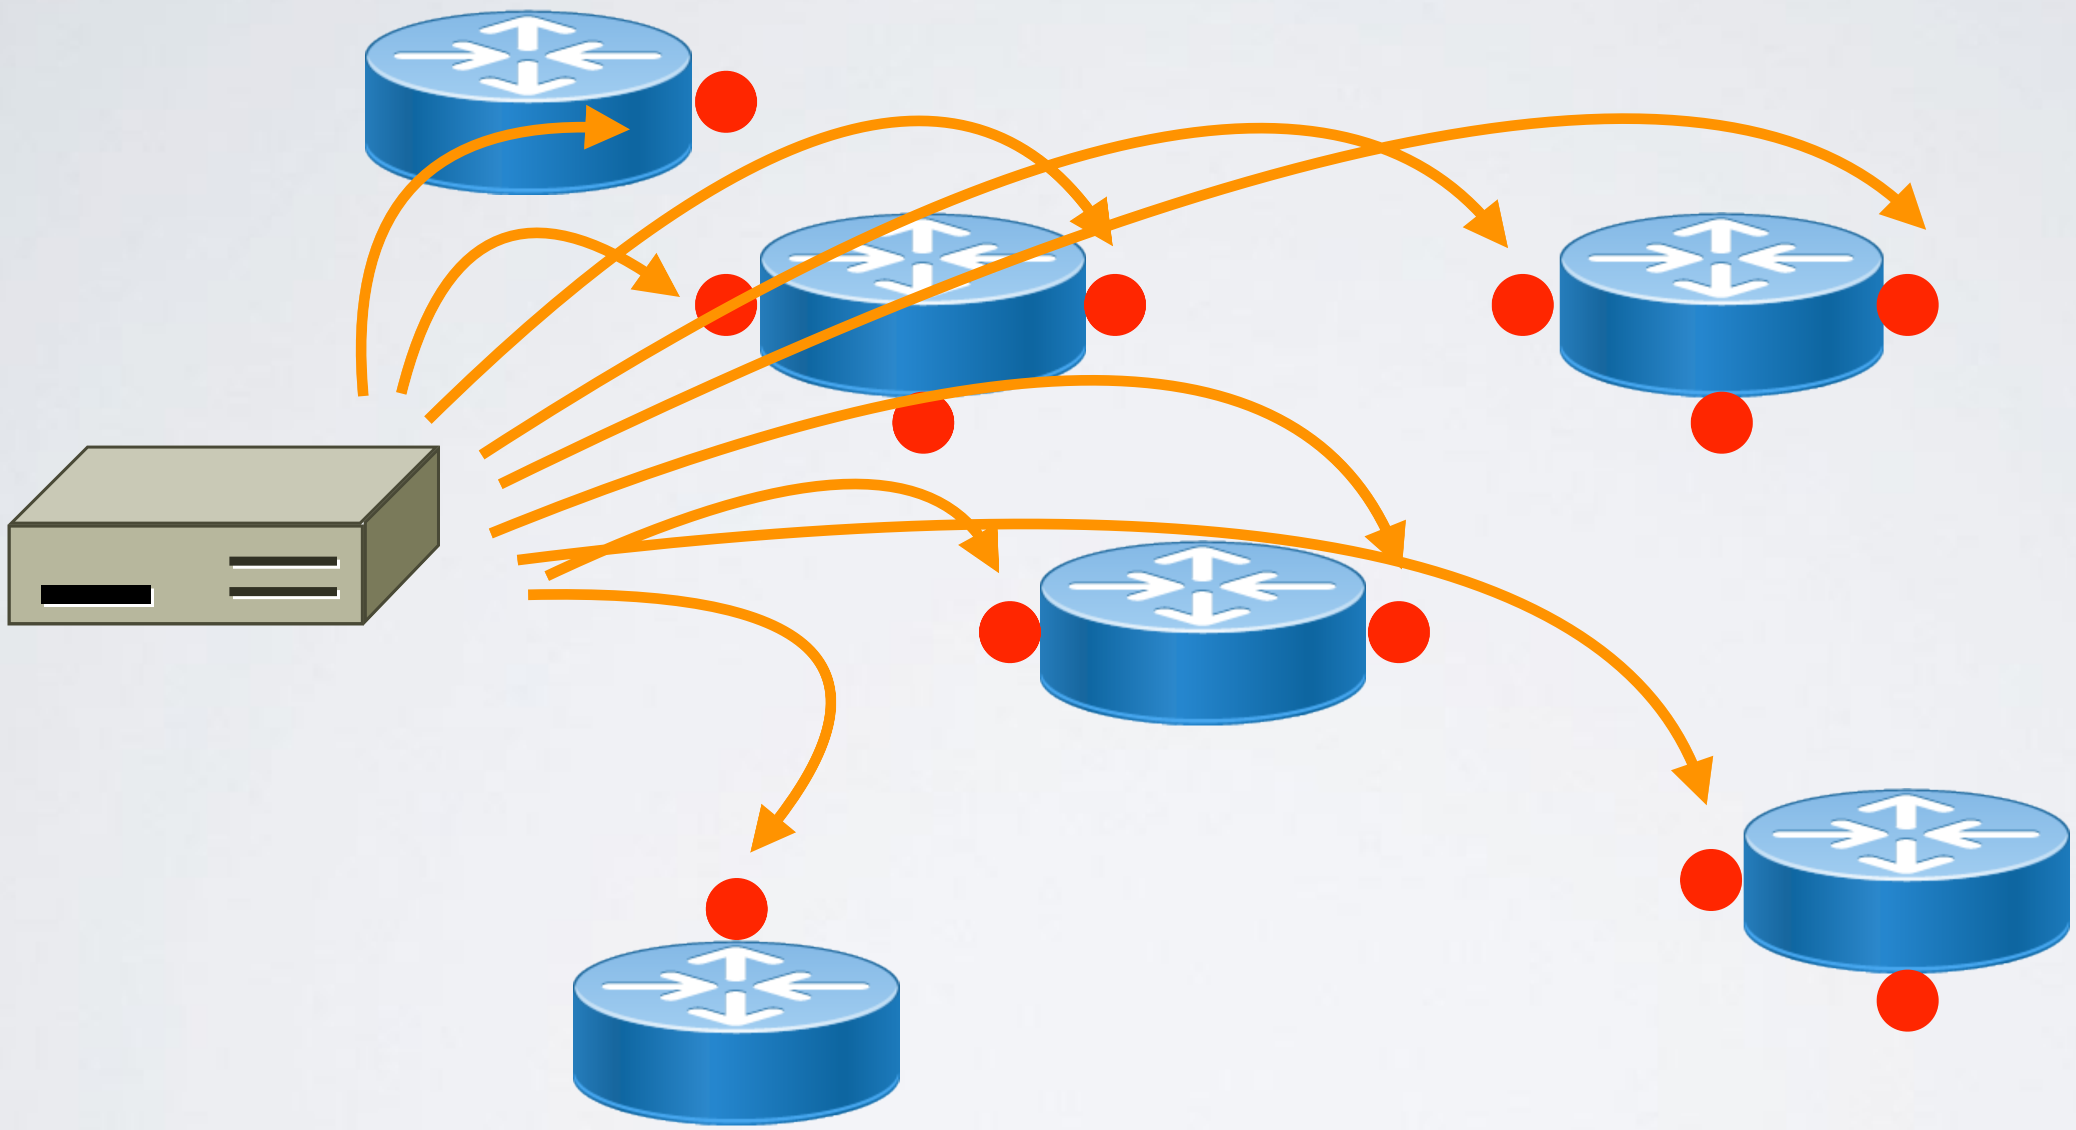

We used a single vantage point located at CAIDA, UC San Diego for the duration of this study

