

The Next Generation of BGP Data Collection Platforms

Thomas Holterbach
University of Strasbourg

RIPE NCC Open House
26 April 2024

Joint work with:

Thomas Alfroy

Thomas Krenc

KC Claffy

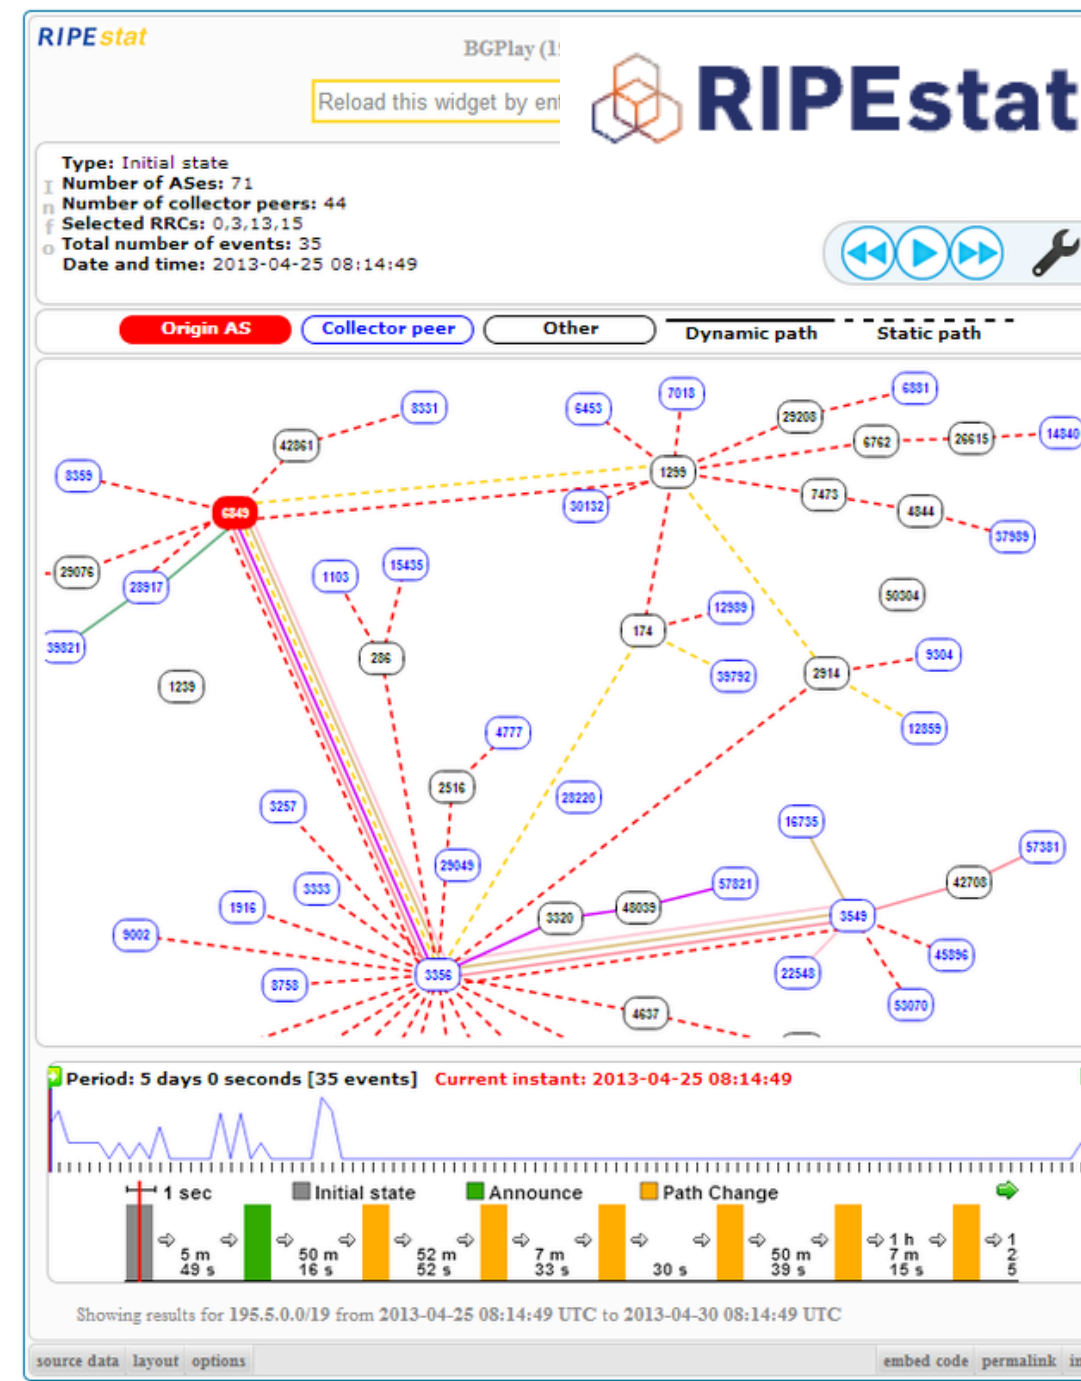
Cristel Pelsser



Map AS topology



Measure connectivity



Detect routing attacks

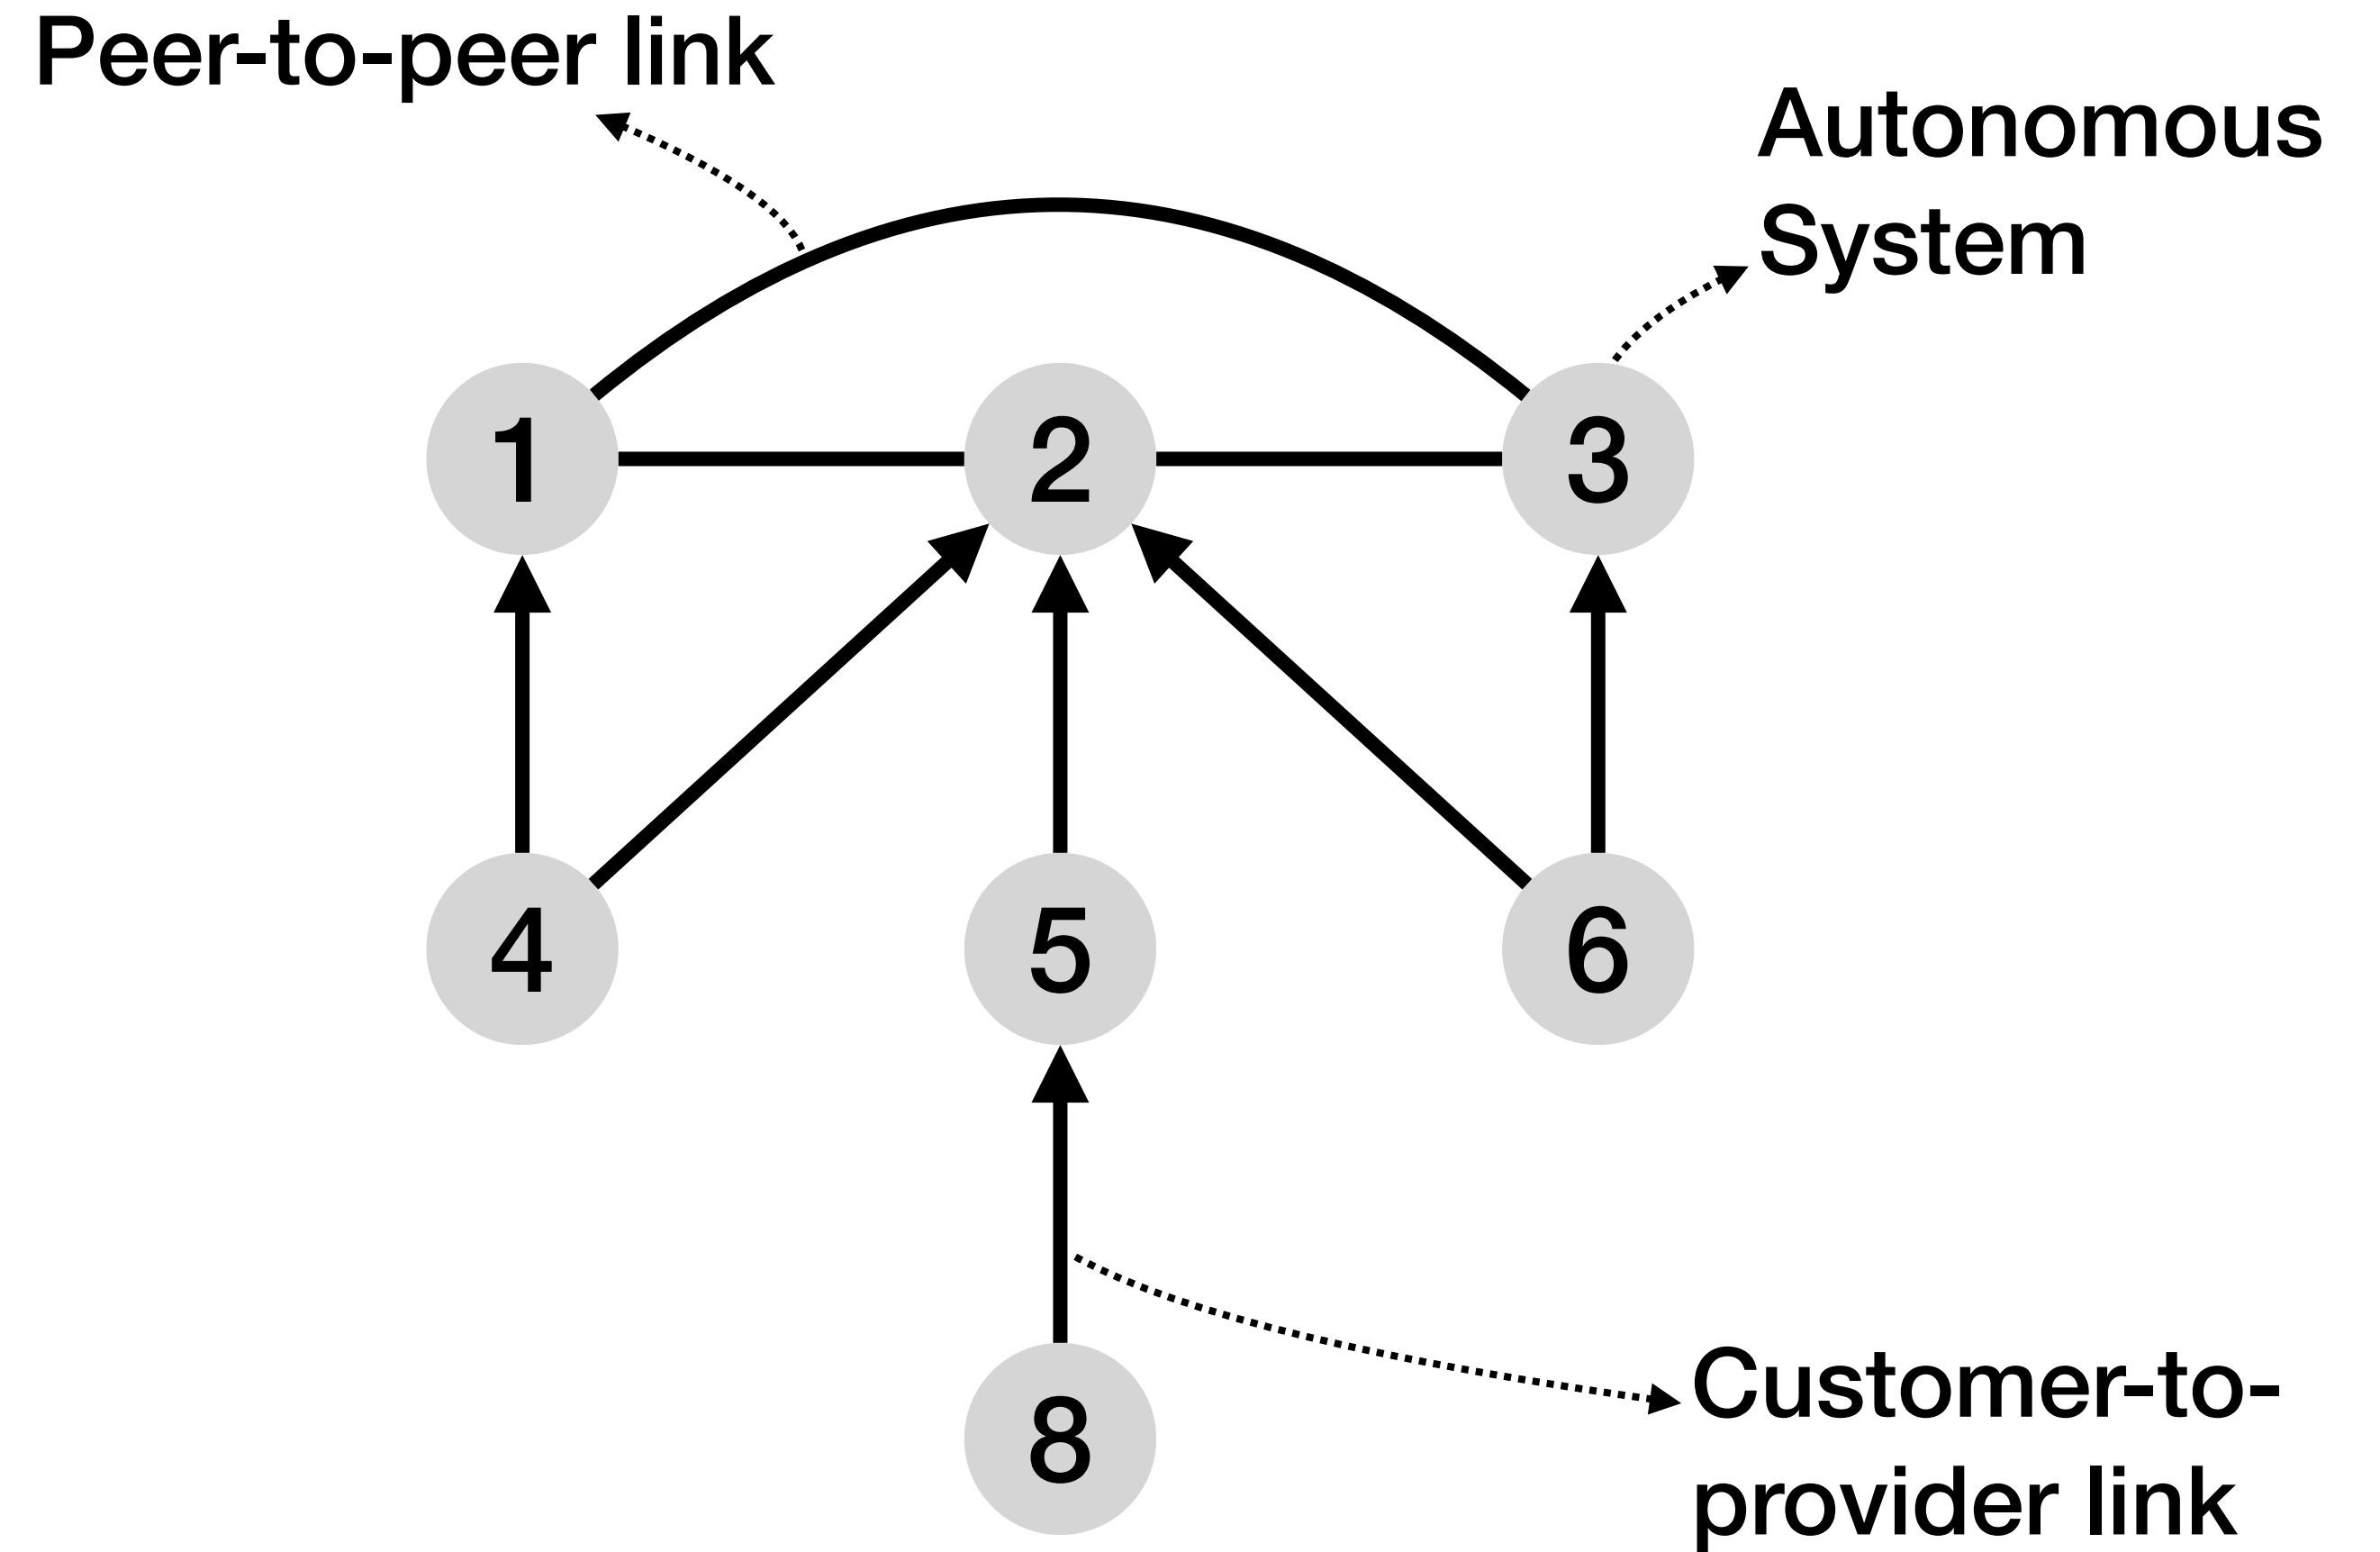
BGP Origin Hijacks Beta Cloudflare Radar