What we did

Central counter: **9290**

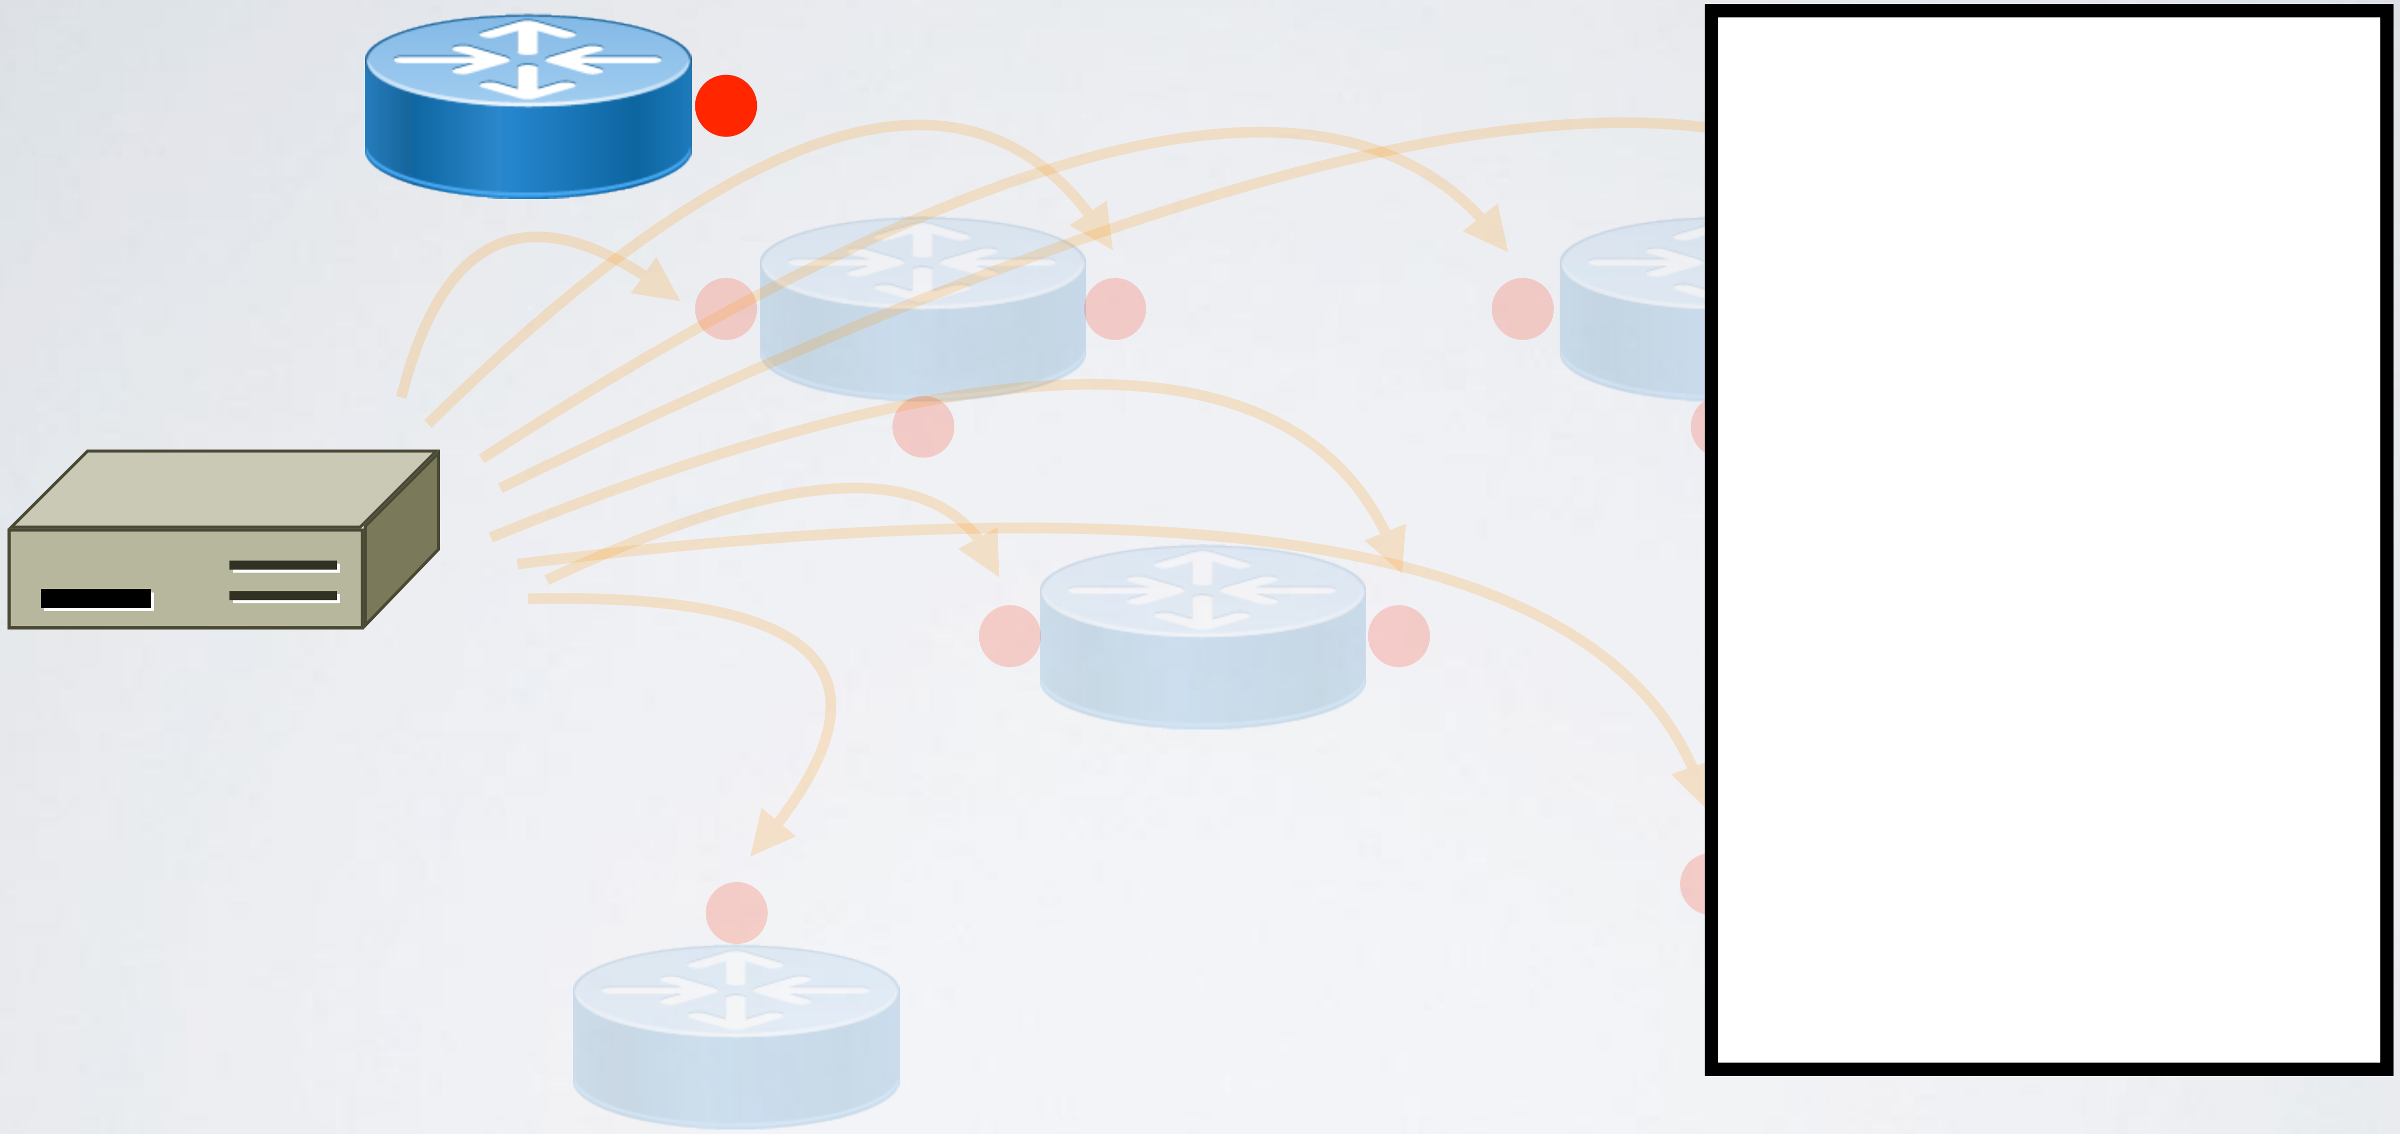

What we did

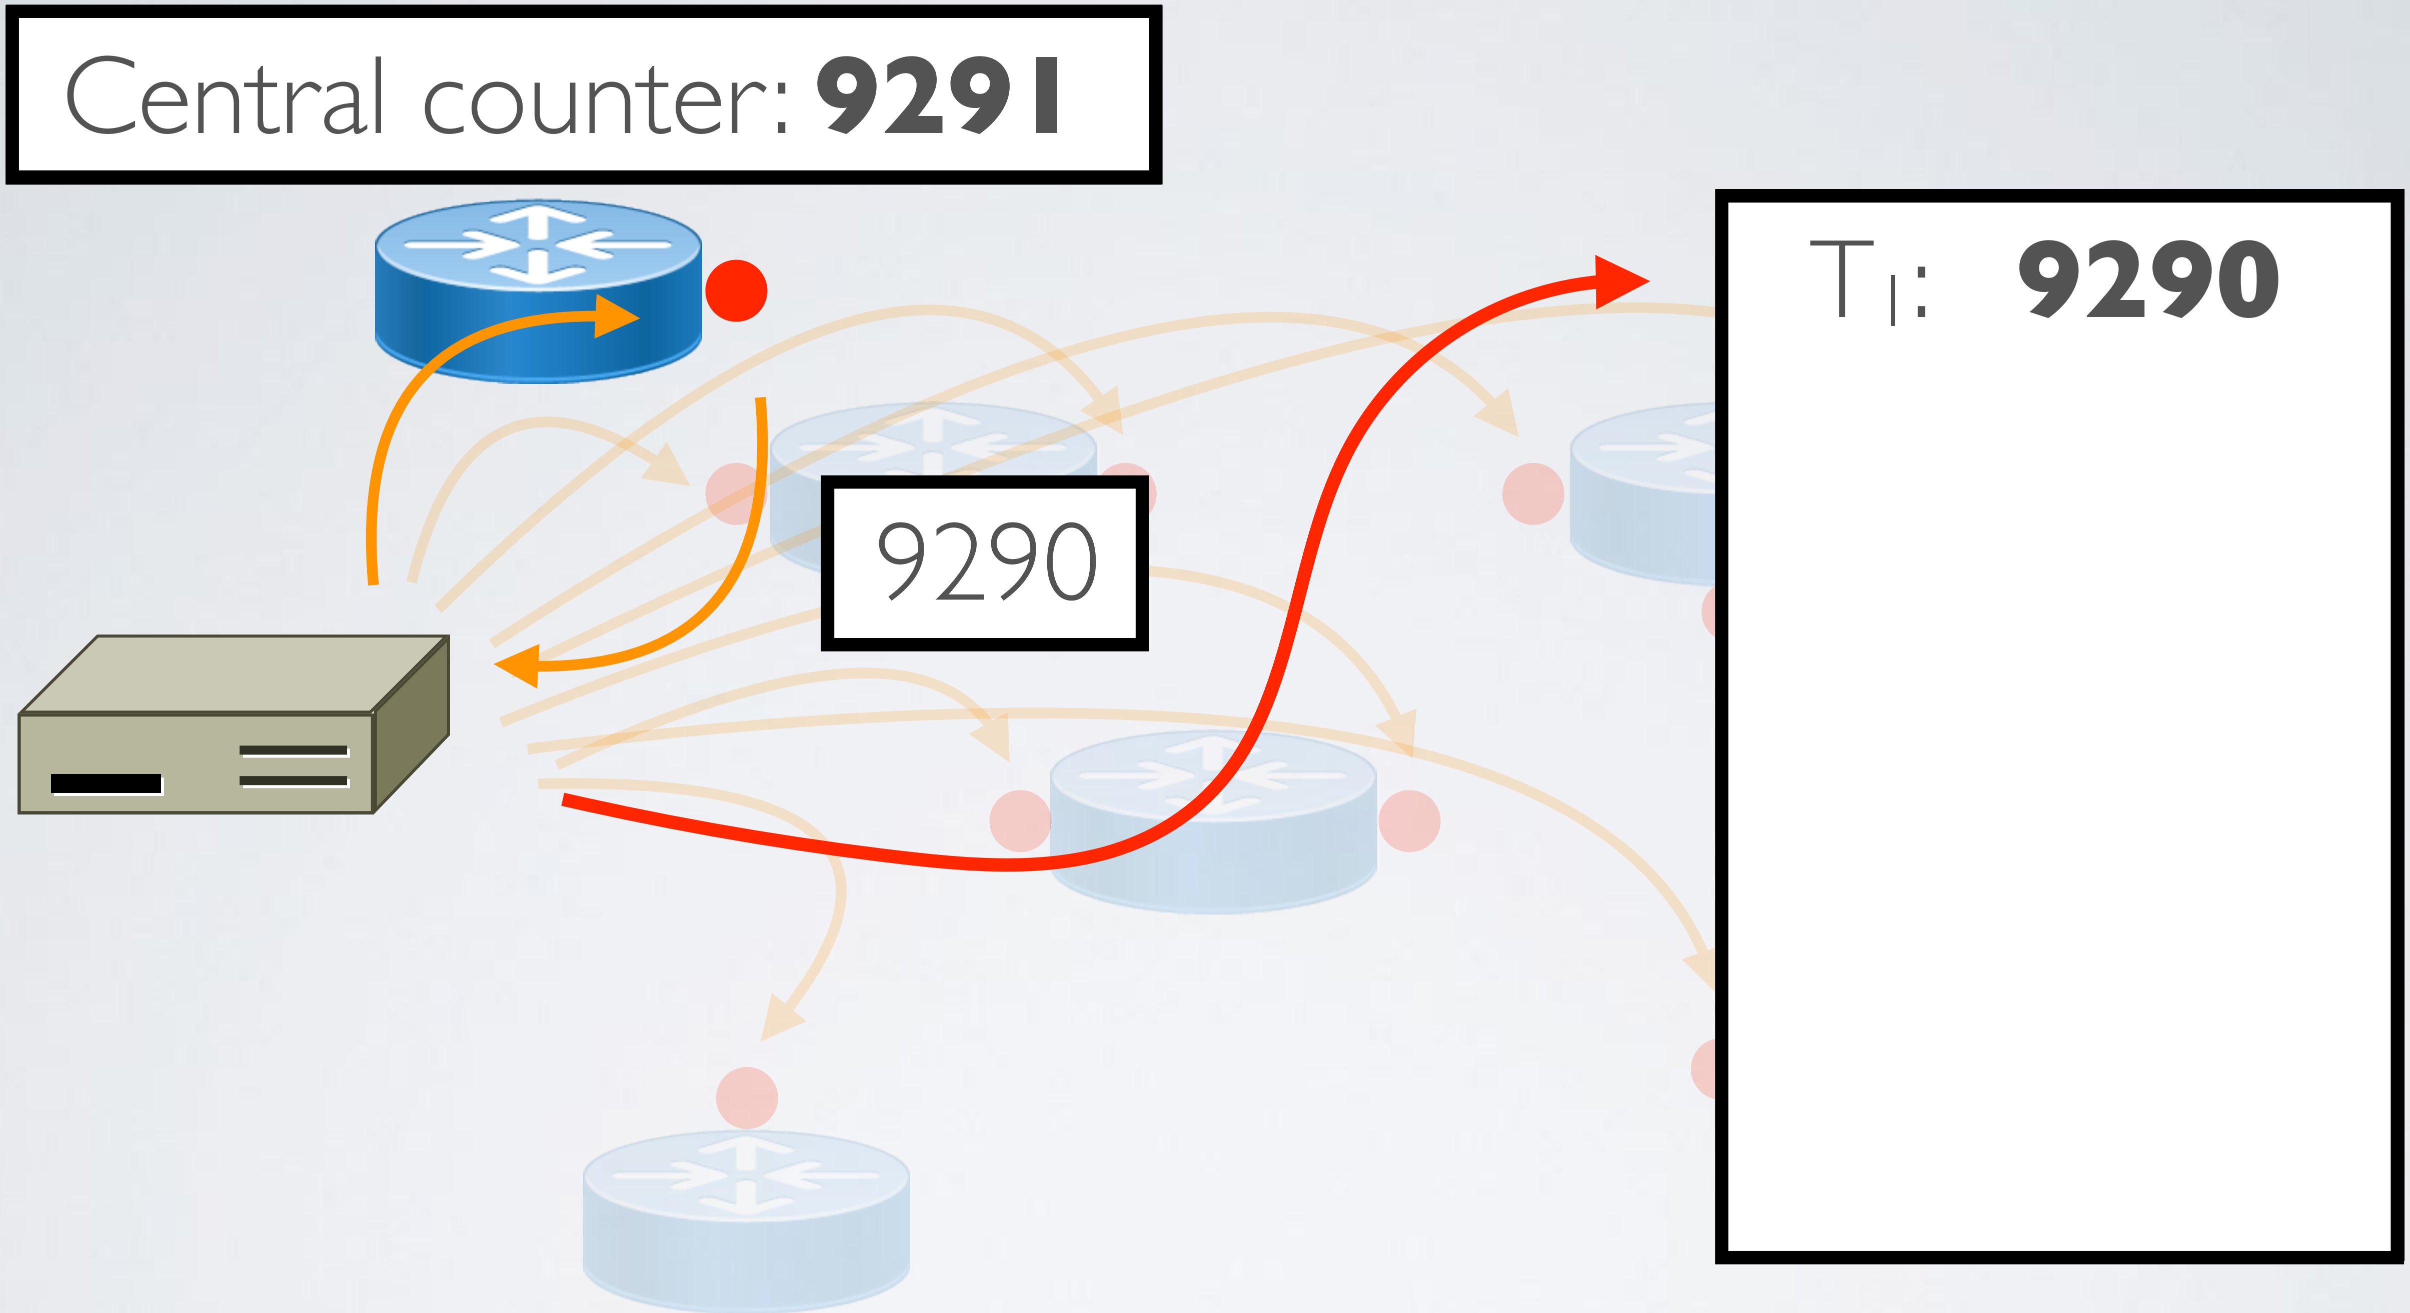

What we did

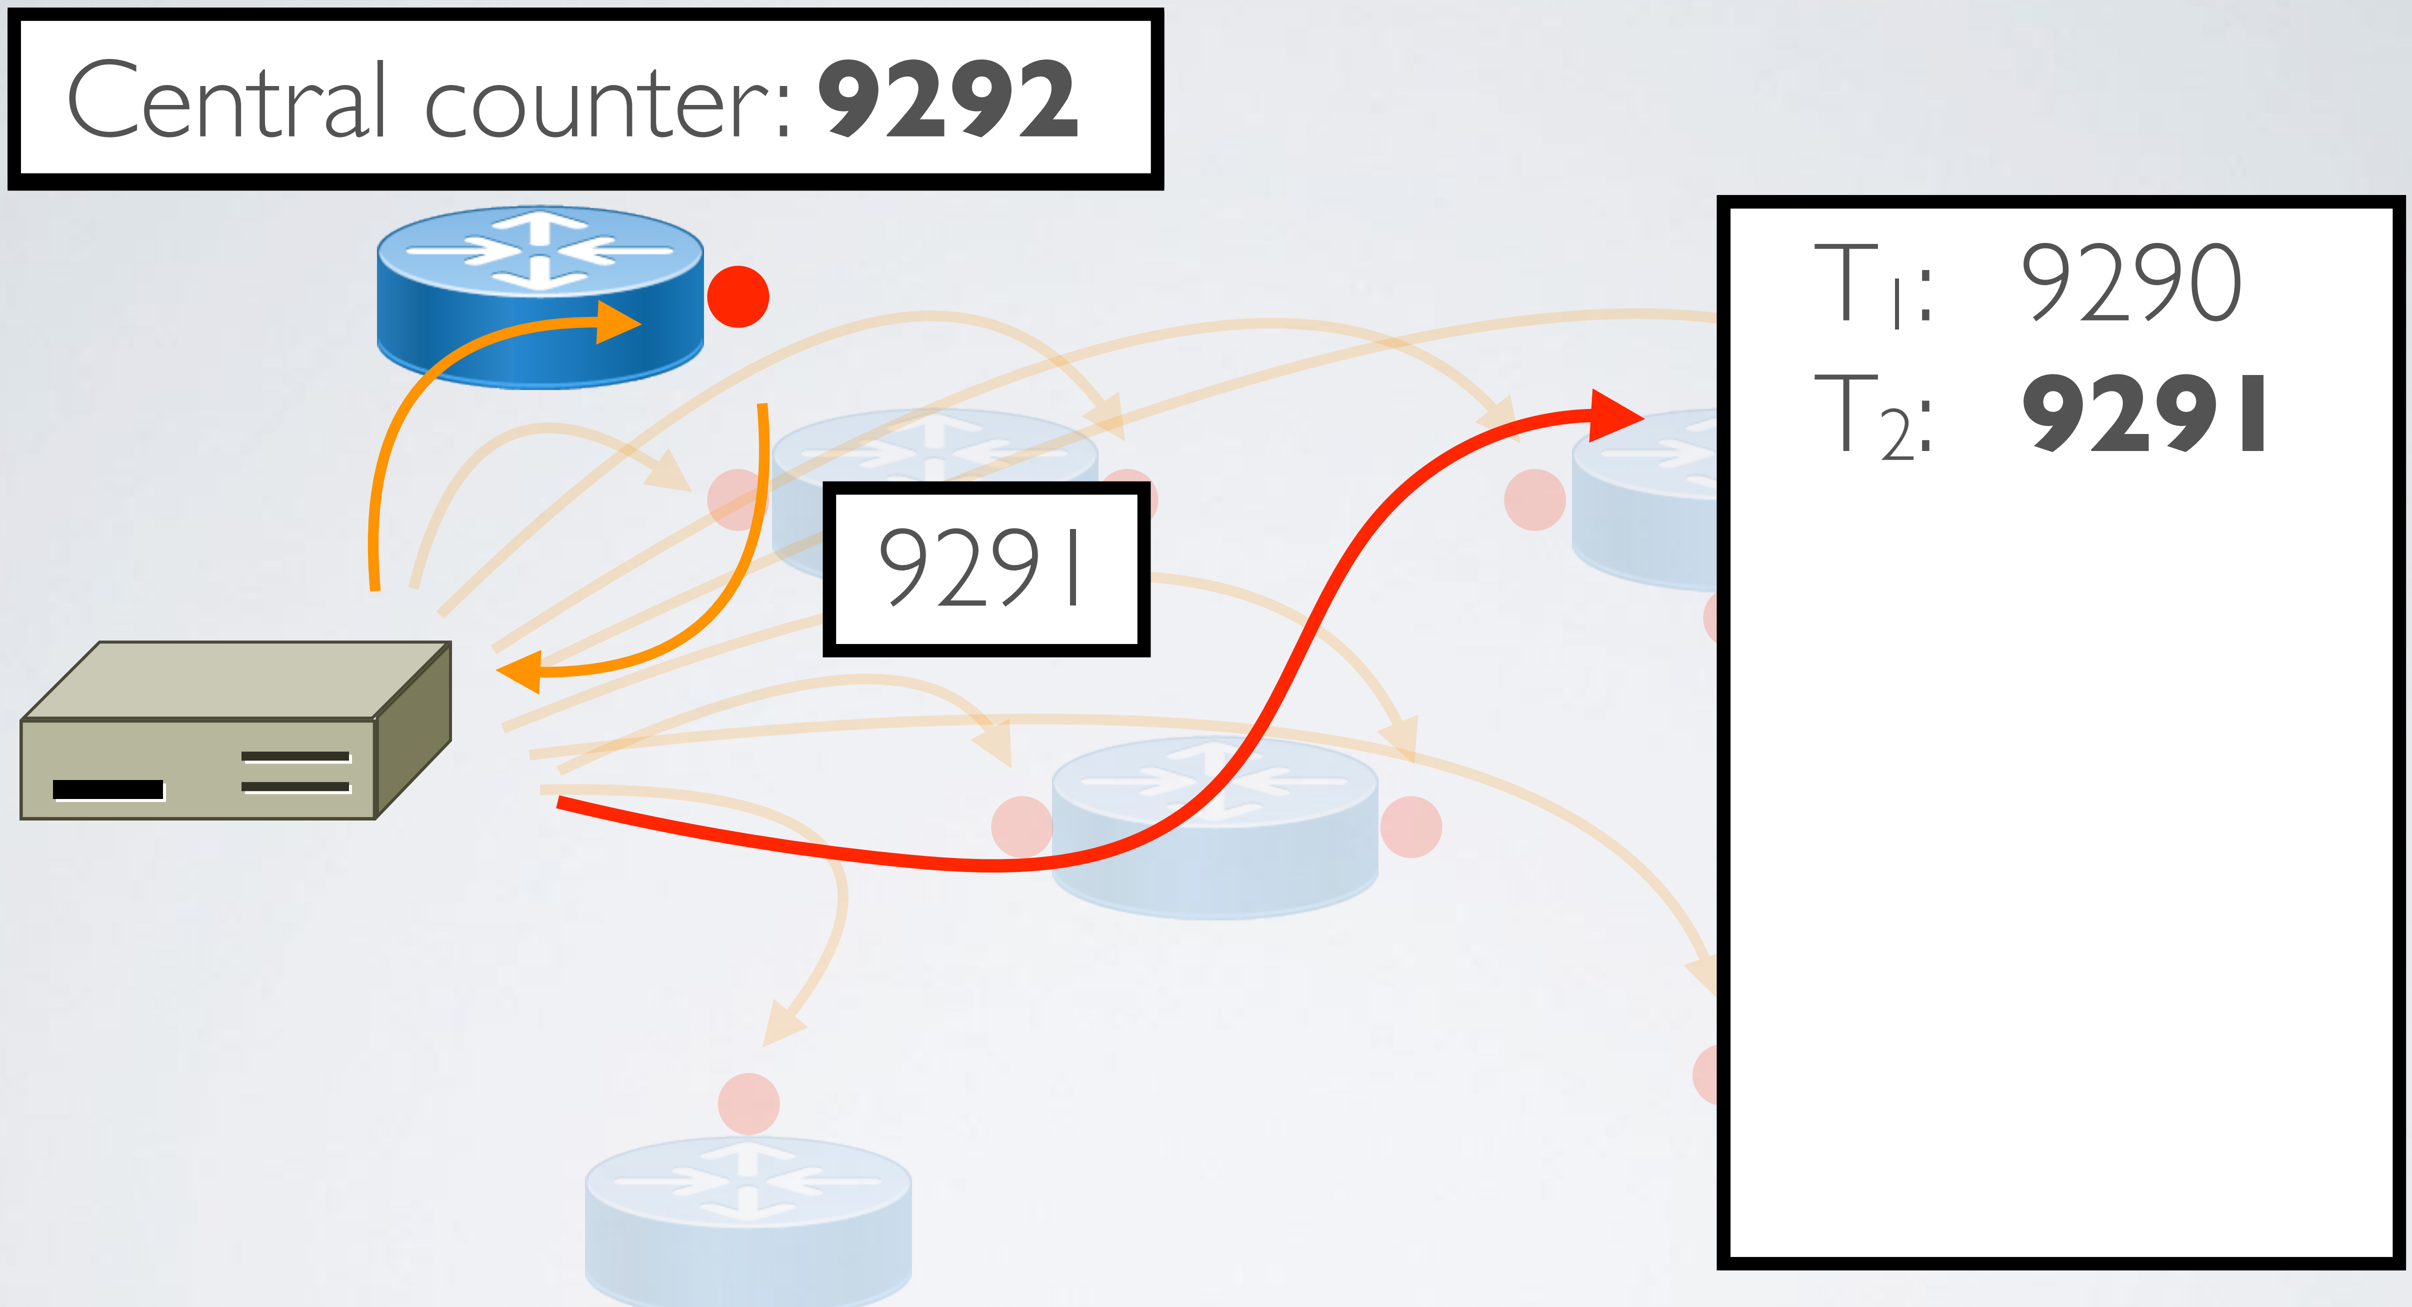

What we did

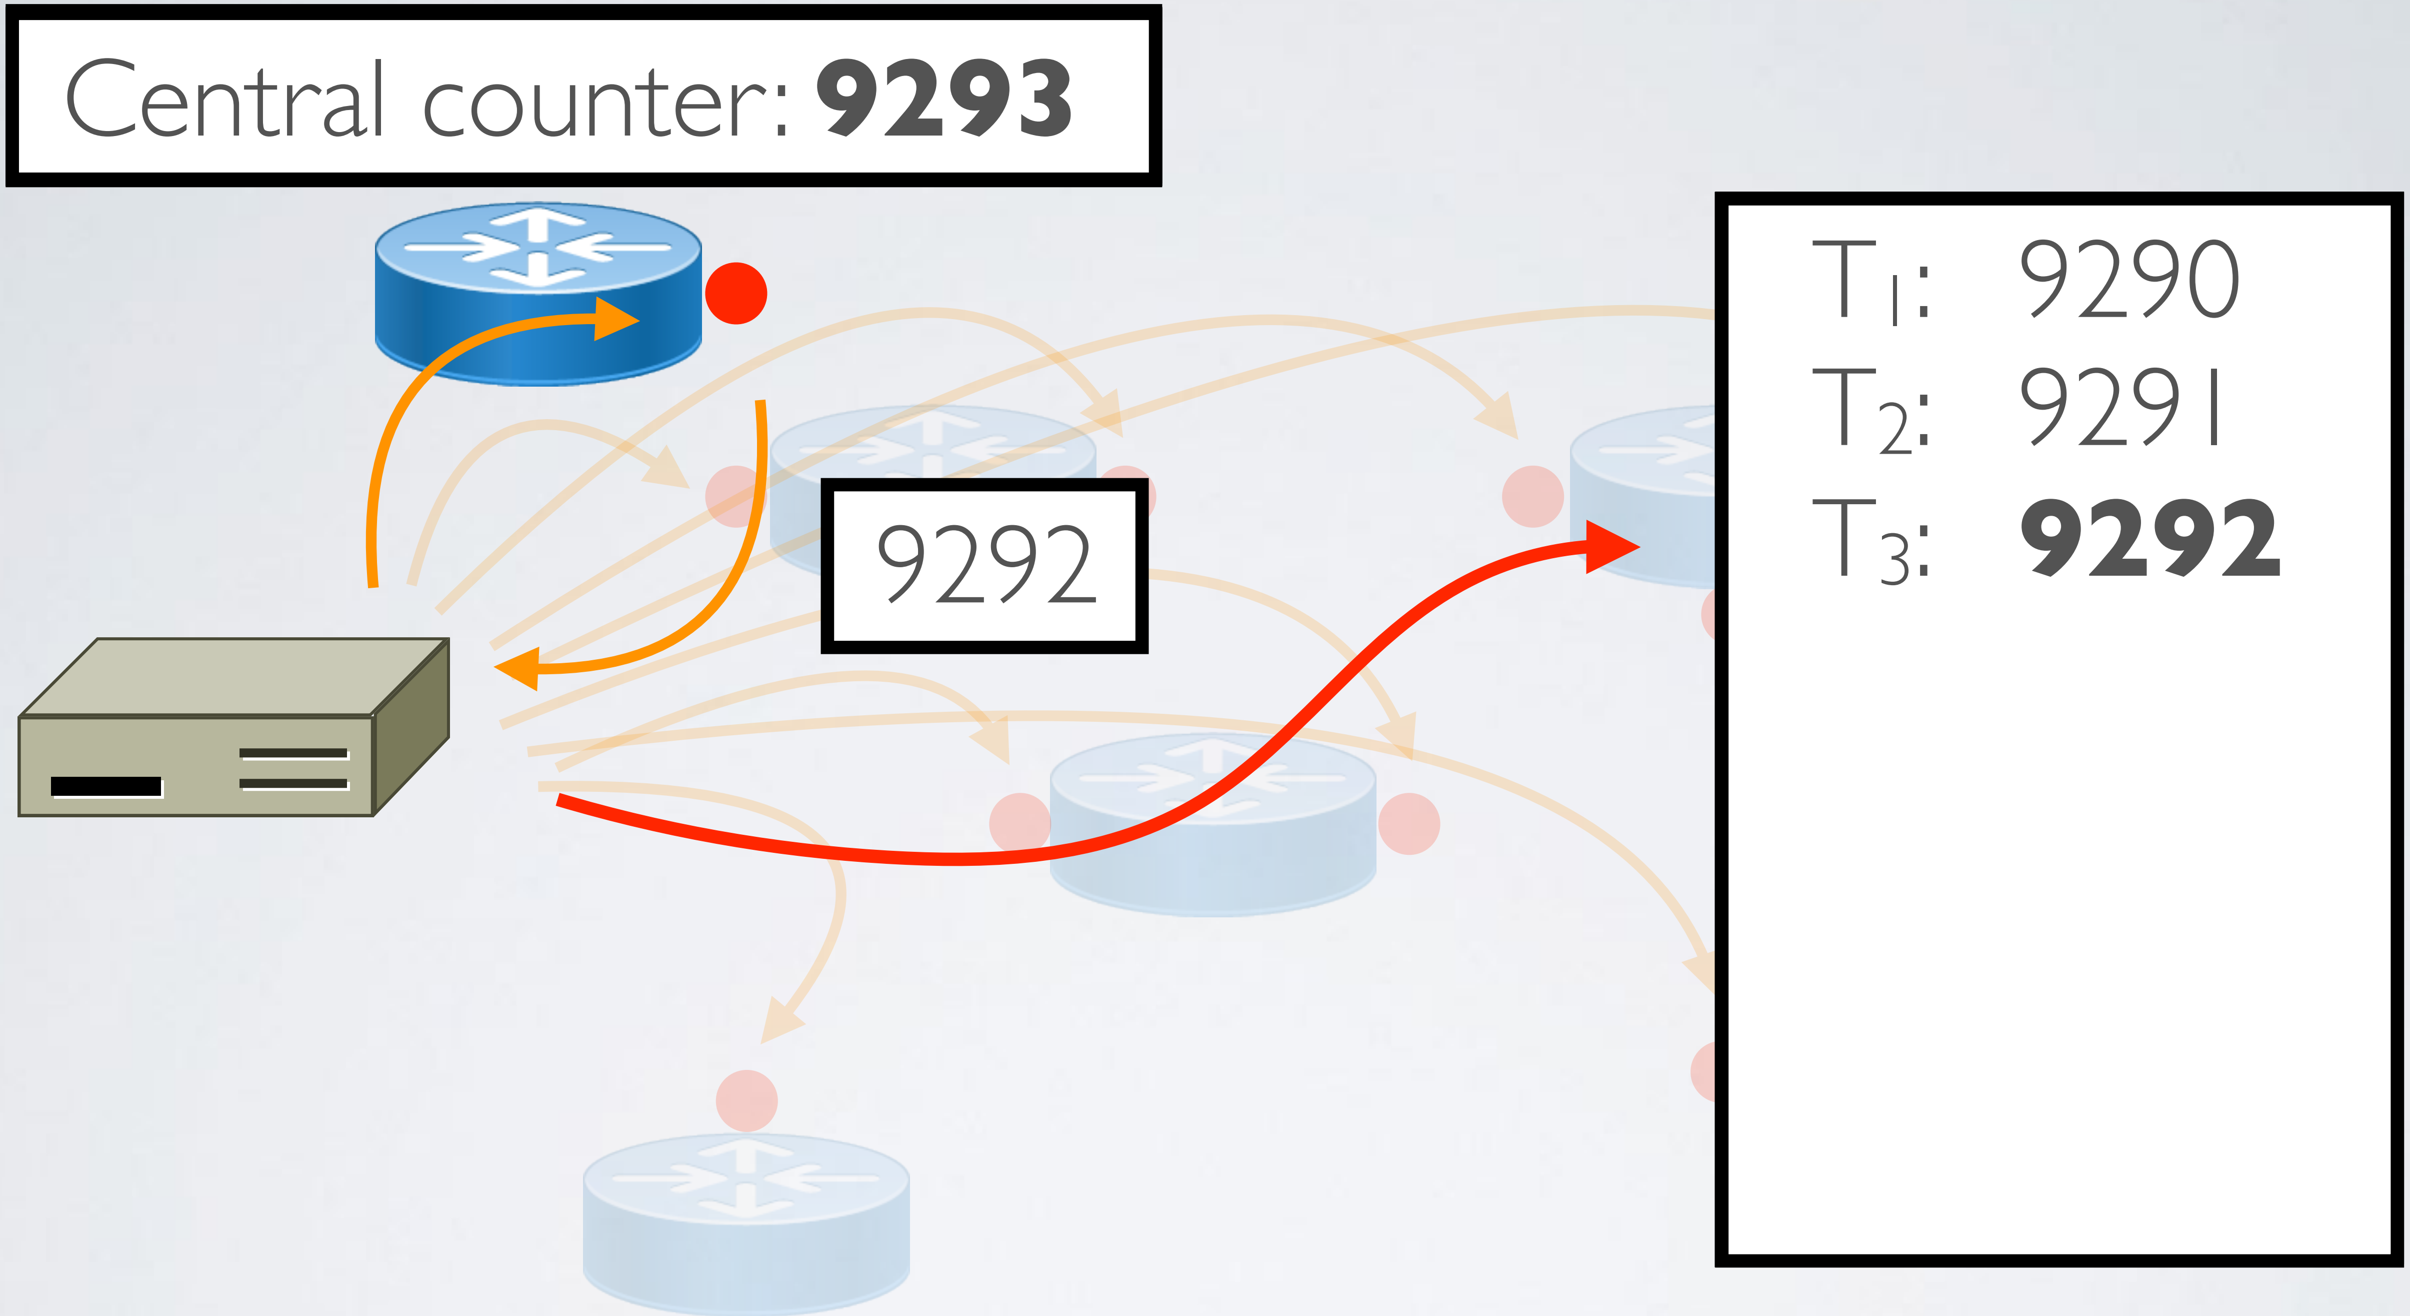

What we did

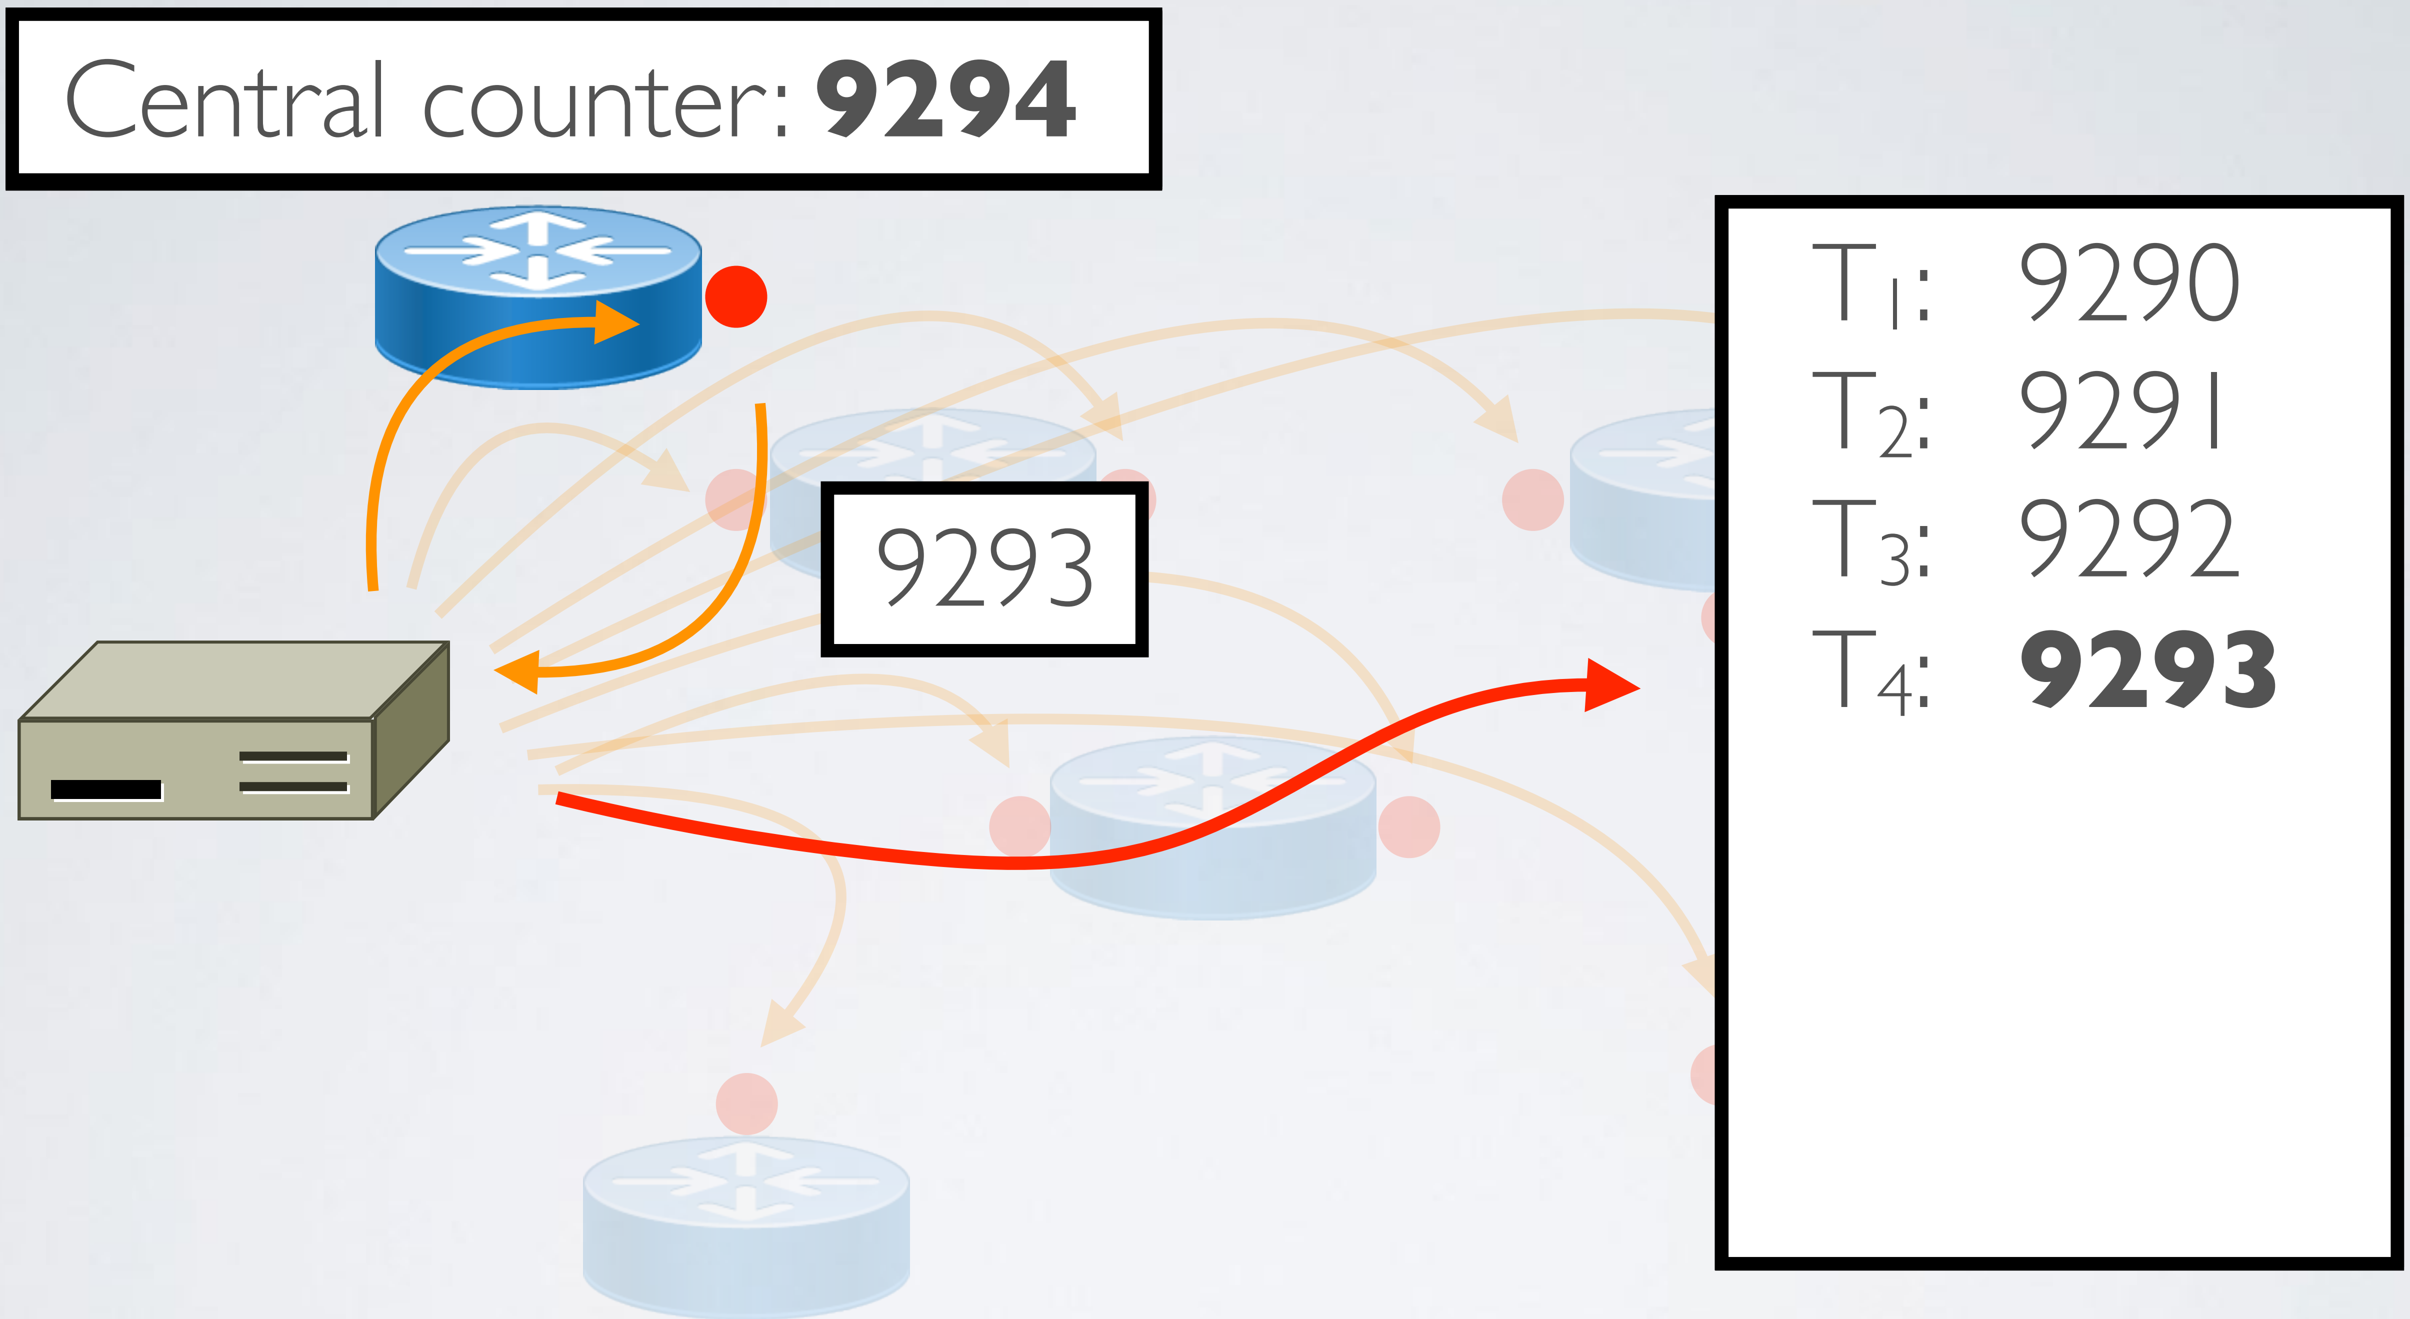

What we did

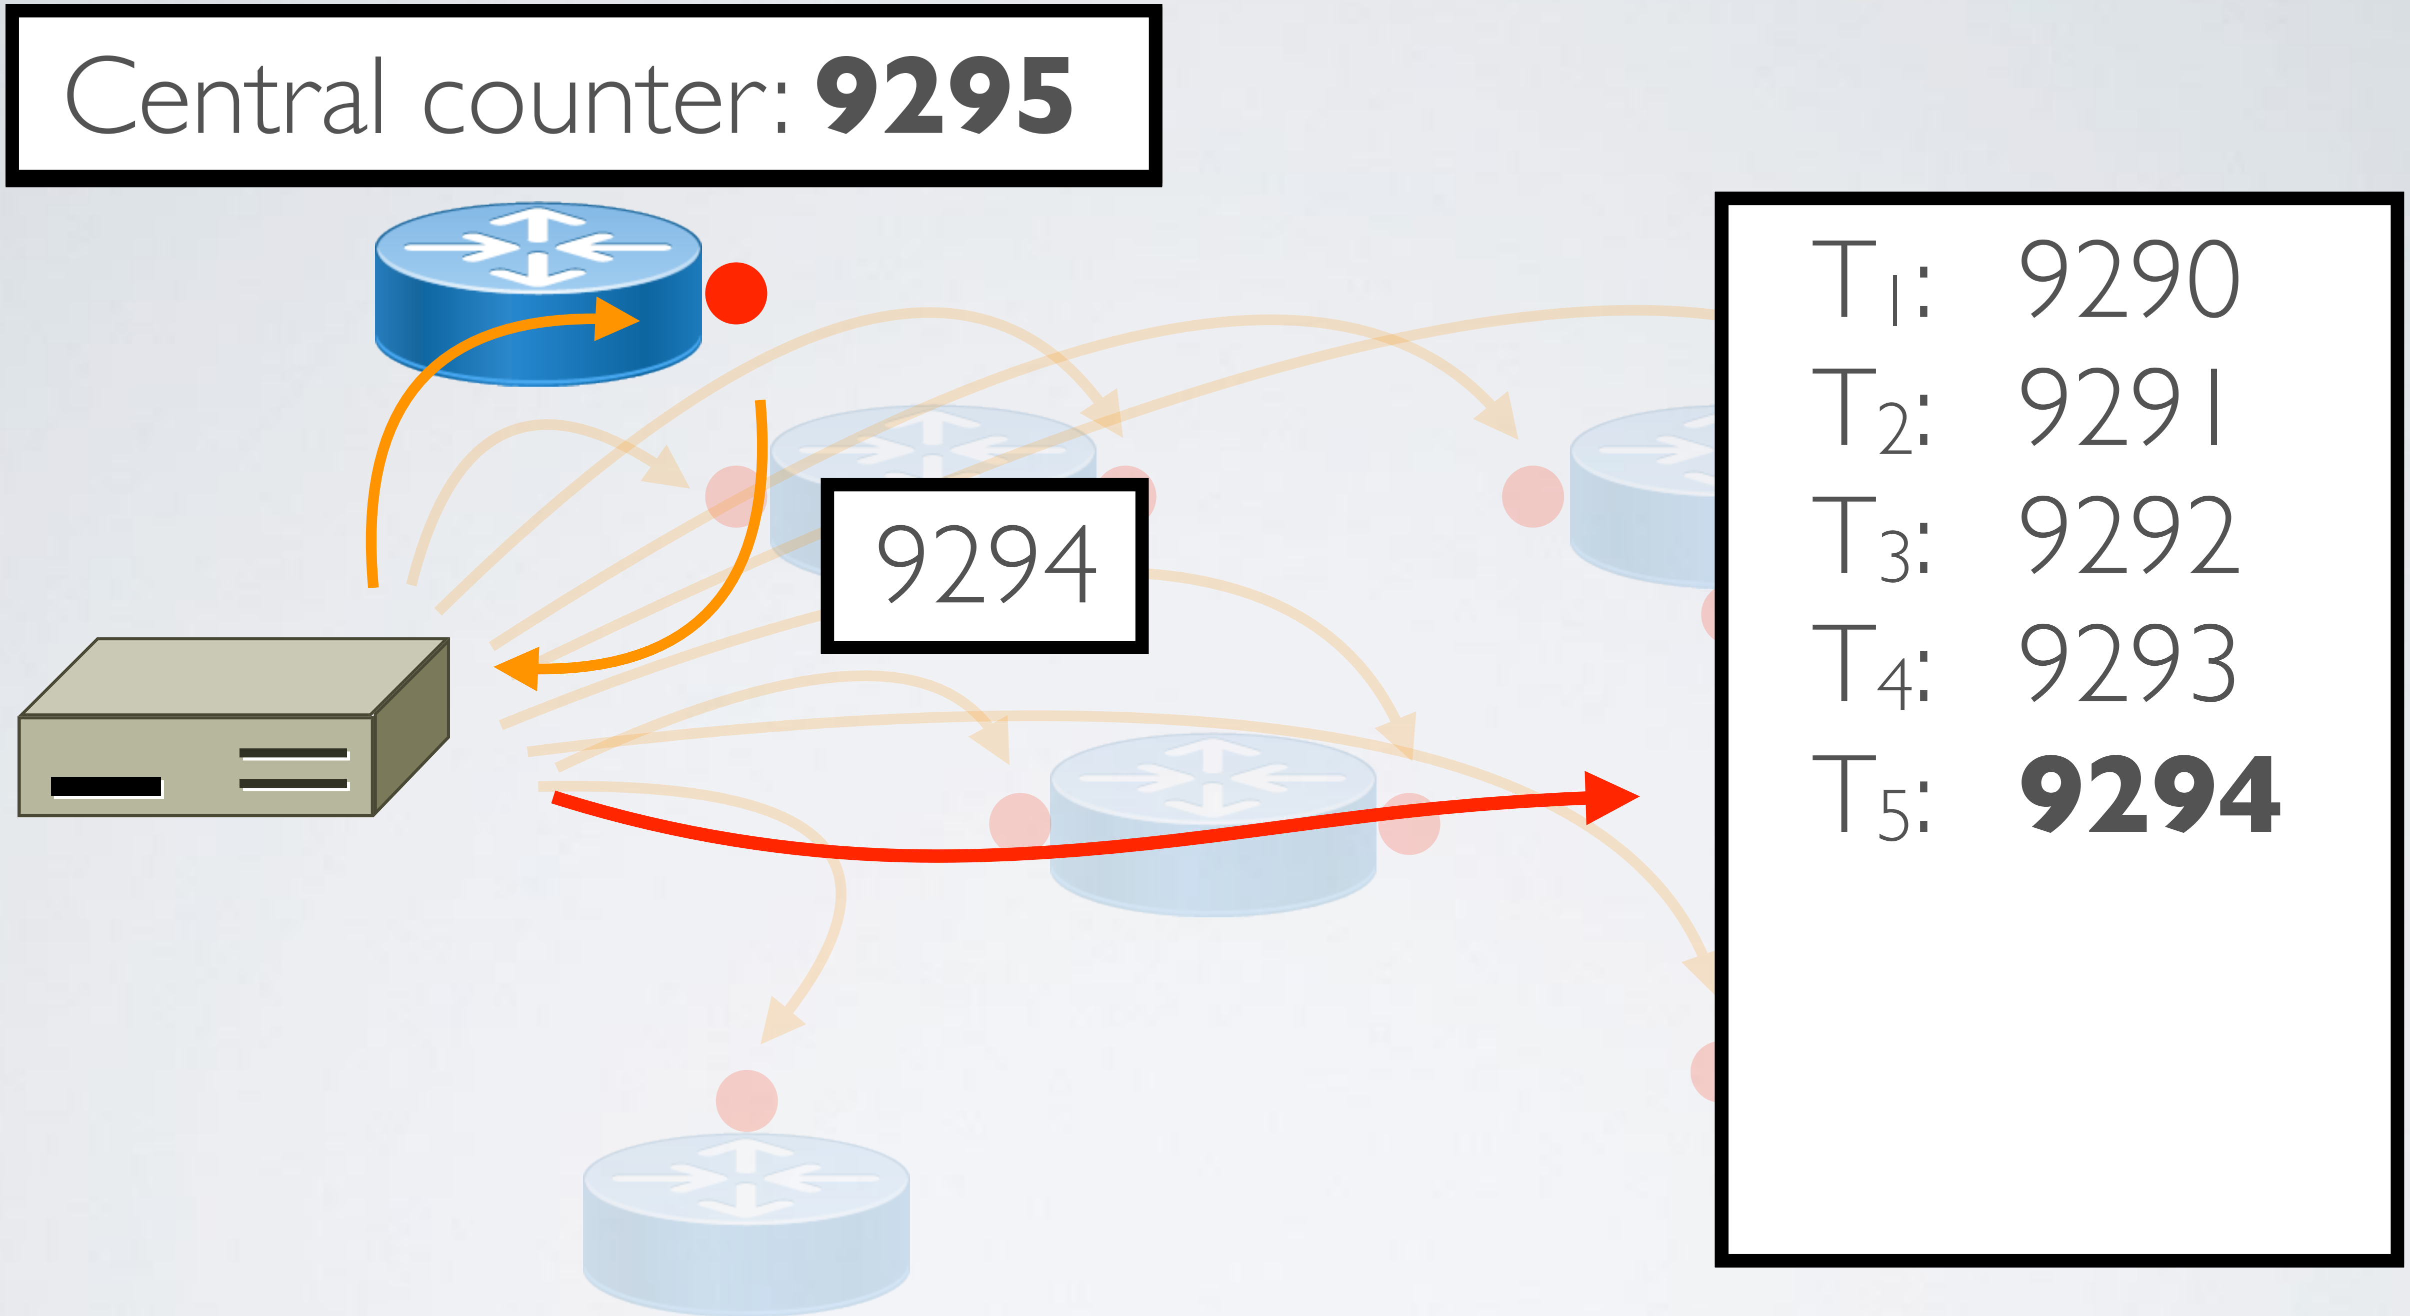

Central counter: **9295**
9294
T_1: 9290
T_2: 9291
T_3: 9292
T_4: 9293
T_5: **9294**