Detected BGP hijack events origina

Detected Origin	Expected Origin(s)
AS263903 (BR)	AS267283 (BR)
AS36962 (ZM)	AS28698 (ZA)
AS28698 (ZA)	AS36962 (ZM)
AS59588 (IQ)	AS204149 (IQ)
AS43754 (IR)	AS60631 (IR)
AS207279 (TR)	AS212238 (GB)
AS262324 (BR)	AS268679 (BR)
AS55330 (AF)	AS58469 (AF)
AS150153	AS150309
AS137432 (BD)	AS140901 (BD)

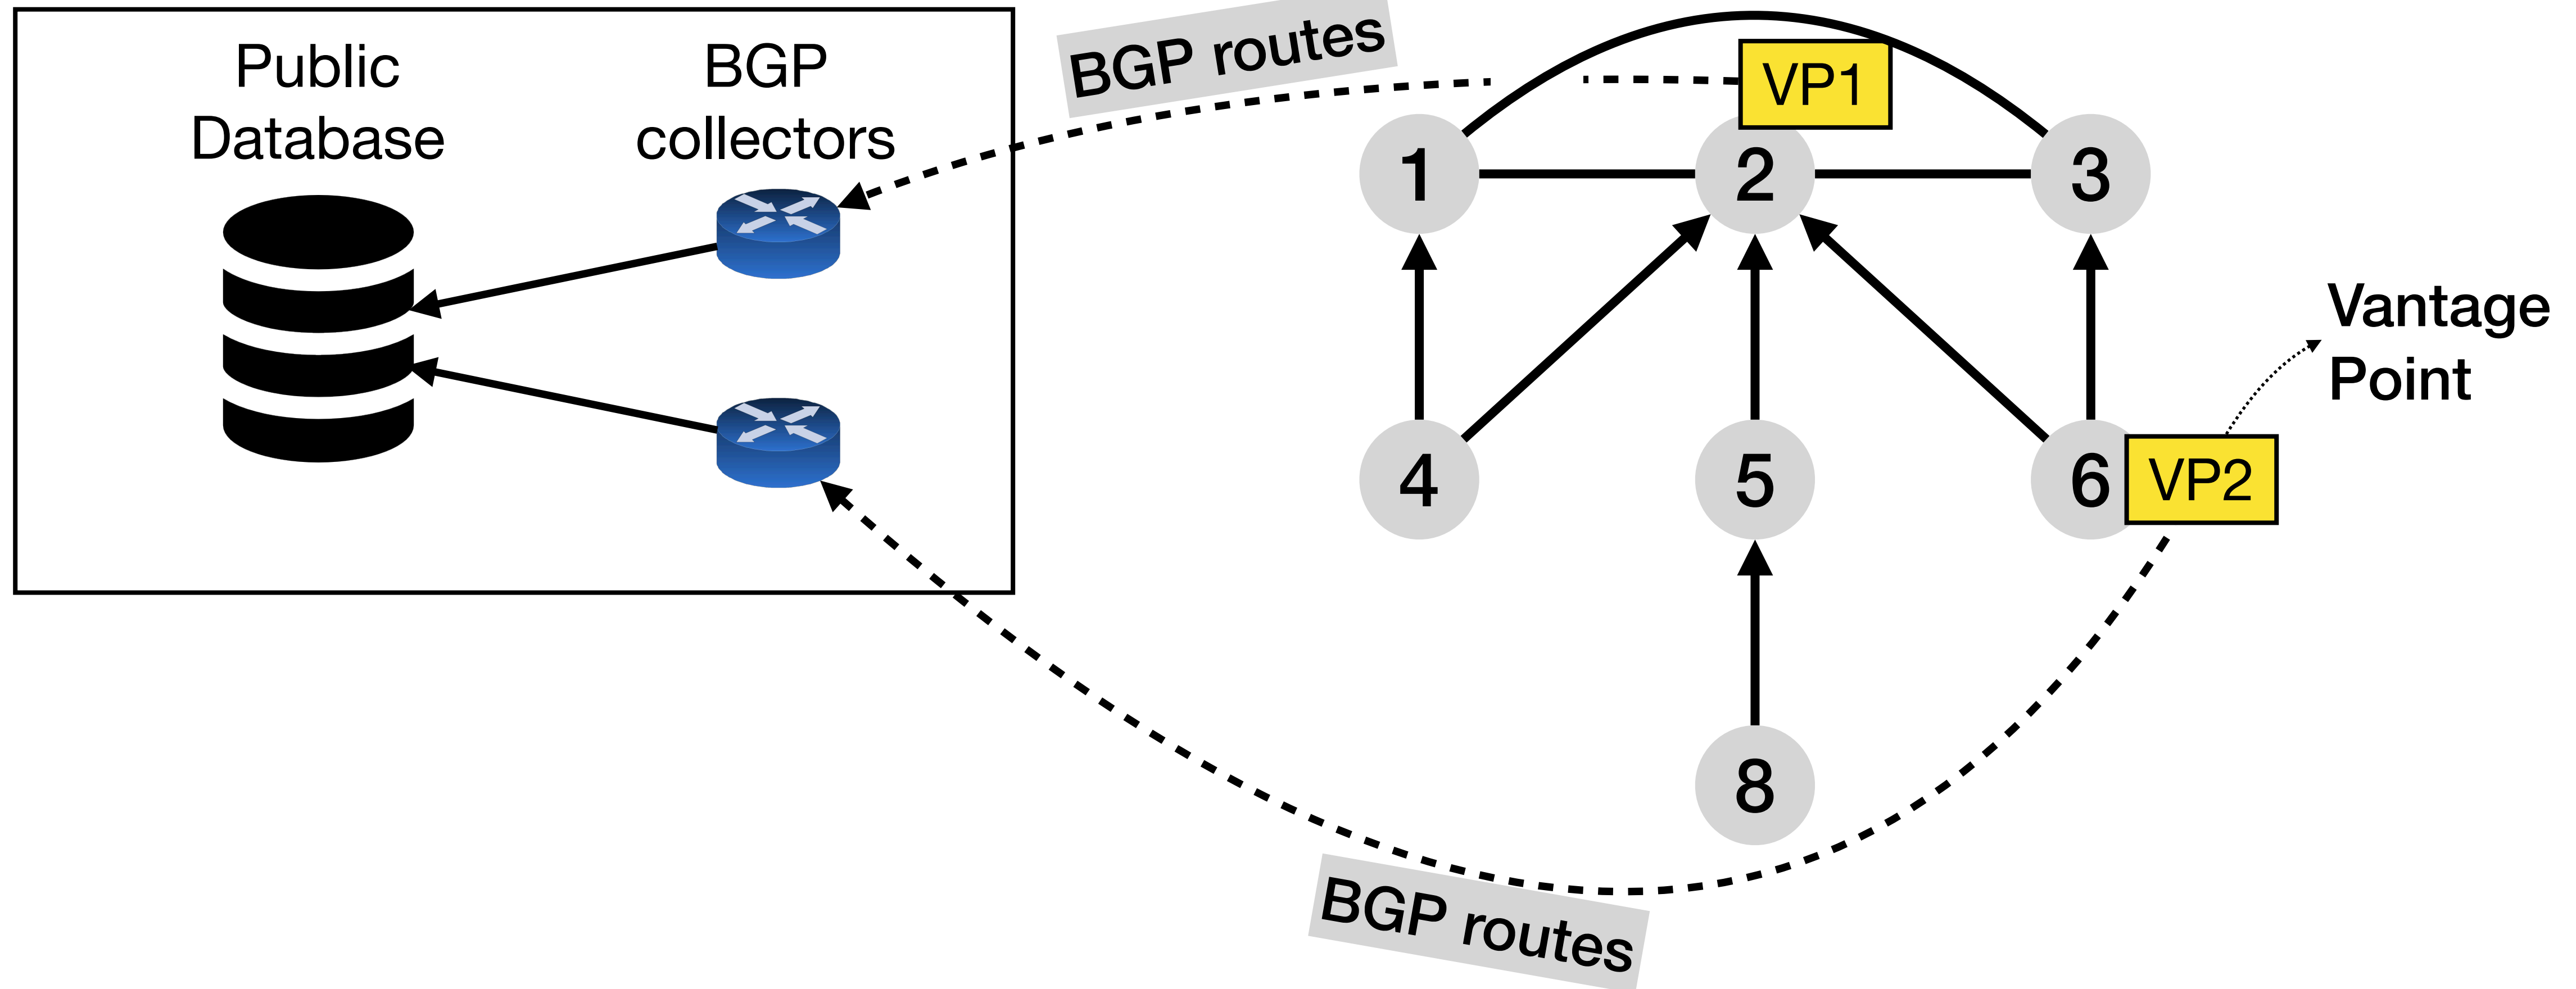
These three use cases (and many more) rely on the collected public BGP routes

Public BGP routes are collected by RIPE RIS and RouteViews



Public BGP routes are collected by RIPE RIS and RouteViews

Collection platform

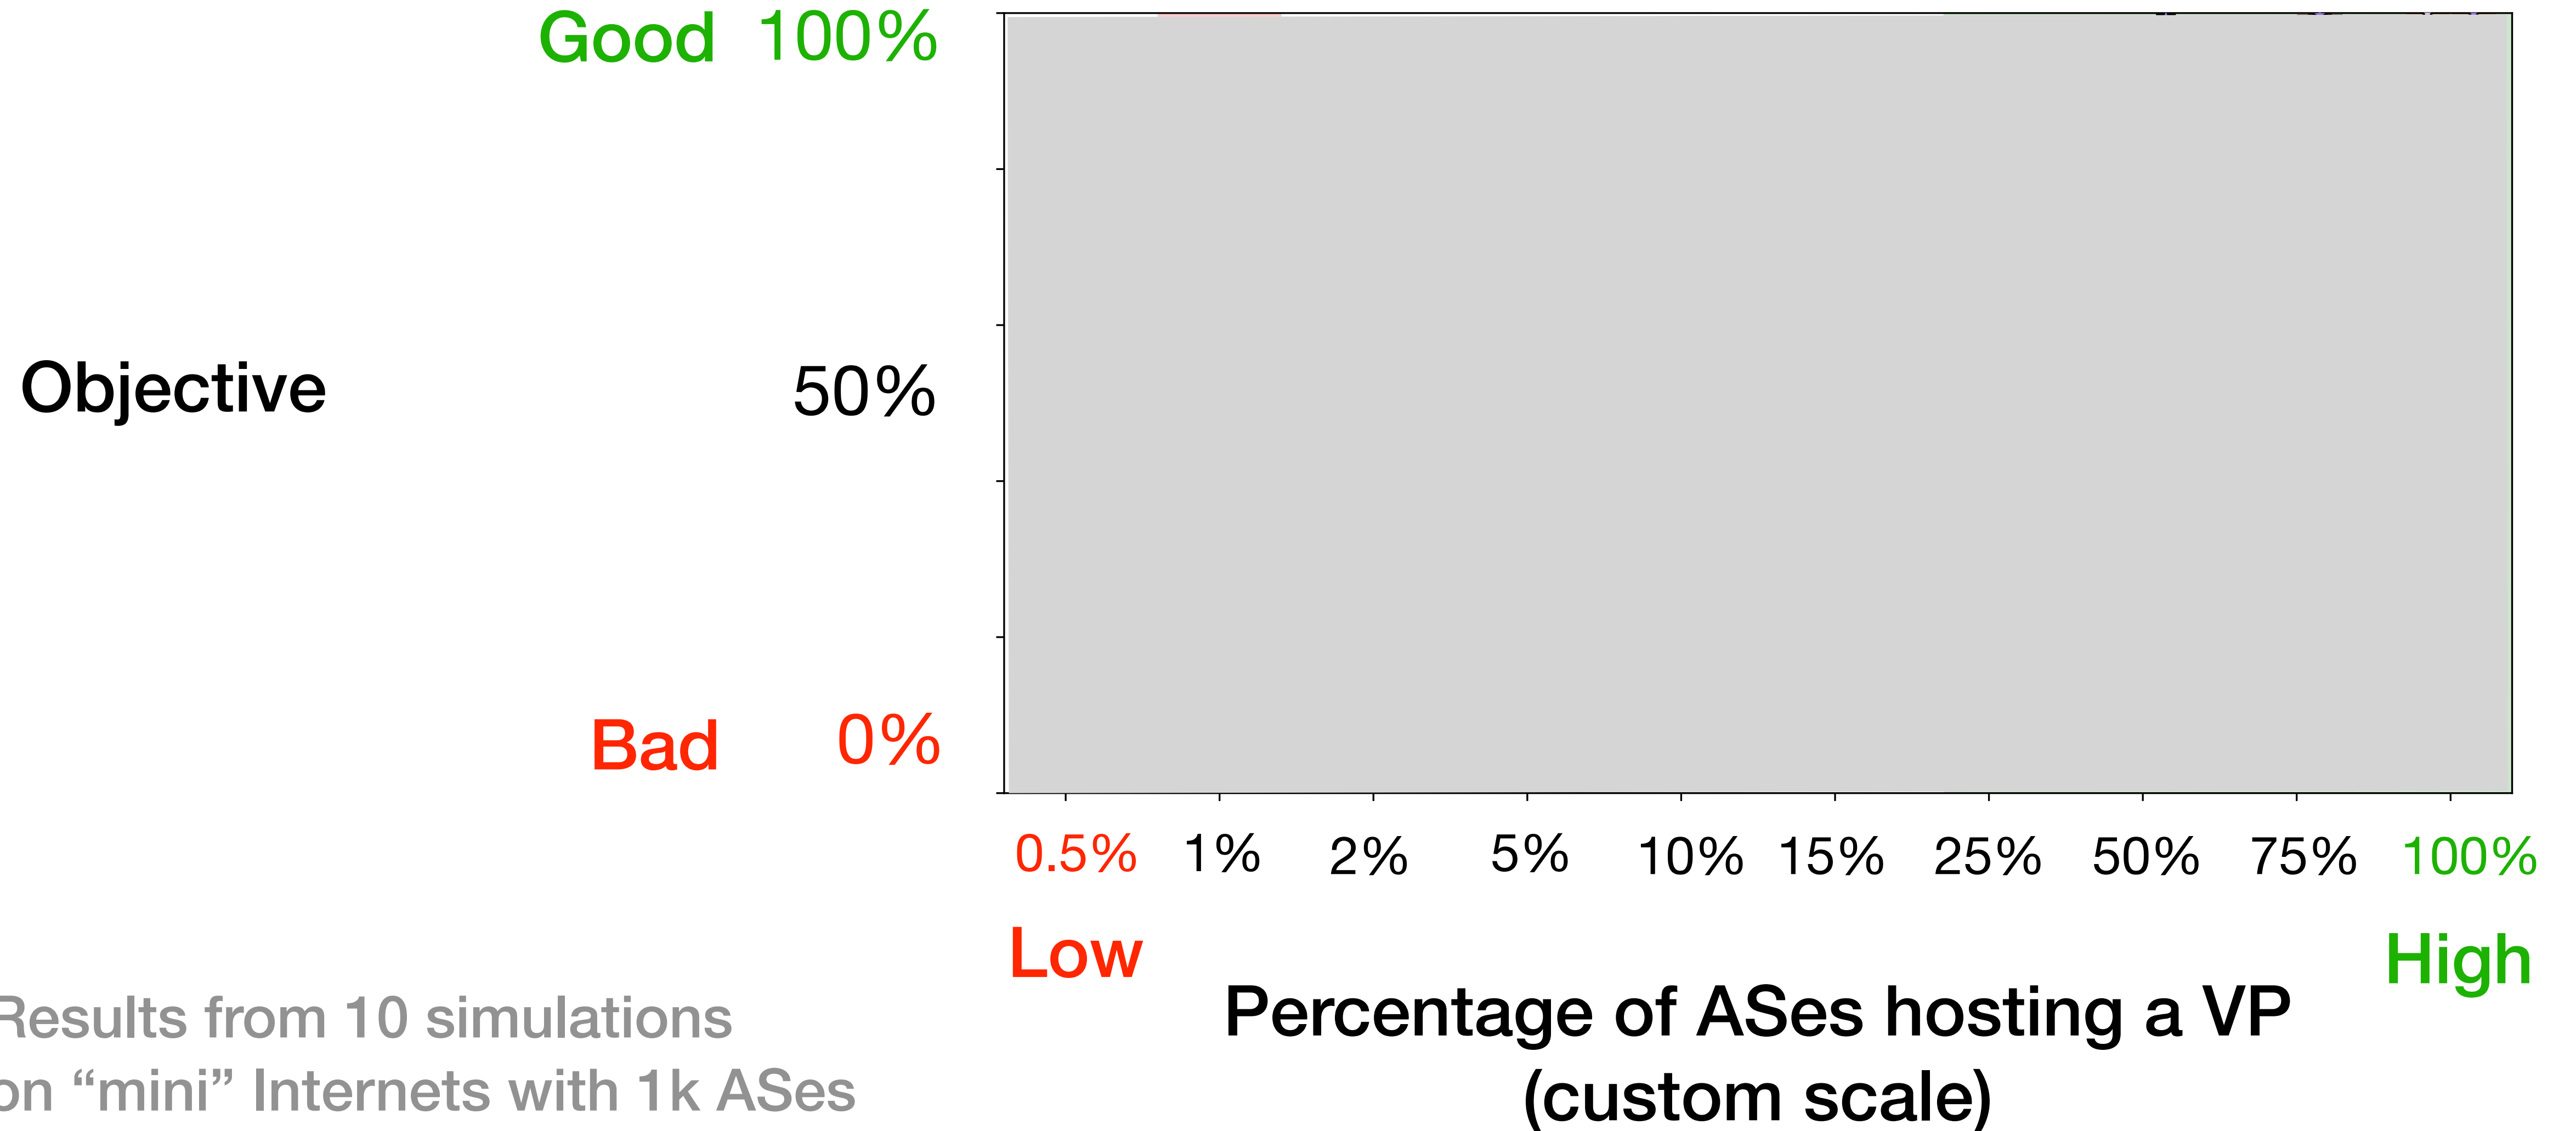


**Three observations motivate reevaluating
how we collect BGP routes**

Three observations motivate reevaluating
how we collect BGP routes

Observation #1: RIPE RIS and RouteViews lack coverage

RIS and RouteViews' low coverage negatively impacts many studies



RIS and RouteViews' low coverage negatively impacts many studies

**RIS+RouteViews
coverage**

Good 100%

Objective

50%

Bad 0%

0.5% 1% 2% 5% 10% 15% 25% 50% 75% 100%

Low

High

**Percentage of ASes hosting a VP
(custom scale)**


Results from 10 simulations
on "mini" Internets with 1k ASes

RIS and RouteViews' low coverage negatively impacts many studies

RIS+RouteViews
coverage

Ideal
coverage

Percentage of
AS links observed



Results from 10 simulations
on “mini” Internets with 1k ASes

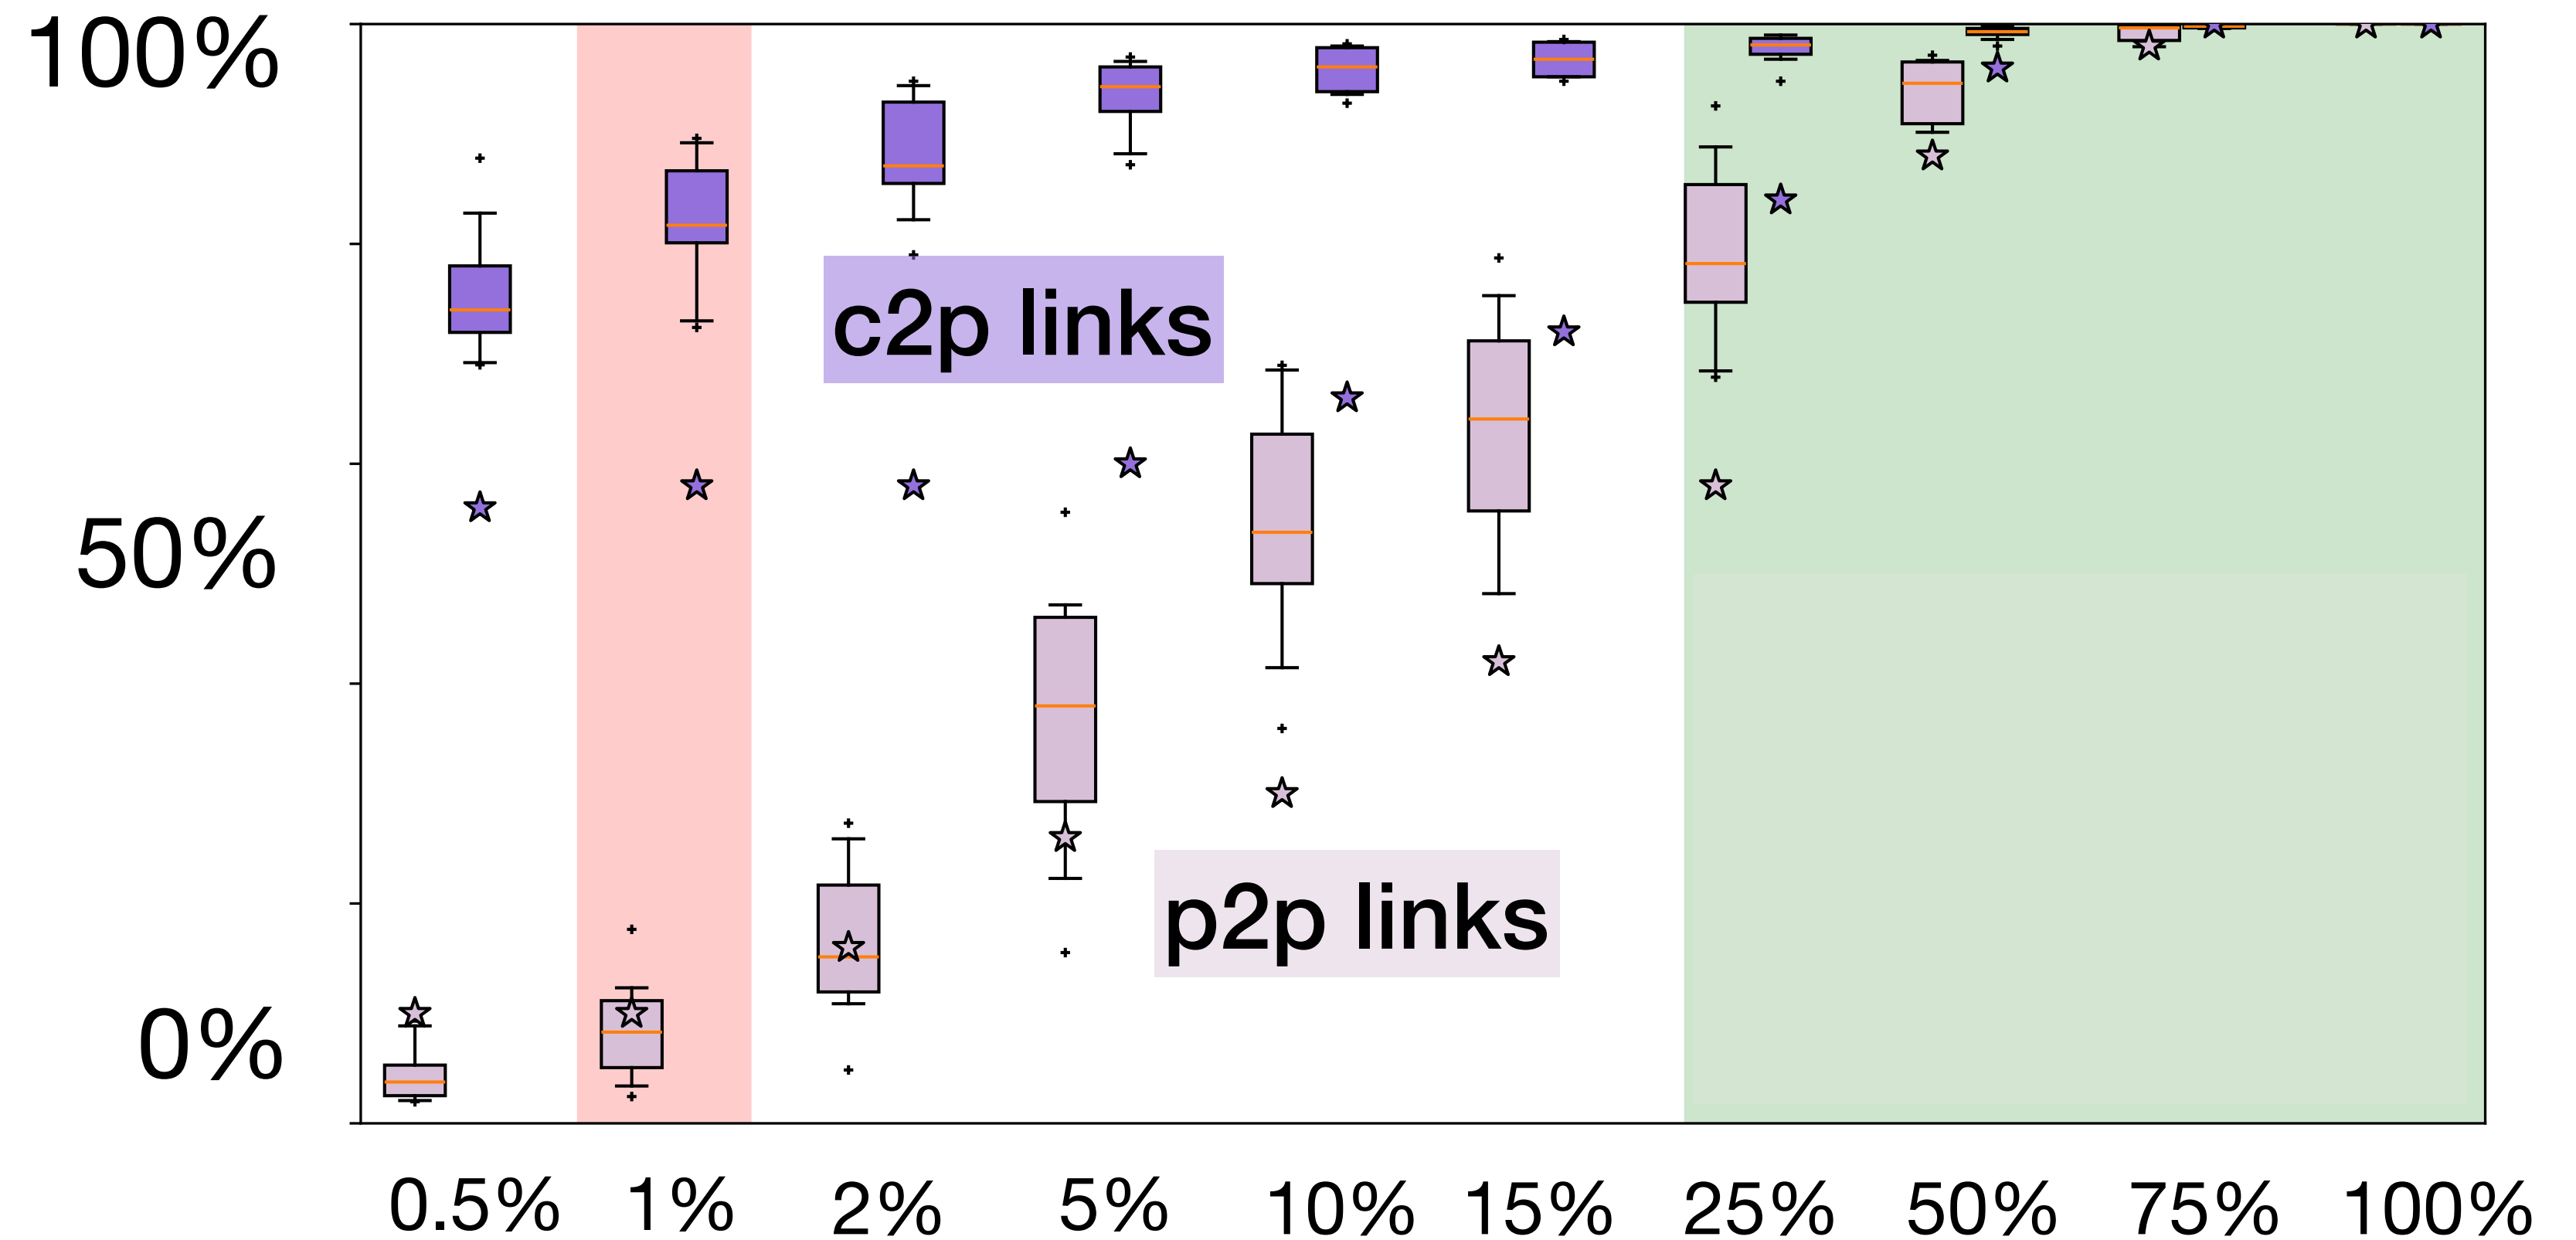
Percentage of ASes hosting a VP
(custom scale)