What we did

Central counter: **1**

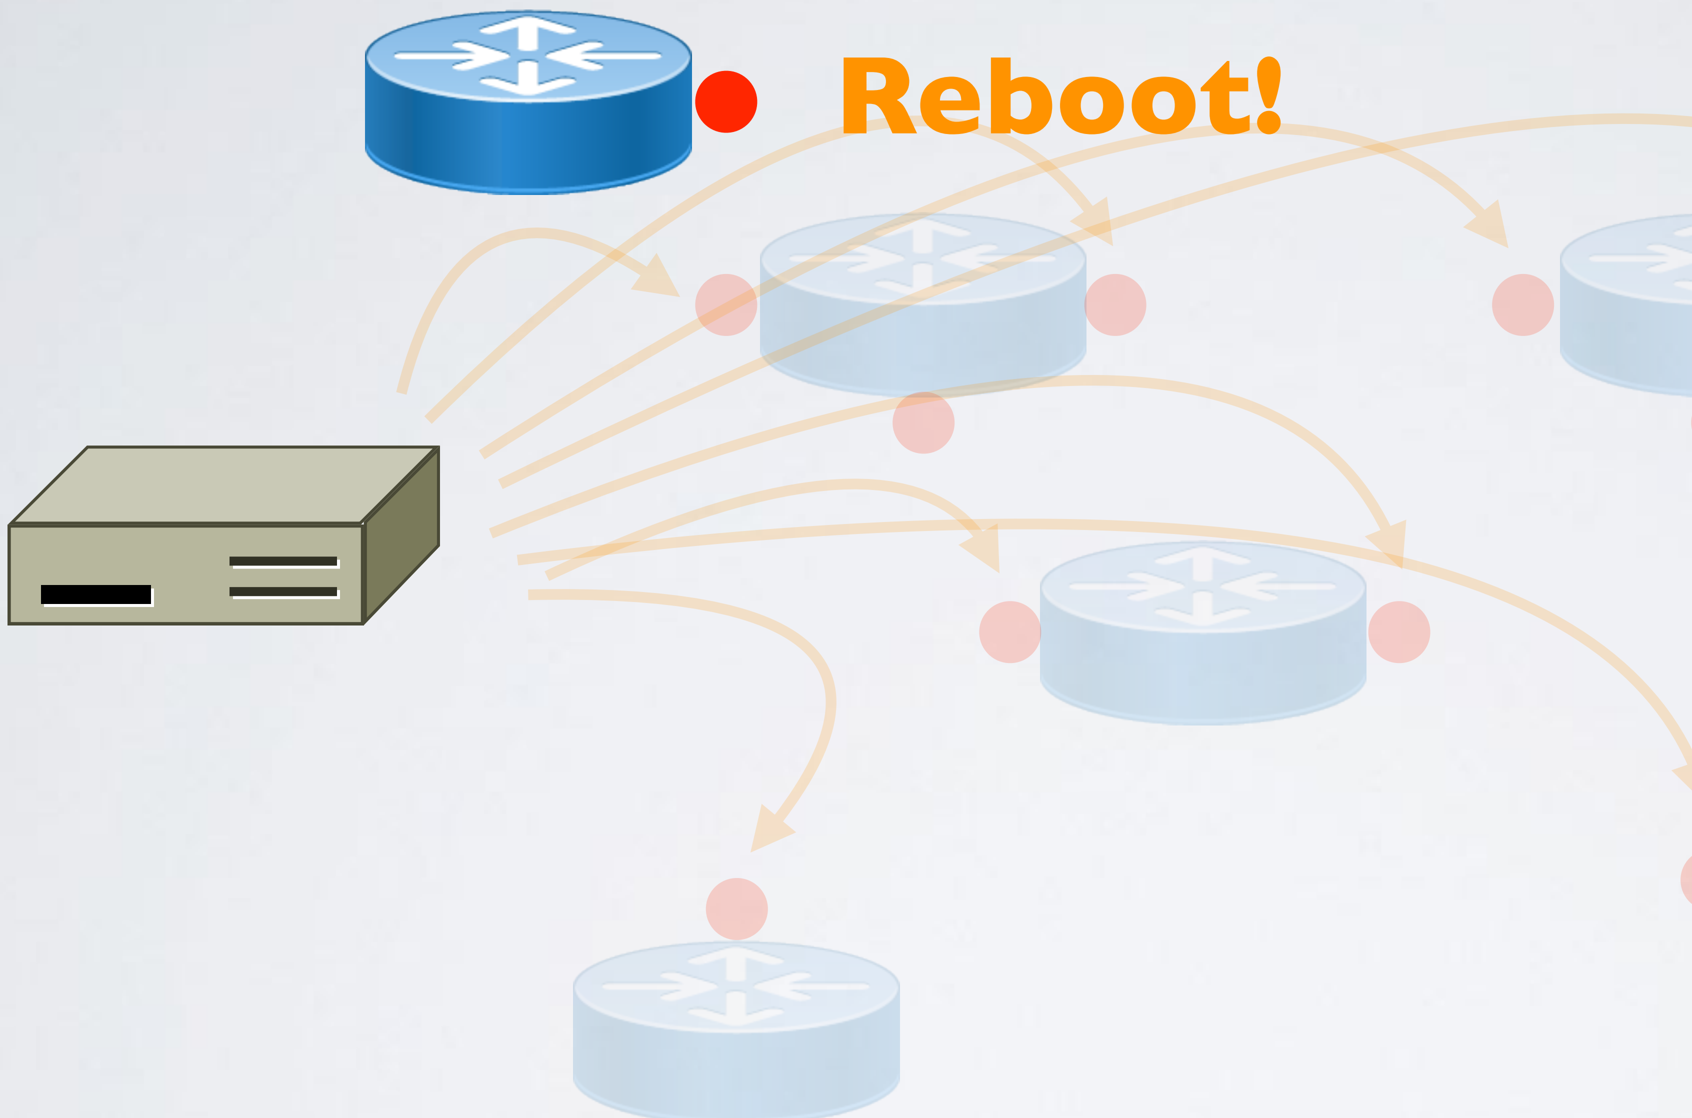

T_1: 9290
T_2: 9291
T_3: 9292
T_4: 9293
T_5: **9294**

What we did

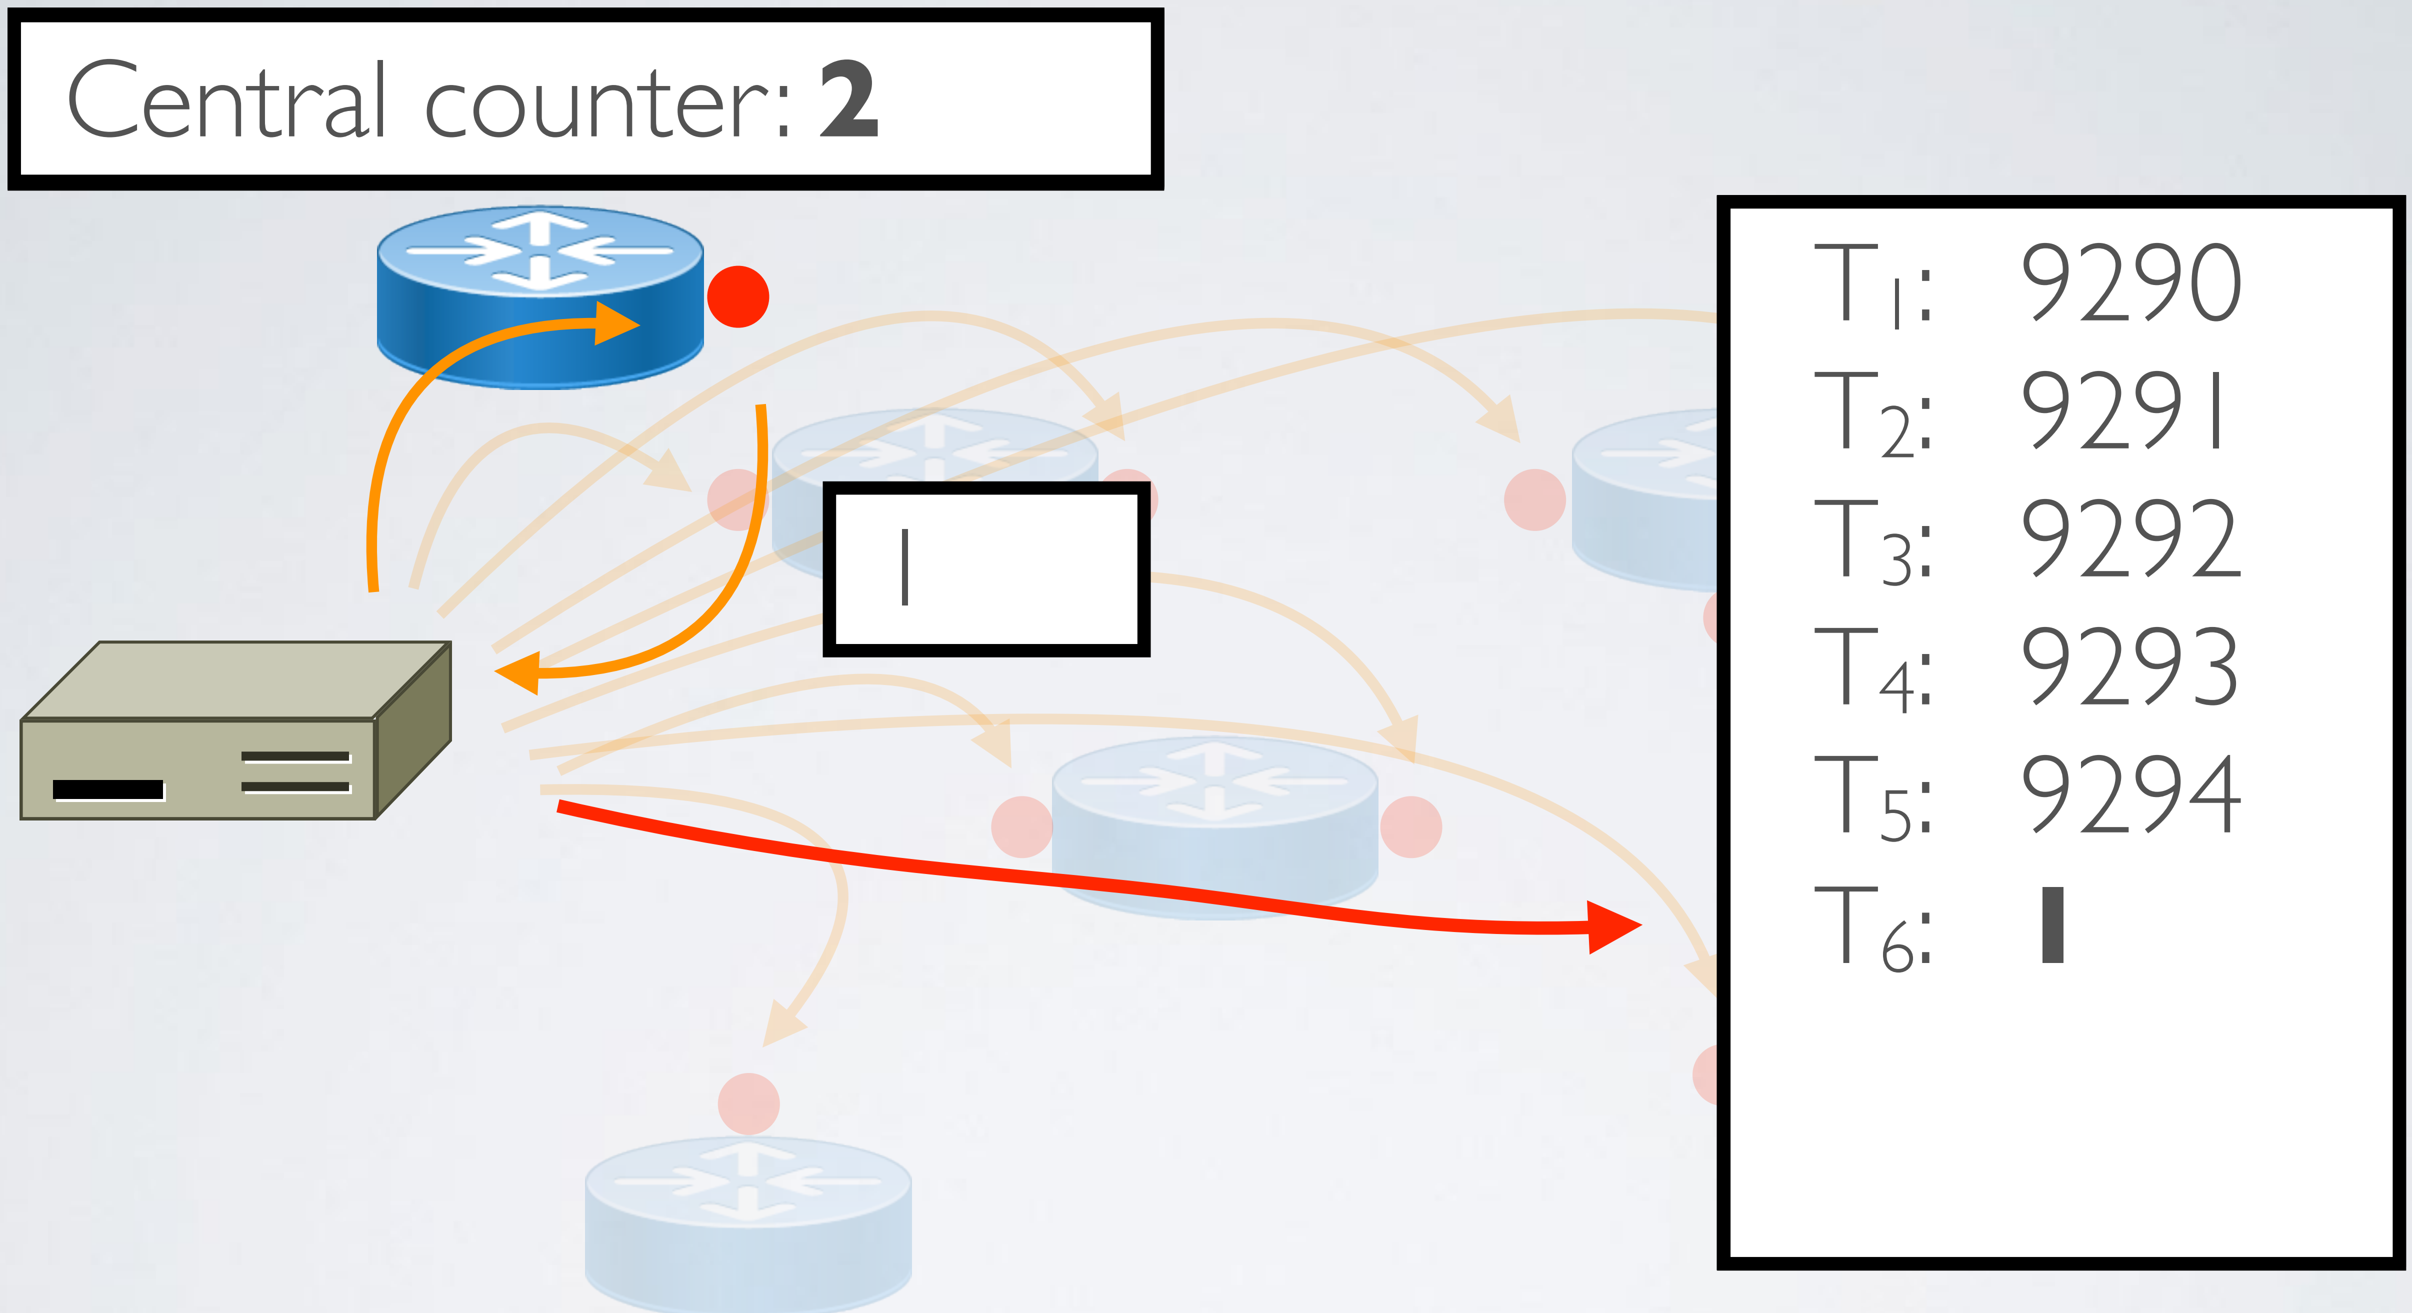

What we did

Central counter: **3**

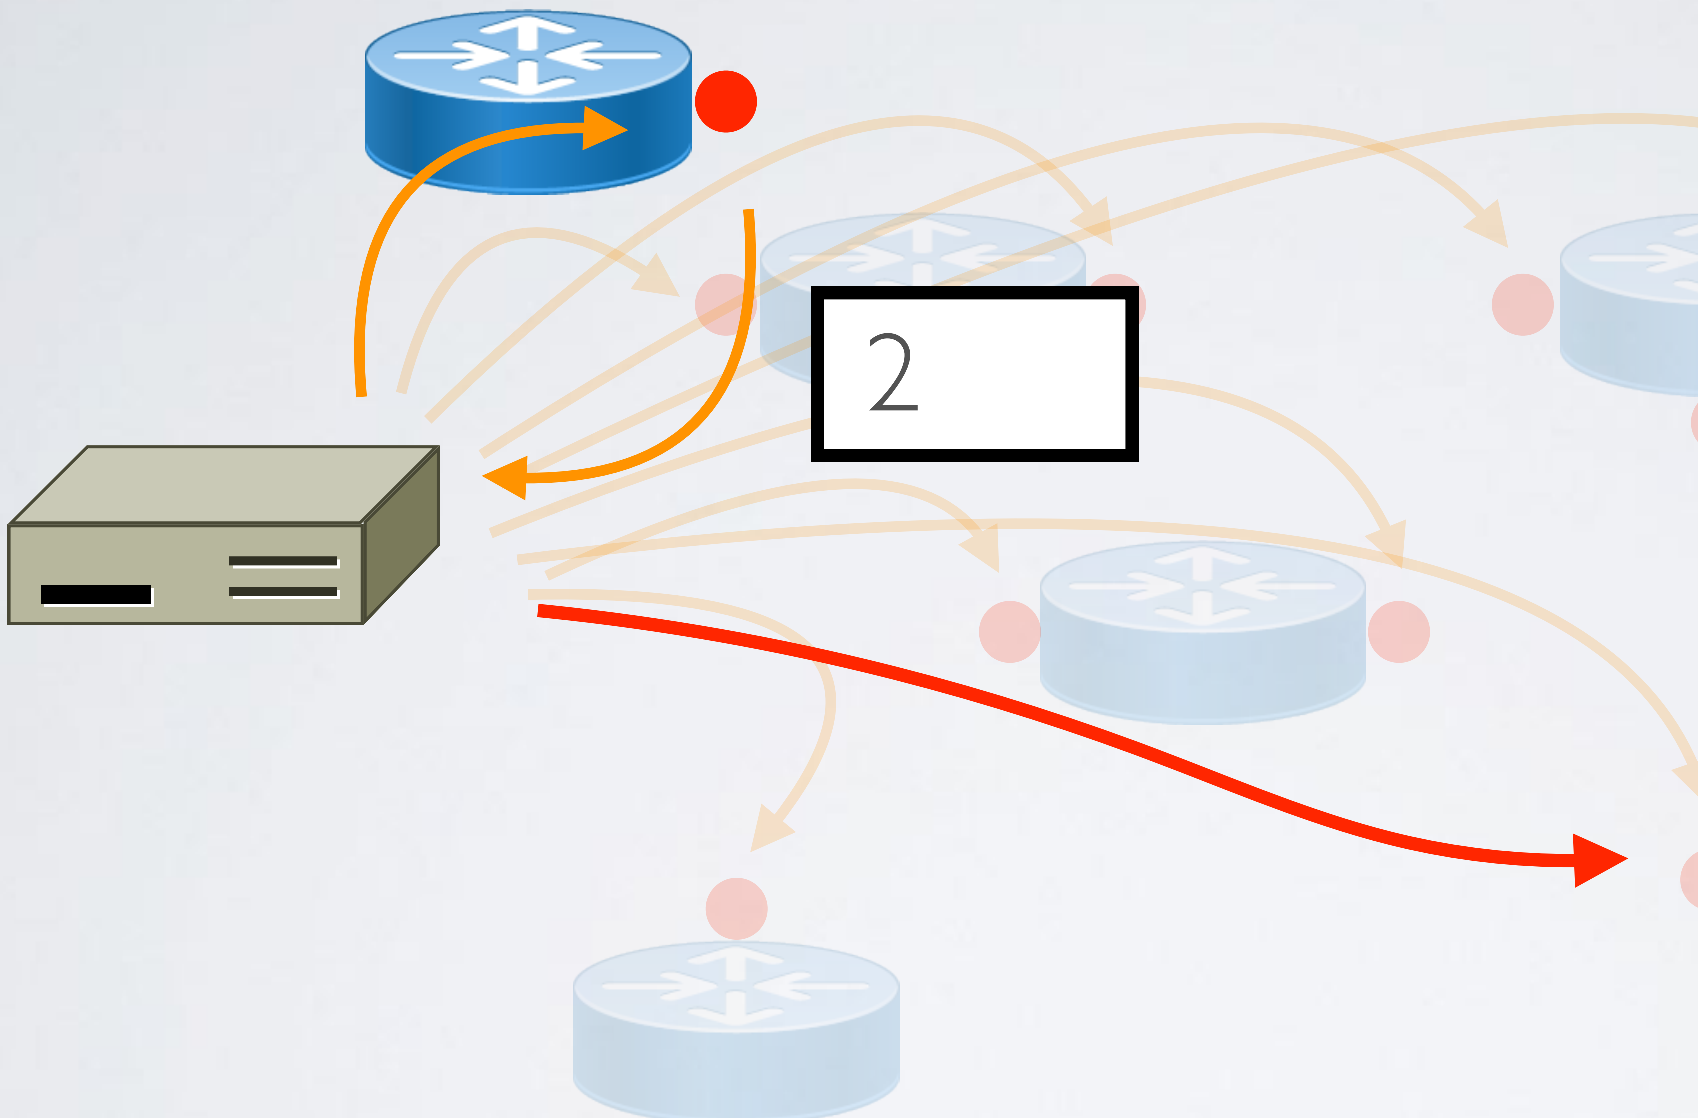

T_1: 9290
T_2: 9291
T_3: 9292
T_4: 9293
T_5: 9294
T_6: 1
T_7: **2**

What we did

Central counter: **4**

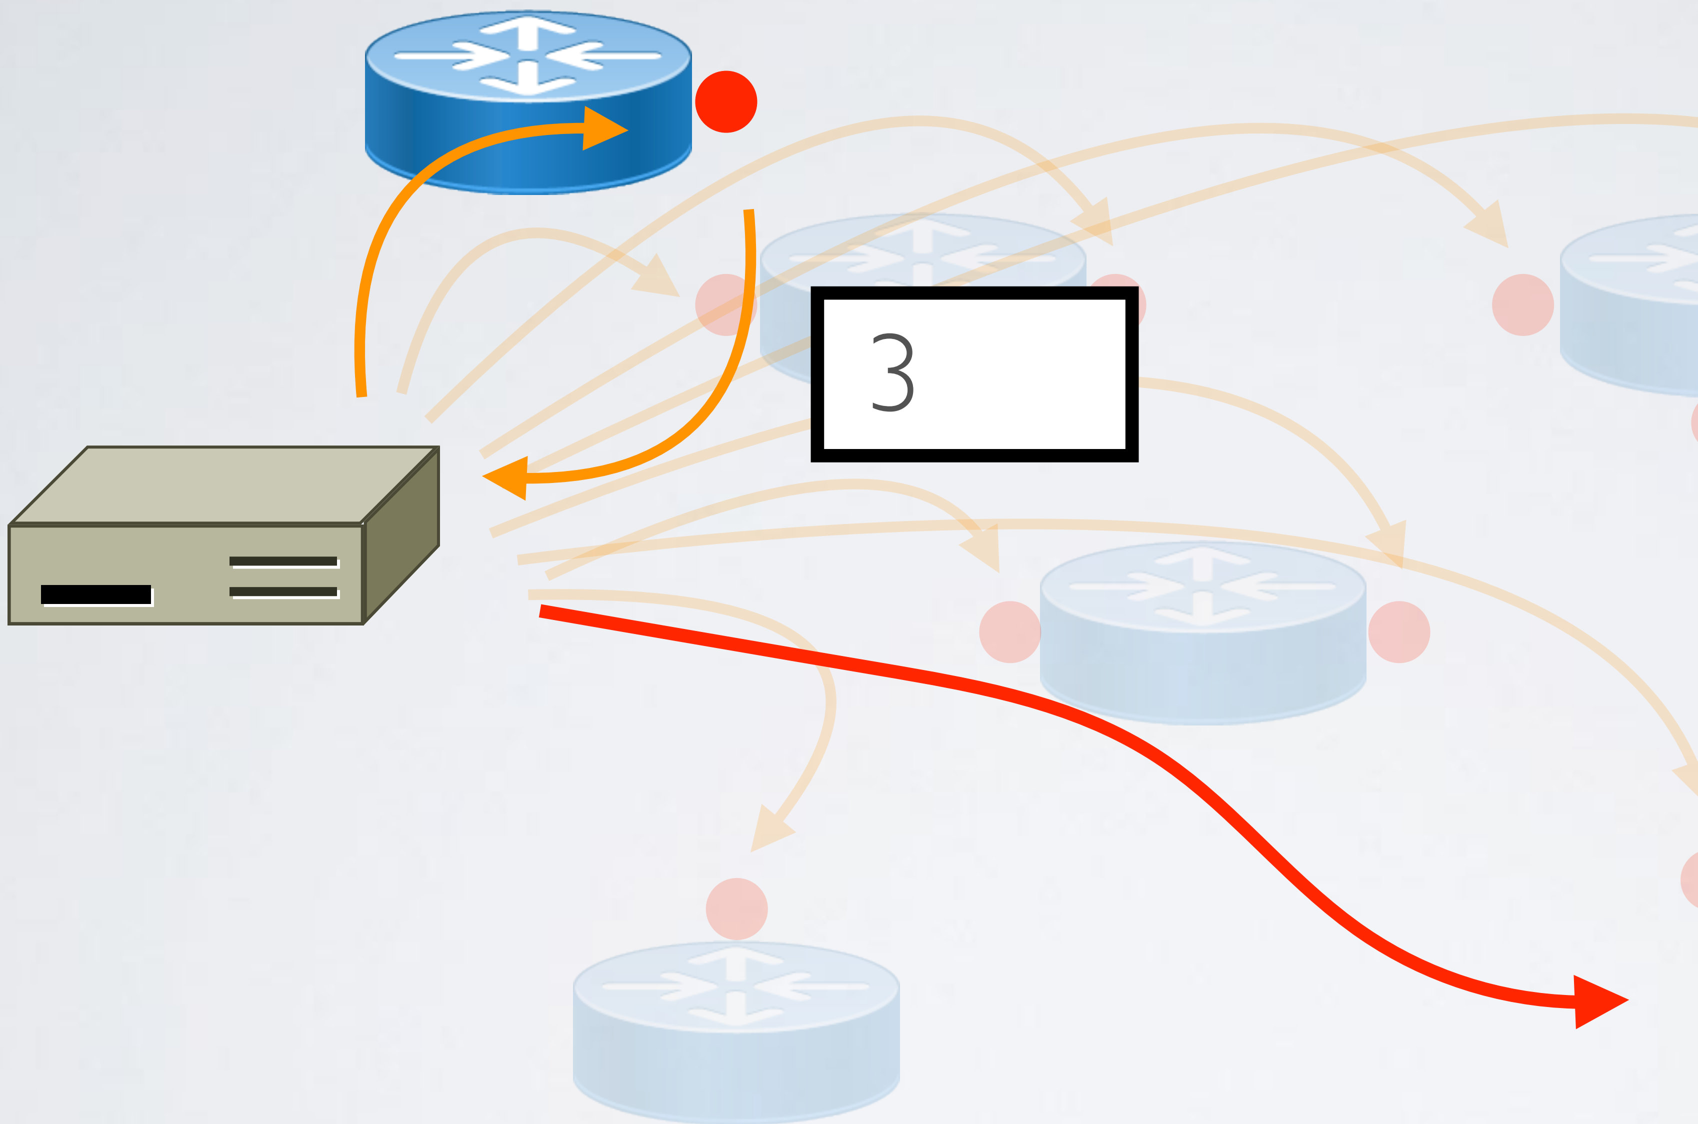

T_1: 9290
T_2: 9291
T_3: 9292
T_4: 9293
T_5: 9294
T_6: 1
T_7: 2
T_8: **3**

What we did

probed router interfaces to infer outage windows using IPID

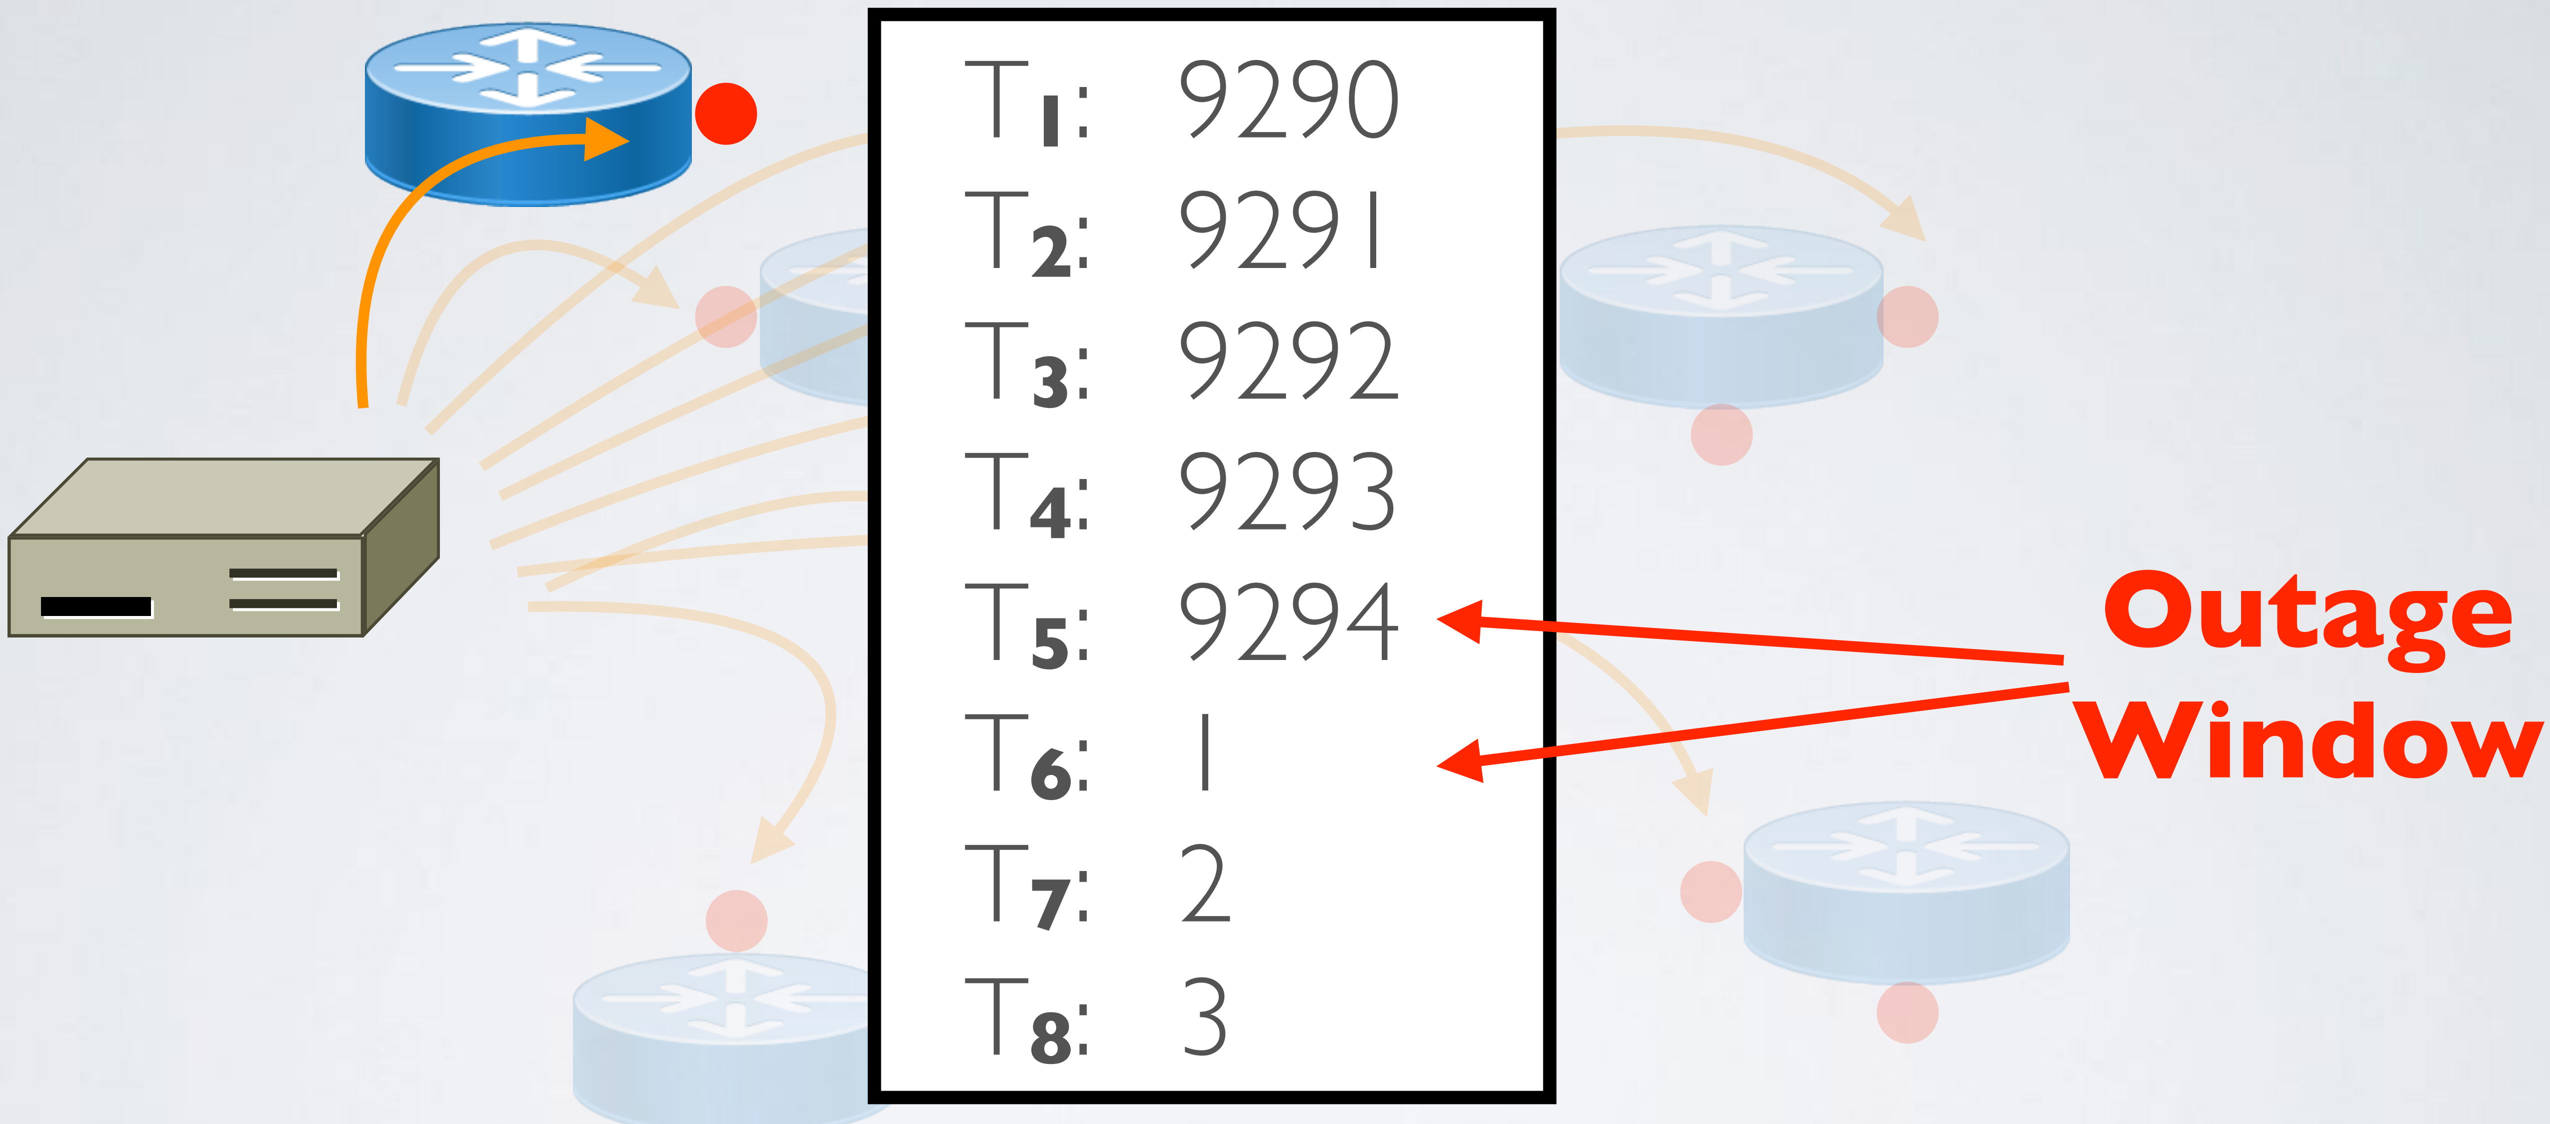

Infer a reboot when time series of values returned from a router is discontinuous, indicating router was restarted

Why IPv6 fragment IDs?

- **IPv4** Fragment IDs:
  - 16 bits, **bursty velocity**: every packet requires unique ID
  - At 100Mbps and 1500 byte packets, Nyquist rate dictates **4 second probing interval**
- **IPv6** Fragment IDs:
  - 32 bits, **low velocity**: IPv6 routers rarely send fragments
  - We average **15 minute probing interval**