RIS and RouteViews' low coverage negatively impacts many studies

**RIS+RouteViews
coverage**

**Ideal
coverage**

**Percentage of
localized* failures**



*We use Feldmann et al.'s algorithm (SIGCOMM'04)

Results from 10 simulations on "mini" Internets with 1k ASes


**Percentage of ASes hosting a VP
(custom scale)**

RIS and RouteViews' low coverage negatively impacts many studies

RIS+RouteViews
coverage

Ideal
coverage

Percentage of
forged-origin hijacks
detected



Results from 10 simulations
on “mini” Internets with 1k ASes

Percentage of ASes hosting a VP
(custom scale)


**Two observations motivate reevaluating
how we collect BGP routes**

Observation #1: RIPE RIS and RouteViews lack coverage

Observation #2: RIPE RIS and RouteViews coverage is flat over time


Despite deploying new VPs, RIS and RouteViews' coverage is flat due the growing size of the Internet

Number of ASes hosting a VP




Despite deploying new VPs, RIS and RouteViews' coverage is flat due the growing size of the Internet

Number of ASes hosting a VP



Percentage of ASes hosting a VP



**Two observations motivate reevaluating
how we collect BGP routes**

Observation #1: RIPE RIS and RouteViews lack coverage


Observation #2: RIPE RIS and RouteViews coverage is flat over time

**Observation #3: Deploying new VPs leads to a unmanageable
number of routes to process**

The number of routes collected increases **quadratically**

The number of routes collected increases **quadratically**

of BGP routes
collected per hour
(*RIS + RouteViews*)




The number of routes collected increases quadratically

of BGP routes collected per hour
(*RIS + RouteViews*)



1TB to process every day!

Data management is challenging for the collection platform



Route Collection at the RIPE NCC - Where are we and where should we go?


Emile Aben — 7 Oct 2020
13 min read

ris events routing

0 ❤️ 1 💬 🔗 📌

Over the past months we've been looking at our Routing Information Service (RIS) and thinking about how to make it best fit for purpose. Ahead of our upcoming RIPE NCC Open House on RIS, this post raises a set of open questions to our community aimed at starting a conversation about how we can keep RIS useful to you.

At the RIPE NCC we run a large route collecting system called RIS. We currently have 21 active route collectors that collect BGP data from roughly 1,300 BGP peering sessions. We know from [our last membership survey](#) that RIS is regarded by our members as a useful and important service. And we



Robert Kistelevi > RIPE NCC Measurement Data Retention... [Want to contribute? Learn how](#)

RIPE NCC Measurement Data Retention Principles

Robert Kistelevi — 22 Nov 2023
Contributors: [Paul de Weerd](#)
6 min read

ris atlas tools community measurements


14 ❤️ 0 💬 🔗 📌

RIPE Atlas and RIPE RIS both provide a wealth of Internet measurement data invaluable to both Internet researchers and network operators alike. But that's not to say that questions about the cost and value of storing this data don't come up from time to time. Here, we open up the discussion on whether change is called for.

RIPE NCC measurement systems, in particular RIPE RIS (Routing Information System) and RIPE Atlas, <https://labs.ripe.net> have been collecting passive and active measurement results (see details below) since 1999 and

Data management is challenging for the collection platform and their users

Do you find the data from RIS and RouteViews expensive to process in terms of computational resources?



Survey conducted among authors of eight top research papers that used BGP data


The Next Generation of BGP Data Collection Platforms




Outline

1. We observe that BGP routes are often redundant

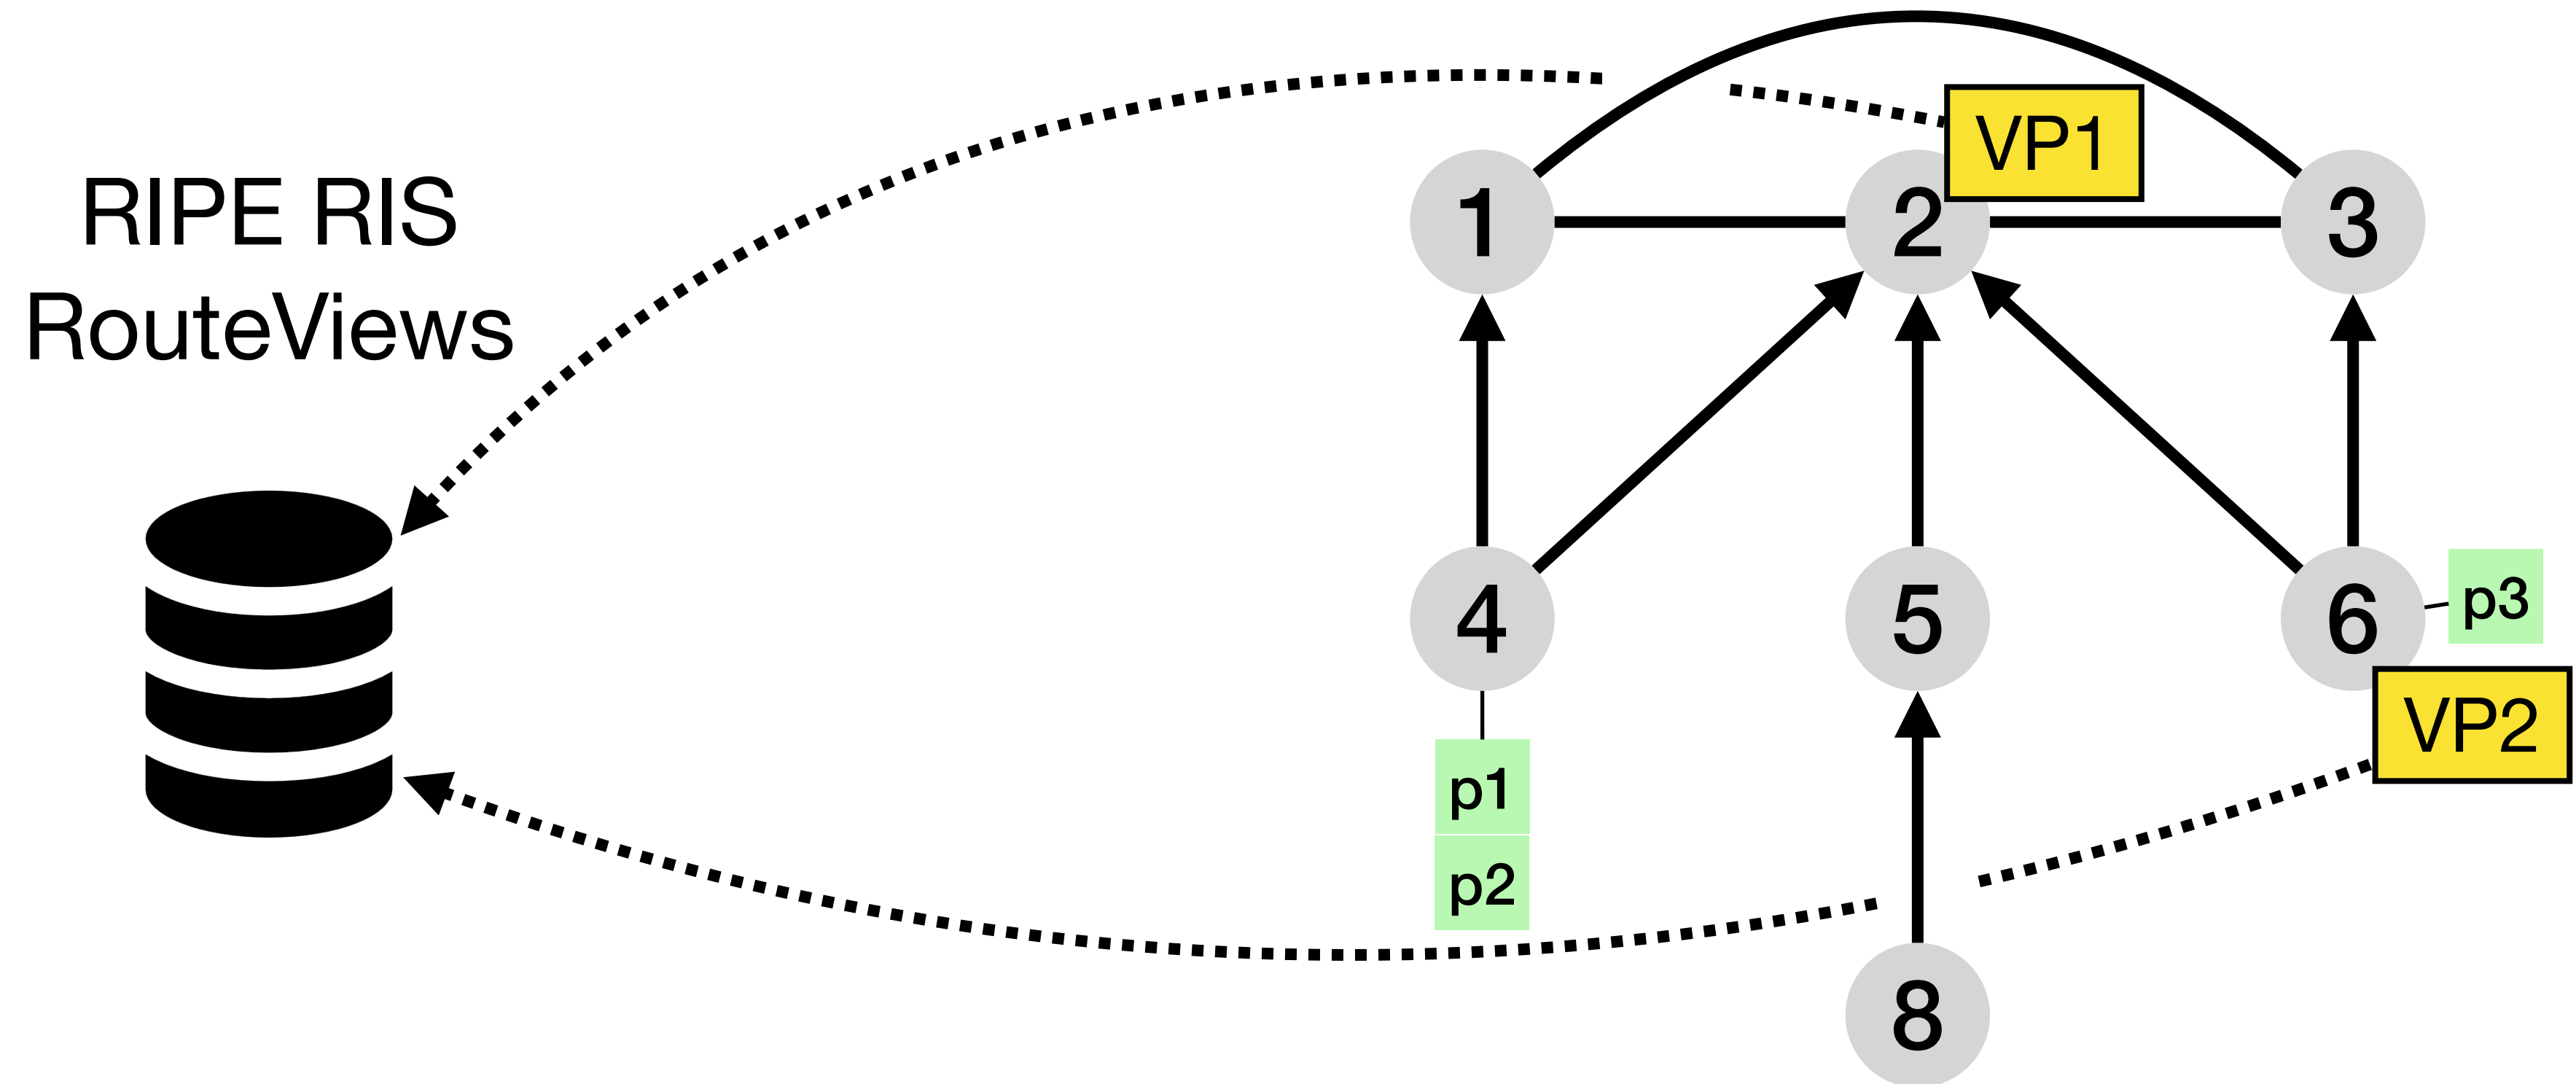
BGP routes can be redundant



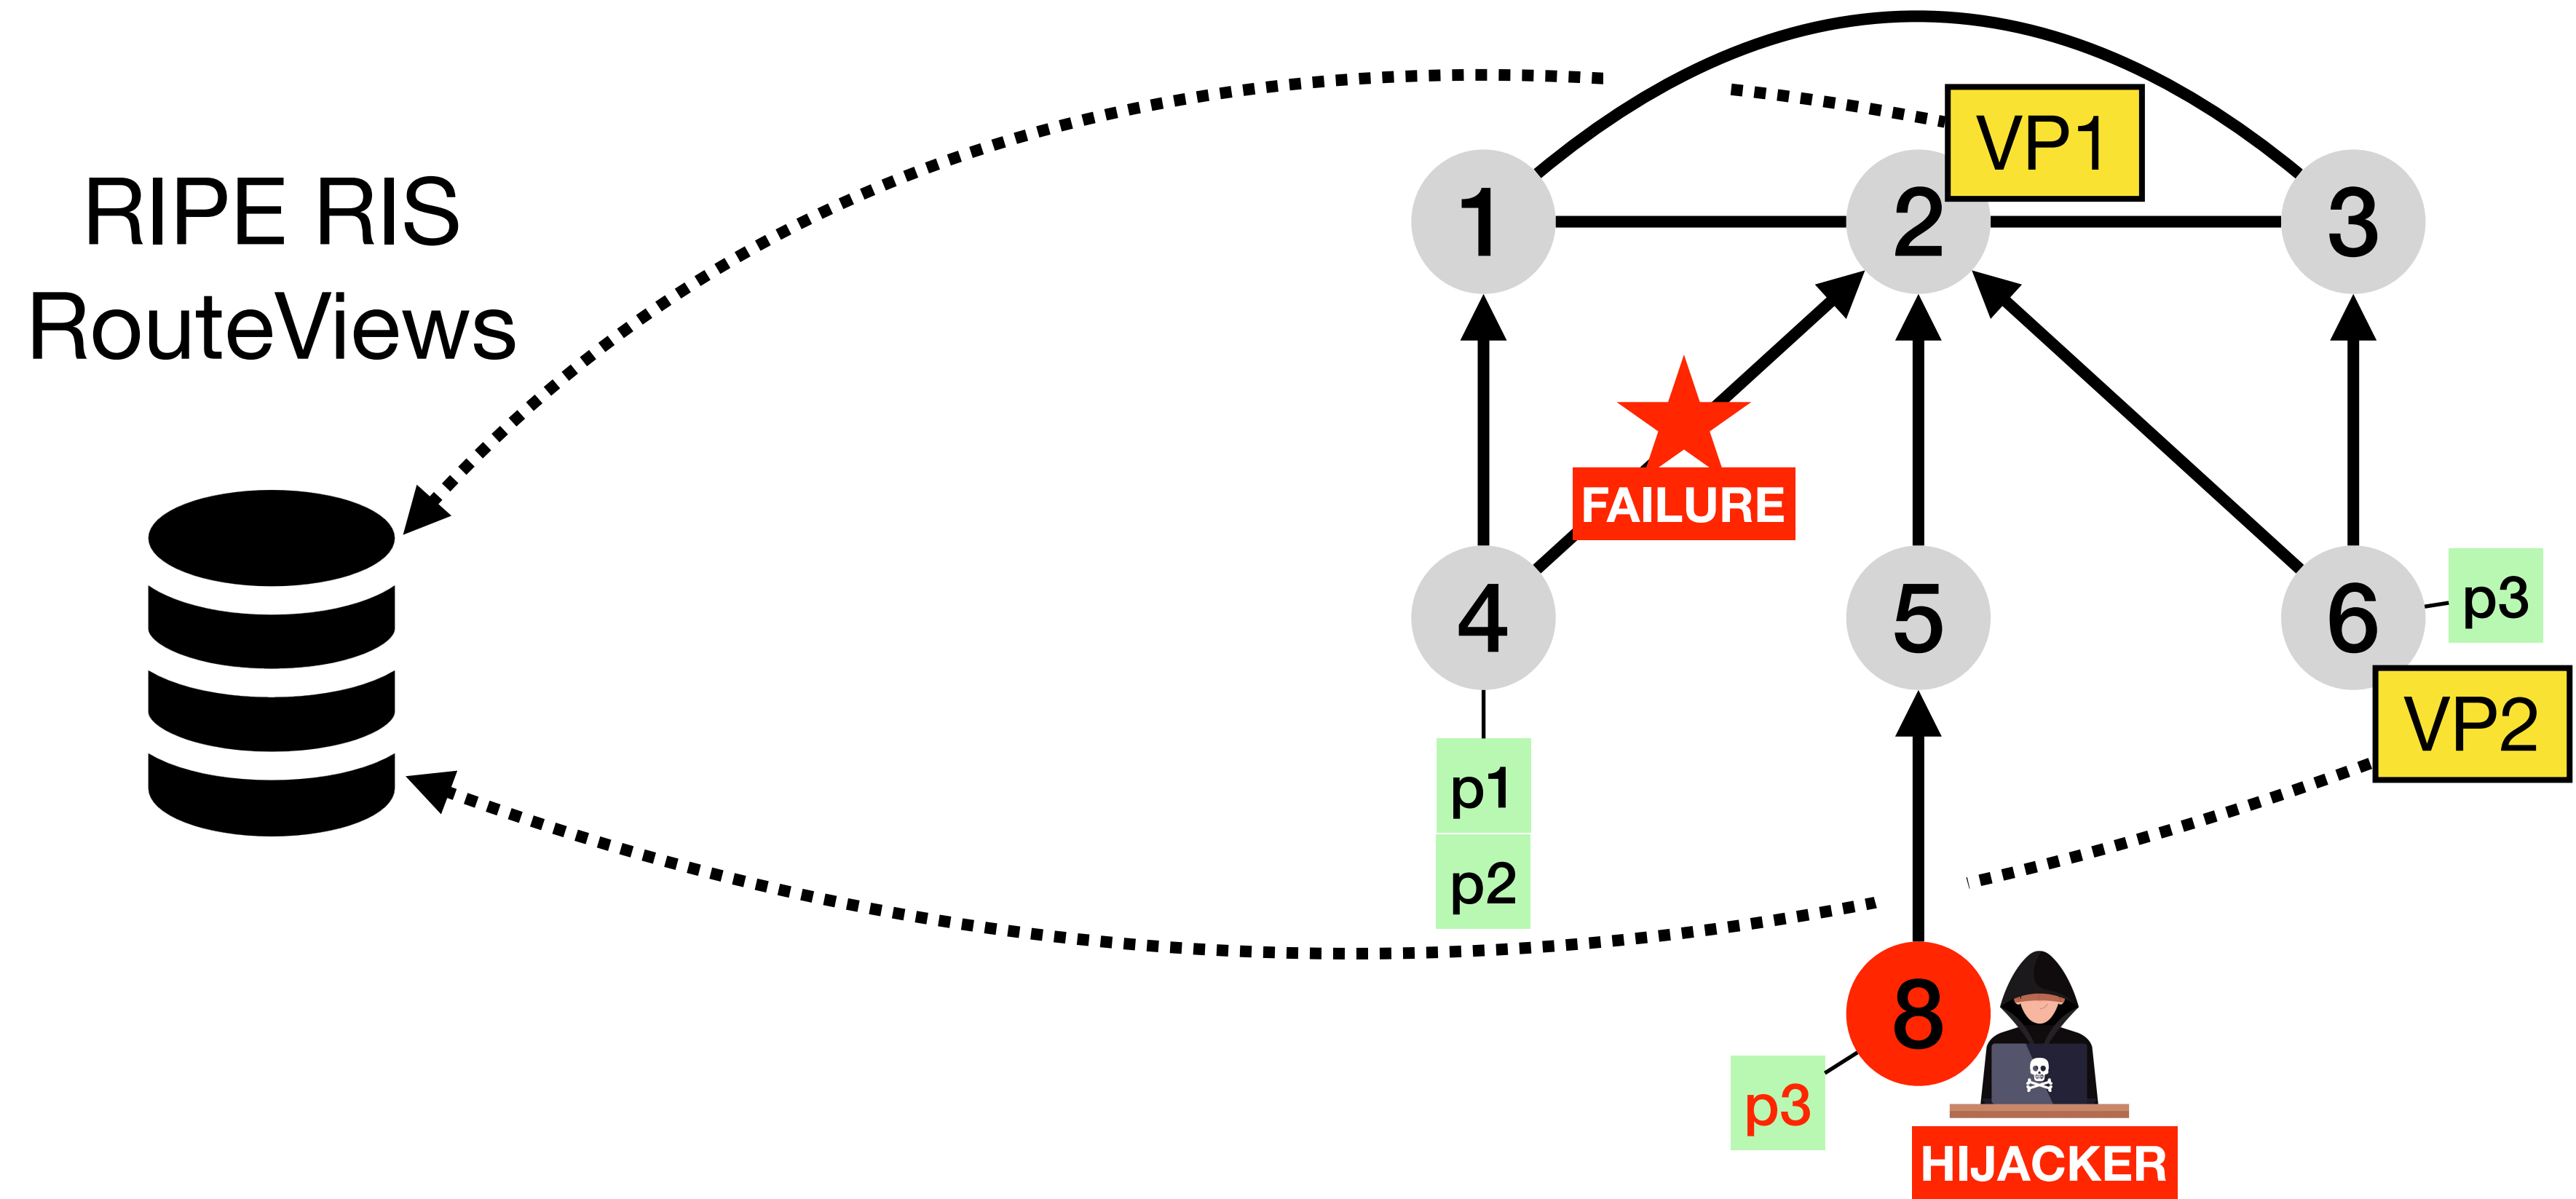
BGP routes can be redundant



BGP routes can be redundant



BGP routes can be redundant





BGP routes can be redundant

Collected routes

VP	prefix	AS path