What we did

correlated routers with prefixes using traceroute paths

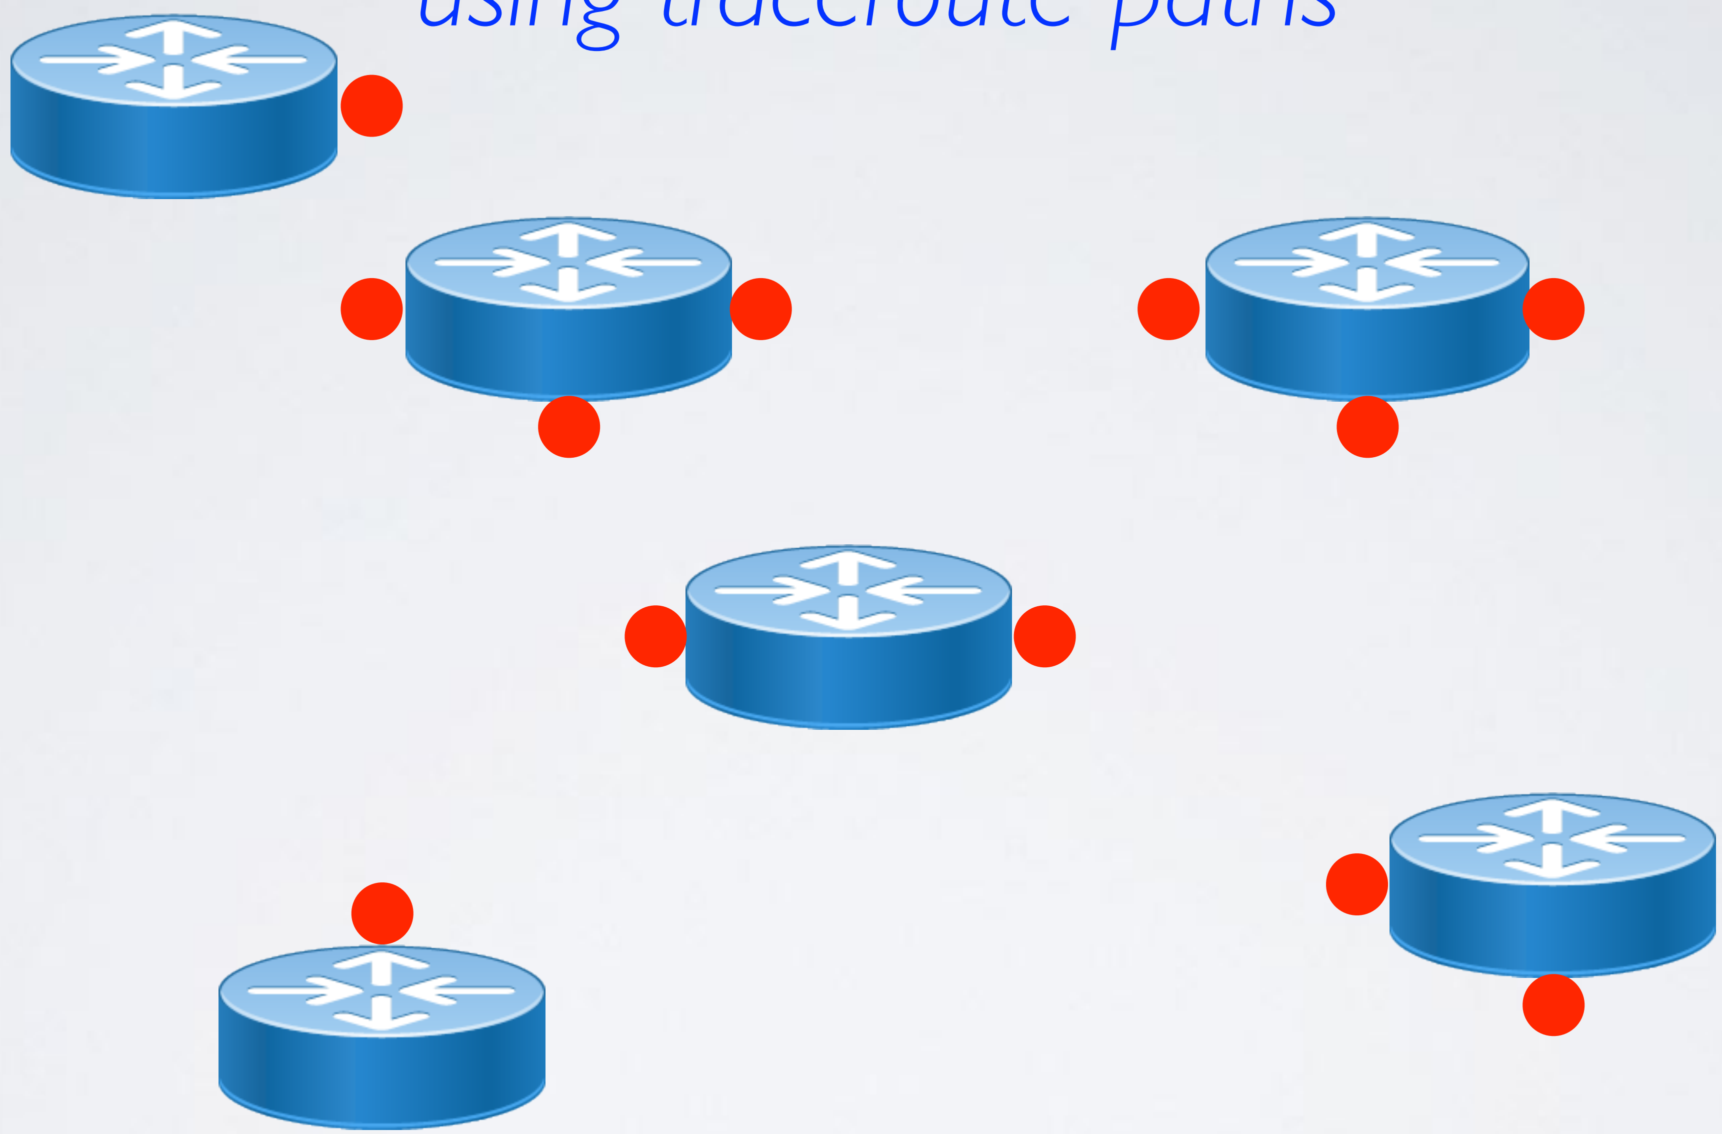

What we did

correlated routers with prefixes using traceroute paths

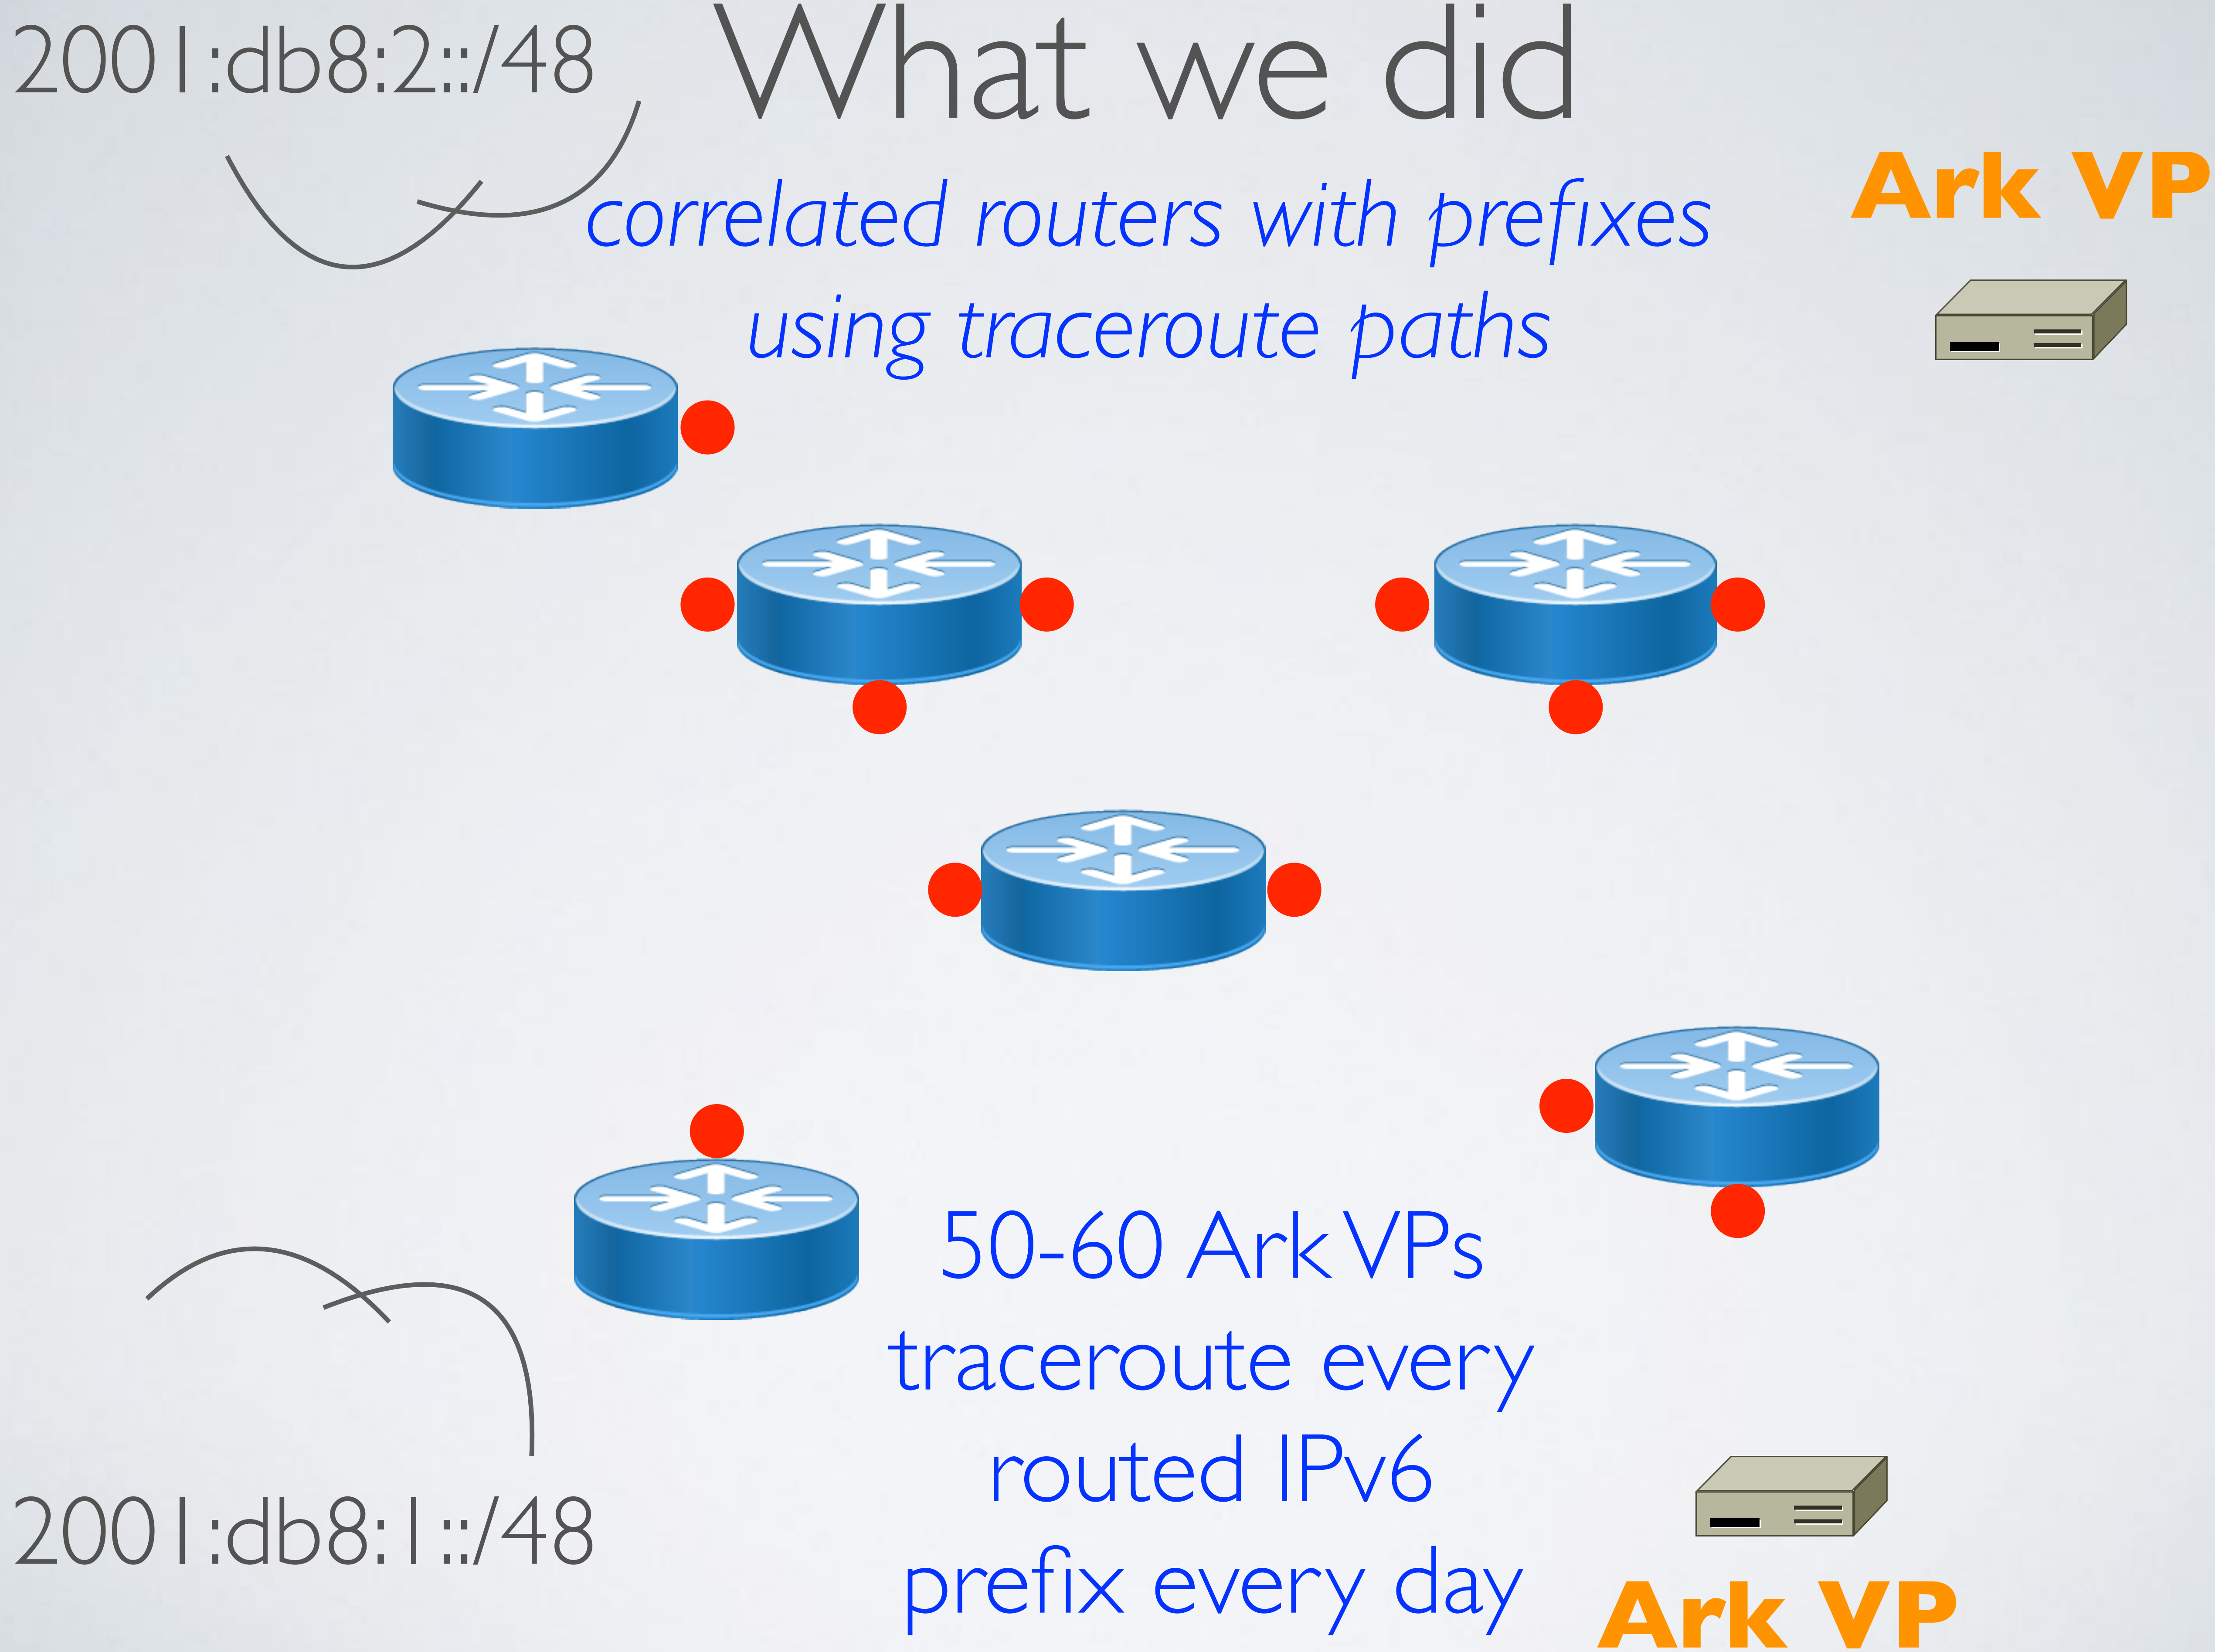

50-60 Ark VPs traceroute every routed IPv6 prefix every day

What we did

correlated routers with prefixes using traceroute paths

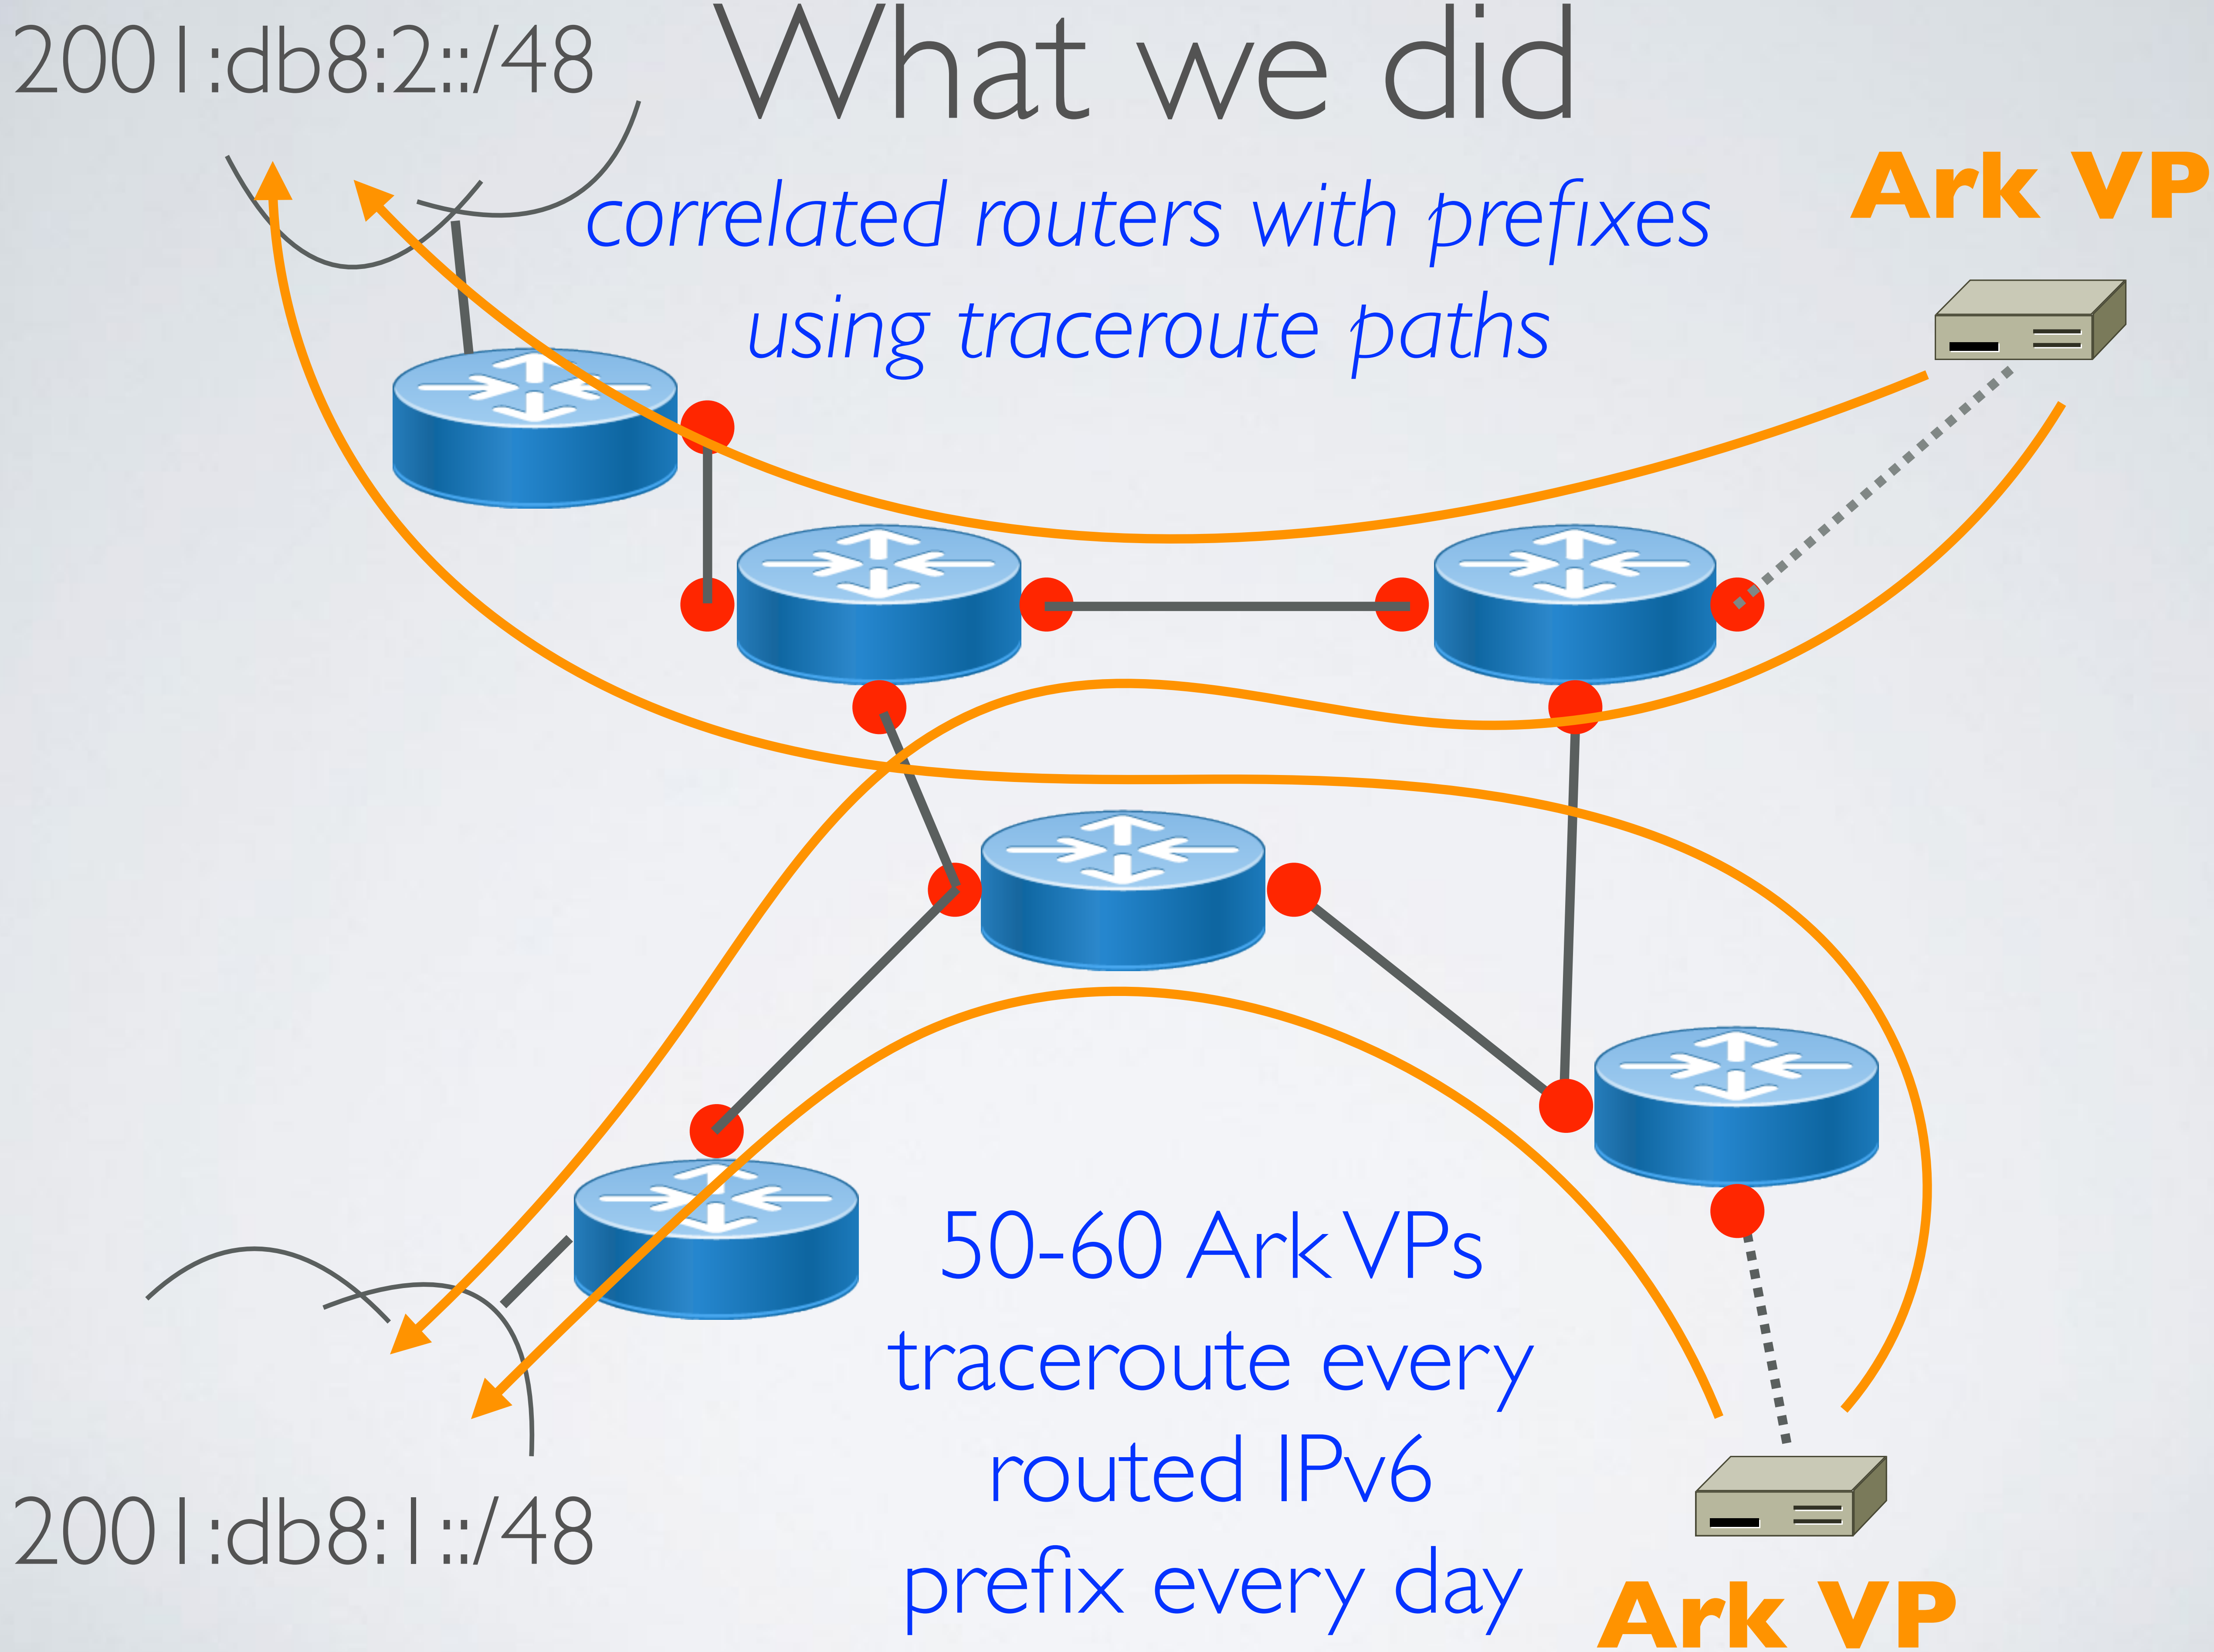

50-60 Ark VPs traceroute every routed IPv6 prefix every day

What we did

computed distance of router from AS announcing network

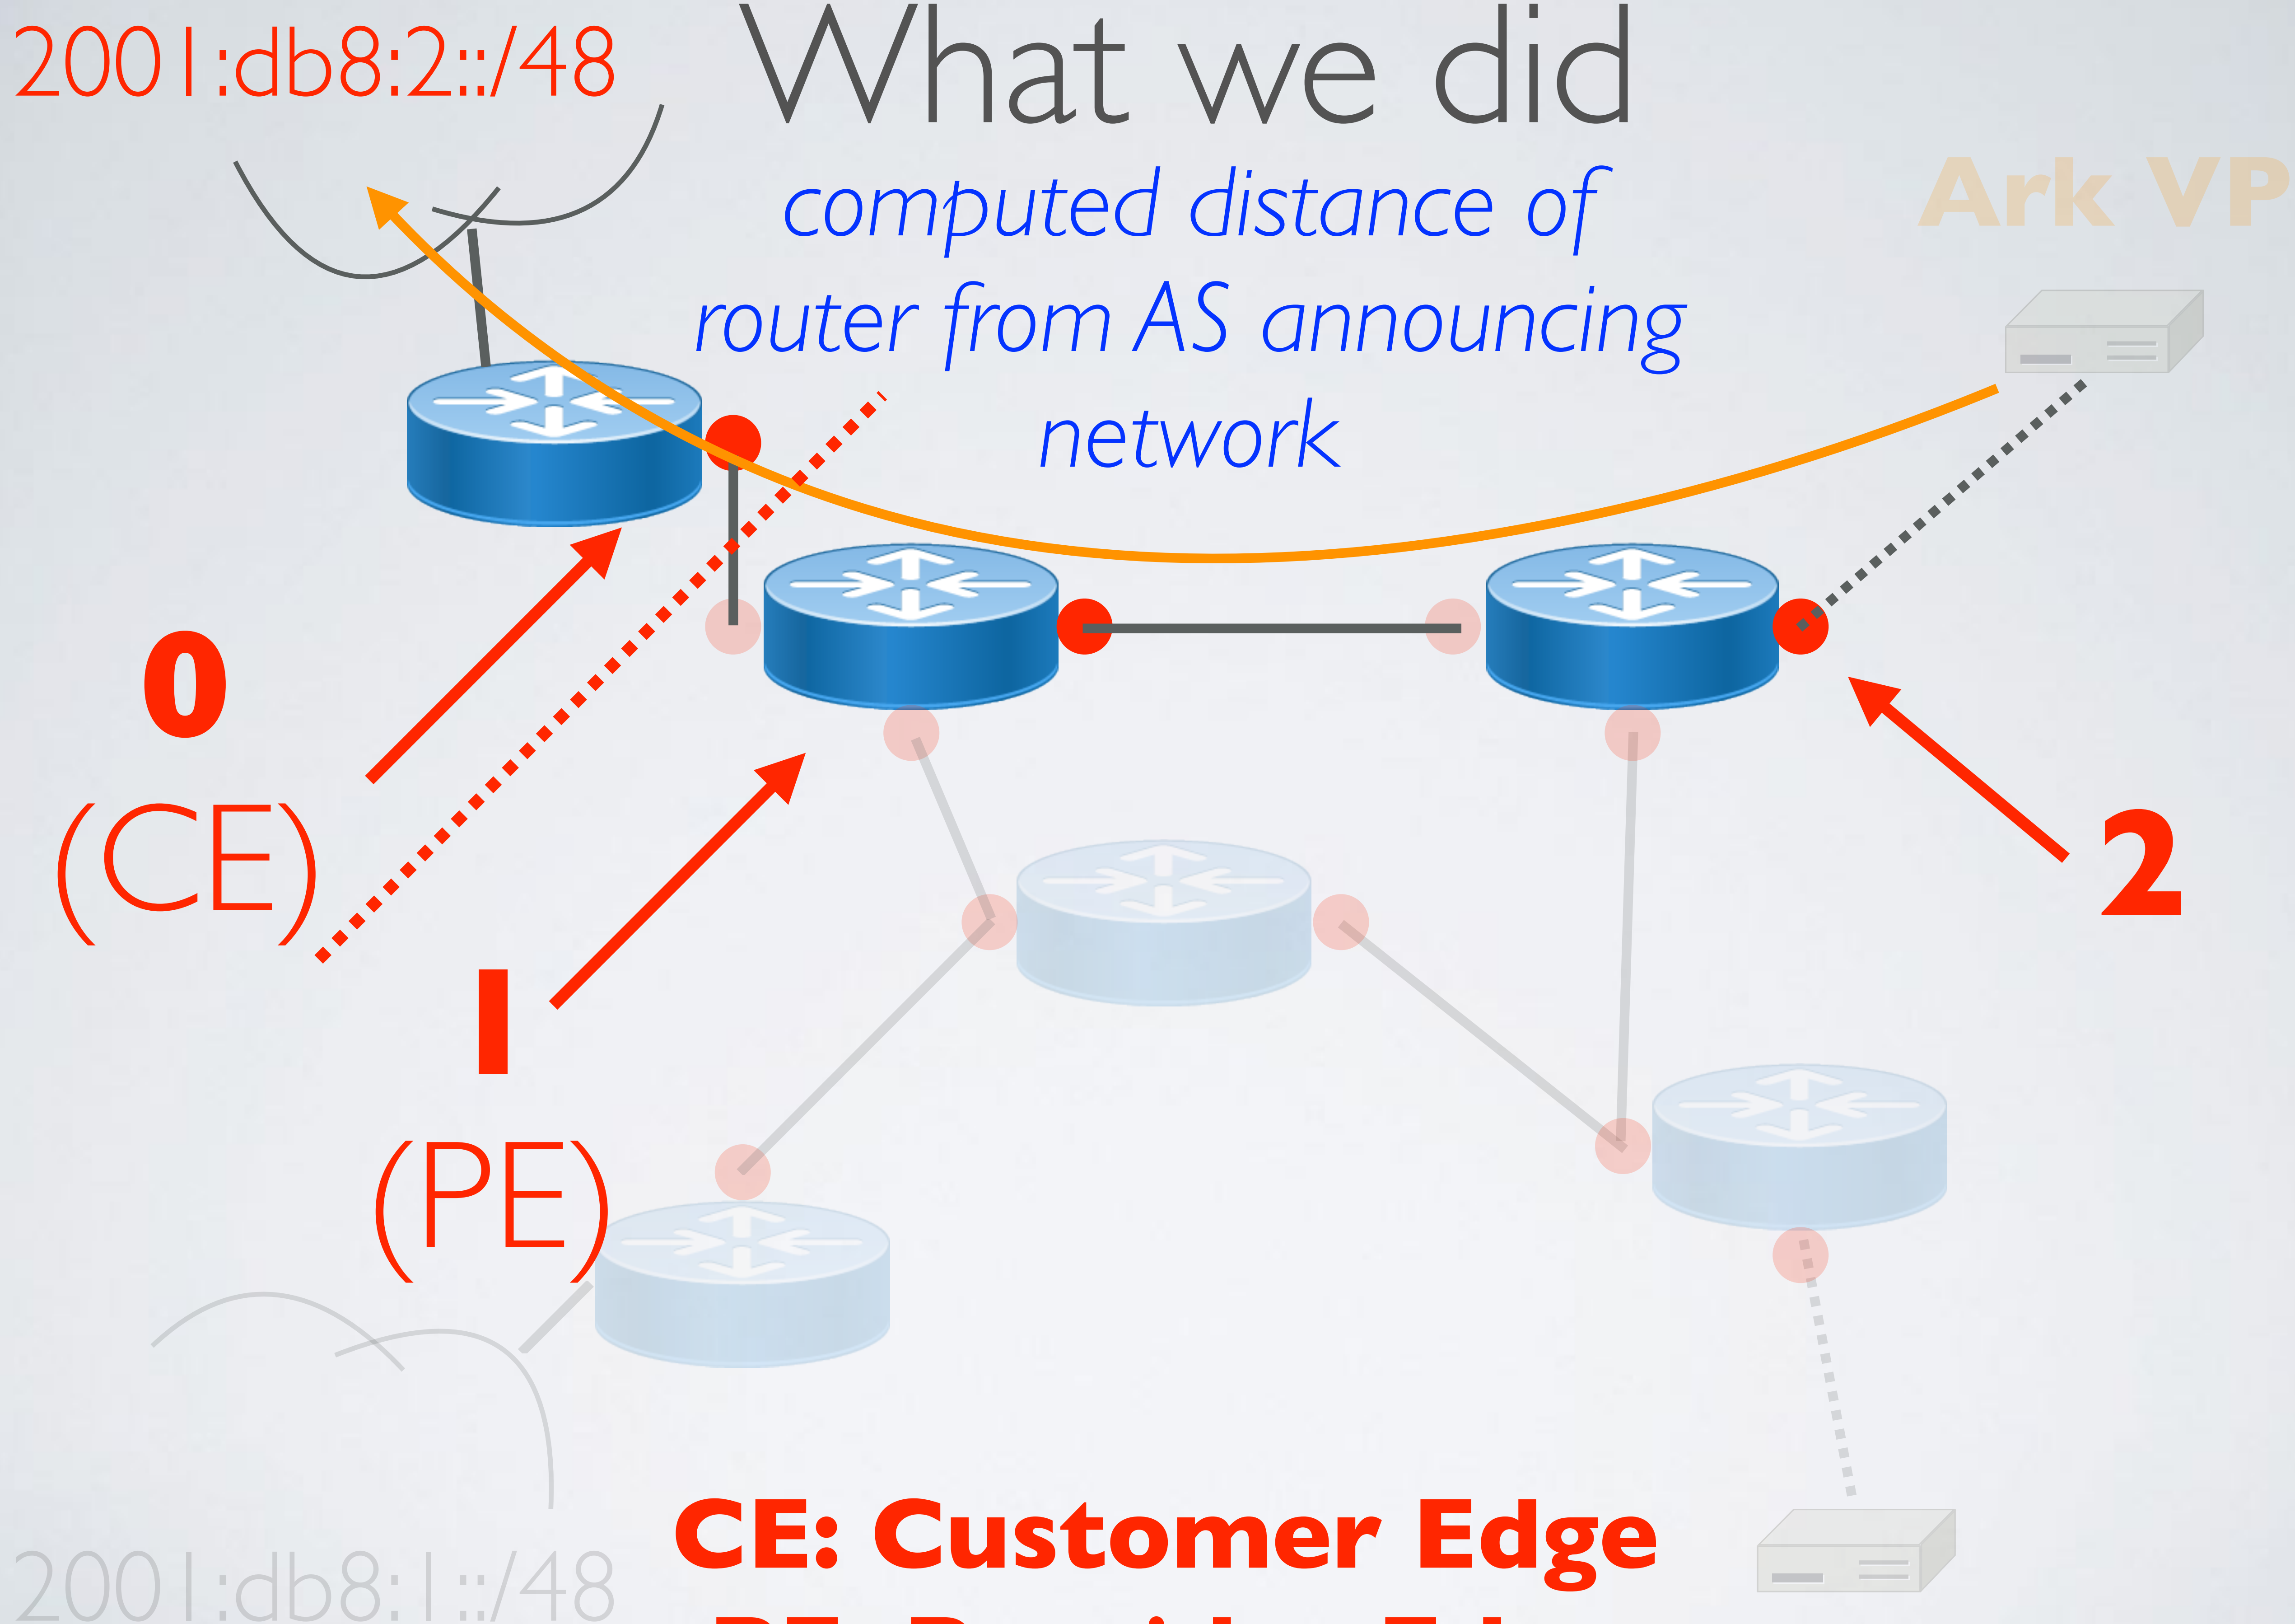

CE: Customer Edge
PE: Provider Edge

What we did

correlated router outage windows with BGP control plane

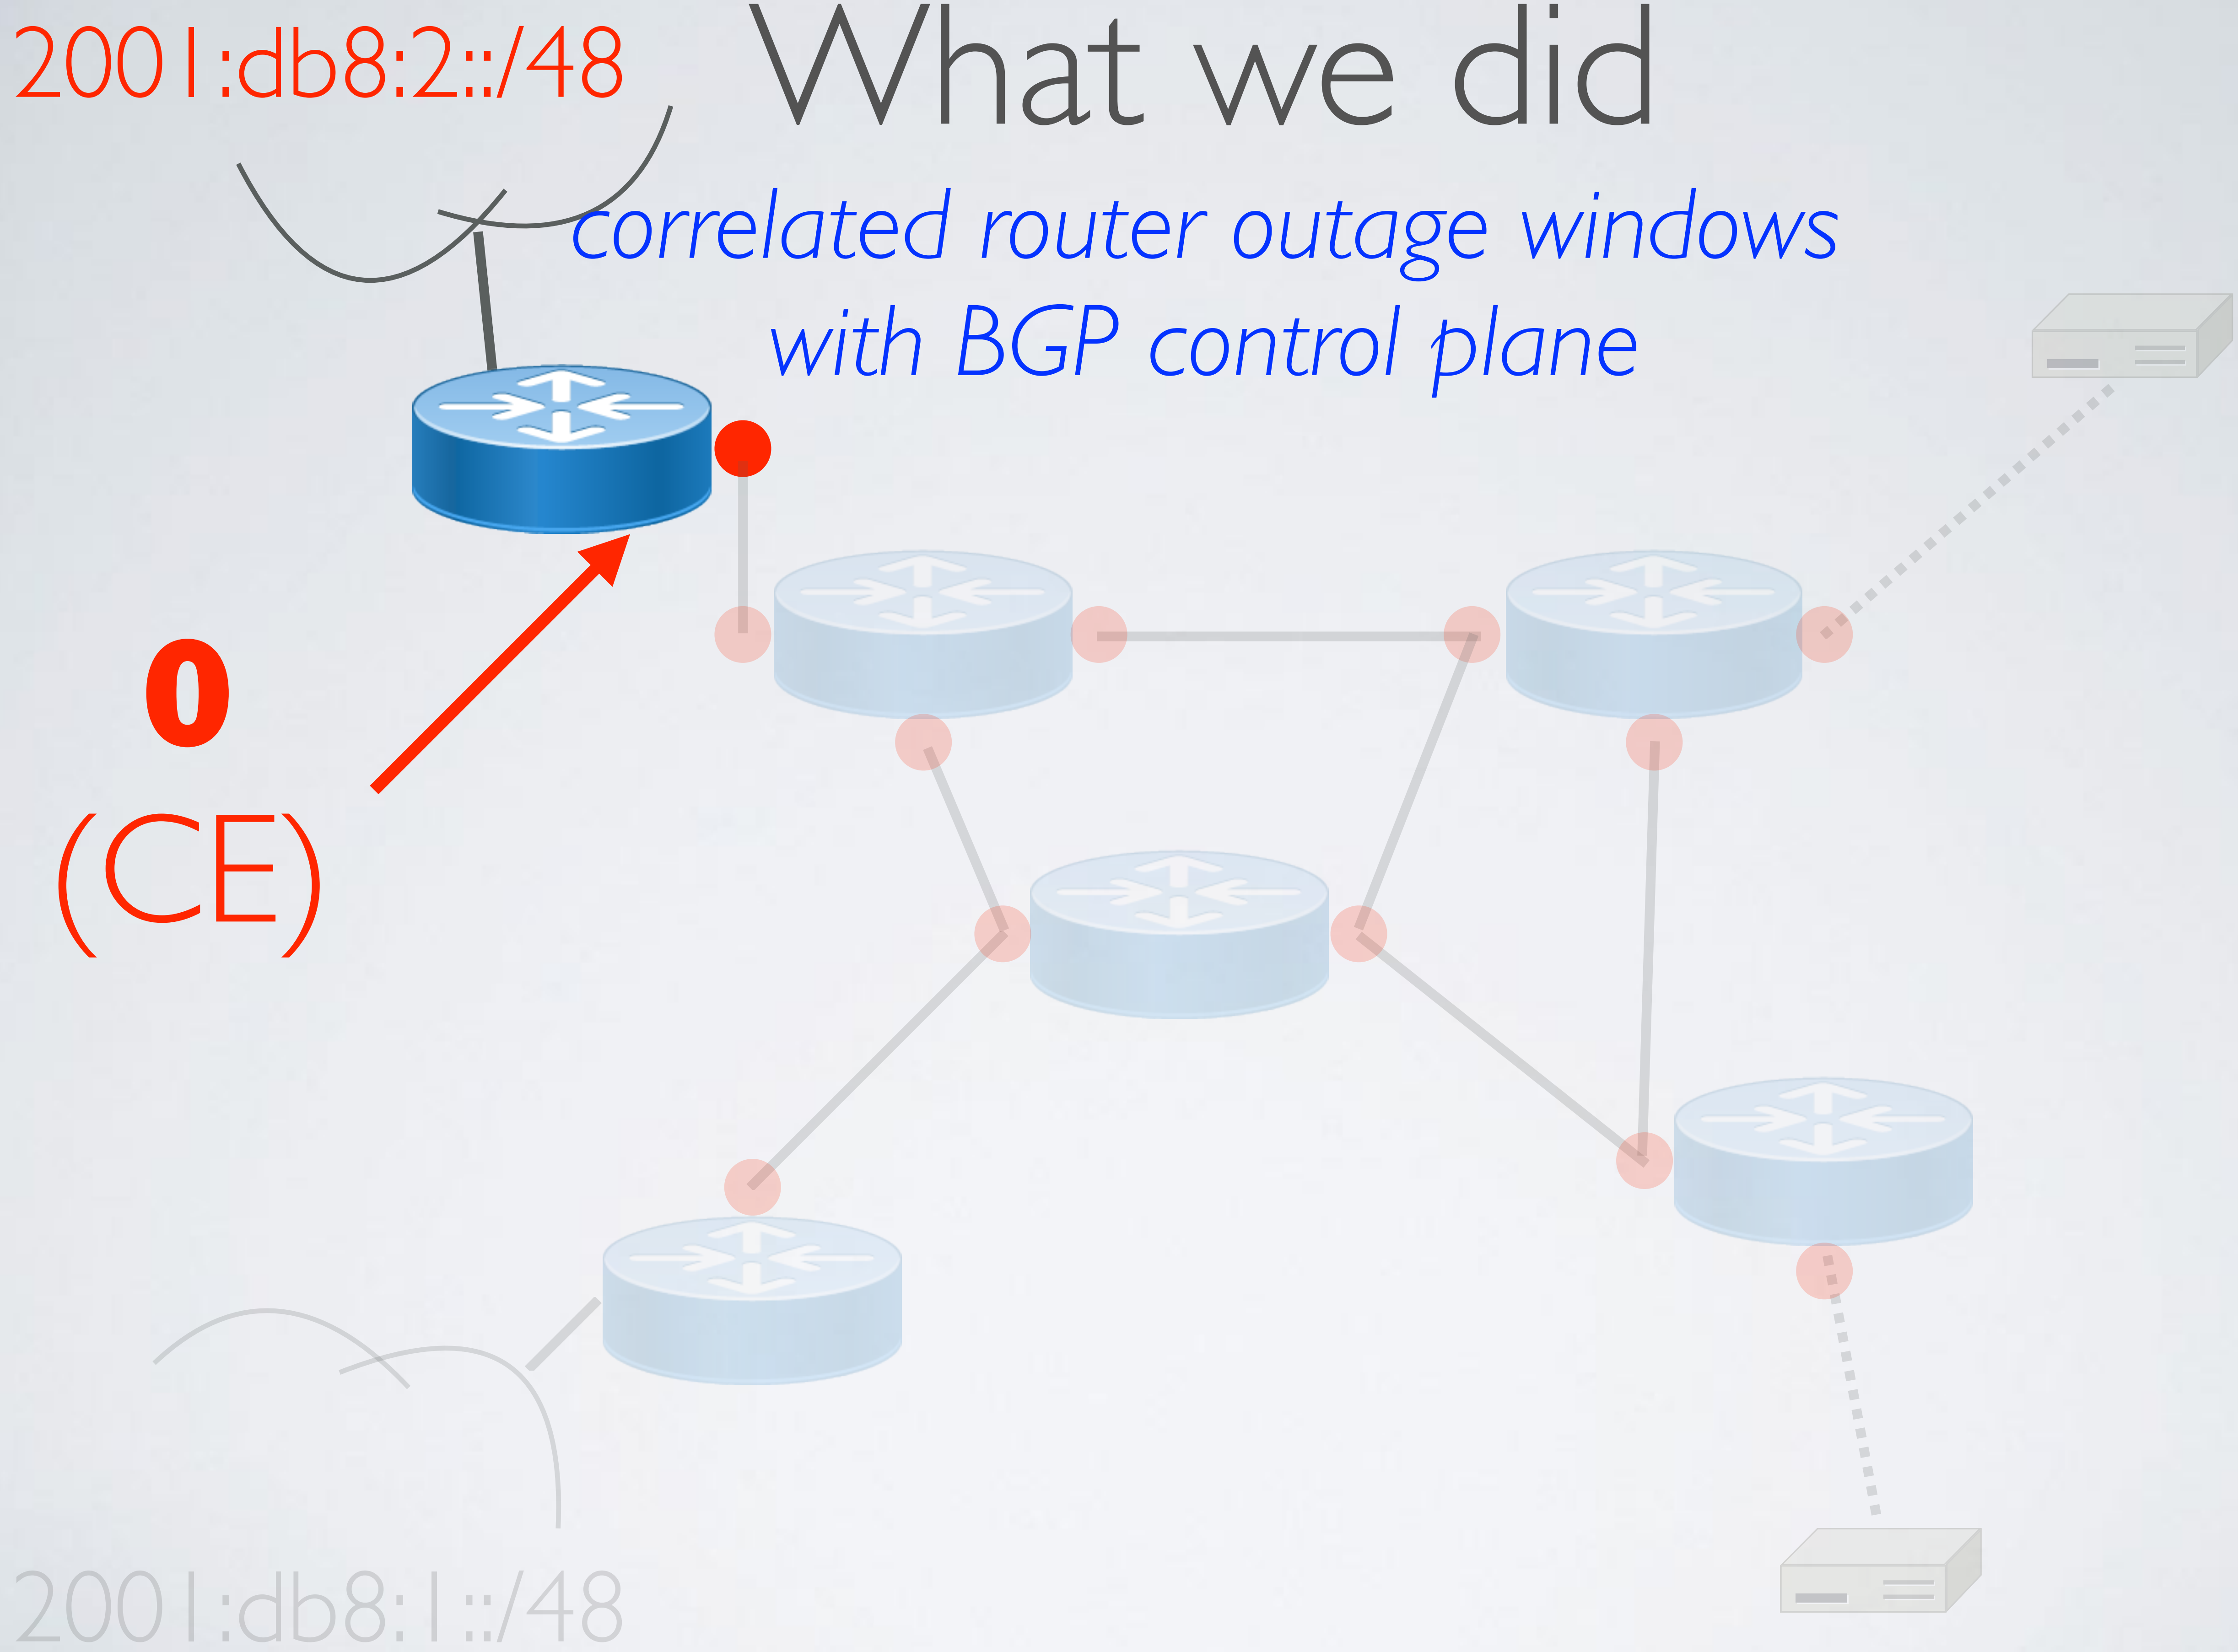

What we did

correlated router outage windows with BGP control plane

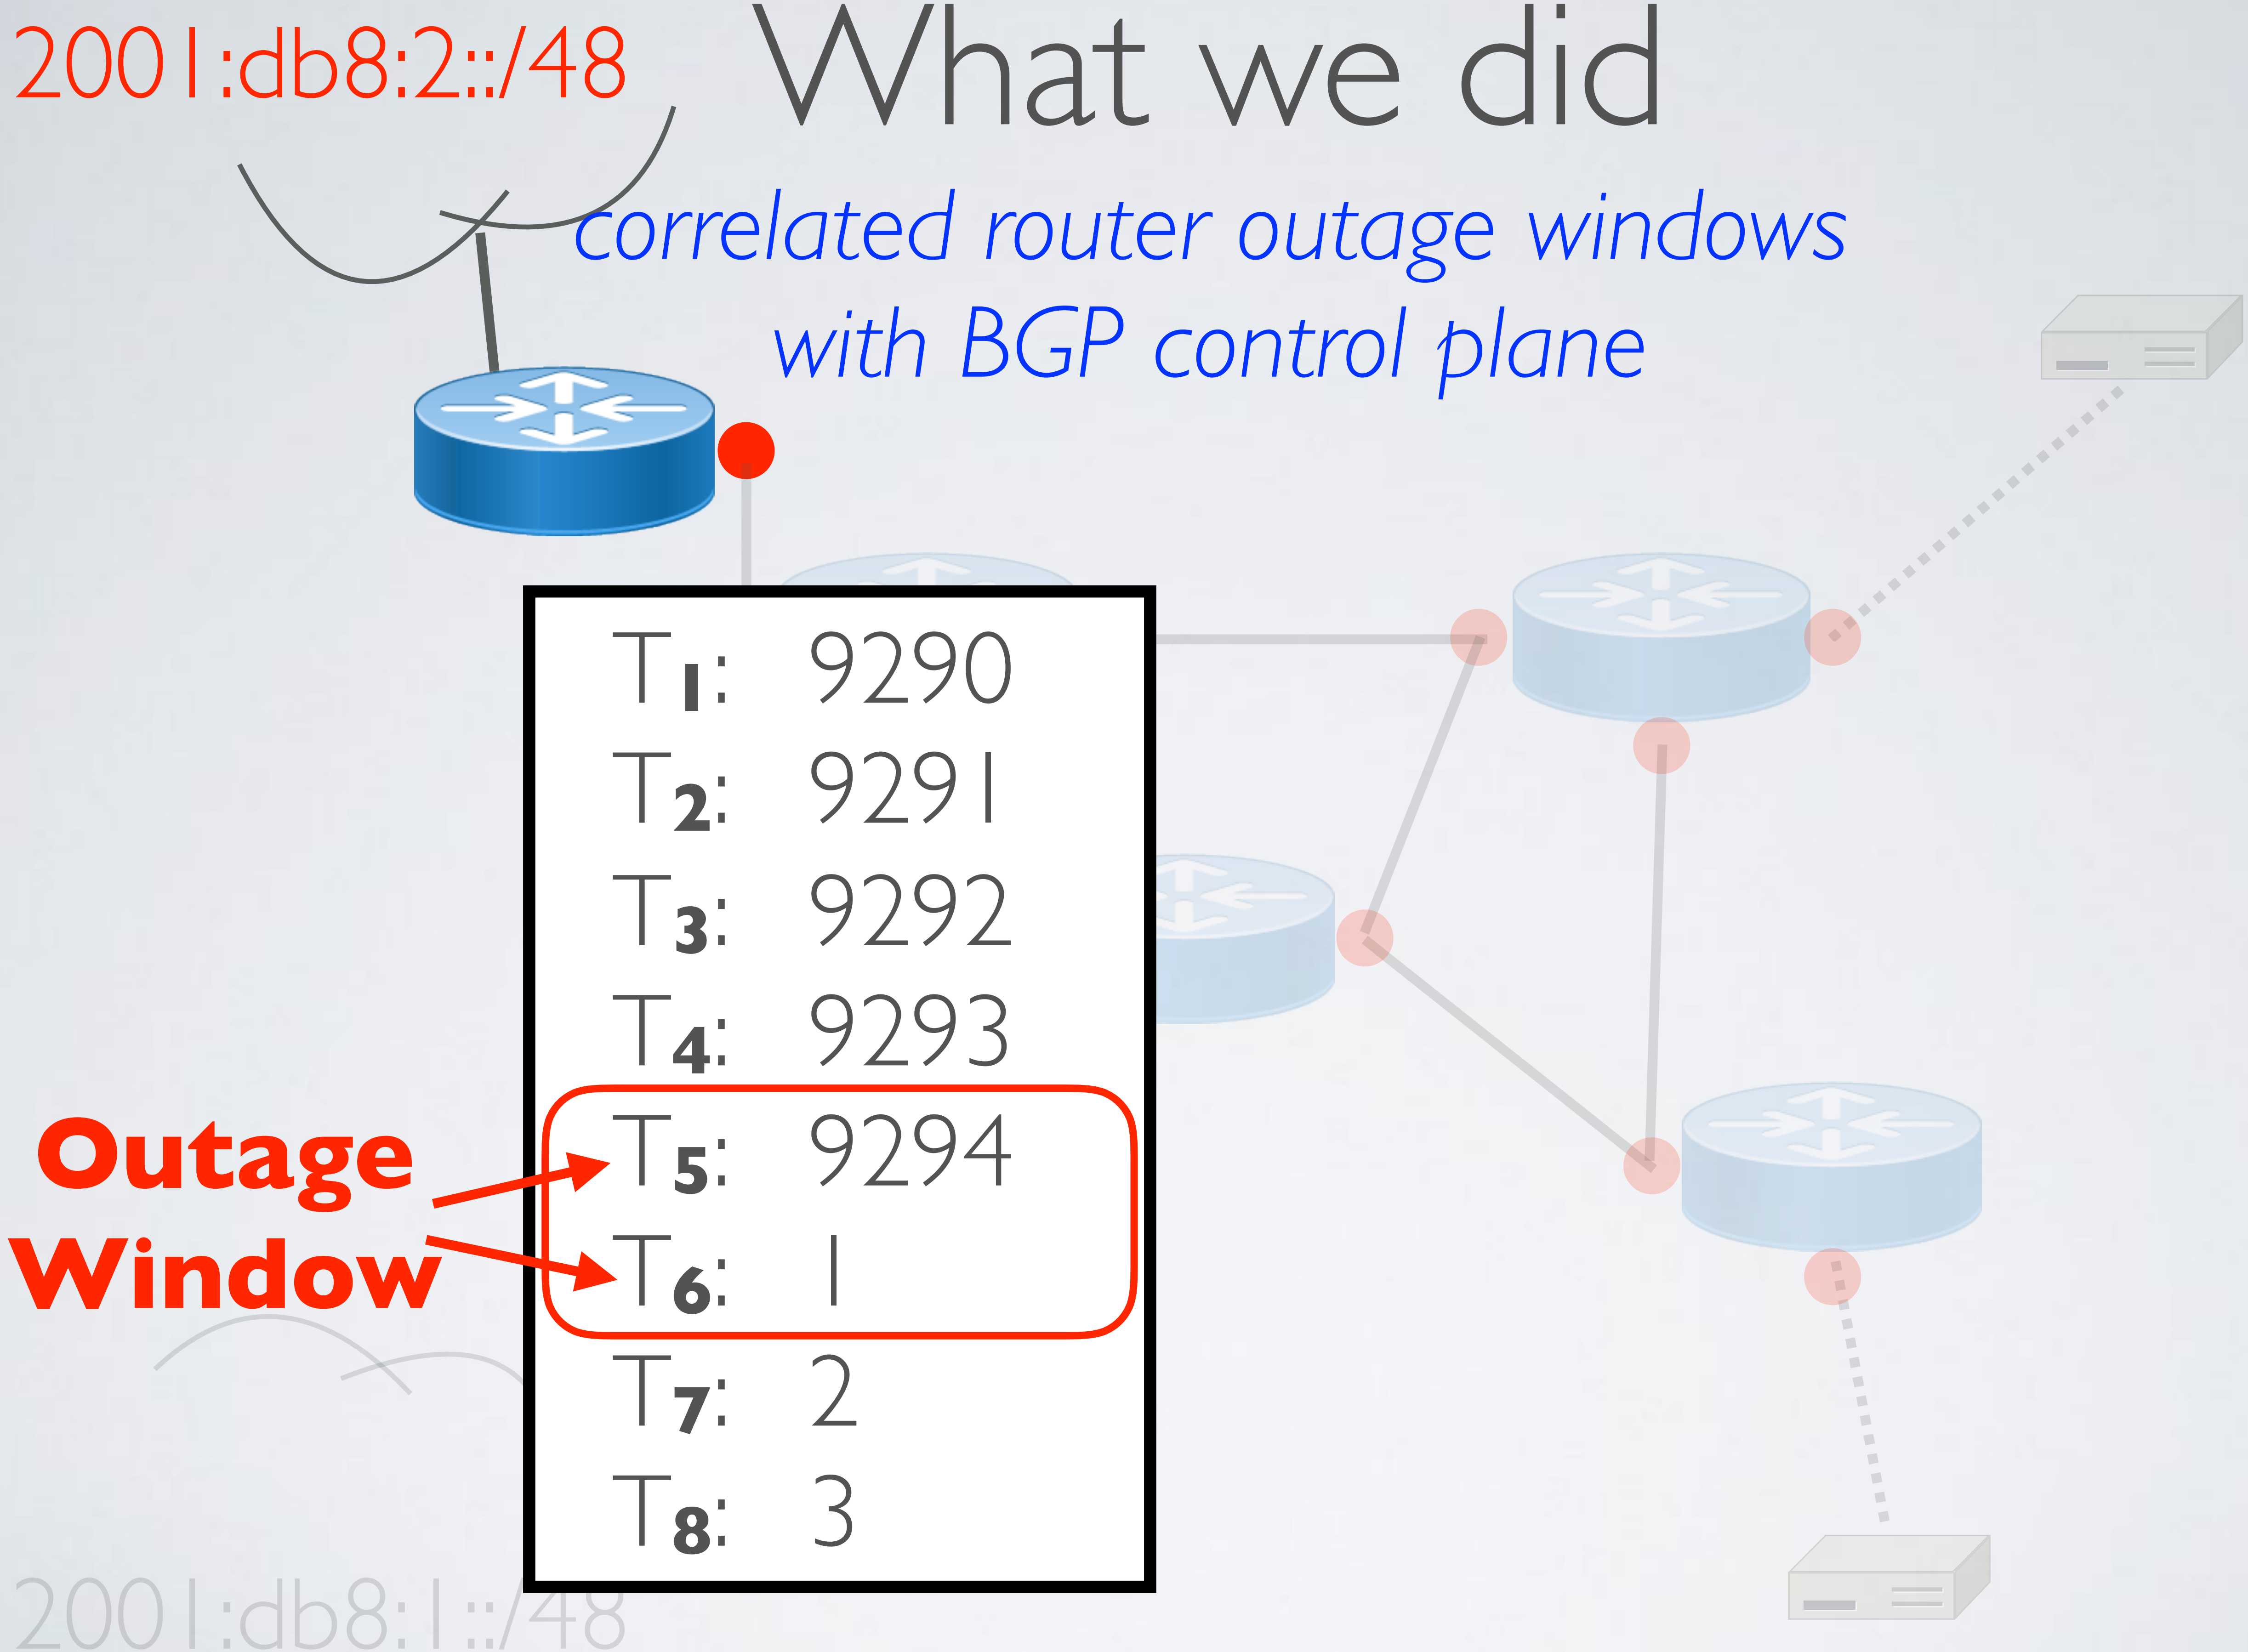

What we did

correlated router outage windows with BGP control plane

2001:db8:2::/48

RouteViews

Outage Window

T_1:	9290
T_2:	9291
T_3:	9292
T_4:	9293
T_5:	9294
T_6:	1
T_7:	2
T_8:	3

2001:db8:2::/48	
$T_{5.2}$:	Peer-1 W
$T_{5.2}$:	Peer-2 W
$T_{5.3}$:	Peer-3 W
$T_{5.3}$:	Peer-4 W
$T_{5.8}$:	Peer-3 A
$T_{5.8}$:	Peer-2 A
$T_{5.8}$:	Peer-1 A
$T_{5.8}$:	Peer-4 A

2001:db8:1::/48

What we did

classified impact on BGP according to observed activity overlapping with inferred outage

- **Complete Withdrawal**: all peers simultaneously withdrew route for at least 70 seconds
  - Single Point of Failure (**SPoF**)
- **Partial Withdrawal**: at least one peer withdrew route for at least 70 seconds, but not all did
- **Churn**: BGP activity for the prefix
- **No Impact**: No observed BGP activity for the prefix

What we did

Data Collection Summary

- Probed IPv6 routers at ~15 minute intervals from 18 Jan 2015 to 30 May 2017 (approx. 2.5 years)
- 149,560 routers allowed reboots to be detected
- We inferred 59,175 (40%) rebooted at least once,750K reboots in total

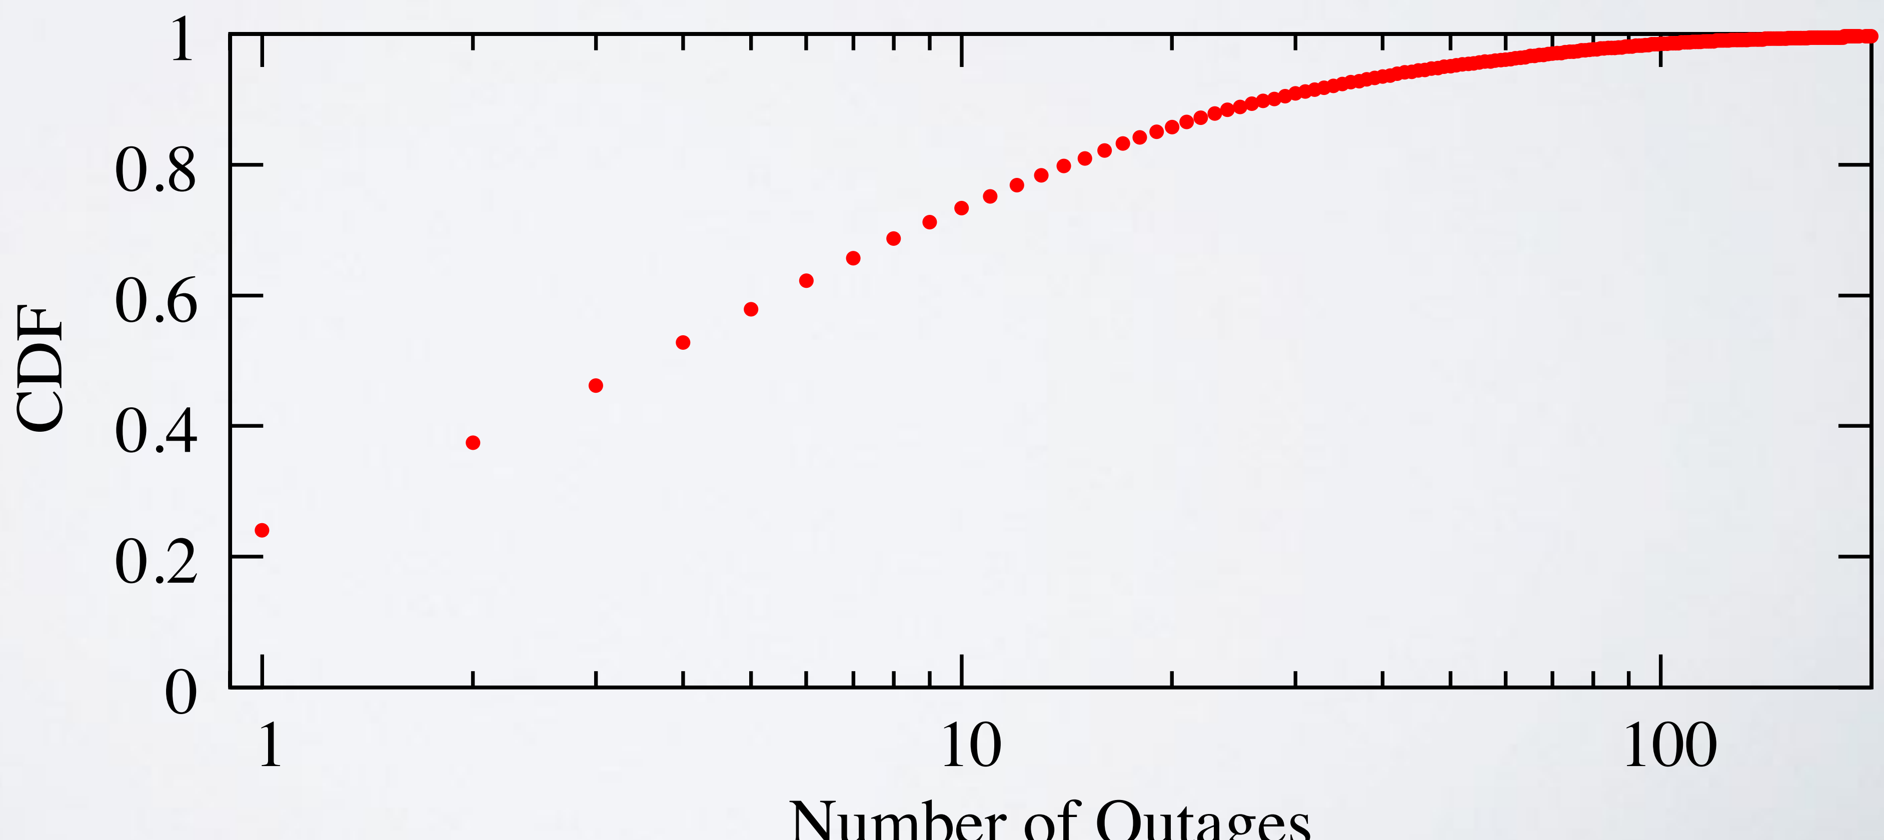

What we found

- **2,385 (4%) of routers** that rebooted (59K) we inferred to be **SPoF** for at least one IPv6 prefix in BGP

- Of SPoF routers, we inferred **59%** to be customer edge router; **8%** provider edge; **29%** within destination AS