RIPE RIS
RouteViews




BGP routes can be redundant

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4

RIPE RIS
RouteViews

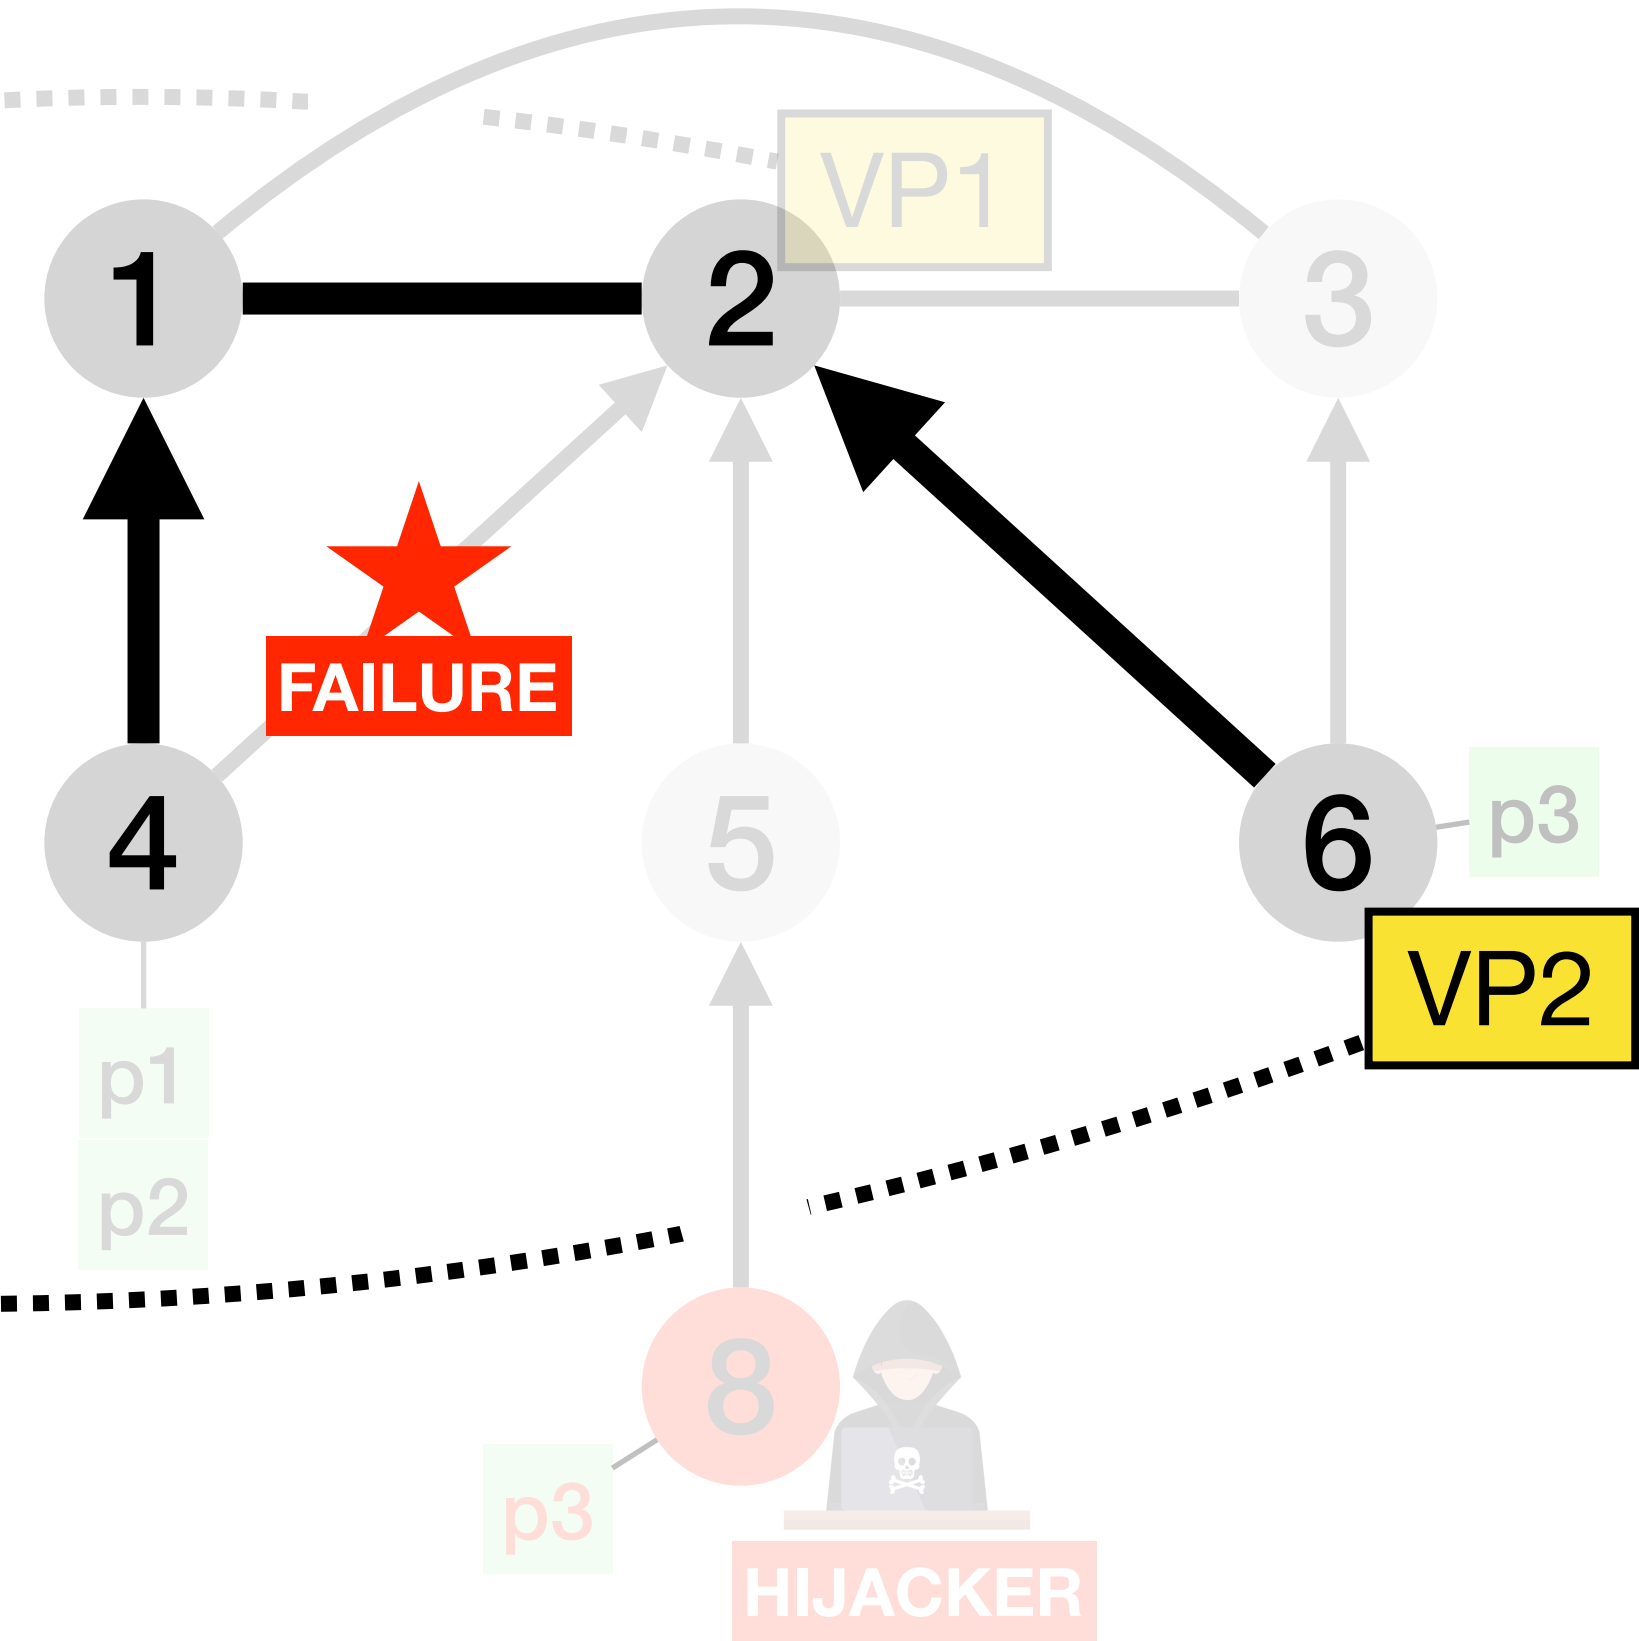


BGP routes can be redundant

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

RIPE RIS
RouteViews



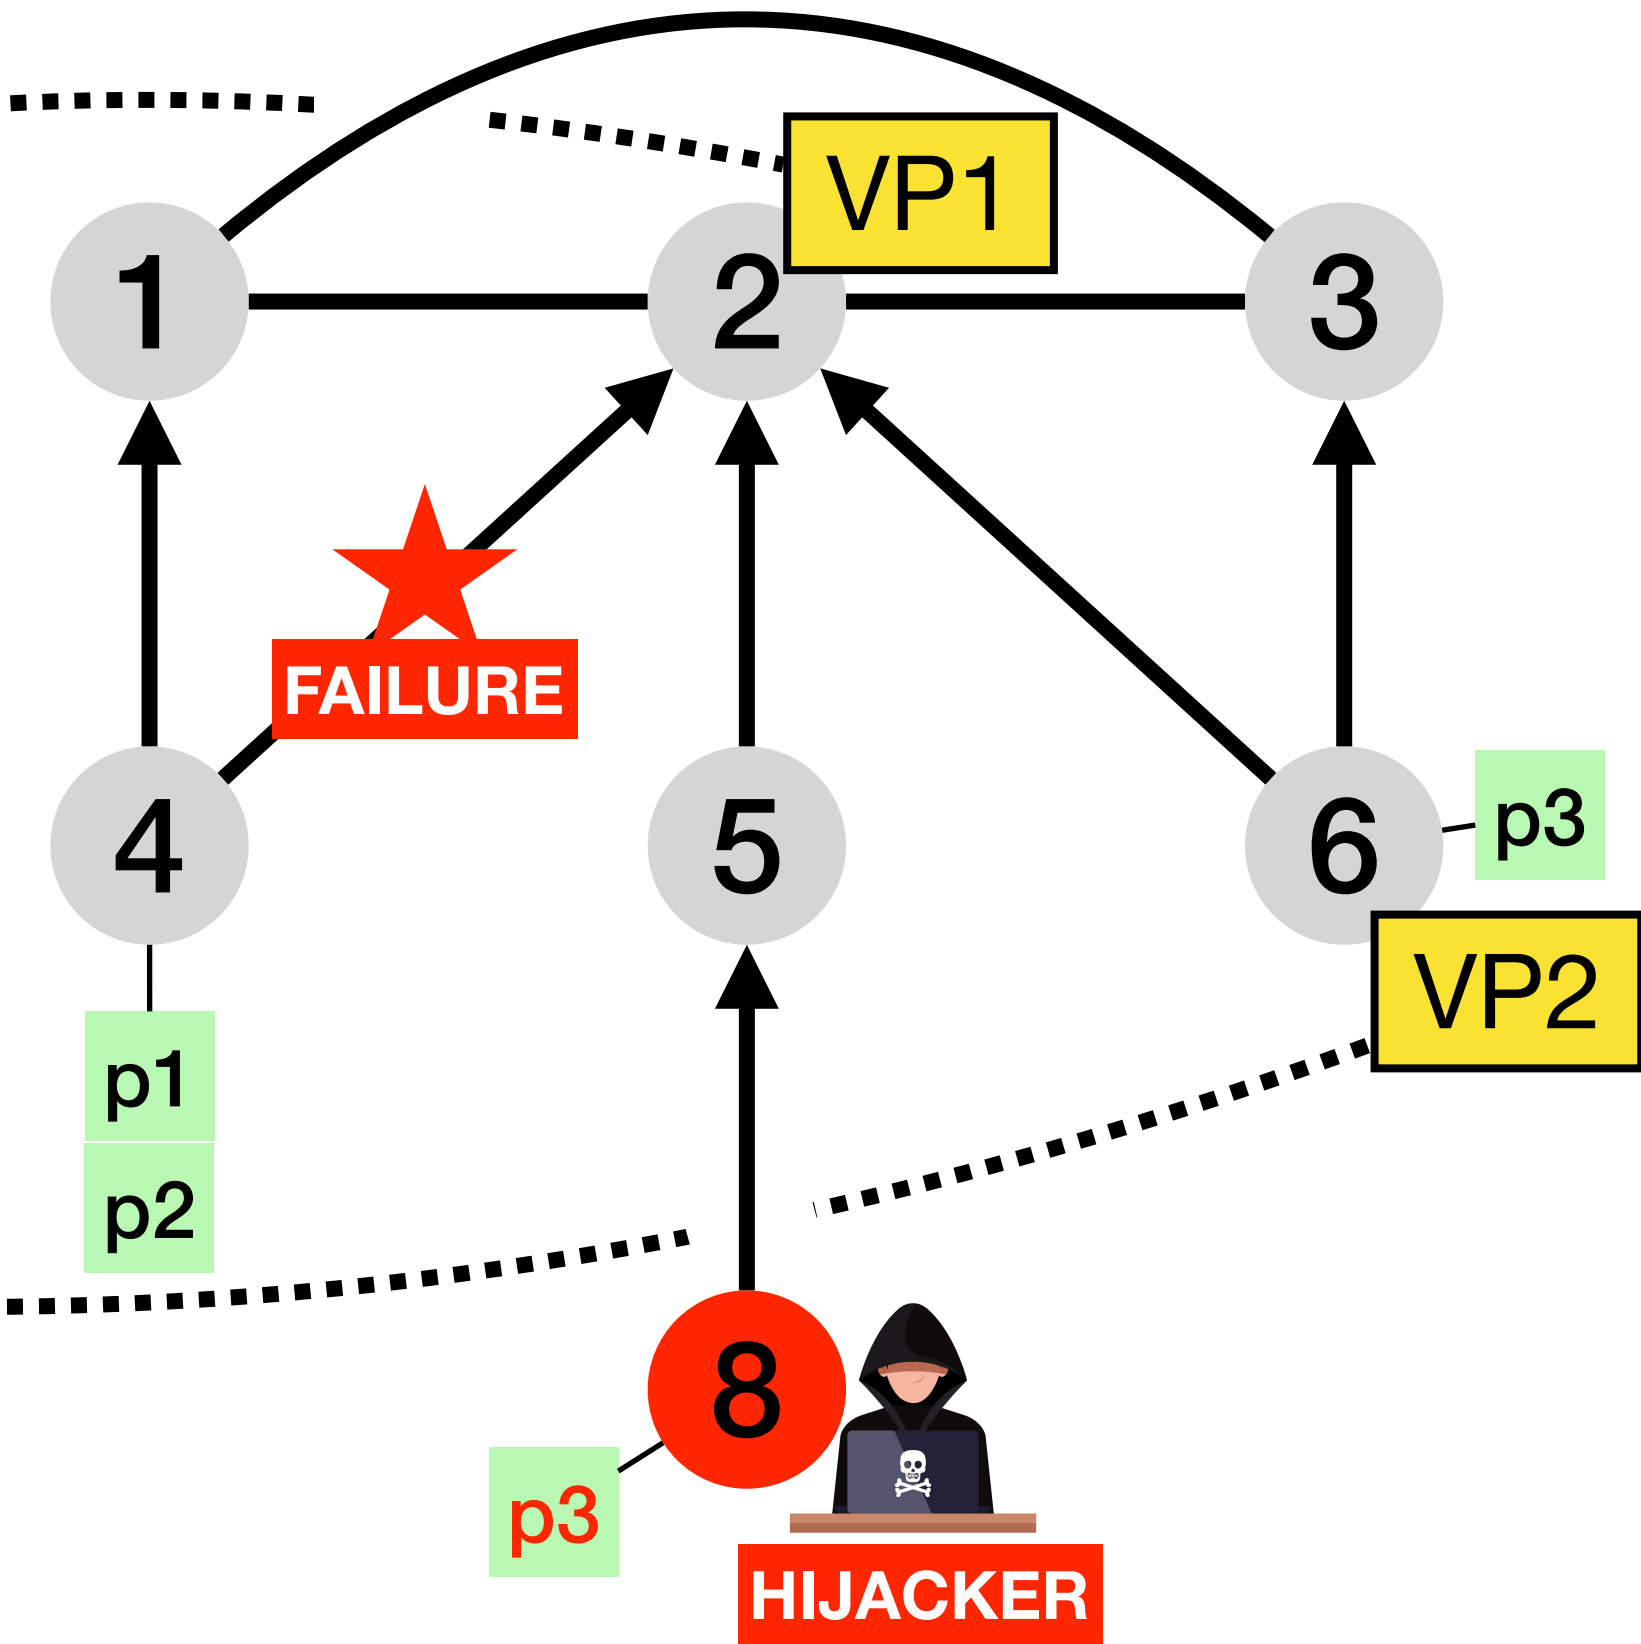
BGP routes can be redundant

Redundant routes

VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

Redundant routes

RIPE RIS