- **No covering prefix for 70%** of withdrawn prefixes

  - During one-week sample, covering prefix presence during withdrawal did not imply data plane reachability

- IPv6 Router reboots **correlated with IPv4** BGP control plane activity

Limitations

- Applicability to IPv4 depends on router being dual-stack
- Requires IPID assigned from a counter
  - Cisco, Huawei, Vyatta, Microtik, HP assign from counter
  - 27.1% responsive for 14 days assigned from counter
- Router outage might end before all peers withdraw route
  - Path exploration + Minimum Route Advertisement Interval (MRAI) + Route Flap Dampening (RFD)
- Complex events: multiple router outages but one detected
  - We observed some complex events and filtered them out

Validation

	Reboots			SPoF		
Network	✓	✗	?	✓	✗	?
US University	7	0	8	7	0	8
US R&E backbone #1	2	0	3	3	2	0
US R&E backbone #2	3	0	1	0	0	4
NZ R&E backbone	11	0	22	4	2	27
Total:	**23**	**0**	**34**	**14**	**4**	**39**

✓ = Validated Inference

✗ = Incorrect Inference

? = Not Validated

Validation

	Reboots			SPoF		
Network	✓	✗	?	✓	✗	?
US University	7	0	8	7	0	8
US R&E backbone #1	2	0	3	3	2	0
US R&E backbone #2	3	0	1	0	0	4
NZ R&E backbone	11	0	22	4	2	27
Total:	**23**	**0**	**34**	**14**	**4**	**39**

Challenging to get validation data: operators often could only tell us about the last reboot

Validation

	Reboots			SPoF		
Network	✓	✗	?	✓	✗	?
US University	7	0	8	7	0	8
US R&E backbone #1	2	0	3	3	2	0
US R&E backbone #2	3	0	1	0	0	4
NZ R&E backbone	11	0	22	4	2	27
Total:	**23**	**0**	**34**	**14**	**4**	**39**

No falsely inferred reboots: we correctly observed the last known reboot of each router

Validation

	Reboots			SPoF		
Network	✓	✗	?	✓	✗	?
US University	7	0	8	7	0	8
US R&E backbone #1	2	0	3	3	2	0
US R&E backbone #2	3	0	1	0	0	4
NZ R&E backbone	11	0	22	4	2	27
Total:	**23**	**0**	**34**	**14**	**4**	**39**

We did not detect some SPoFs

Data Collection Summary

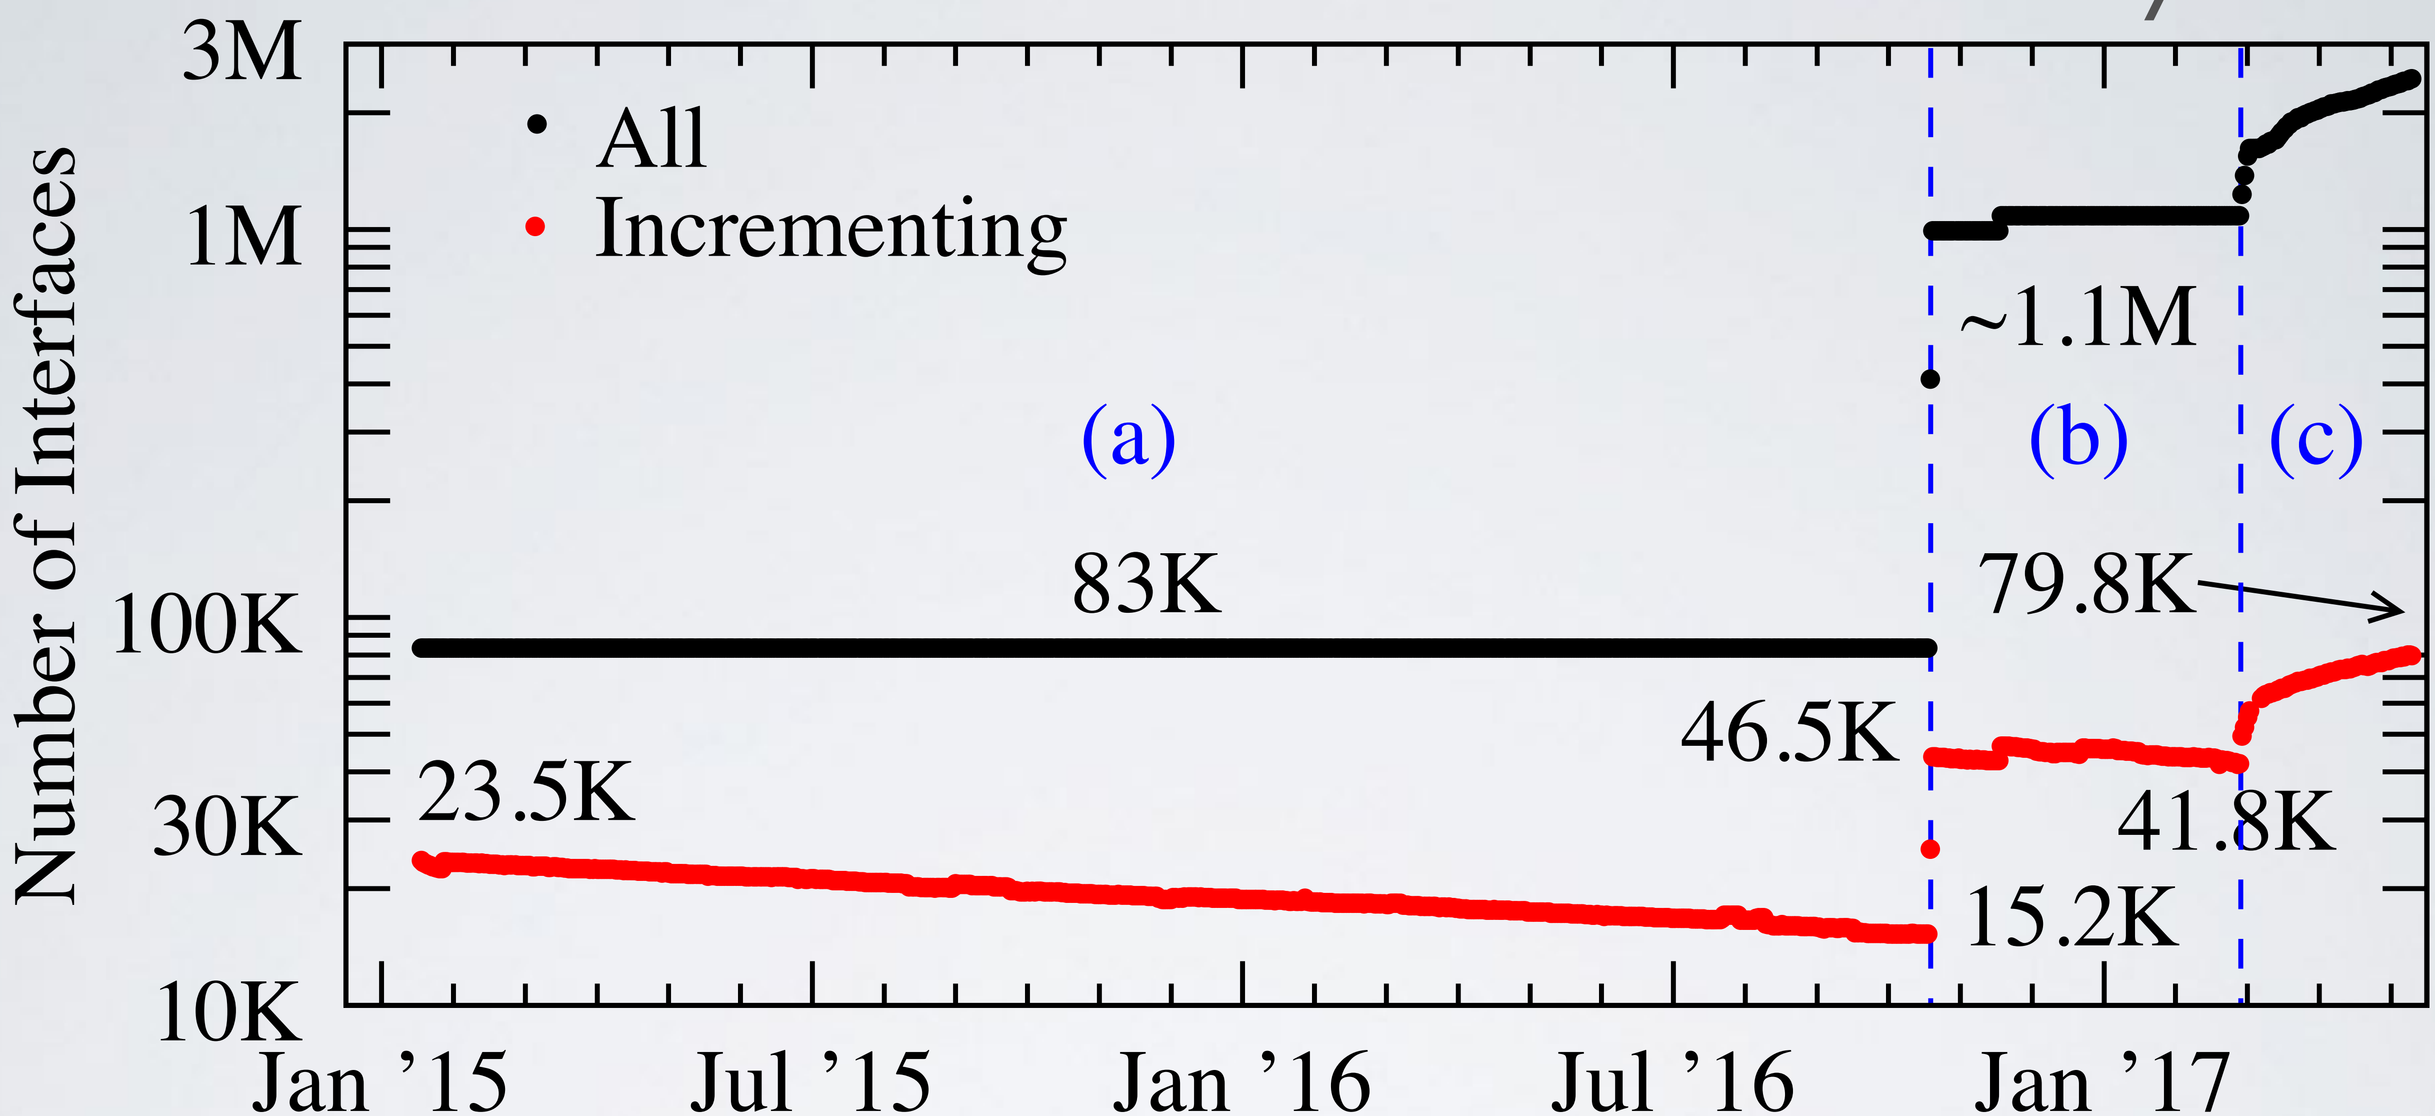

	PPS	List	Unresponsive
(a)	100	Static 83K	12-24 hours
(b)	225	Static 1.1M	12-24 hours
(c)	200	Dynamic, ~2.4M	7-14 days

Correlating BGP/router outages

Control: six hours prior to inferred outages, Feb 2015

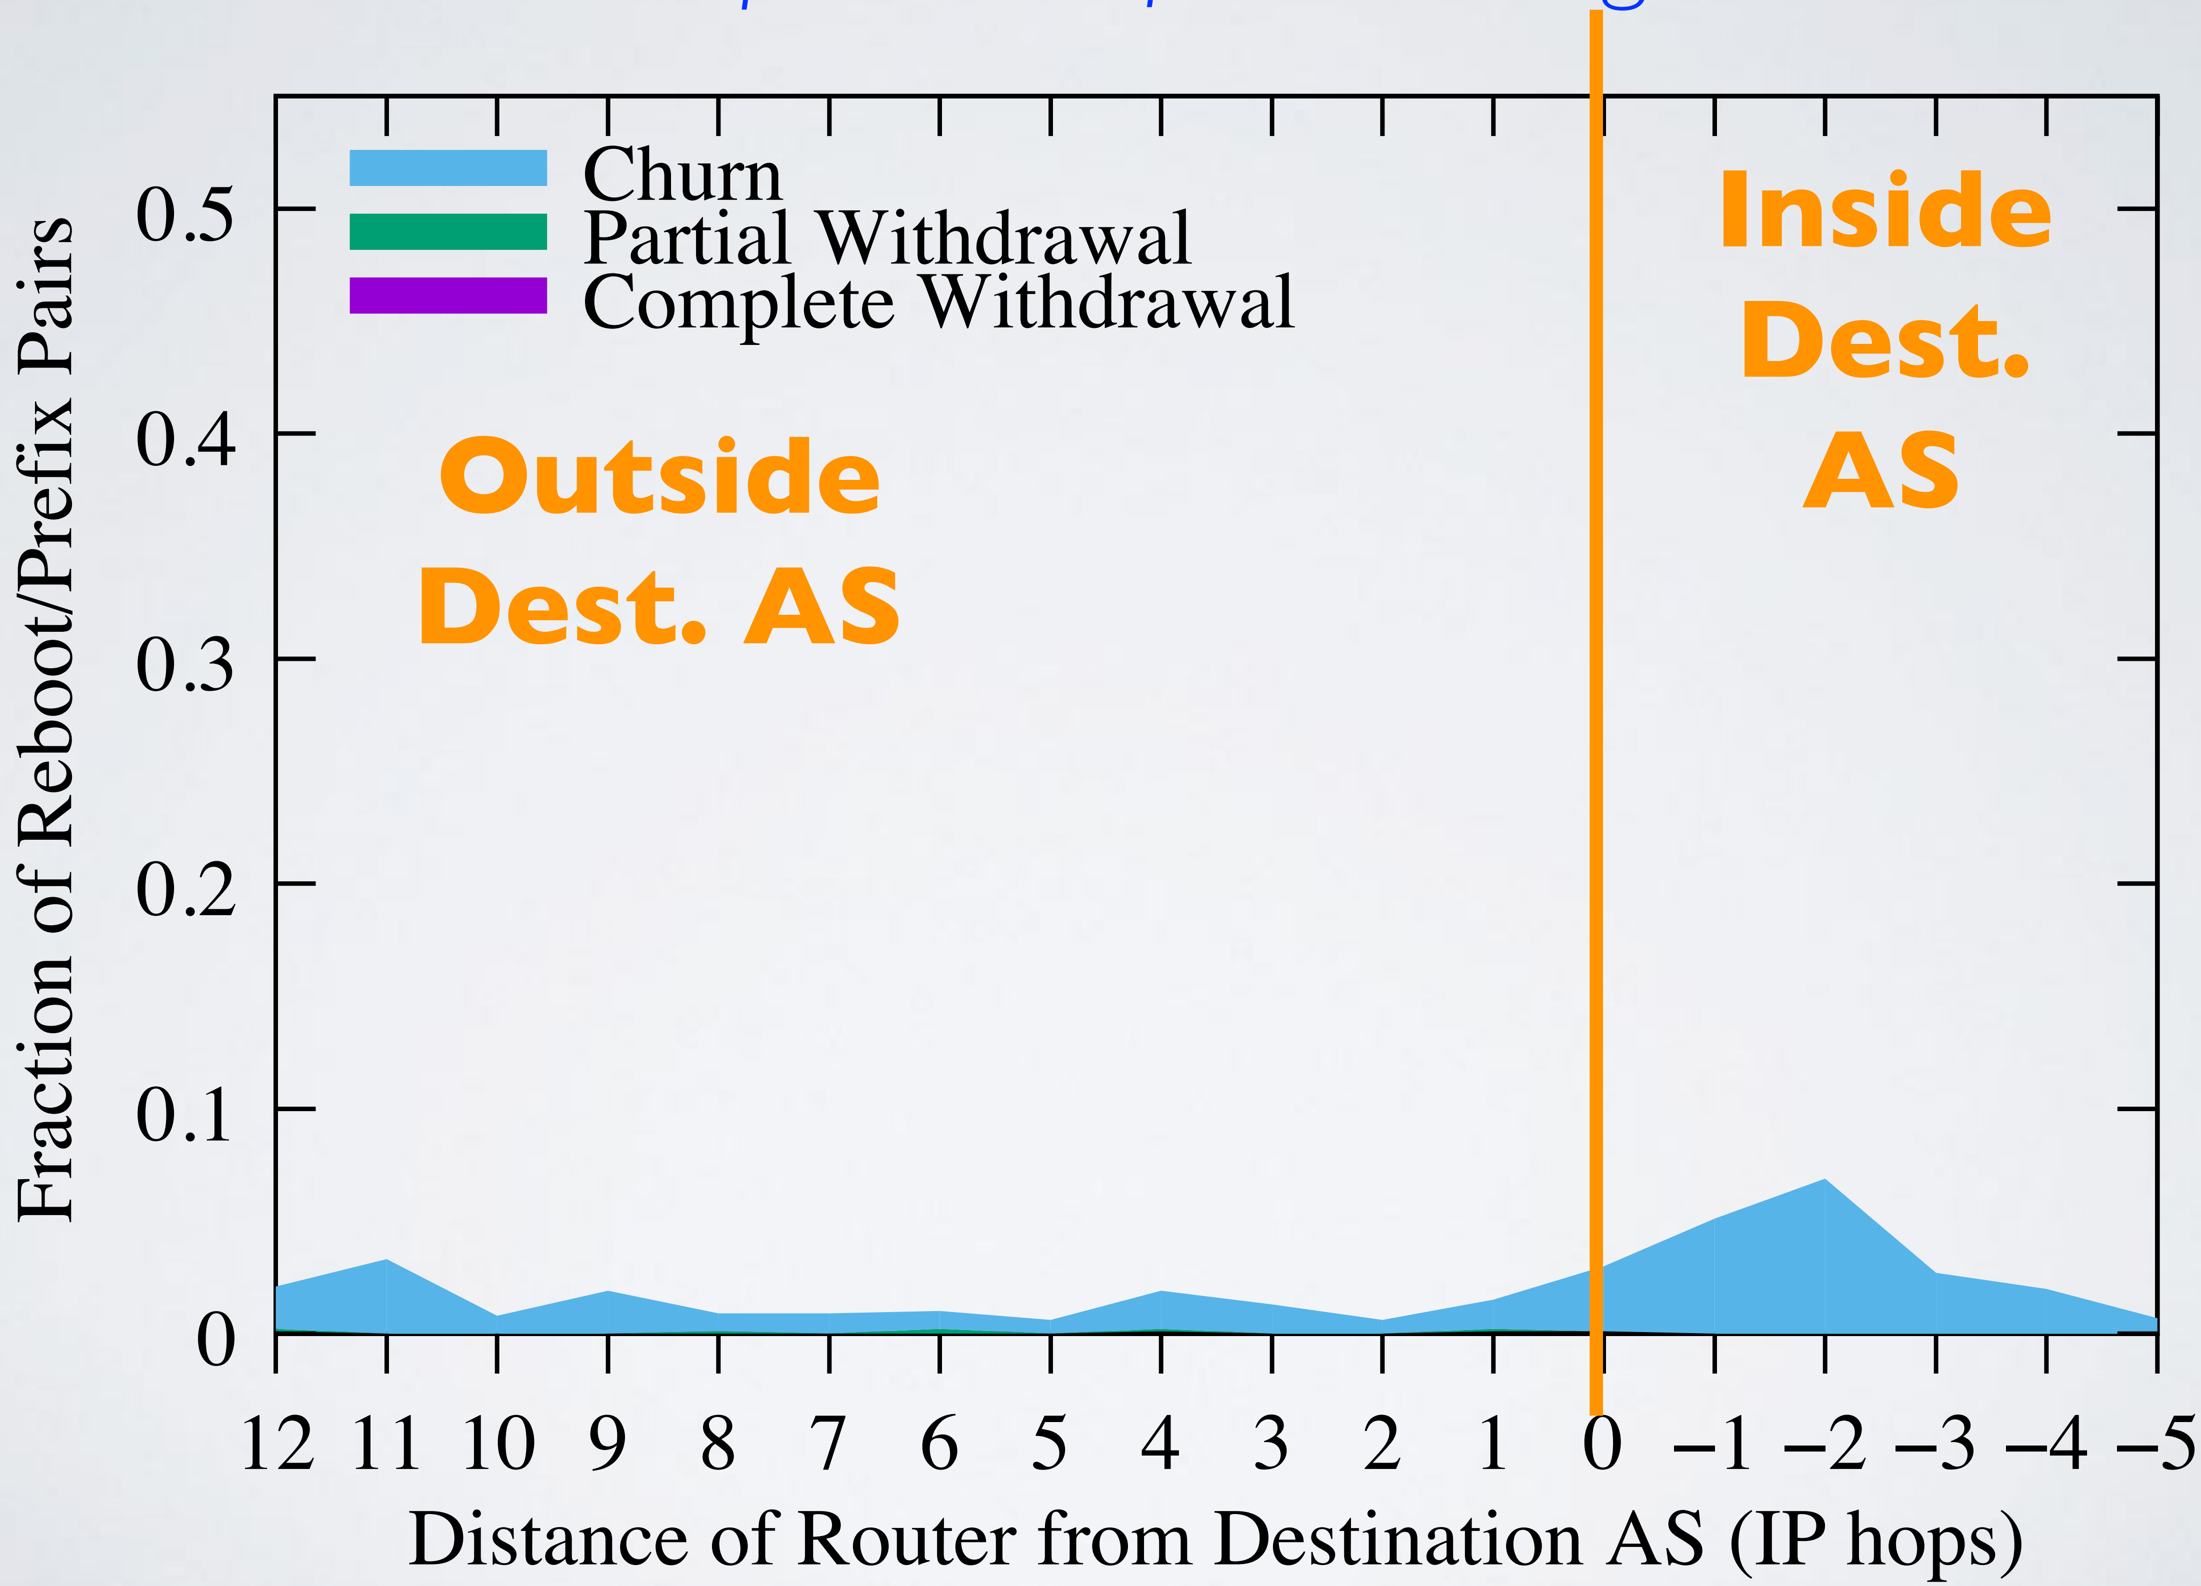

Correlating BGP/router outages

During the inferred outages, Feb 2015

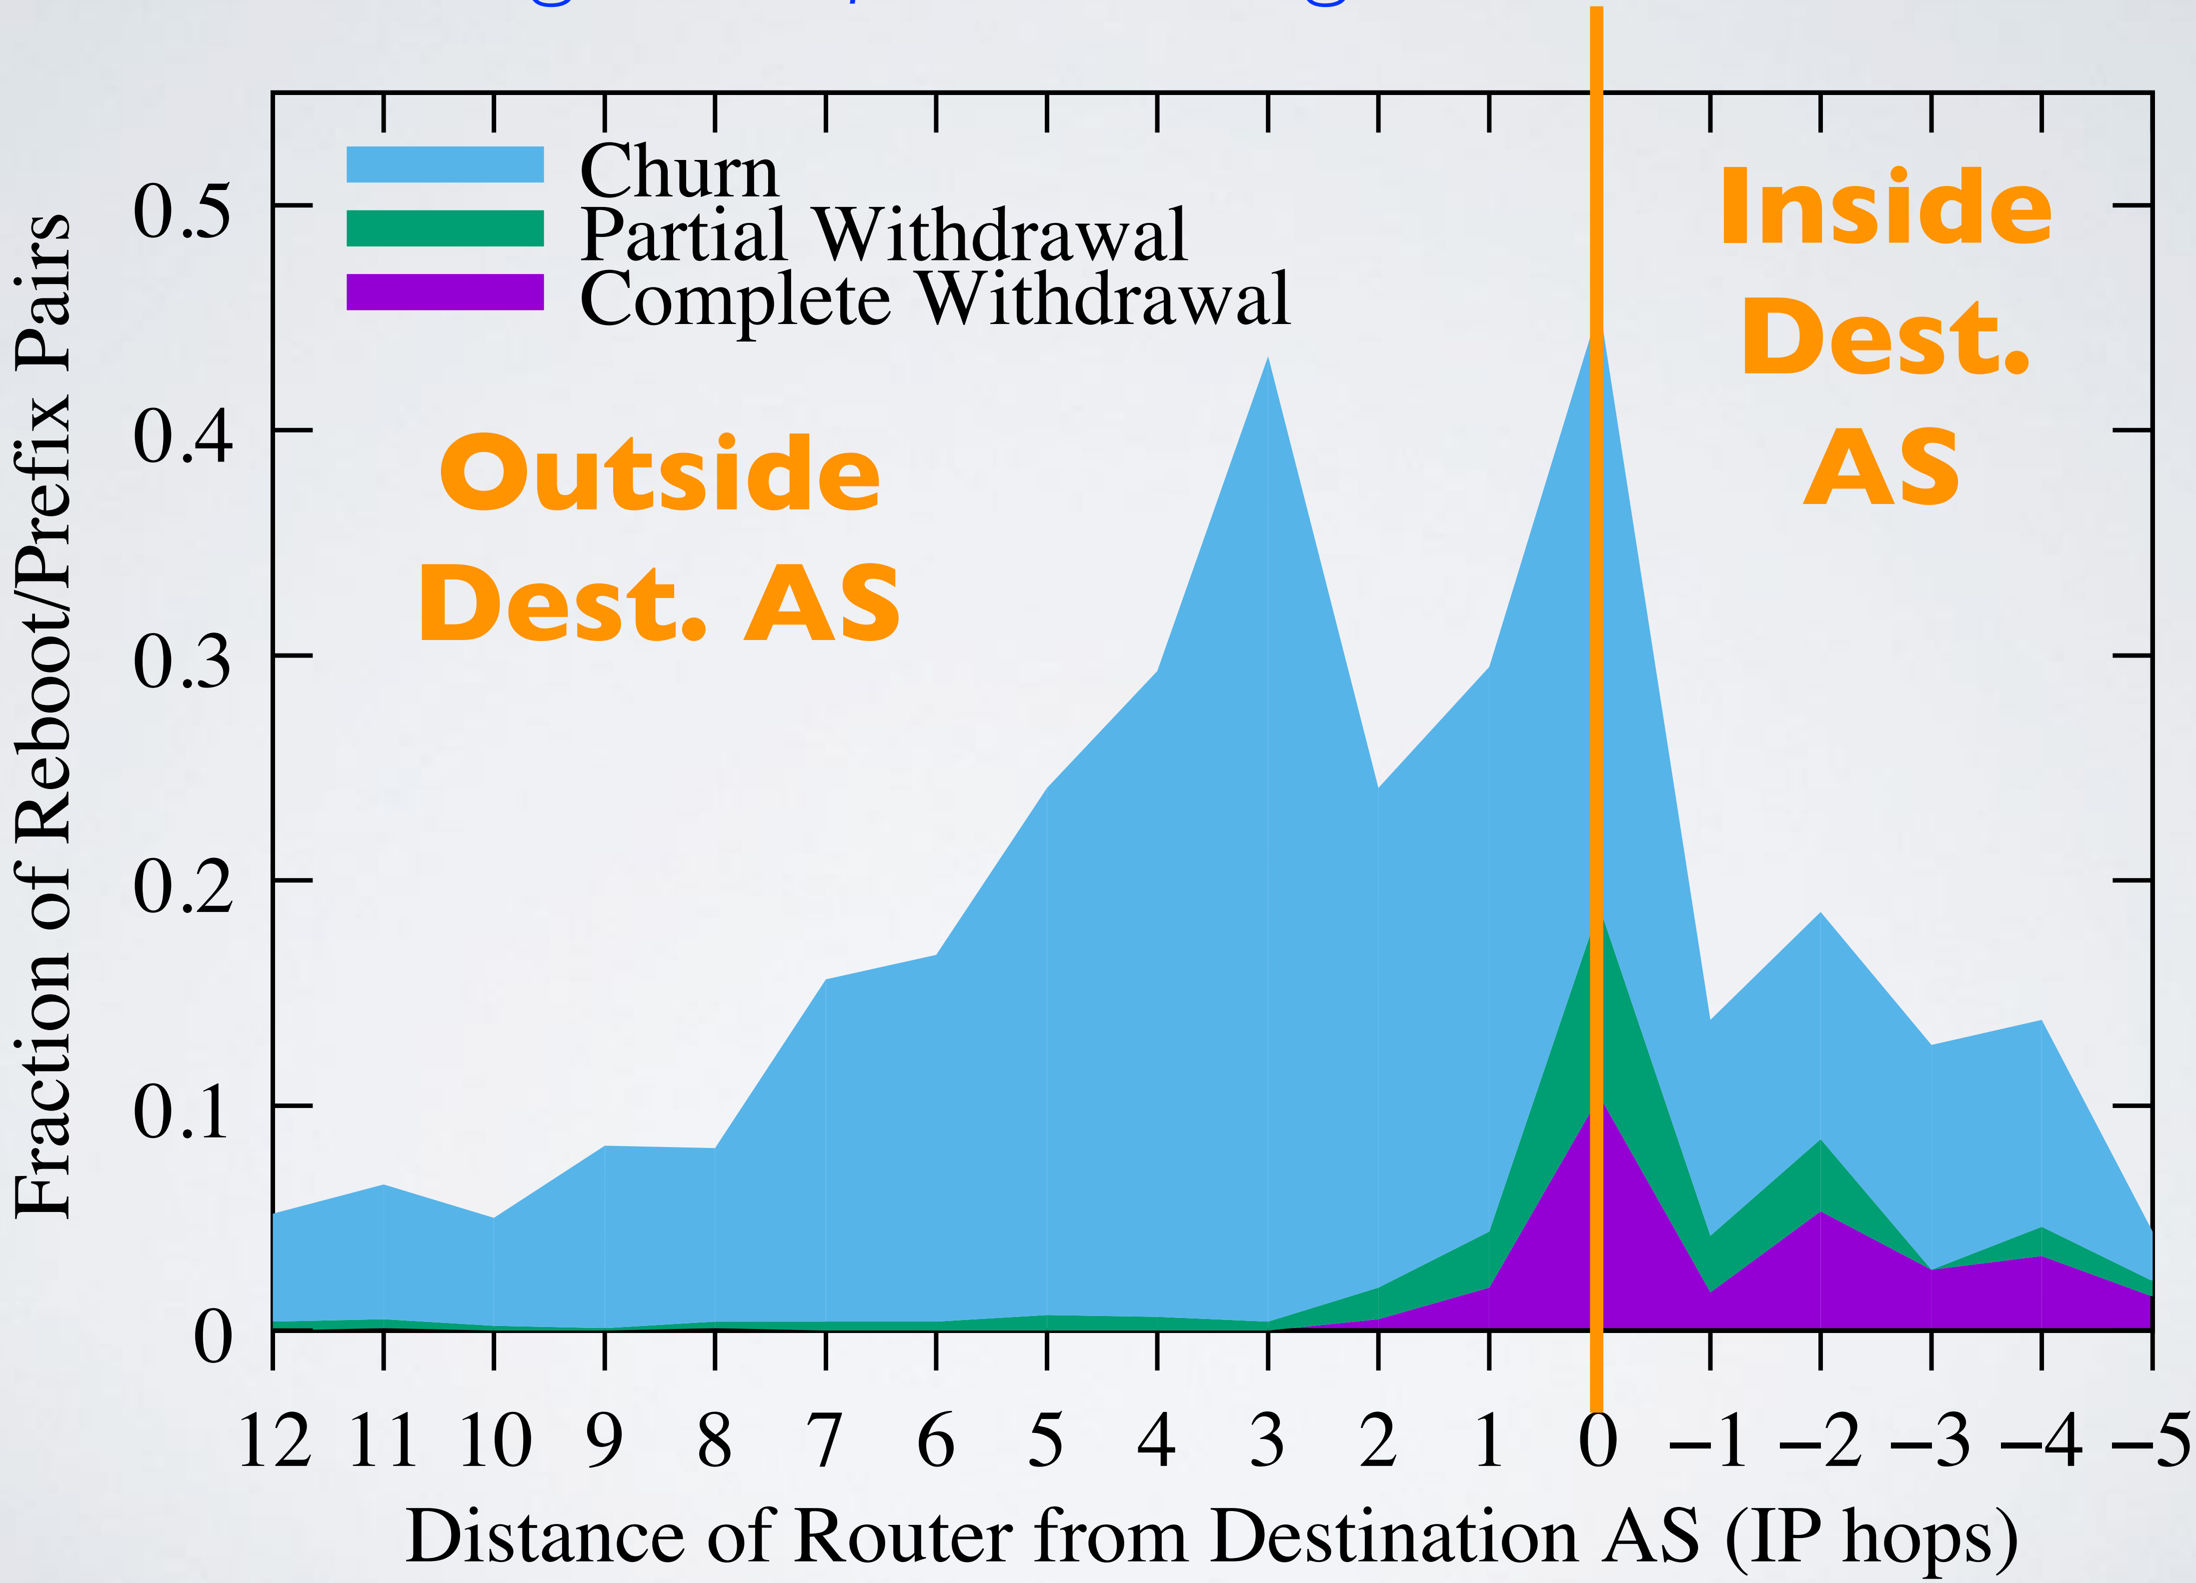

BGP Prefix Withdrawals: SPoF

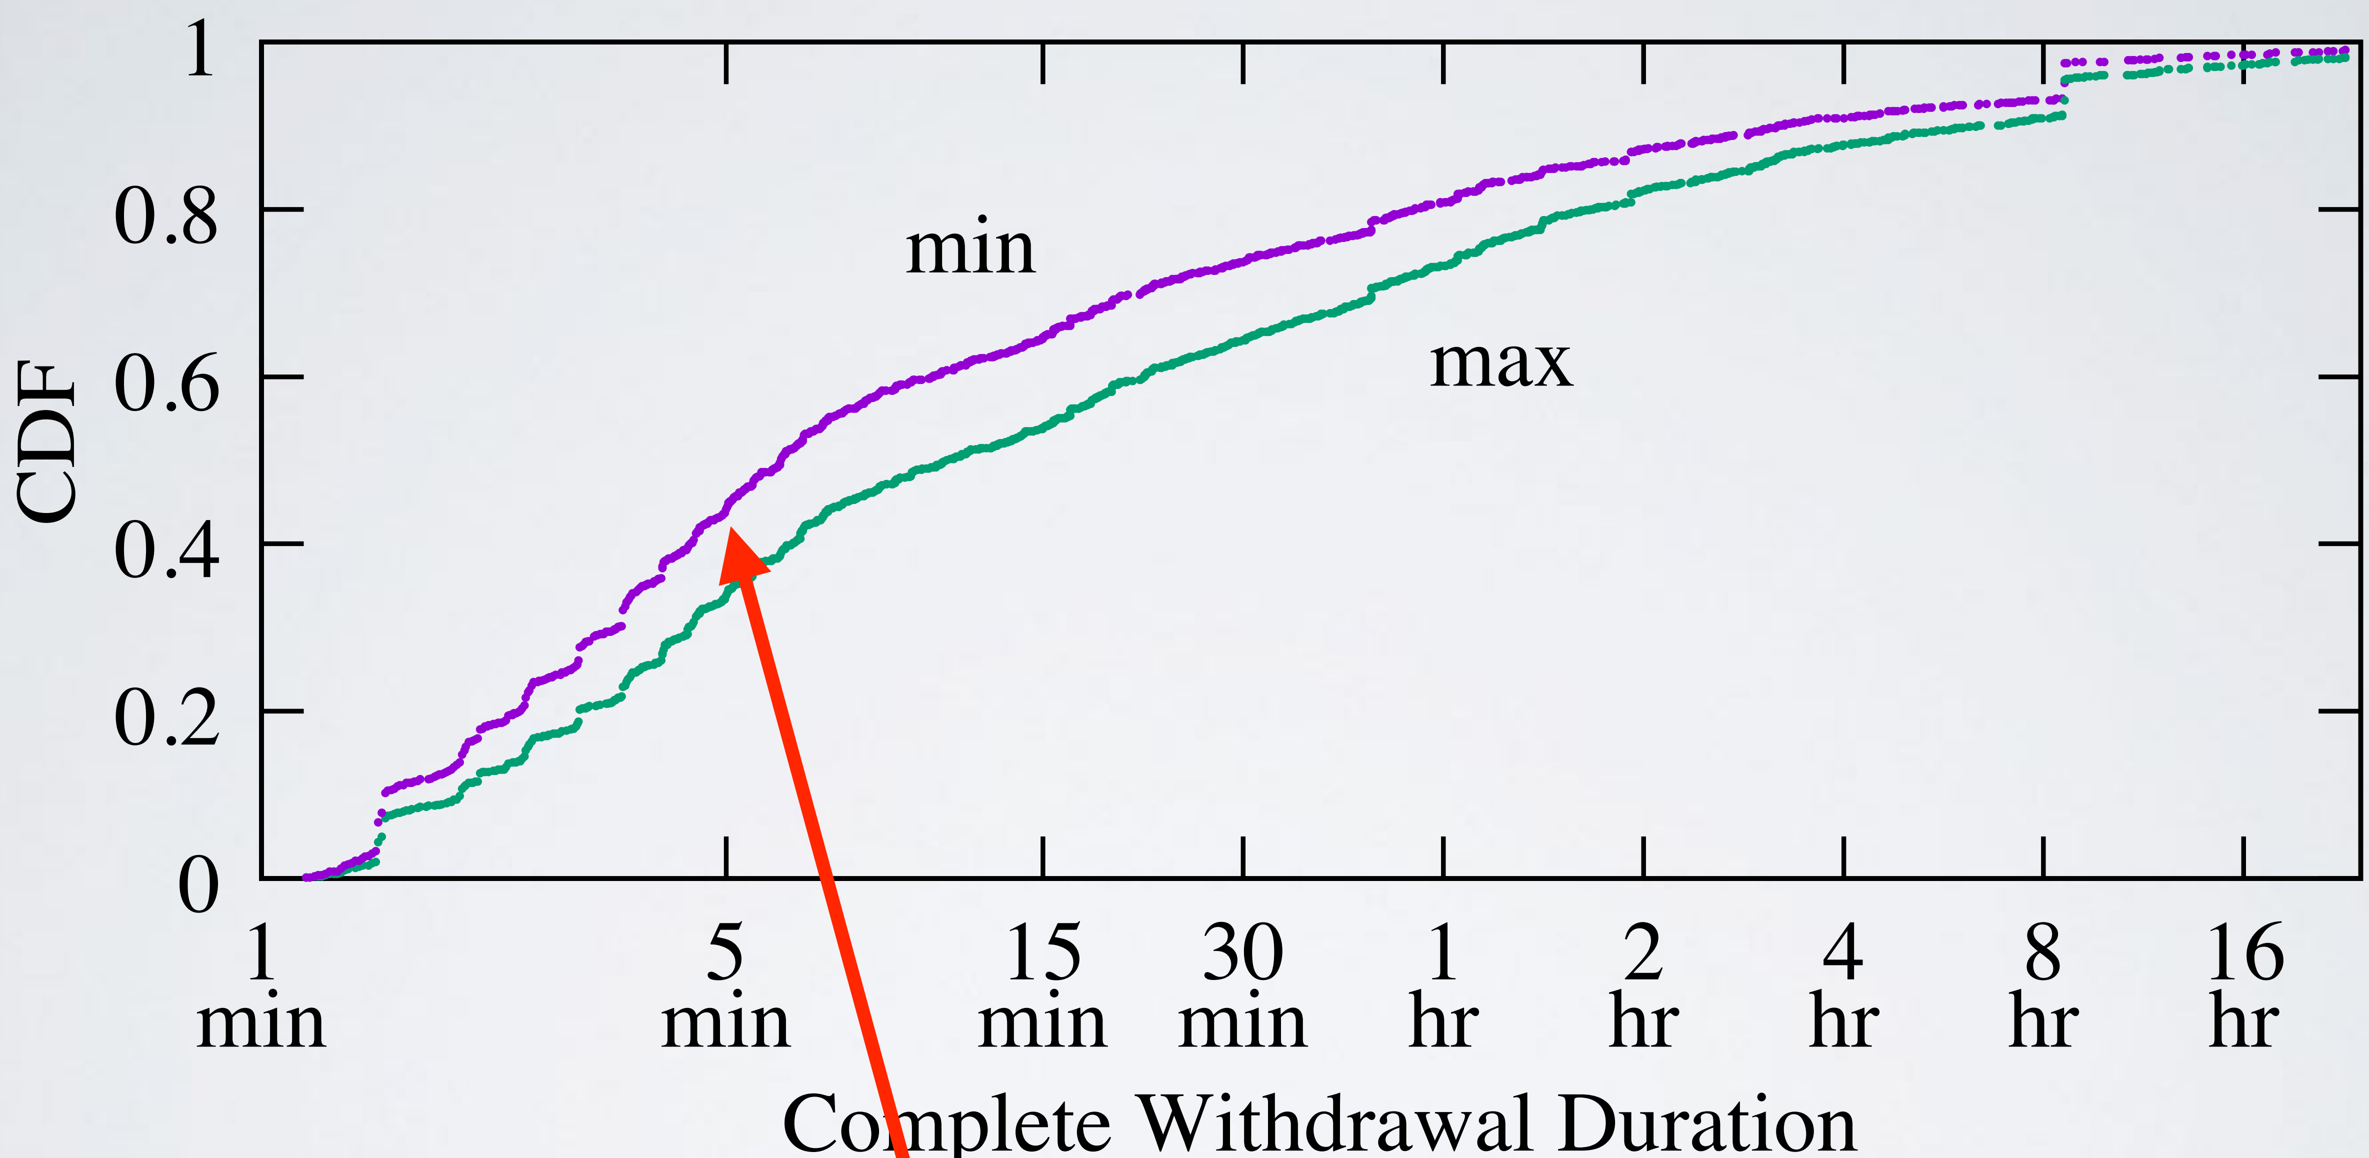

44% less than 5 minutes, suggestive of router maintenance or router crash

SPoF prefixes mostly single homed

Especially SPoFs outside destination AS, as expected

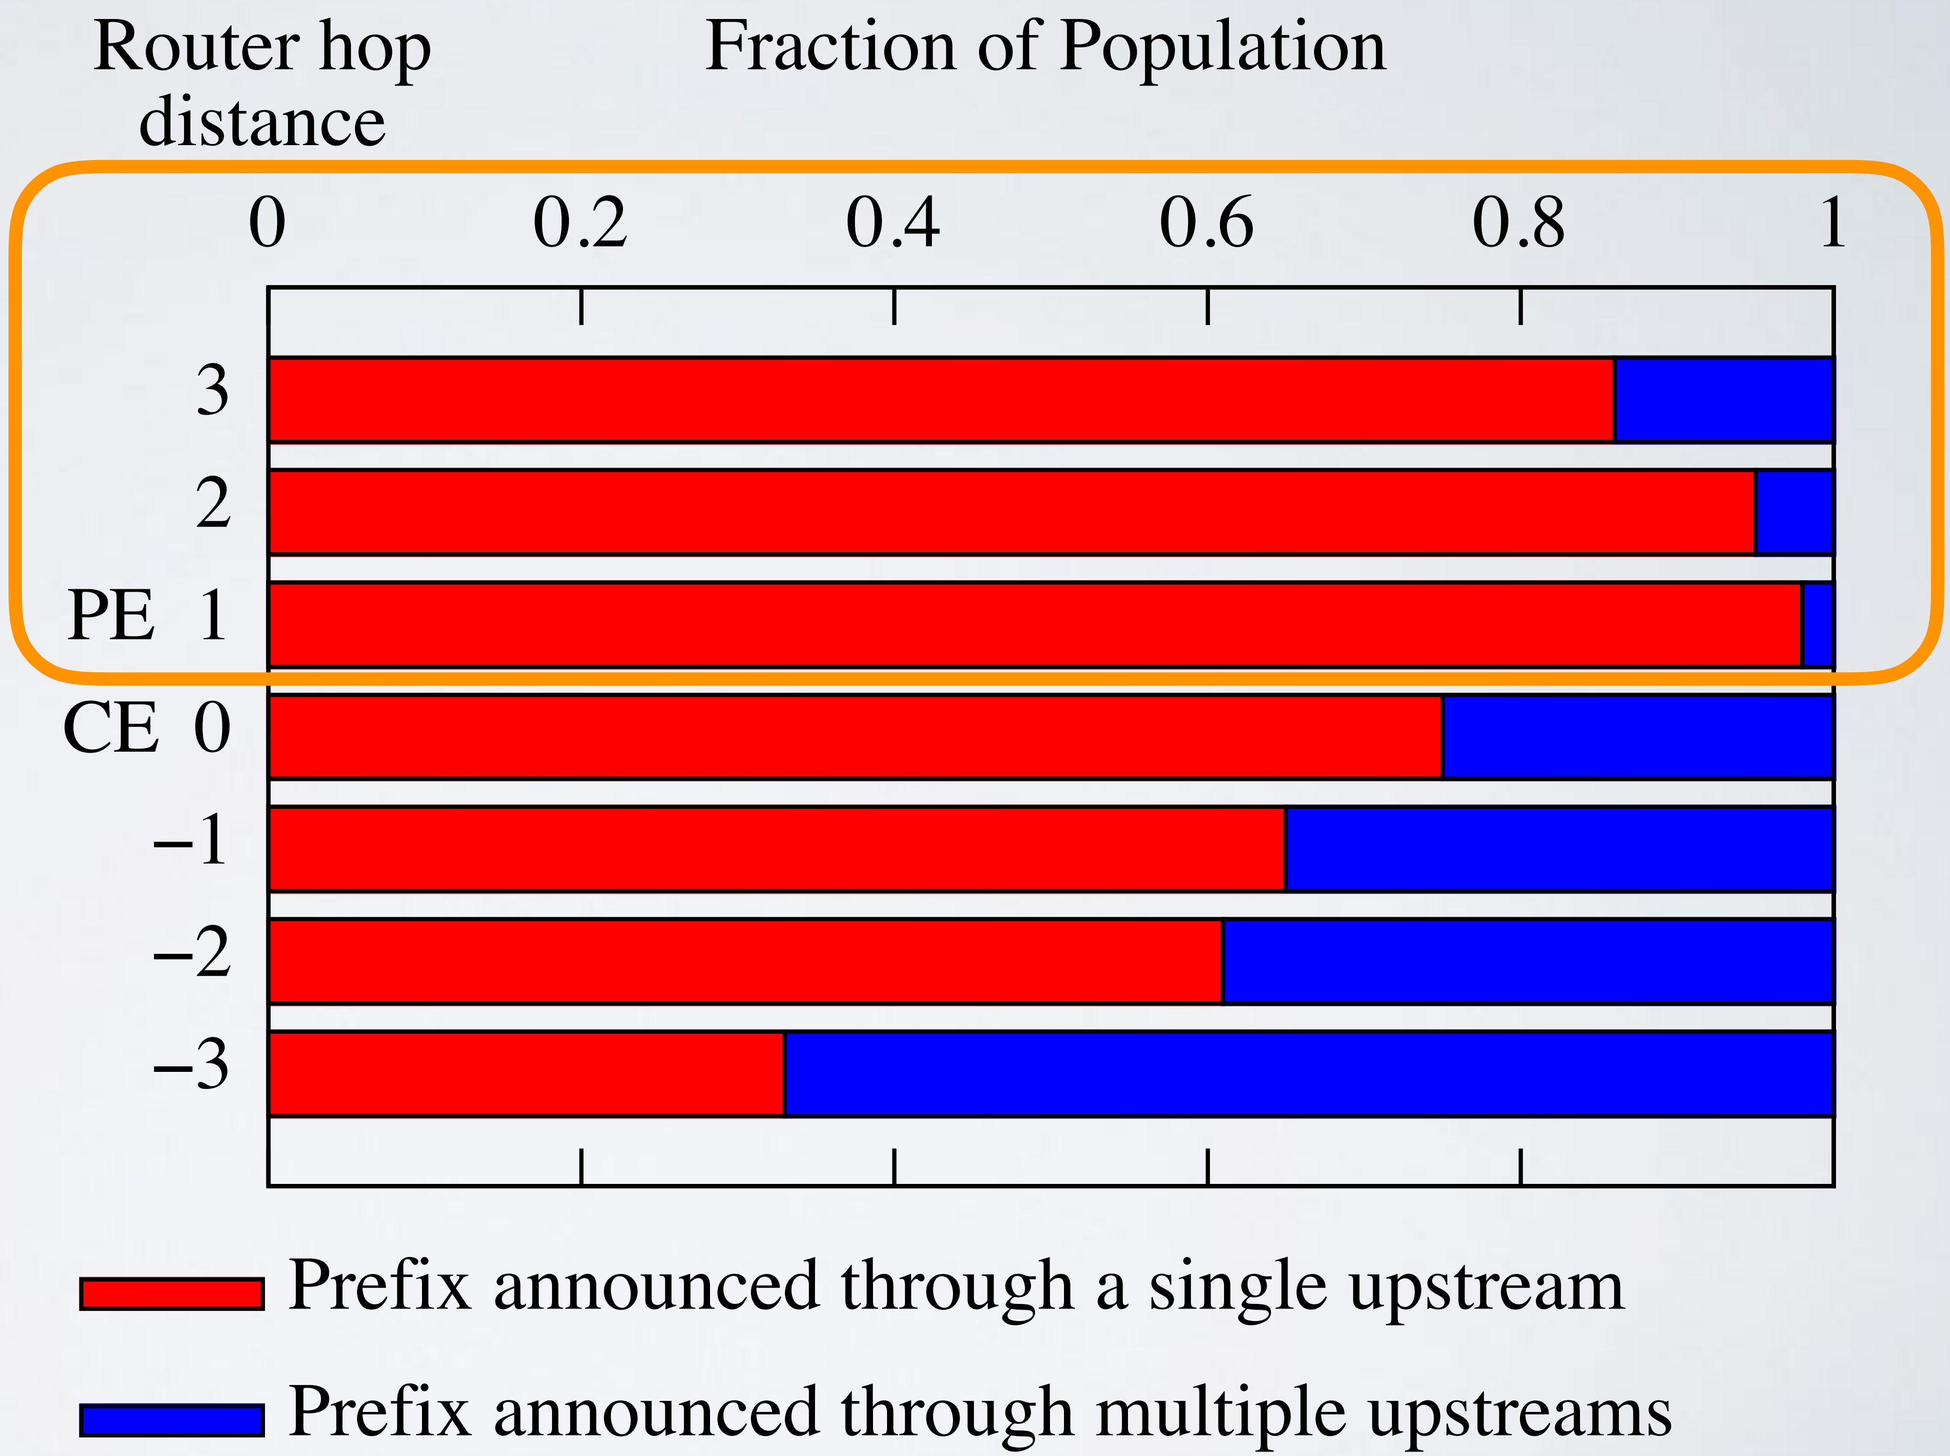

Impact on IPv4 prefixes in BGP

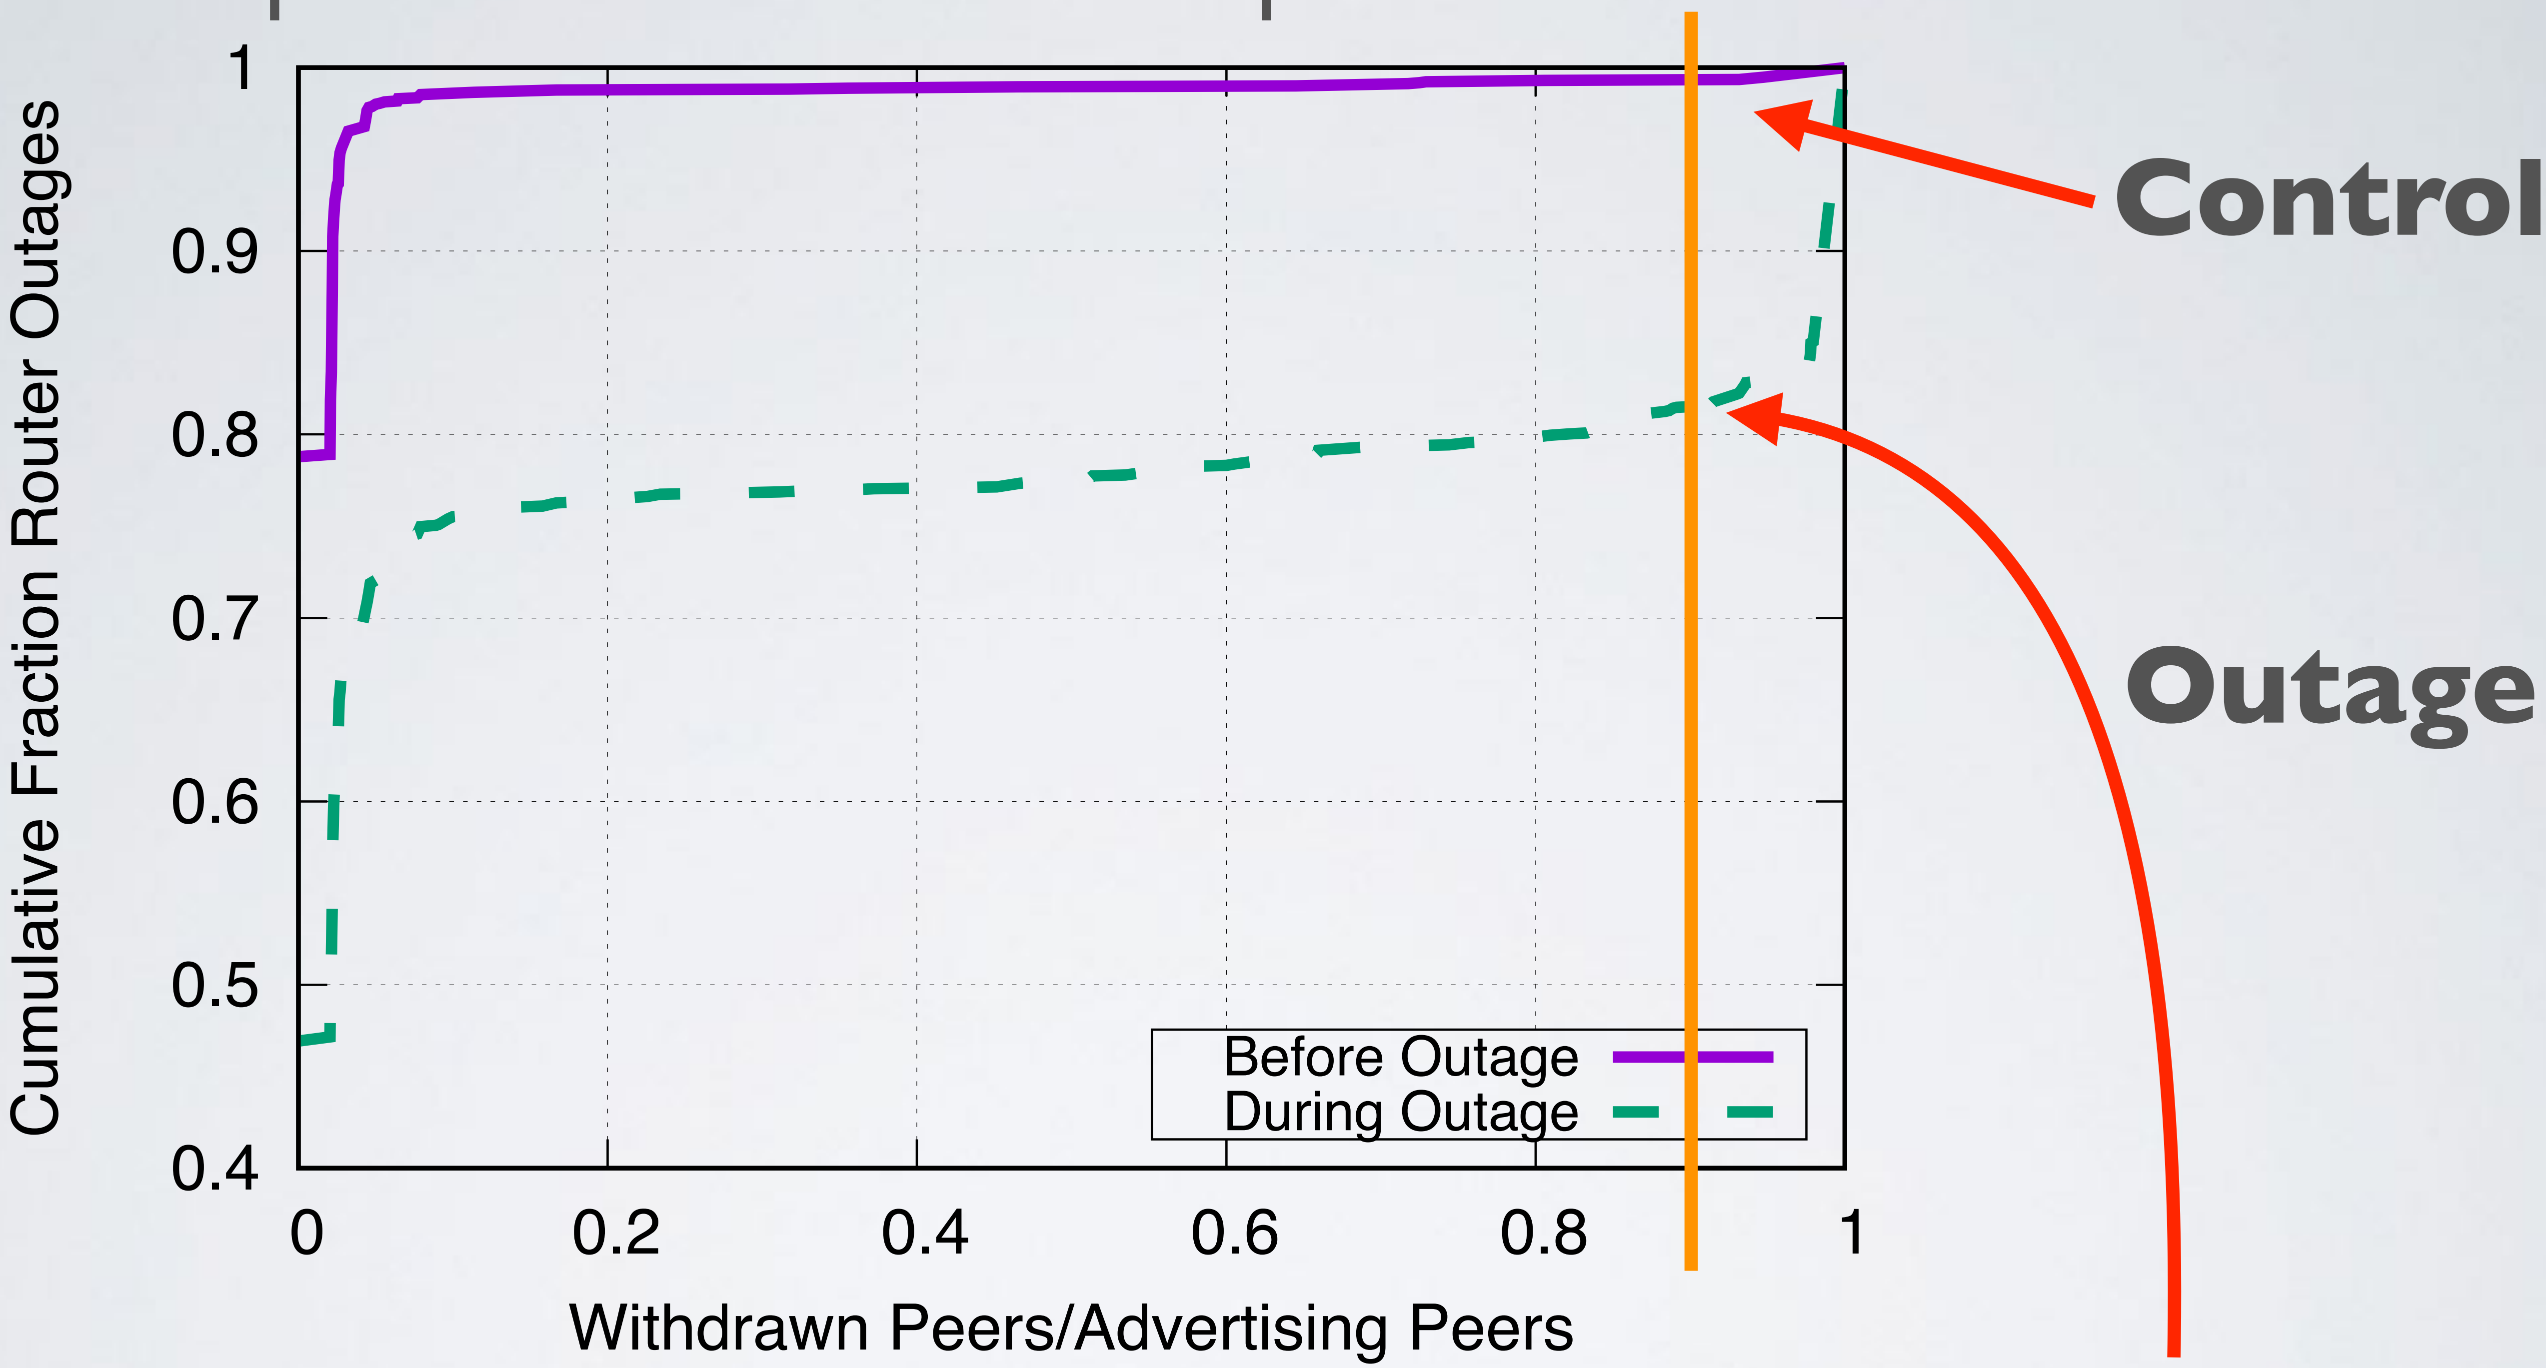

We examined IPv4 prefixes for 5% sample of reboots.
19% of correlated IPv4 prefixes withdrawn
by at least 90% of peers during router outage window.

Summary

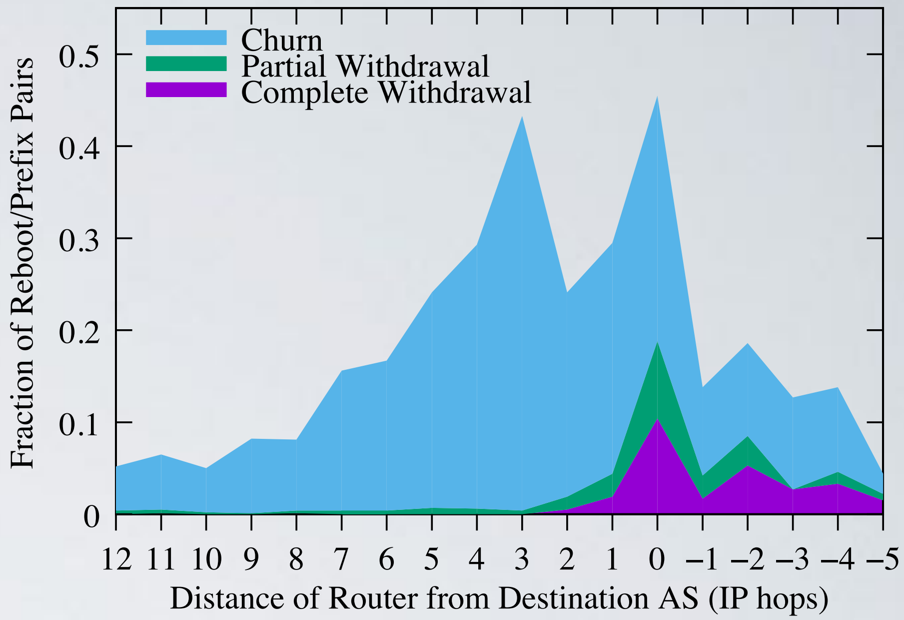

- Step towards root-cause analysis of inter-domain routing outages and events
  - Explore applicability of method to measurement of other critical Internet infrastructure: DNS, Web, Email
- In our 2.5 year sample of 59K routers that rebooted
  - 4% (2.3K) were SPoF
  - SPoF were mostly confined to the edge: 59% customer edge
- We released our code as part of scamper

https://www.caida.org/tools/measurement/scamper/

Backup Slides

Impact on IPv4 Services

censys.io April 2017

Active Hosts	**39,107**
HTTP	25,592
HTTPS	16,321
SSH	11,277
DNS	7,922
SMTP	7,383
IMAP	5,127

Web

Email

We examined IPv4 prefixes for 5% sample of reboots where at least 90% of peers during router outage window.