RouteViews



Redundant BGP routes are not so useful

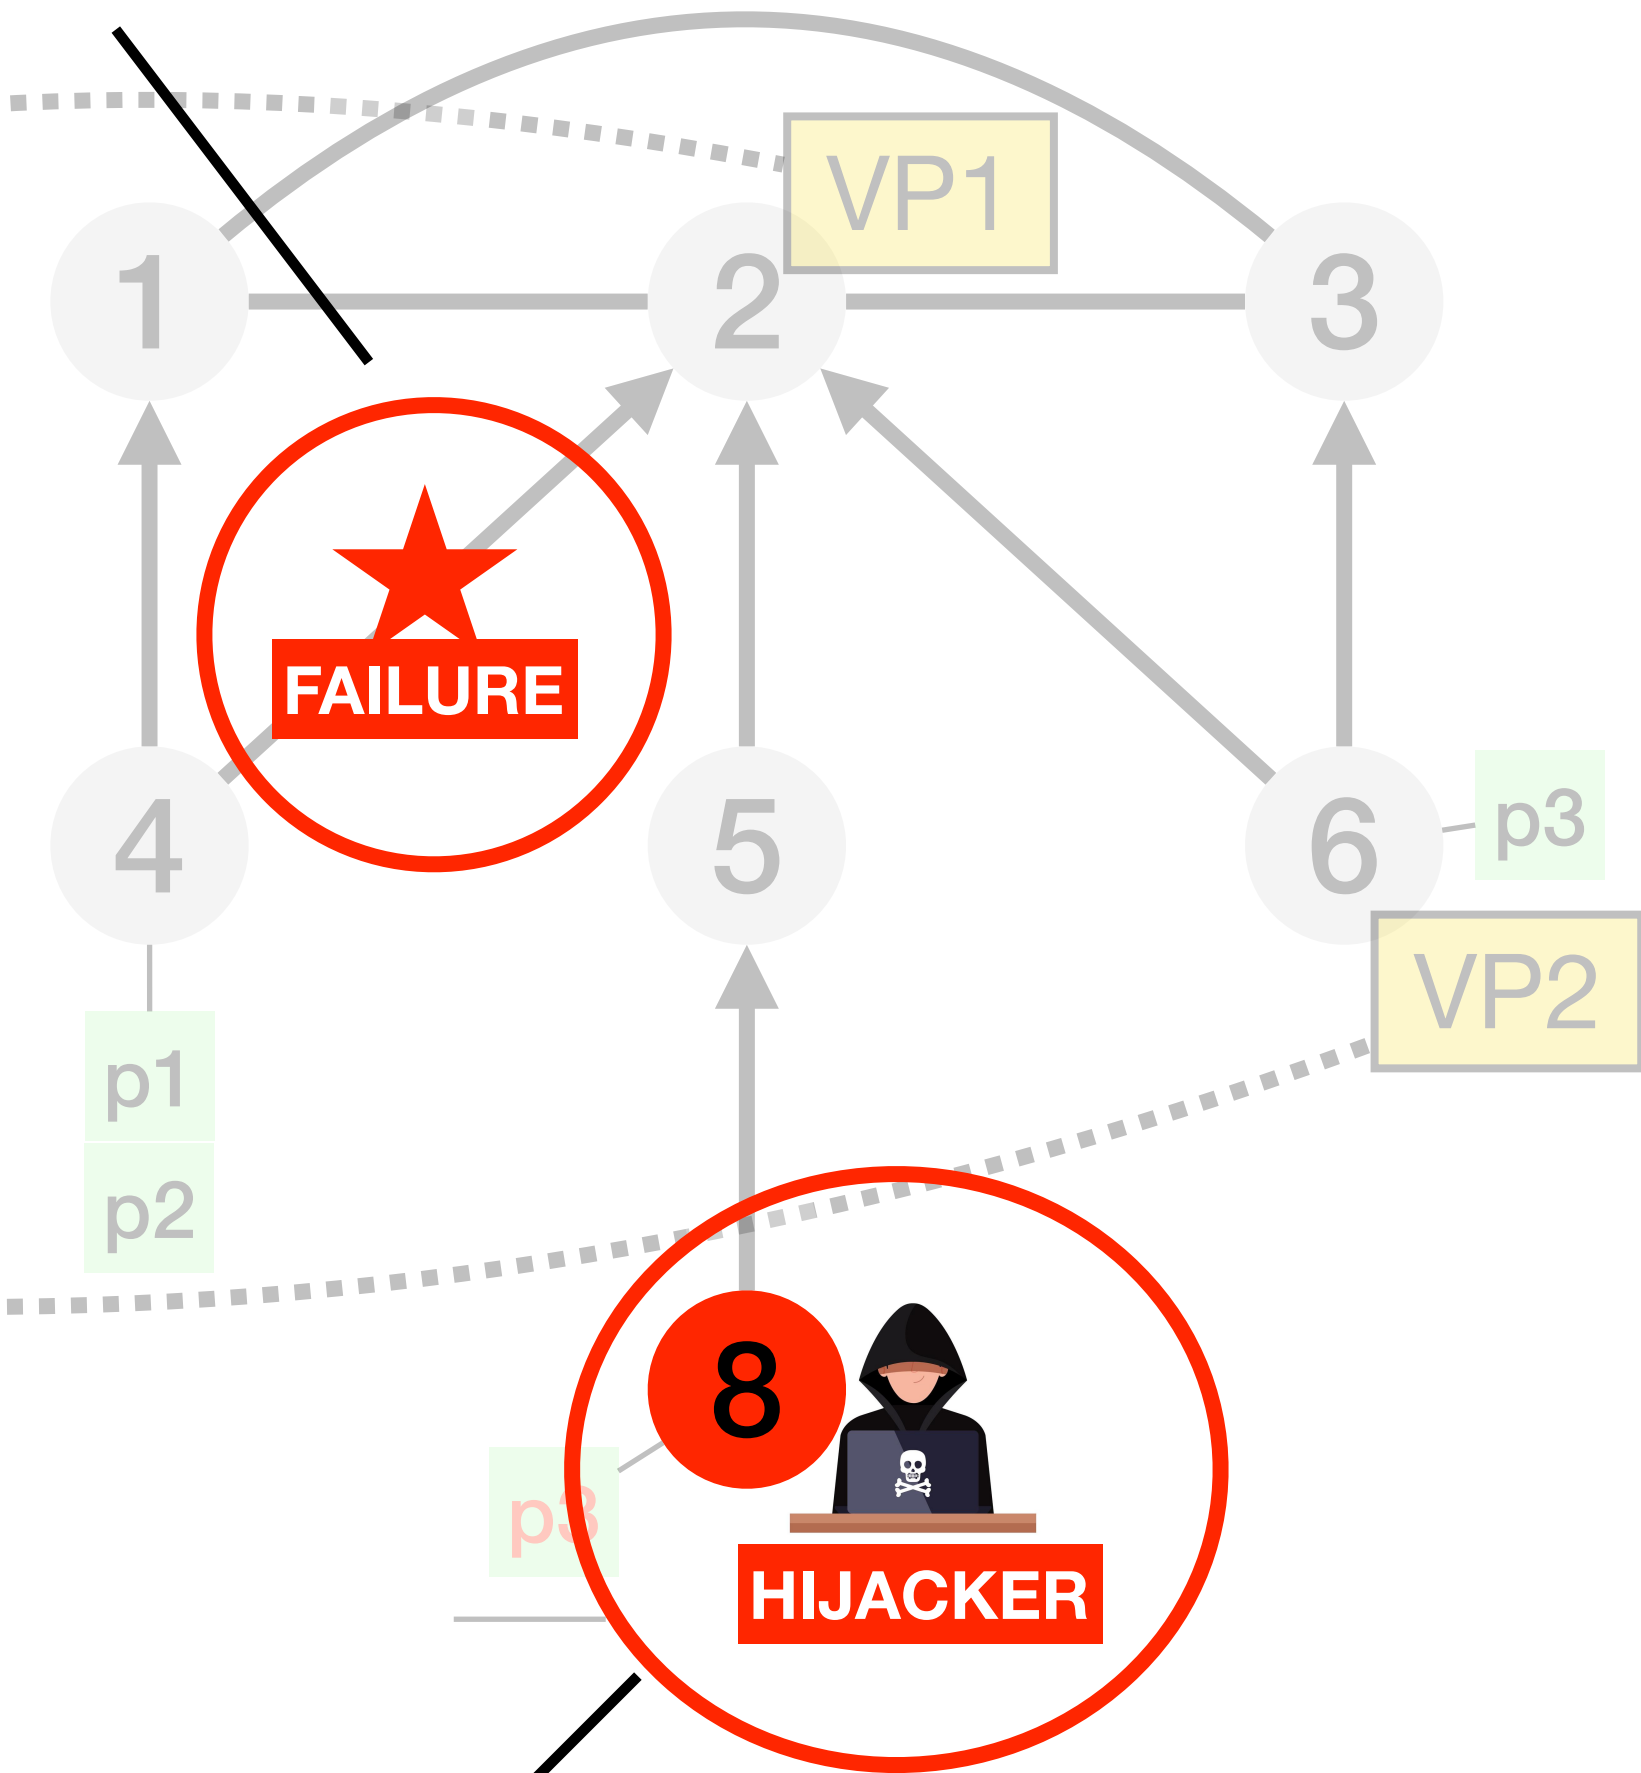
The failure is visible in one direction only

Redundant routes

VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4

Redundant routes

RIPE RIS
RouteViews




The hijack is not detected

Outline

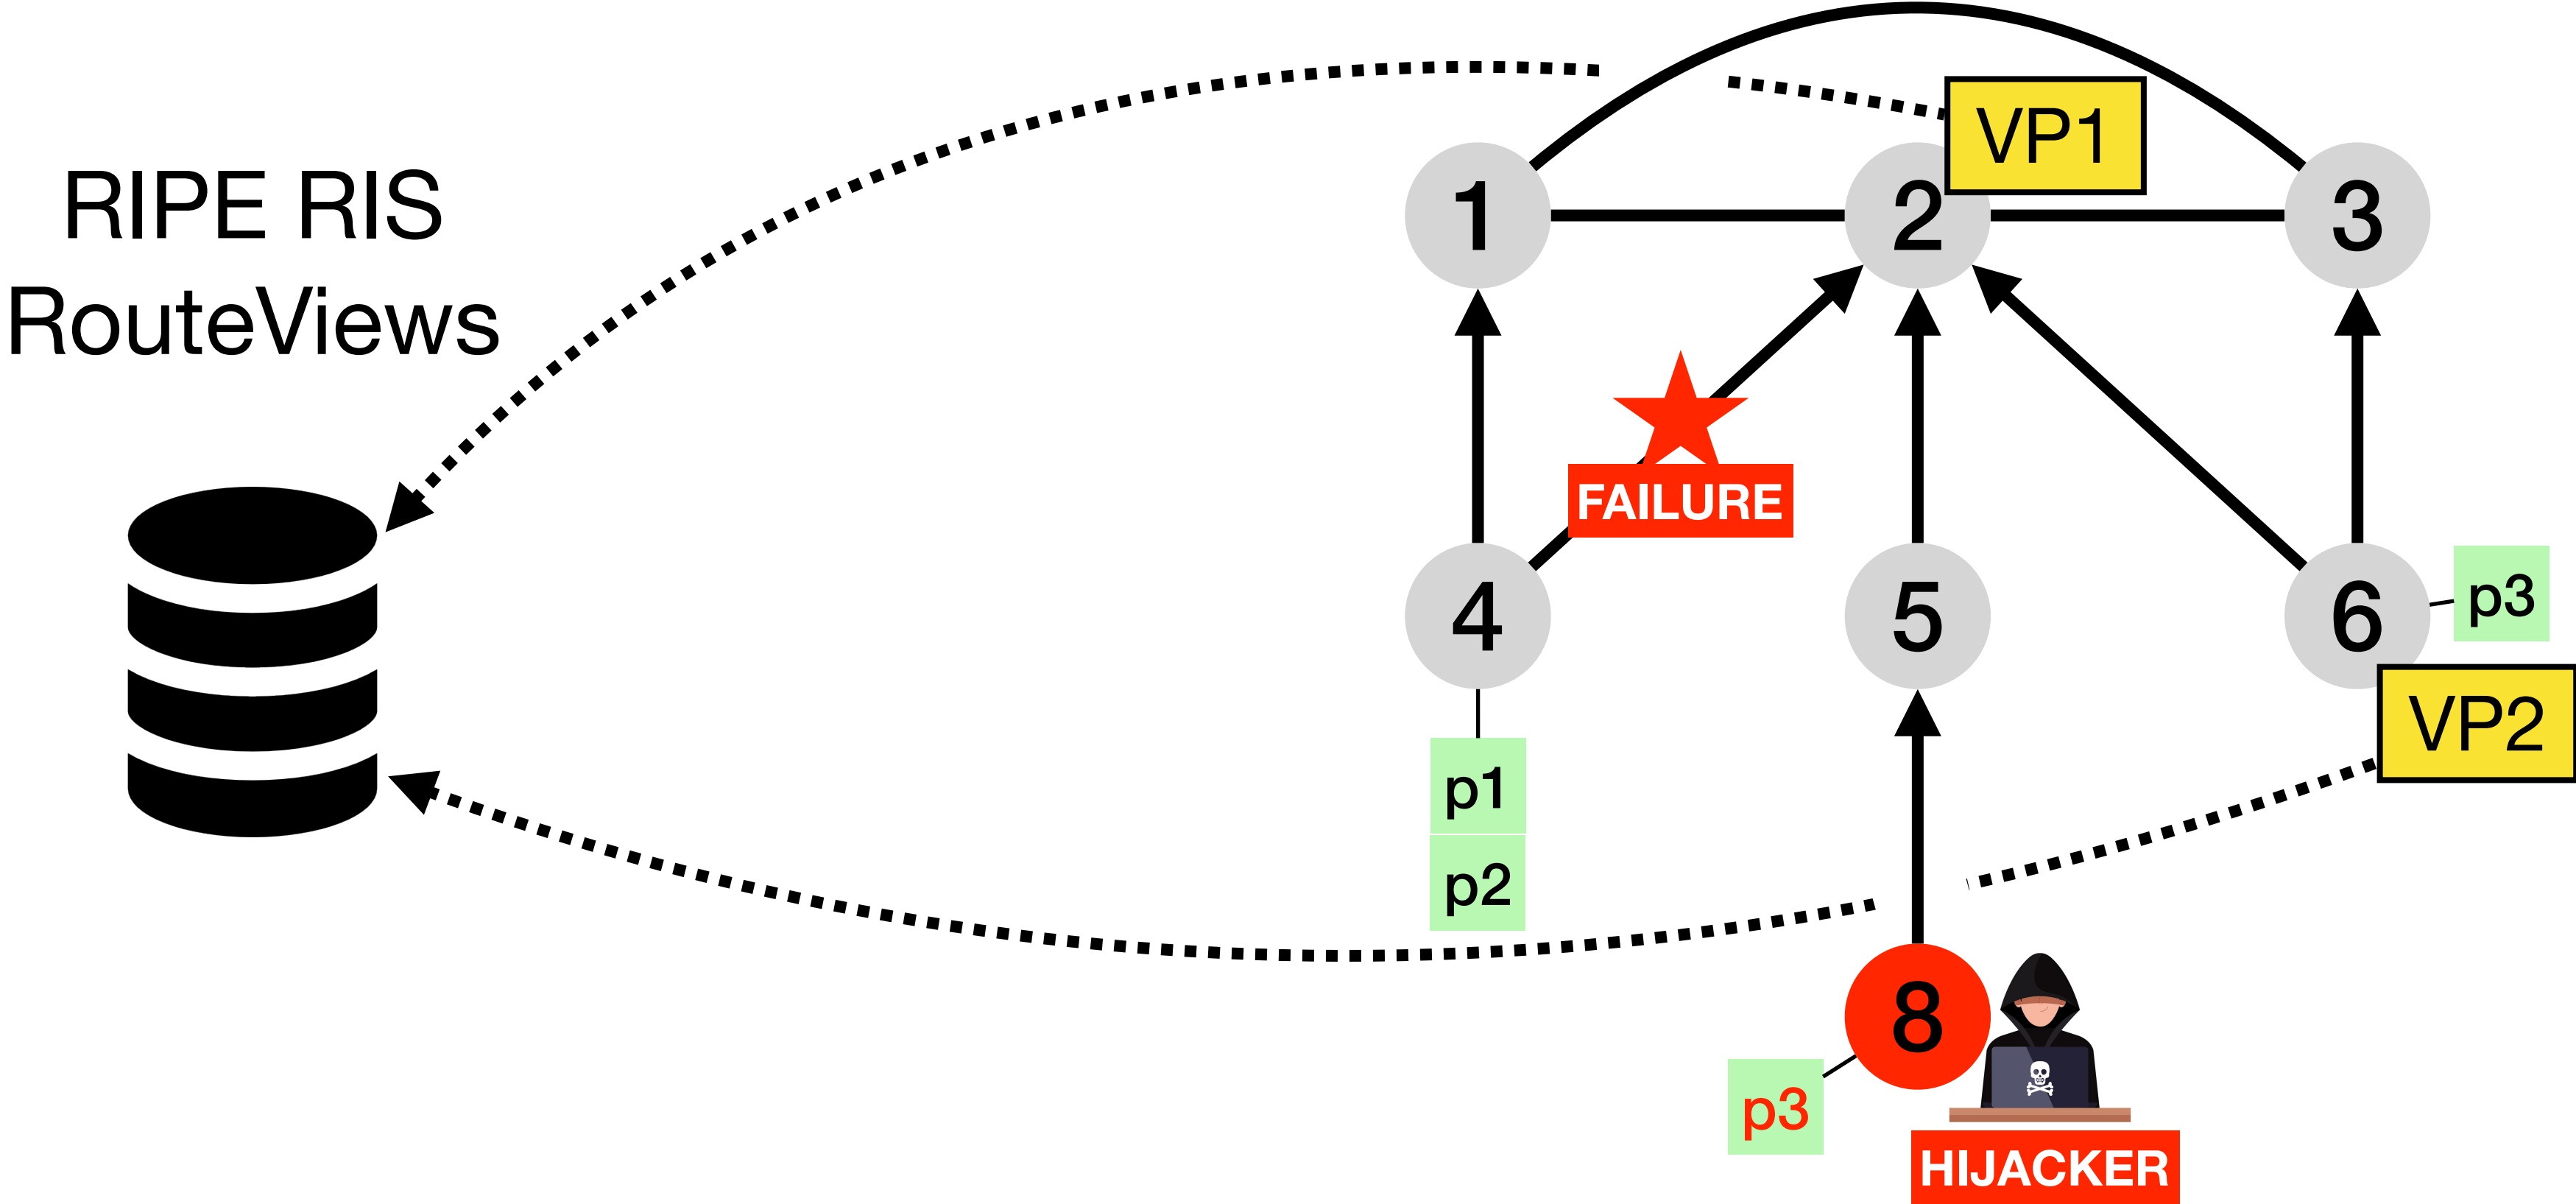
1. We observe that BGP routes are often redundant
2. Redundant BGP routes enable an **overshoot-and-discard** collection scheme