Partial Withdrawals

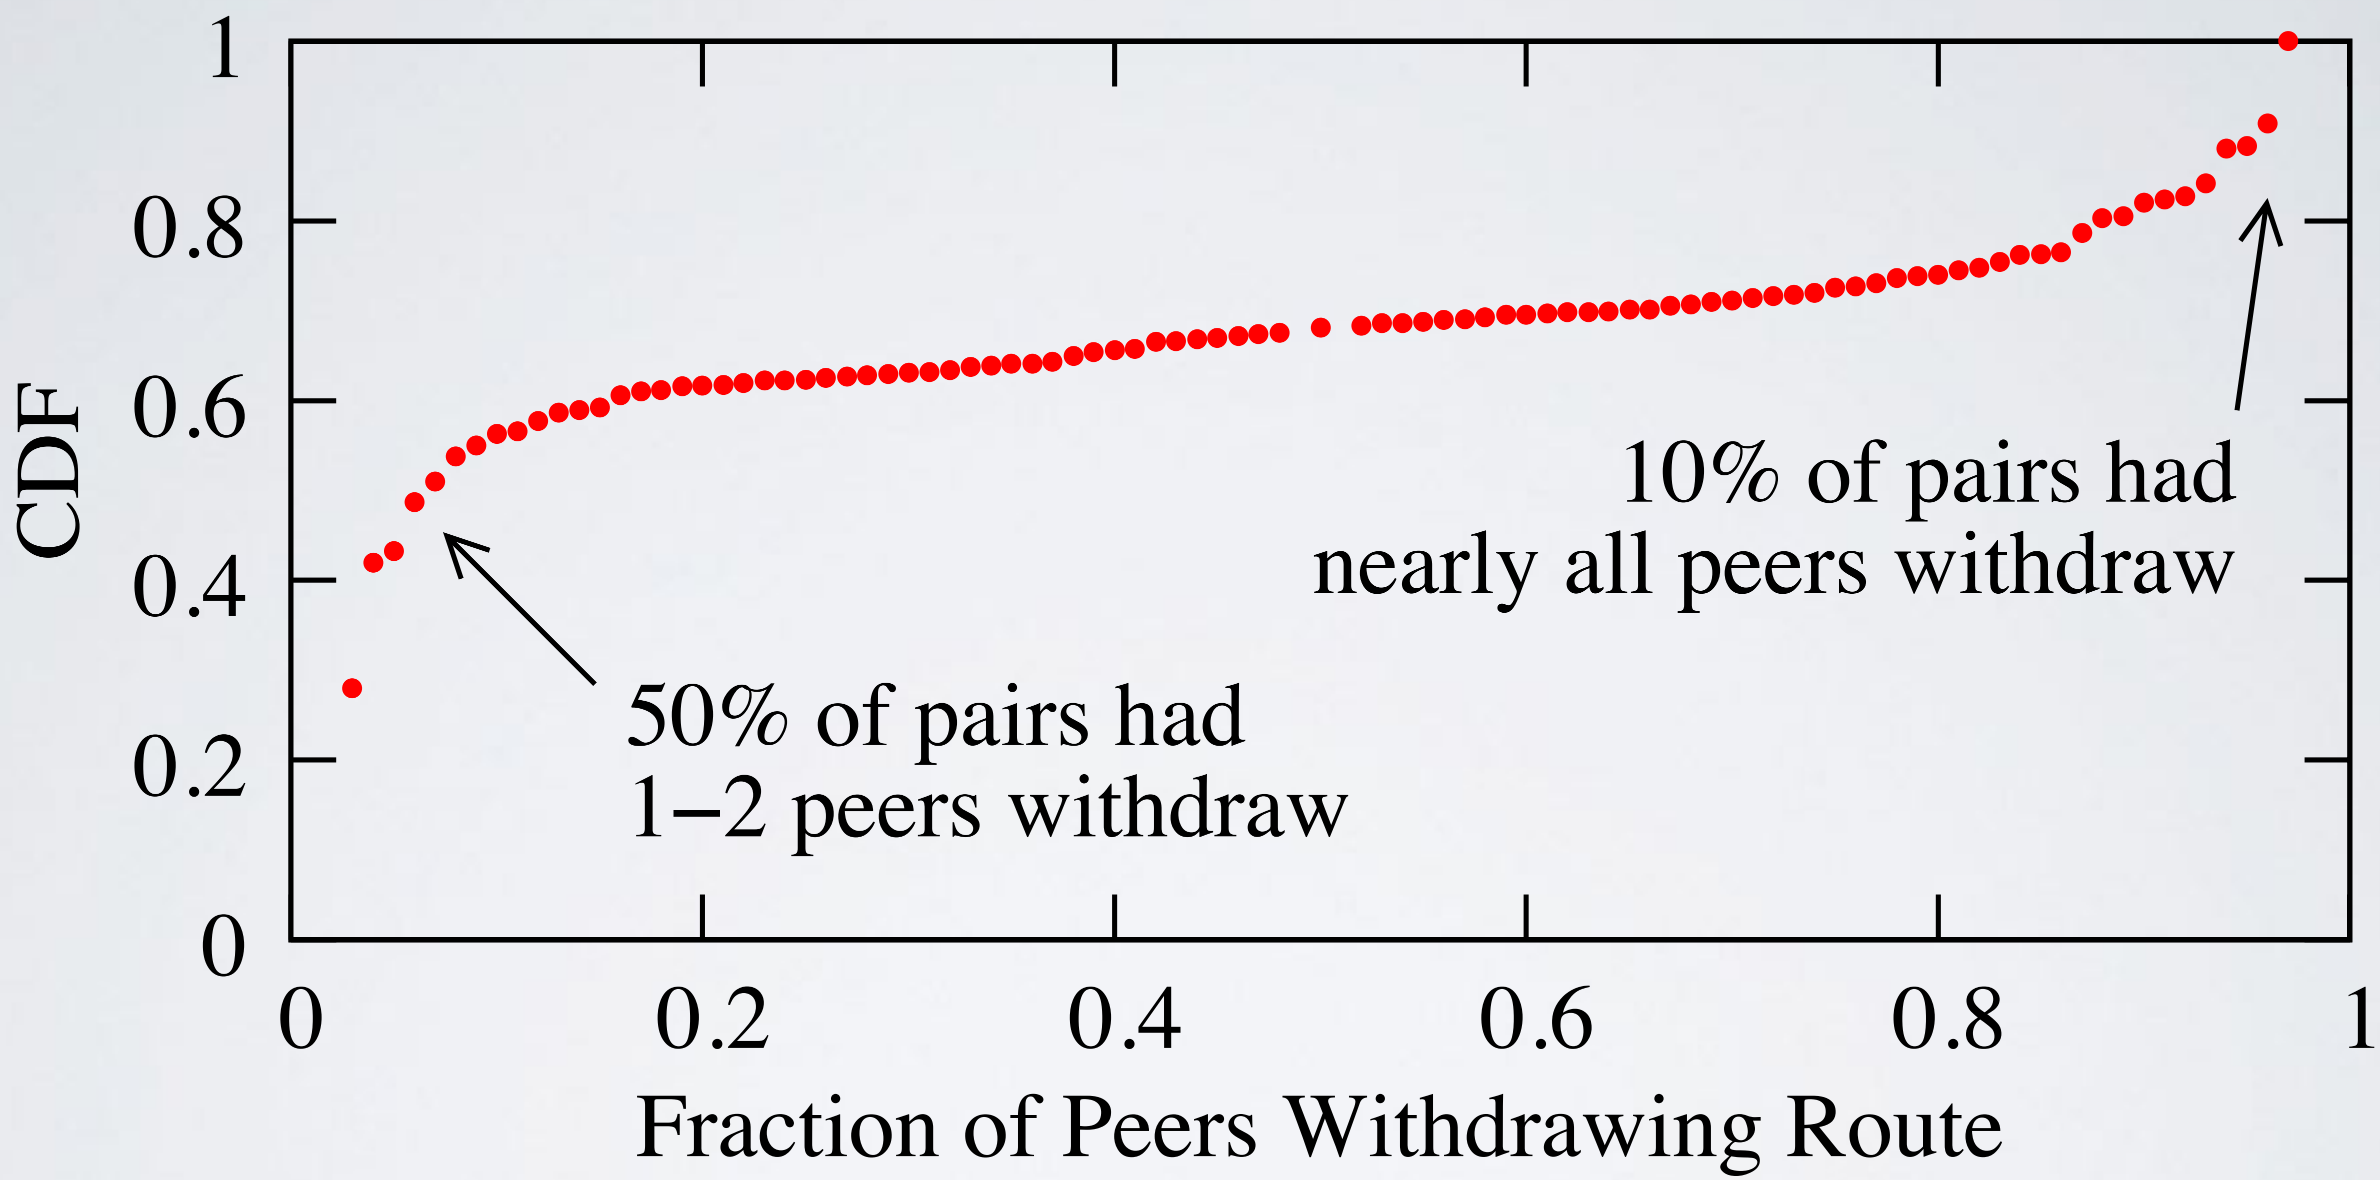

50% of pairs had 1-2 peers withdraw prefix
10% of pairs had nearly all peers withdraw prefix

Degrees of ASes monitored

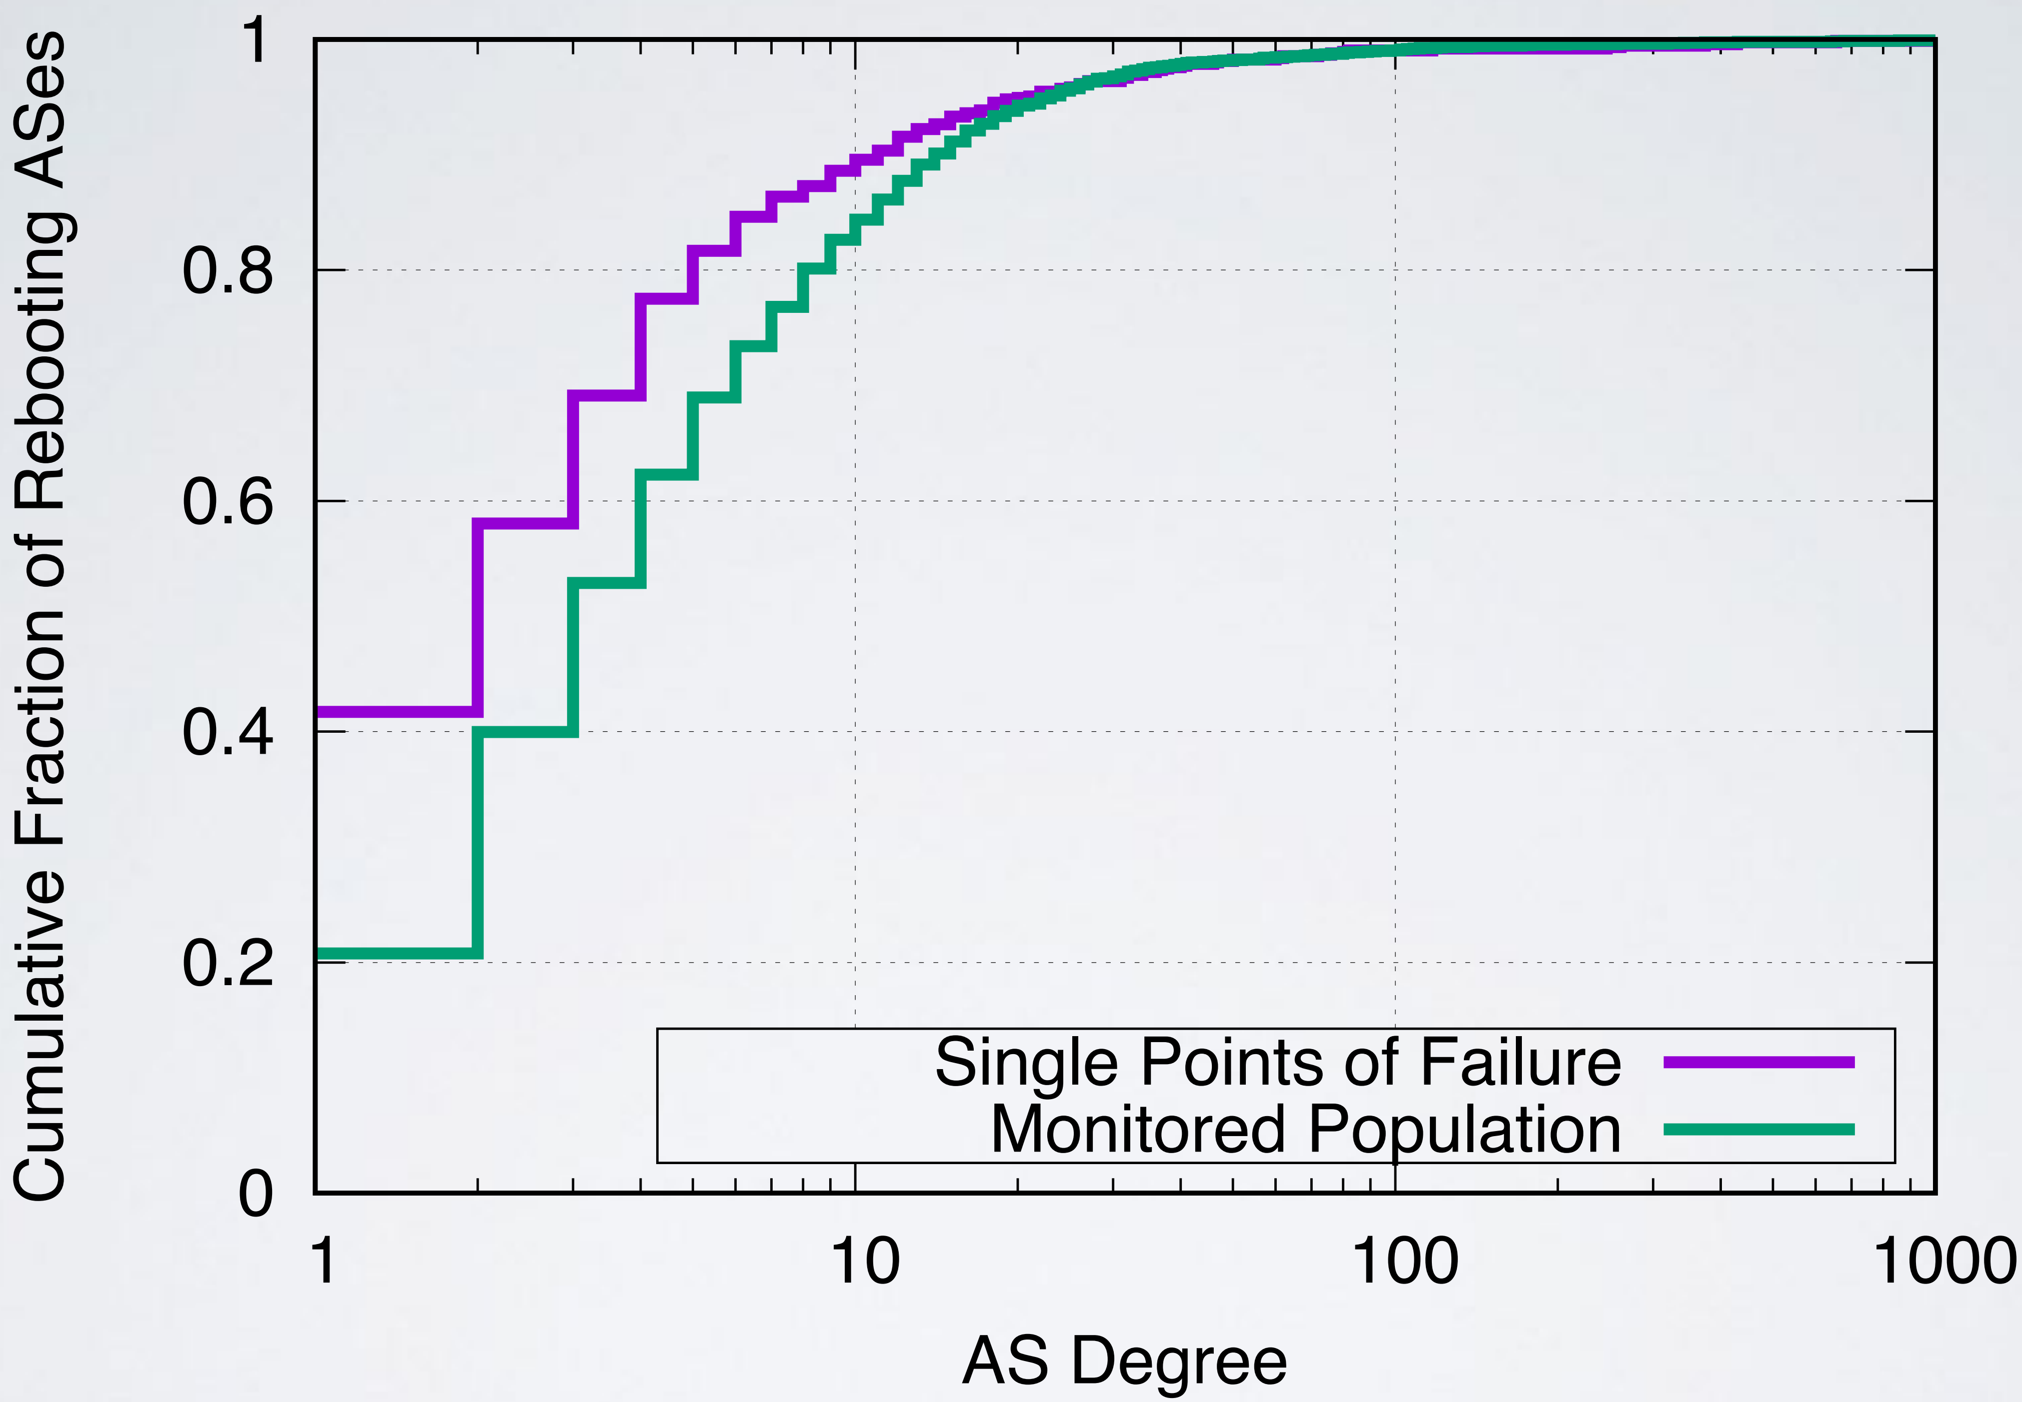

ASes that were inferred to have a SPoF were disproportionately low-degree ASes

Activity for IPv4 prefixes in BGP

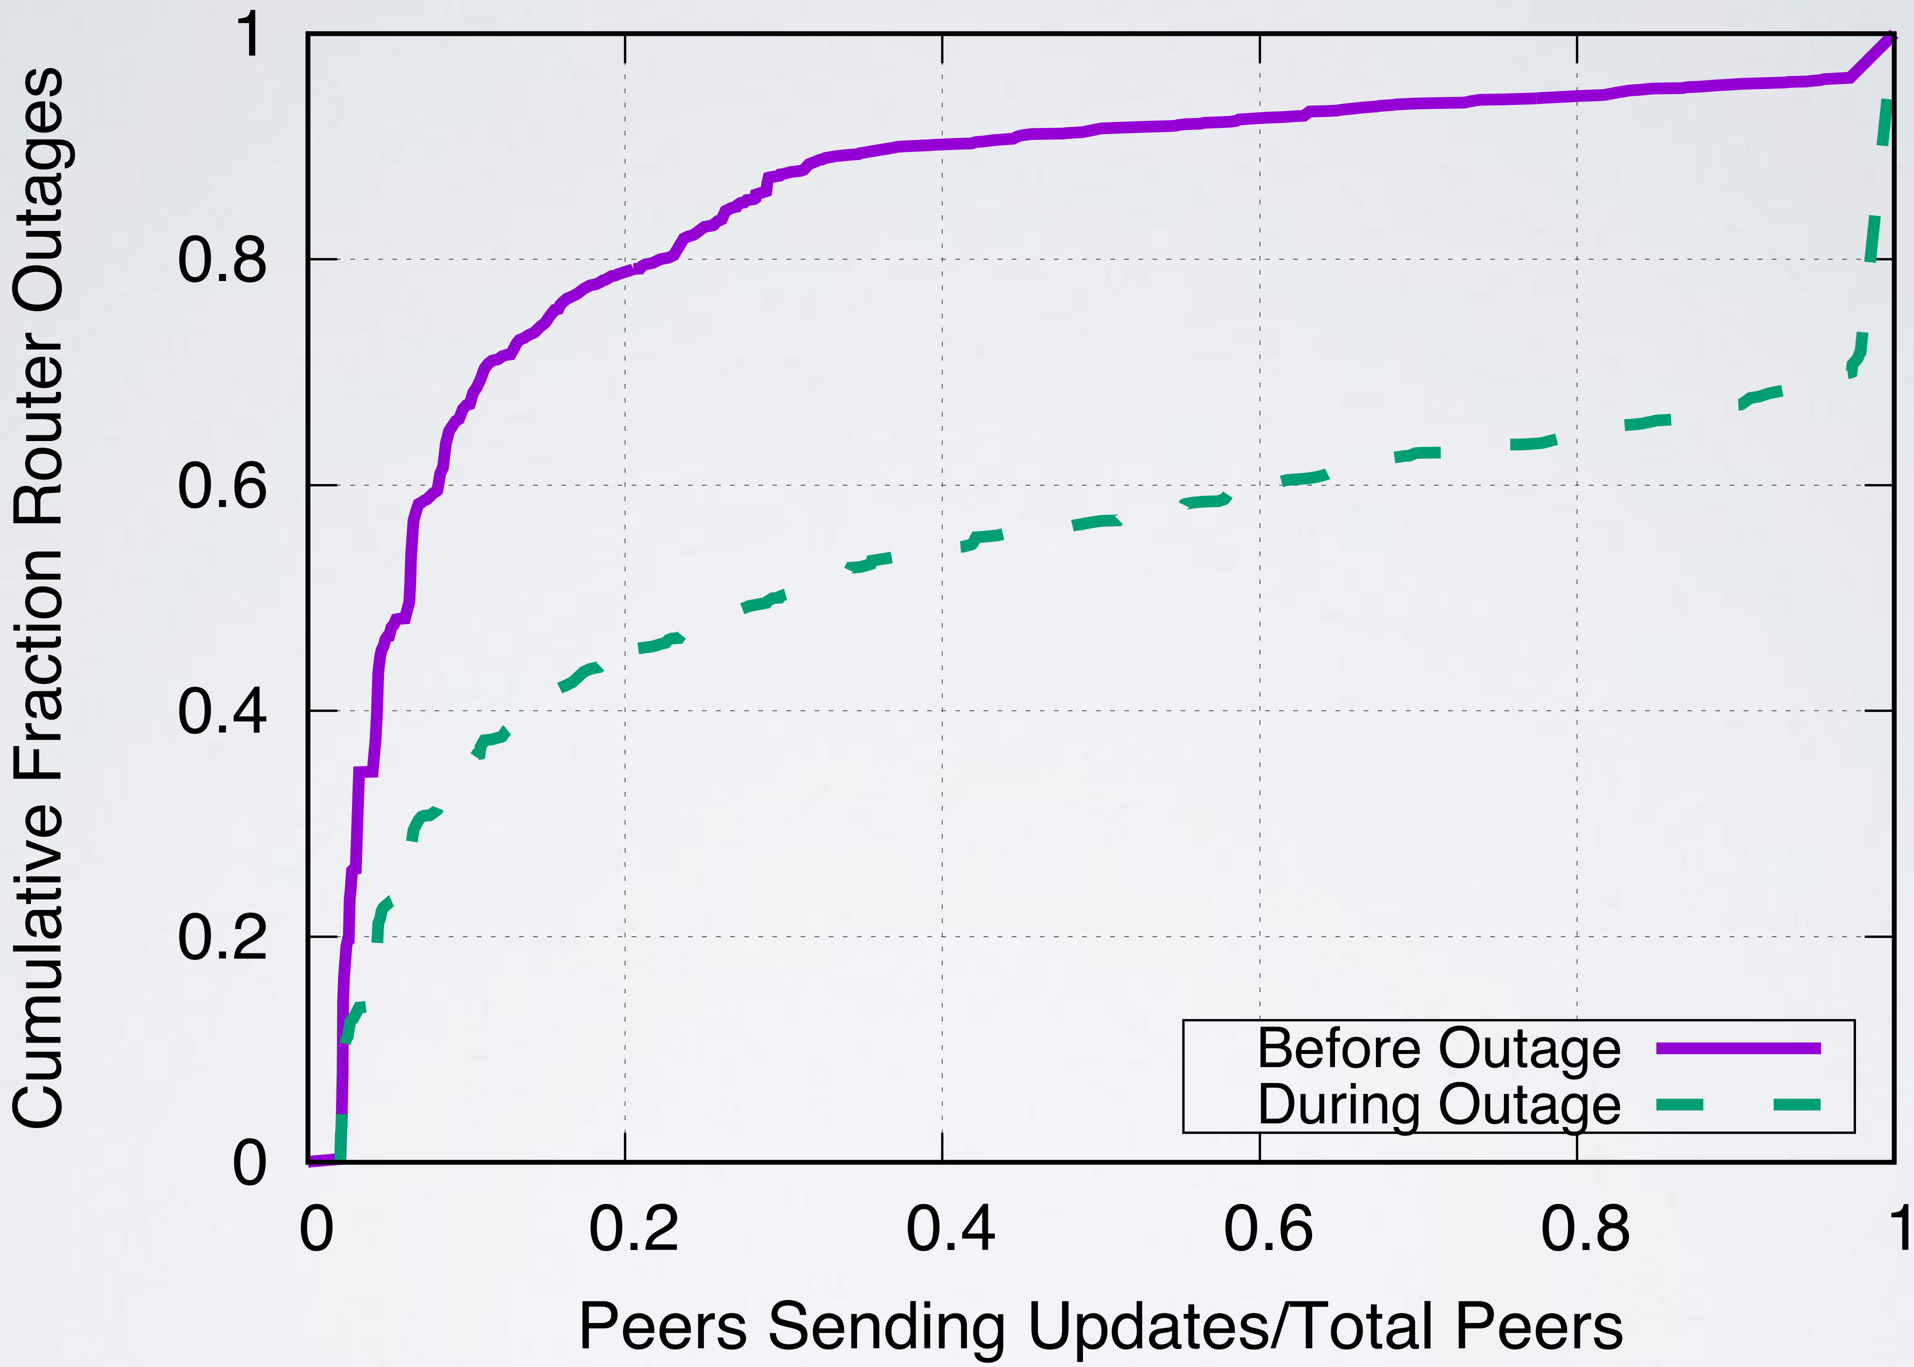

At least 70% of peers reported BGP activity on IPv4 prefixes for 50% of the inferred router outages

Reboot Window Durations

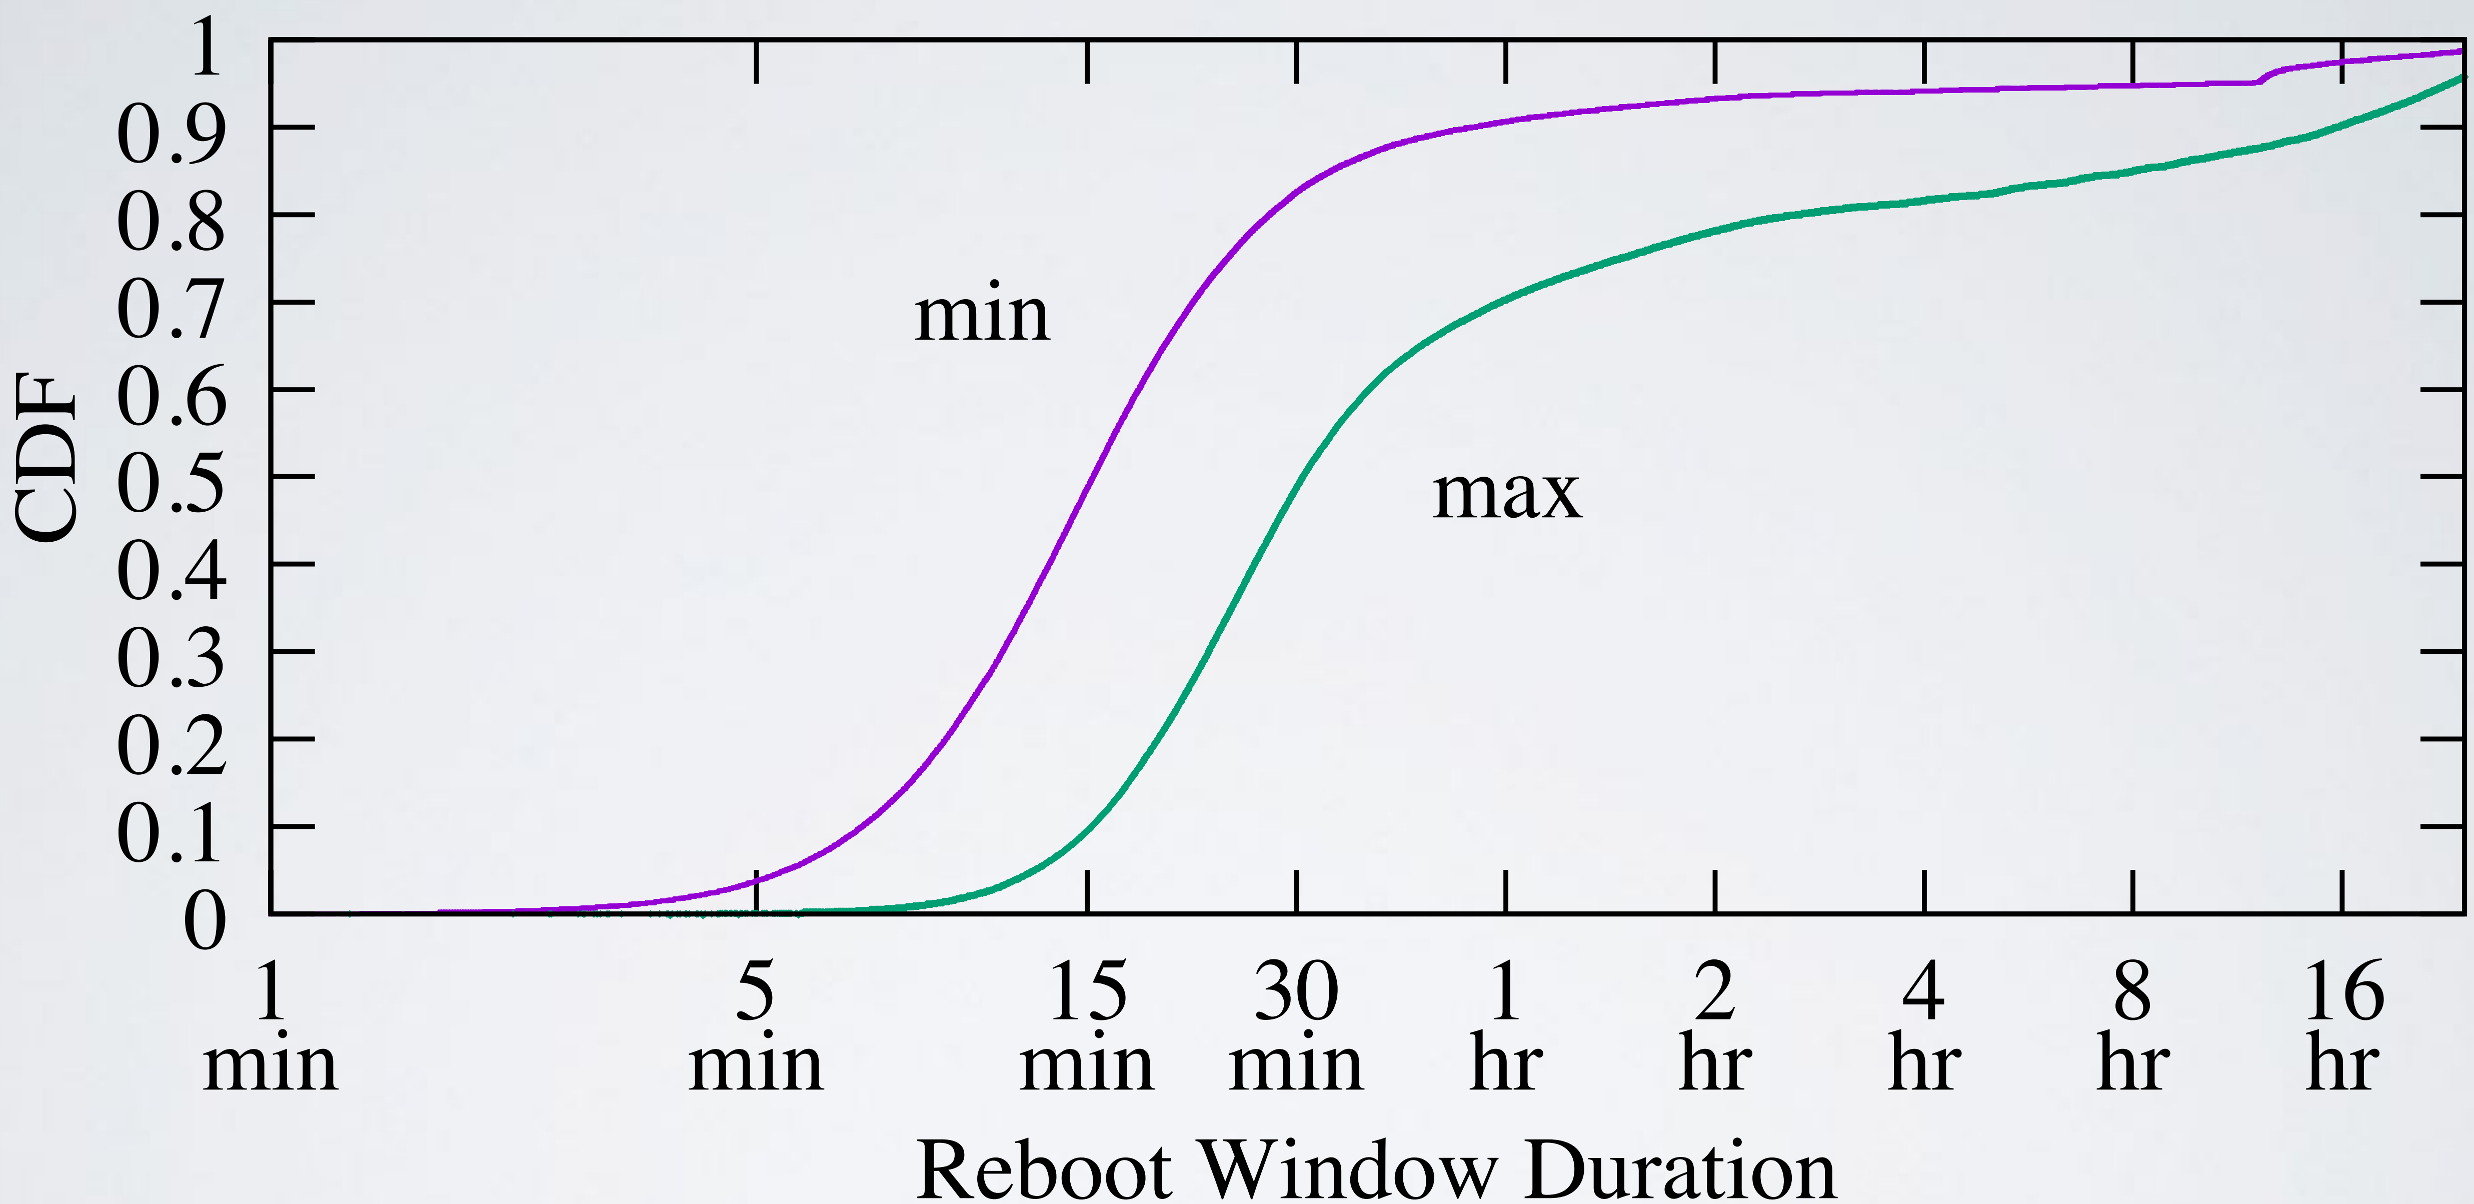

Half the maximum reboot lengths were less than 30 minutes (~two probing rounds)

Router + BGP outage correlation

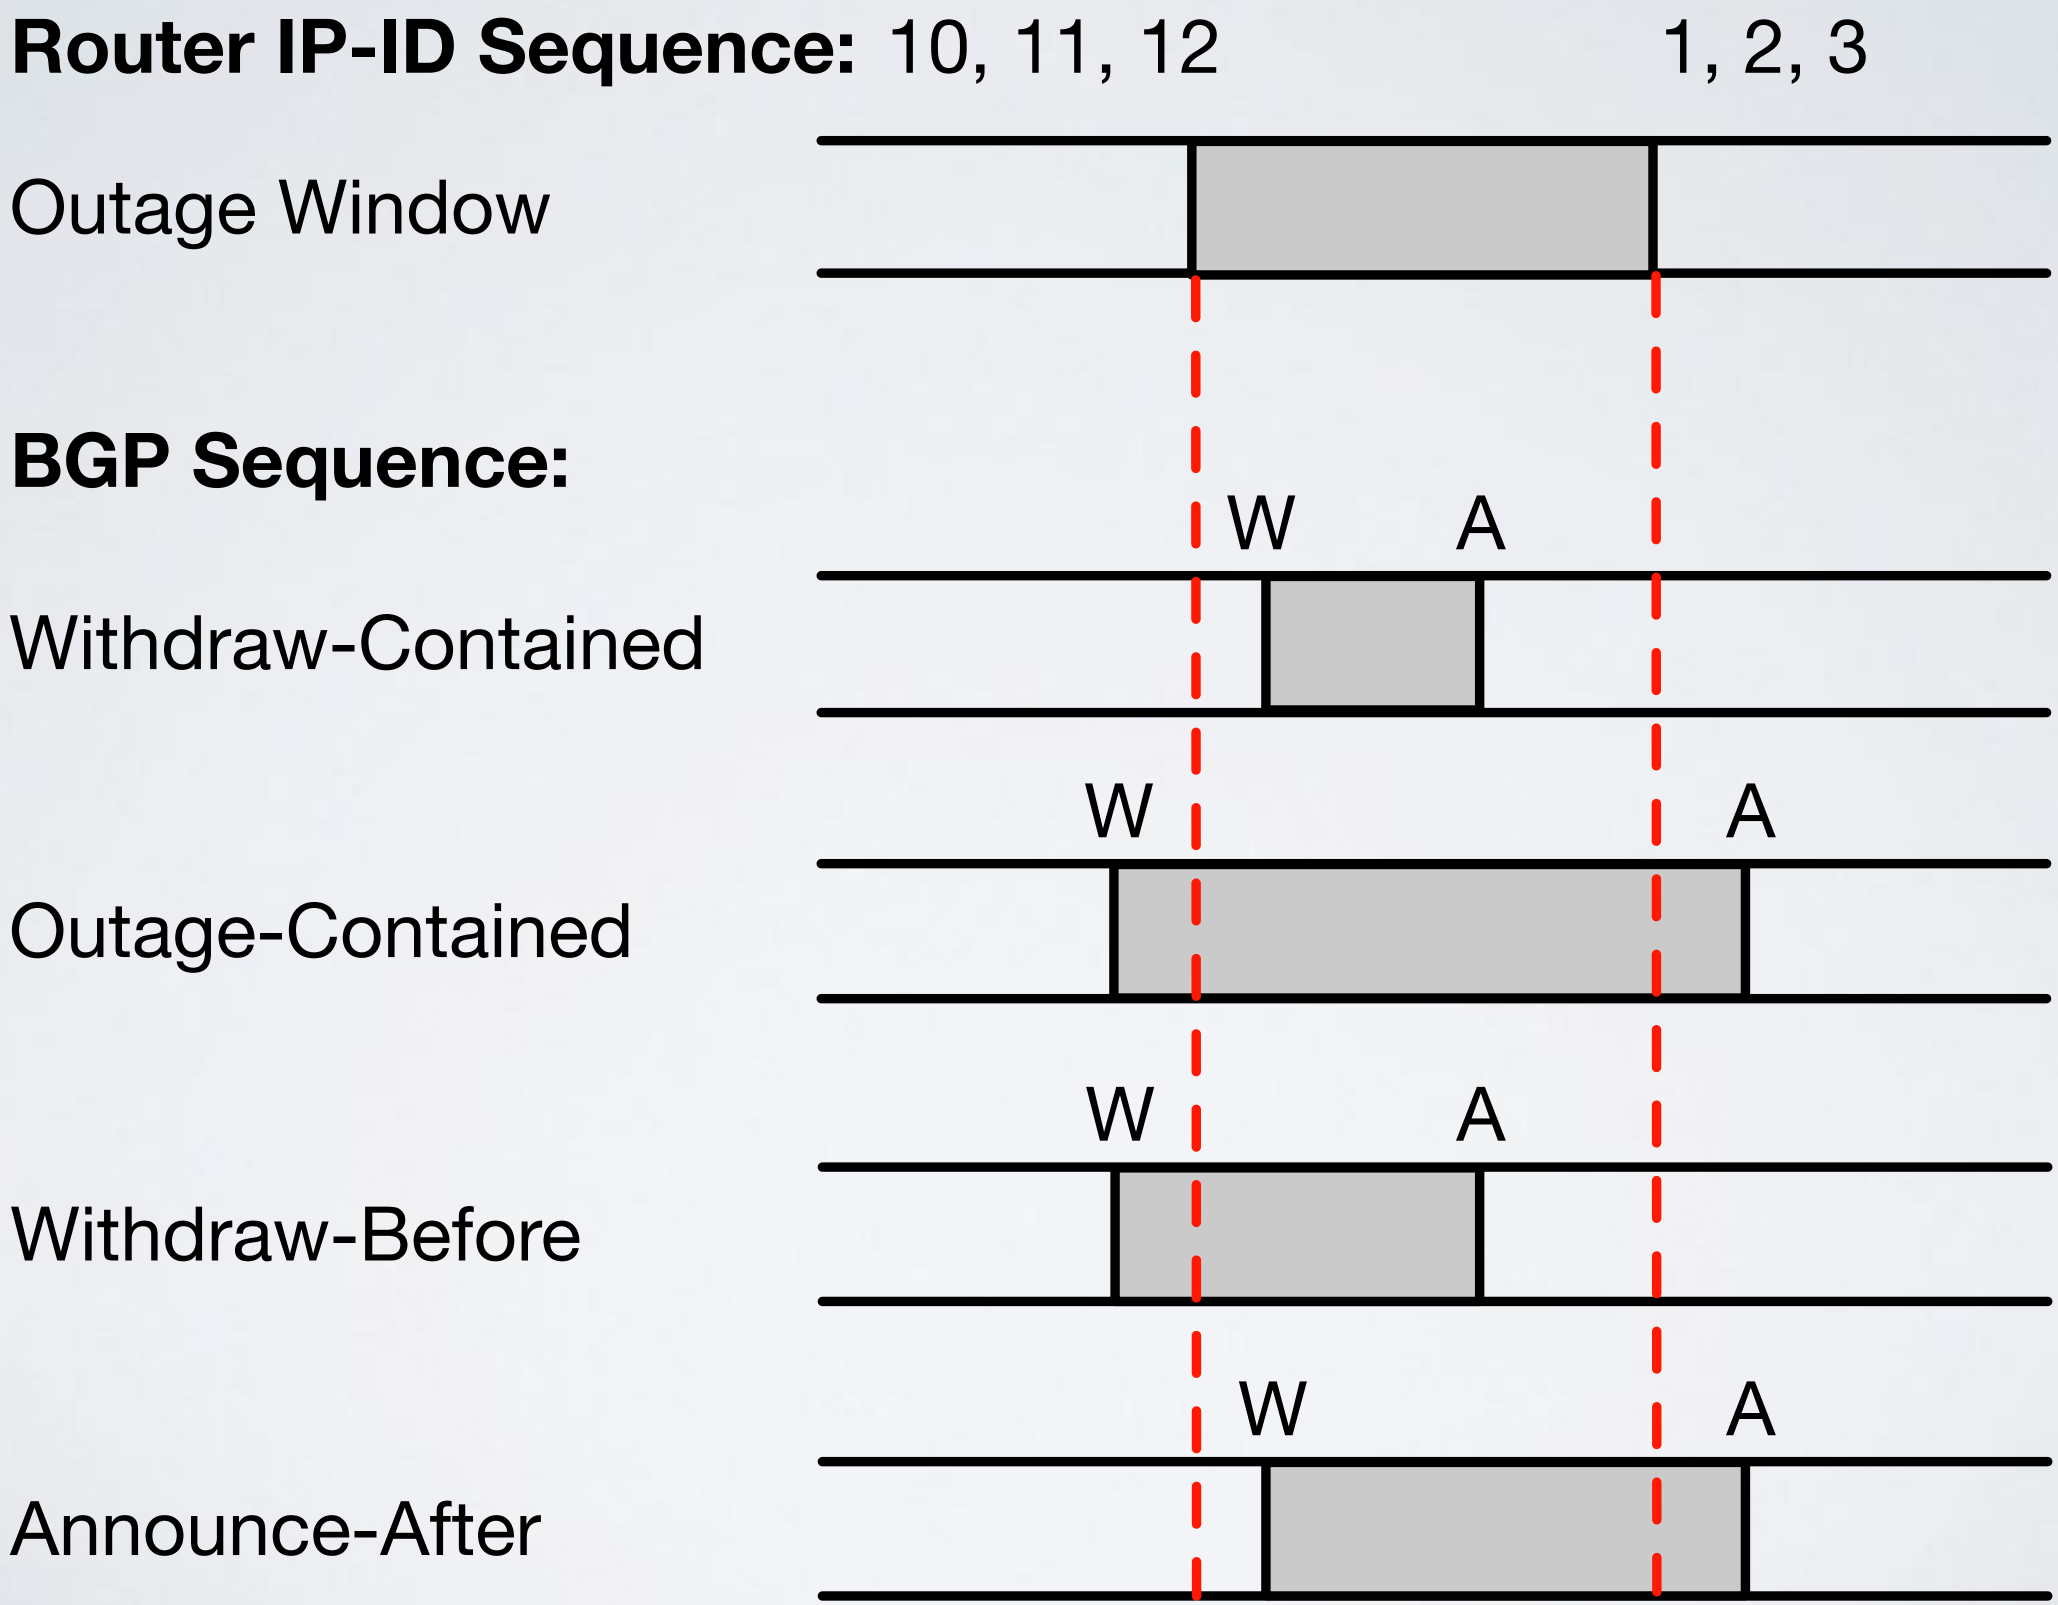

Data processing pipeline

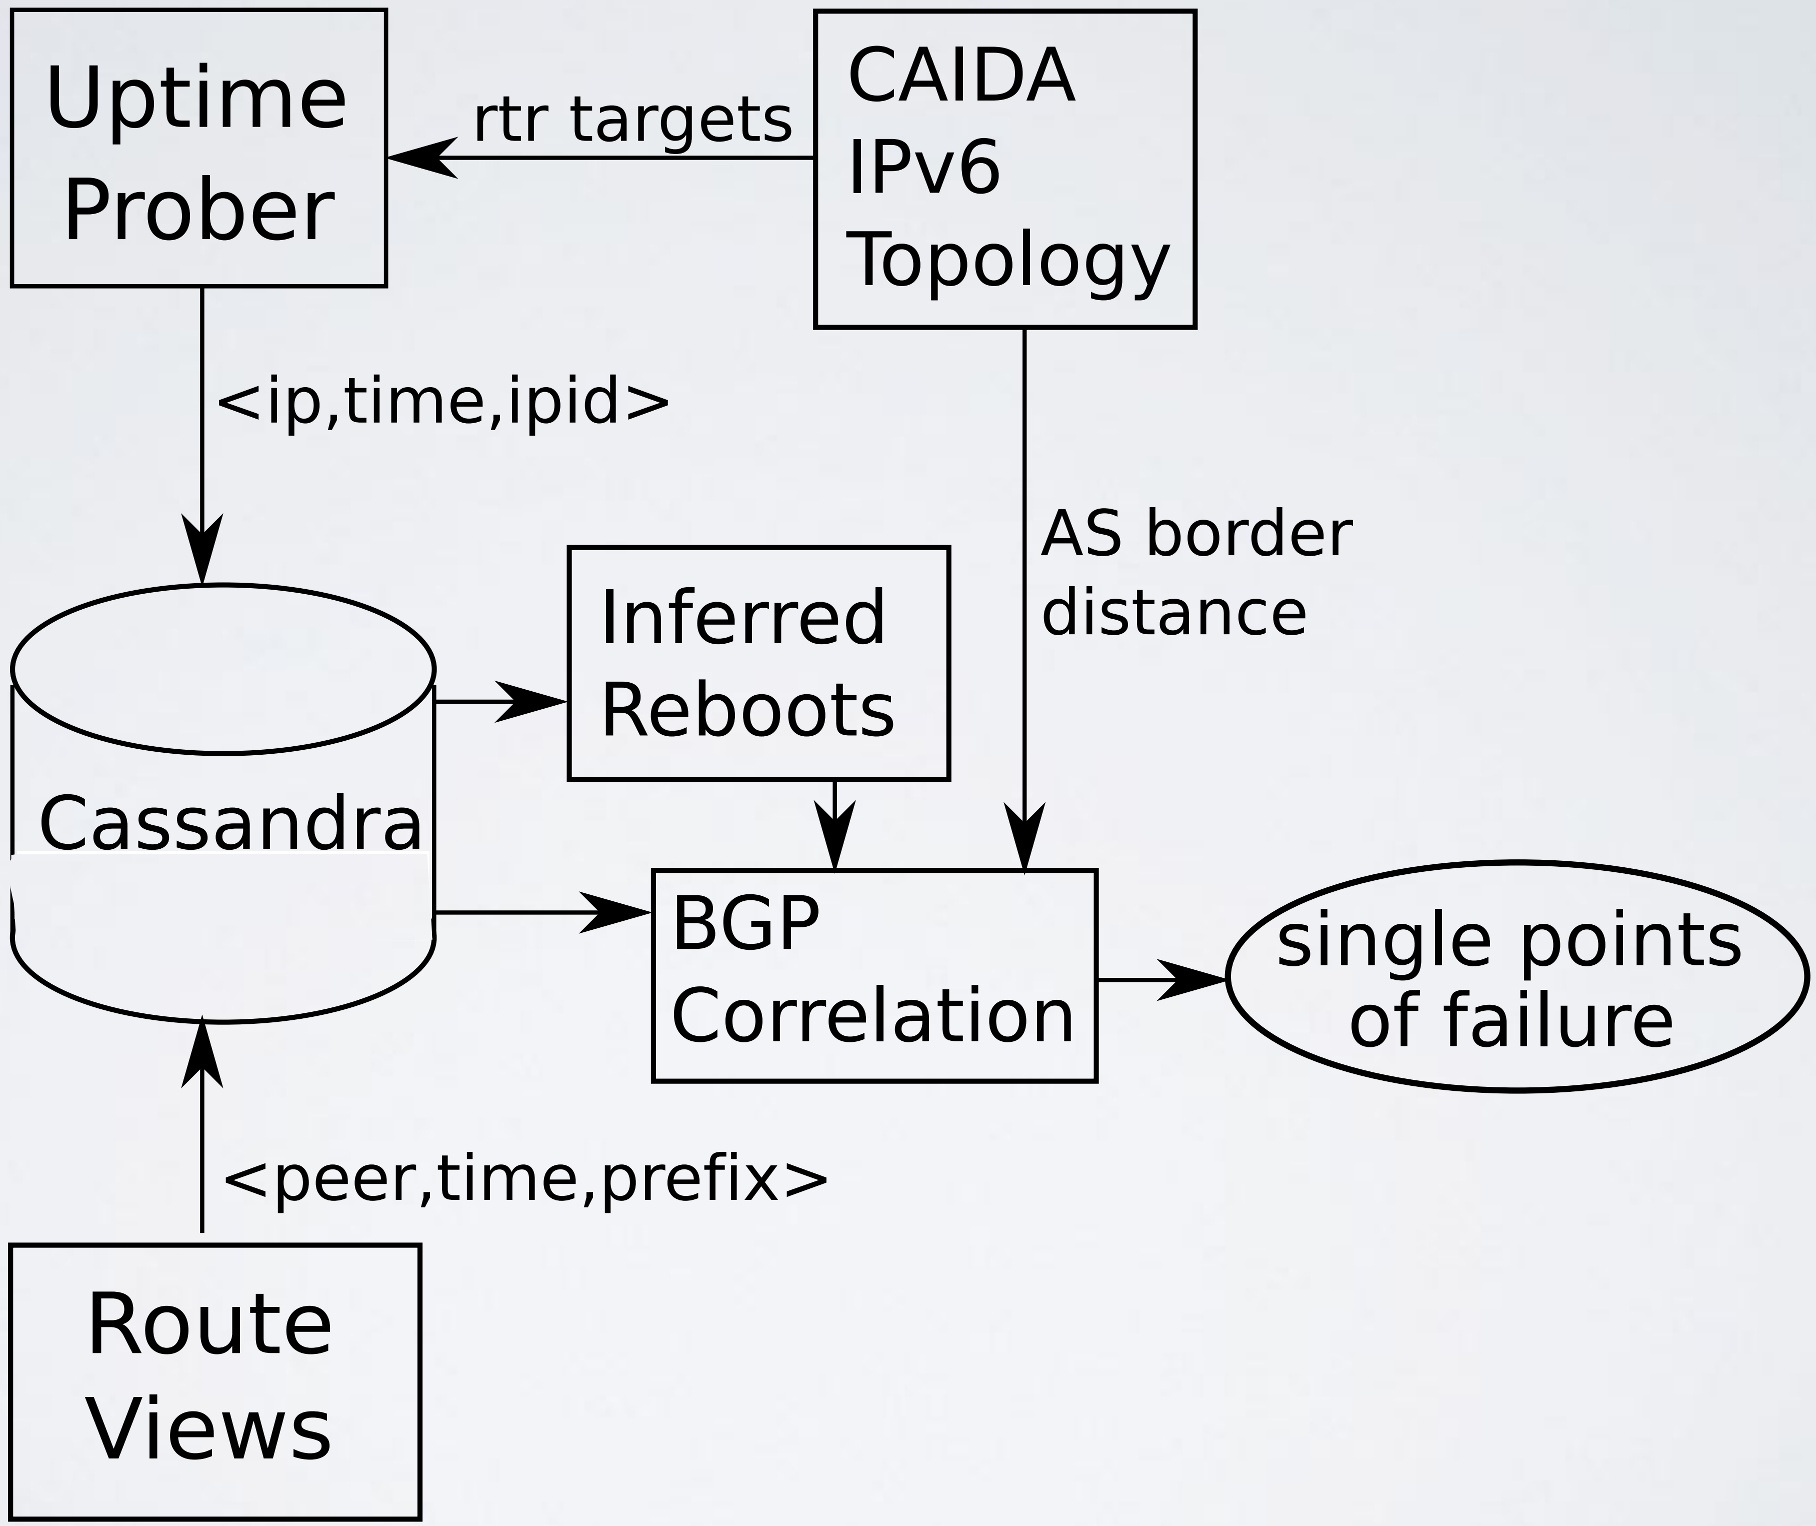

Inferring router position

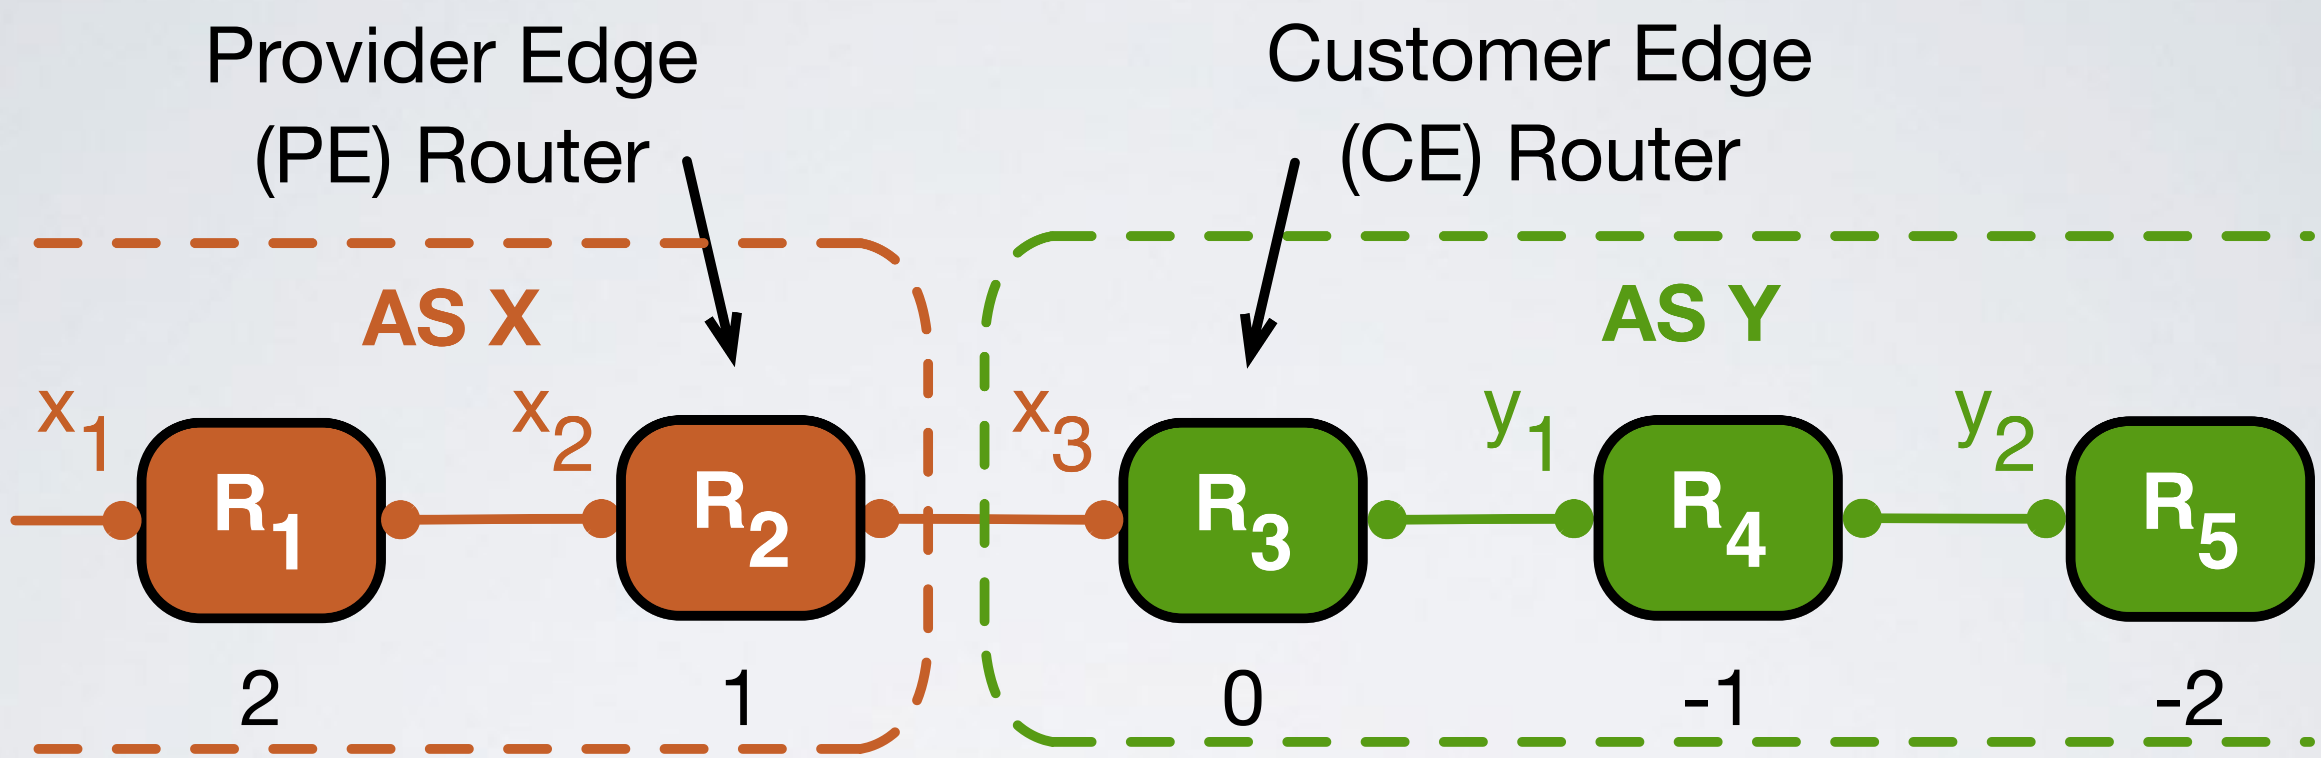

(a) interface addresses routed by Y appear in traceroute

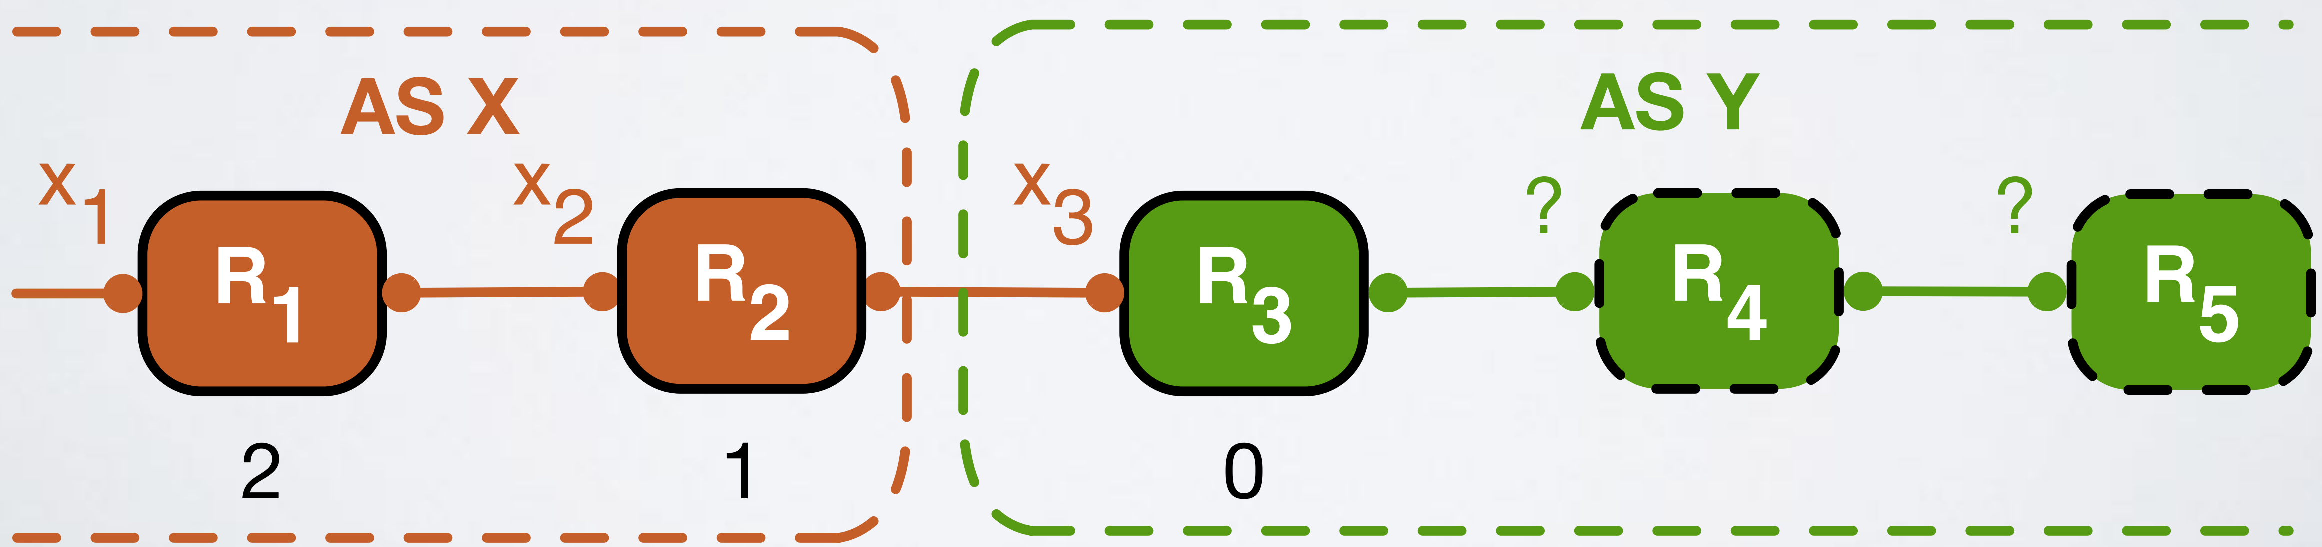

(b) no interface addresses routed by Y appear in traceroute

Data Collection Summary

	18 Jan '15 18 Oct '16 (a)	18 Oct '16 24 Feb '17 (b)	24 Feb '17 30 May '17 (c)
Probing rate	100 pps	225 pps	200 pps
Interfaces	83K seen Dec '14	1.1M seen Jun to Oct '16	Dynamic. 2.4M in May '17
Responsive	every round ~15 mins	every round ~15 mins	every round ~15 mins
Unresponsive	12-24 hours	12-24 hours	7-14 days

Why IPv6 fragment IDs?

IPv4 ID values are 16 bits with bursty velocity as every packet requires a unique value.

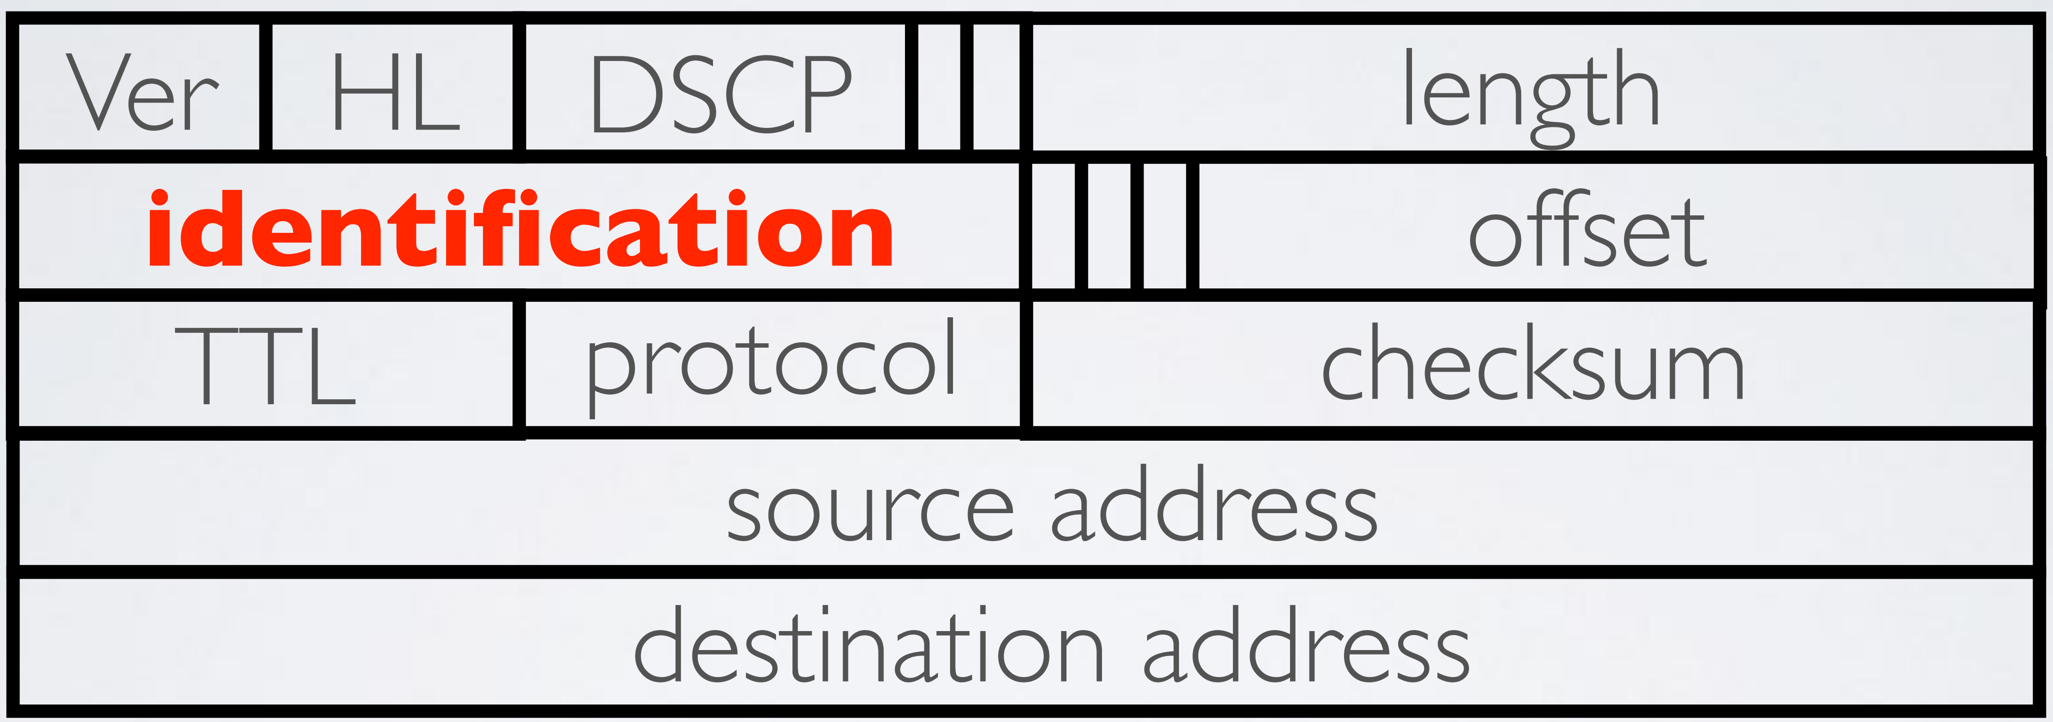

At 100Mbps and 1500 byte packets.
Nyquist rate dictates a 4 second probing interval

Why IPv6 fragment IDs?

IPv6 ID values are 32 bits with low velocity as systems rarely send fragmented packets.

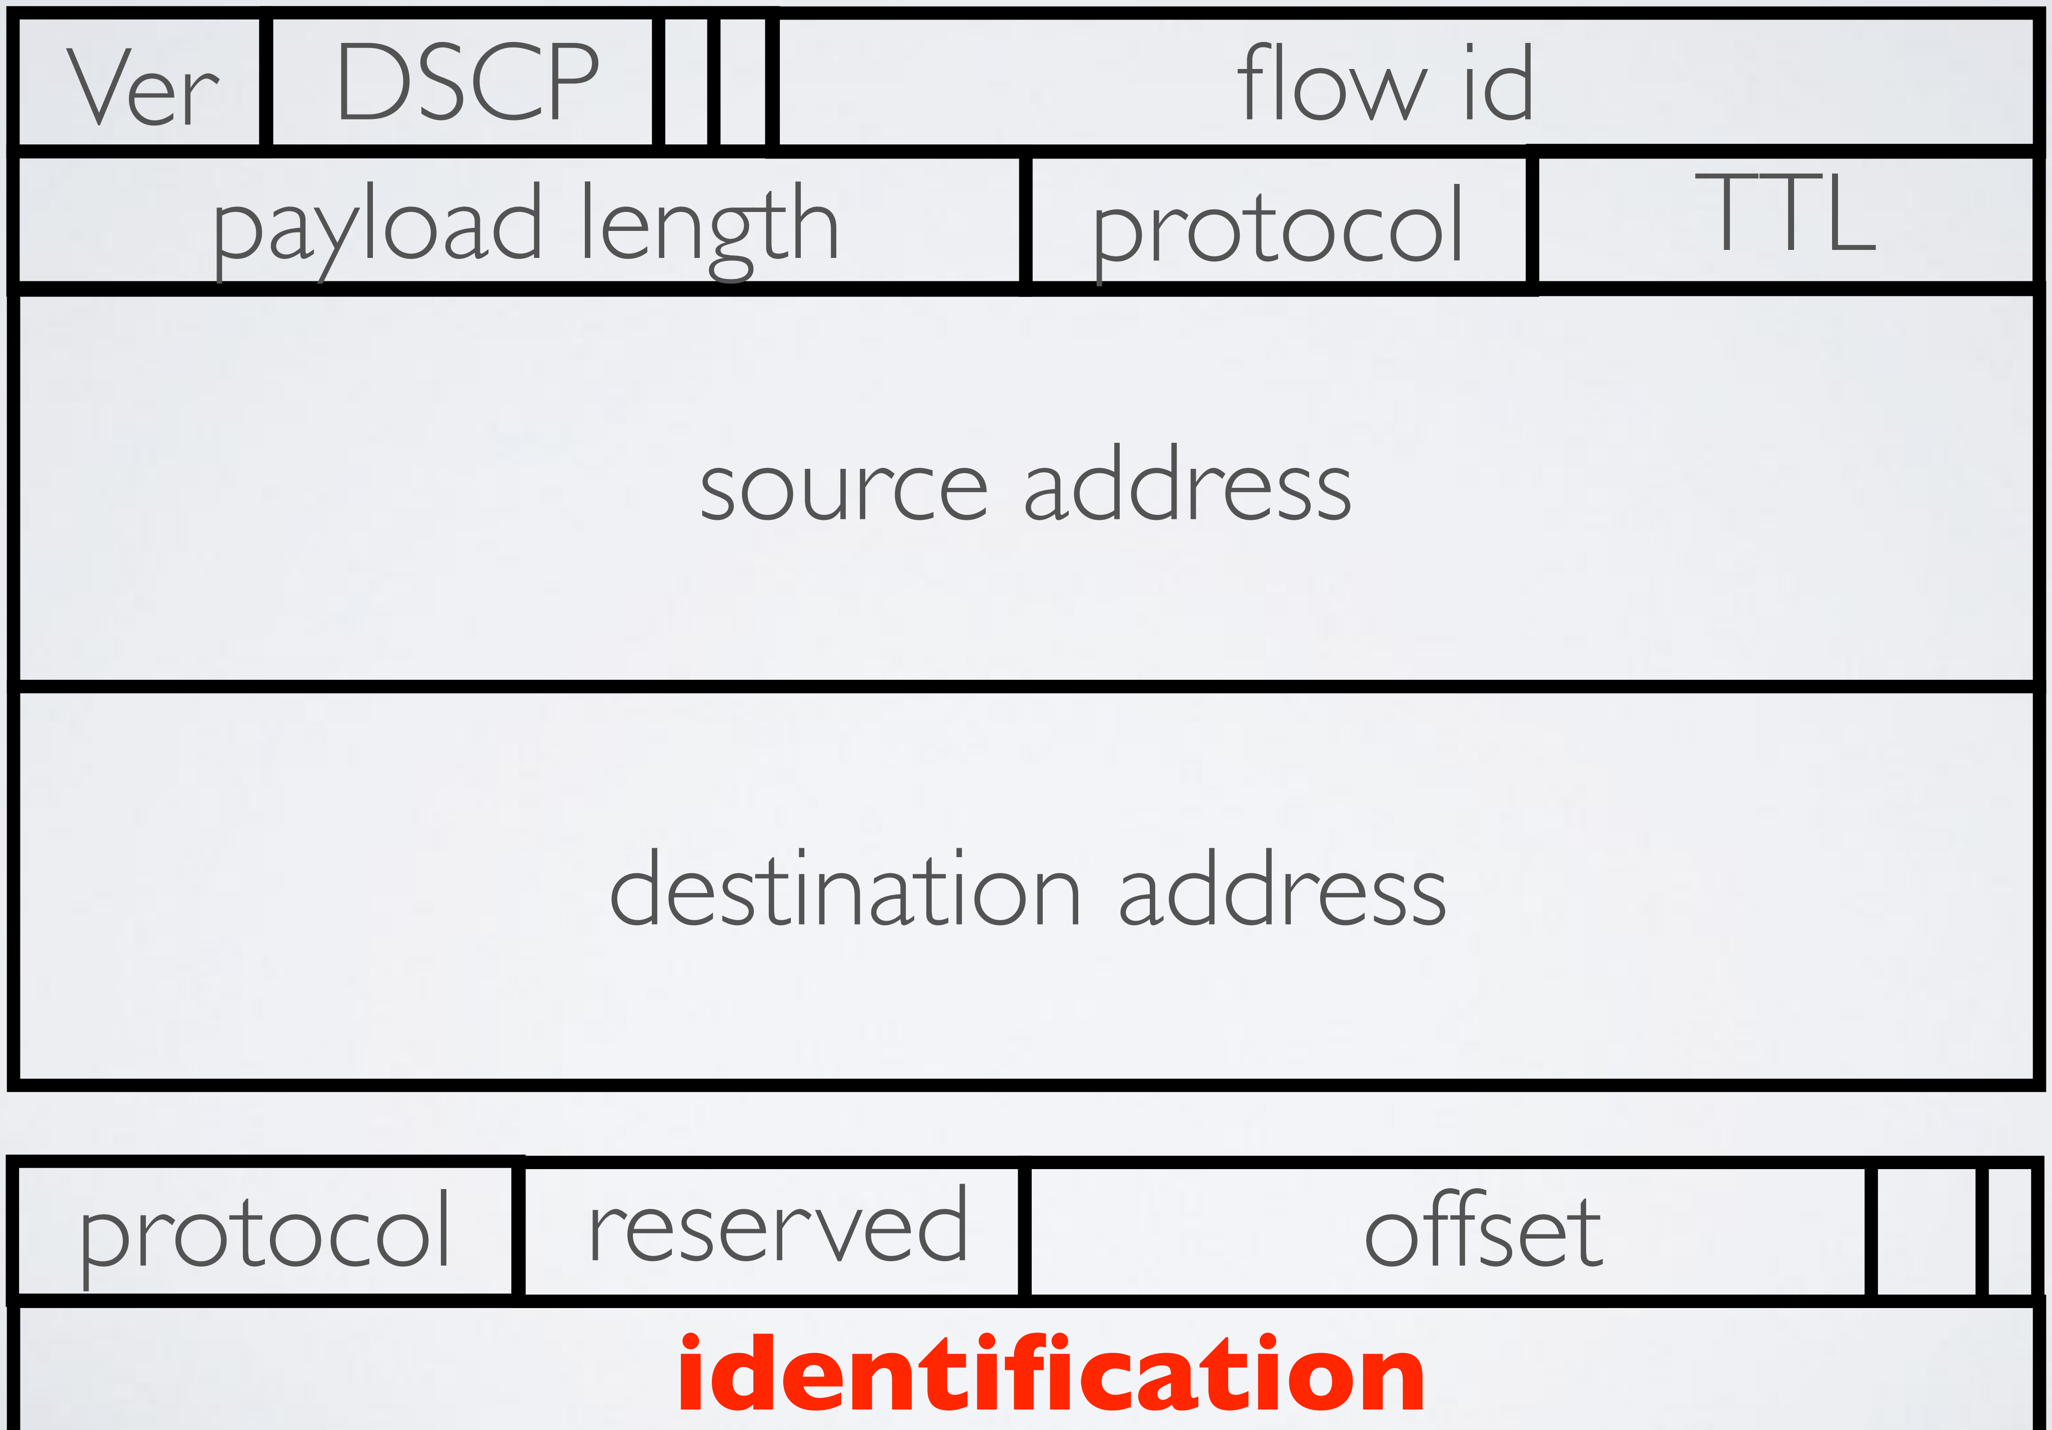

Soliciting IPv6 Fragment IDs