The “overshoot-and-discard” data collection paradigm

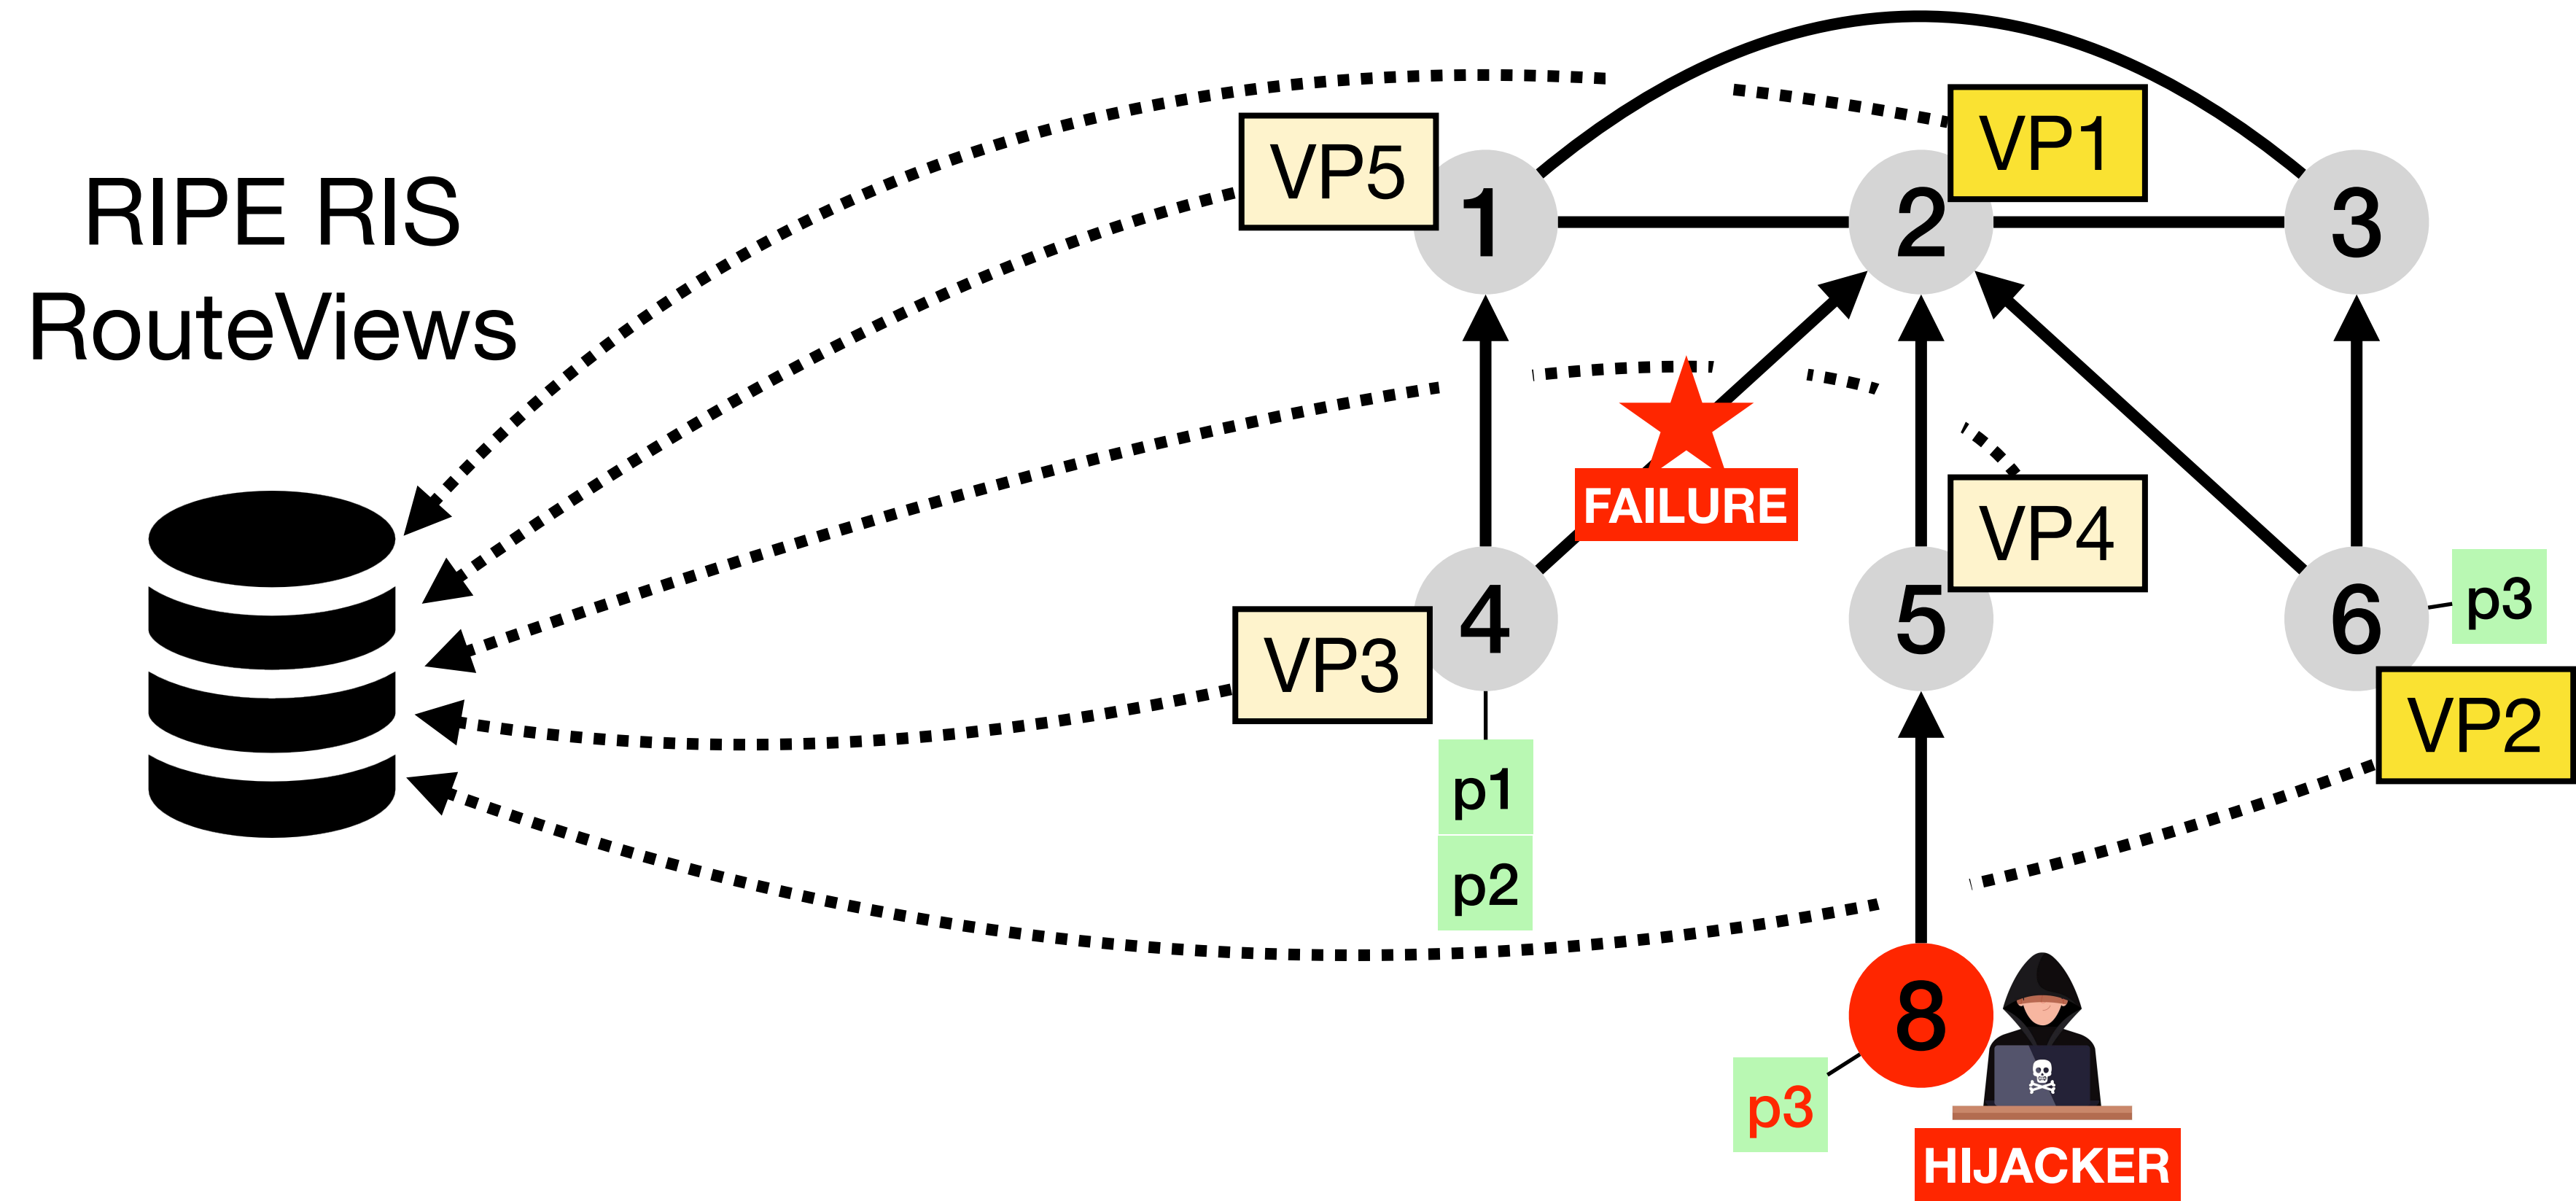


Overshoot: We collect data from as many VPs as possible
To prevent missing important information

Overshoot: deploying as many VPs as possible



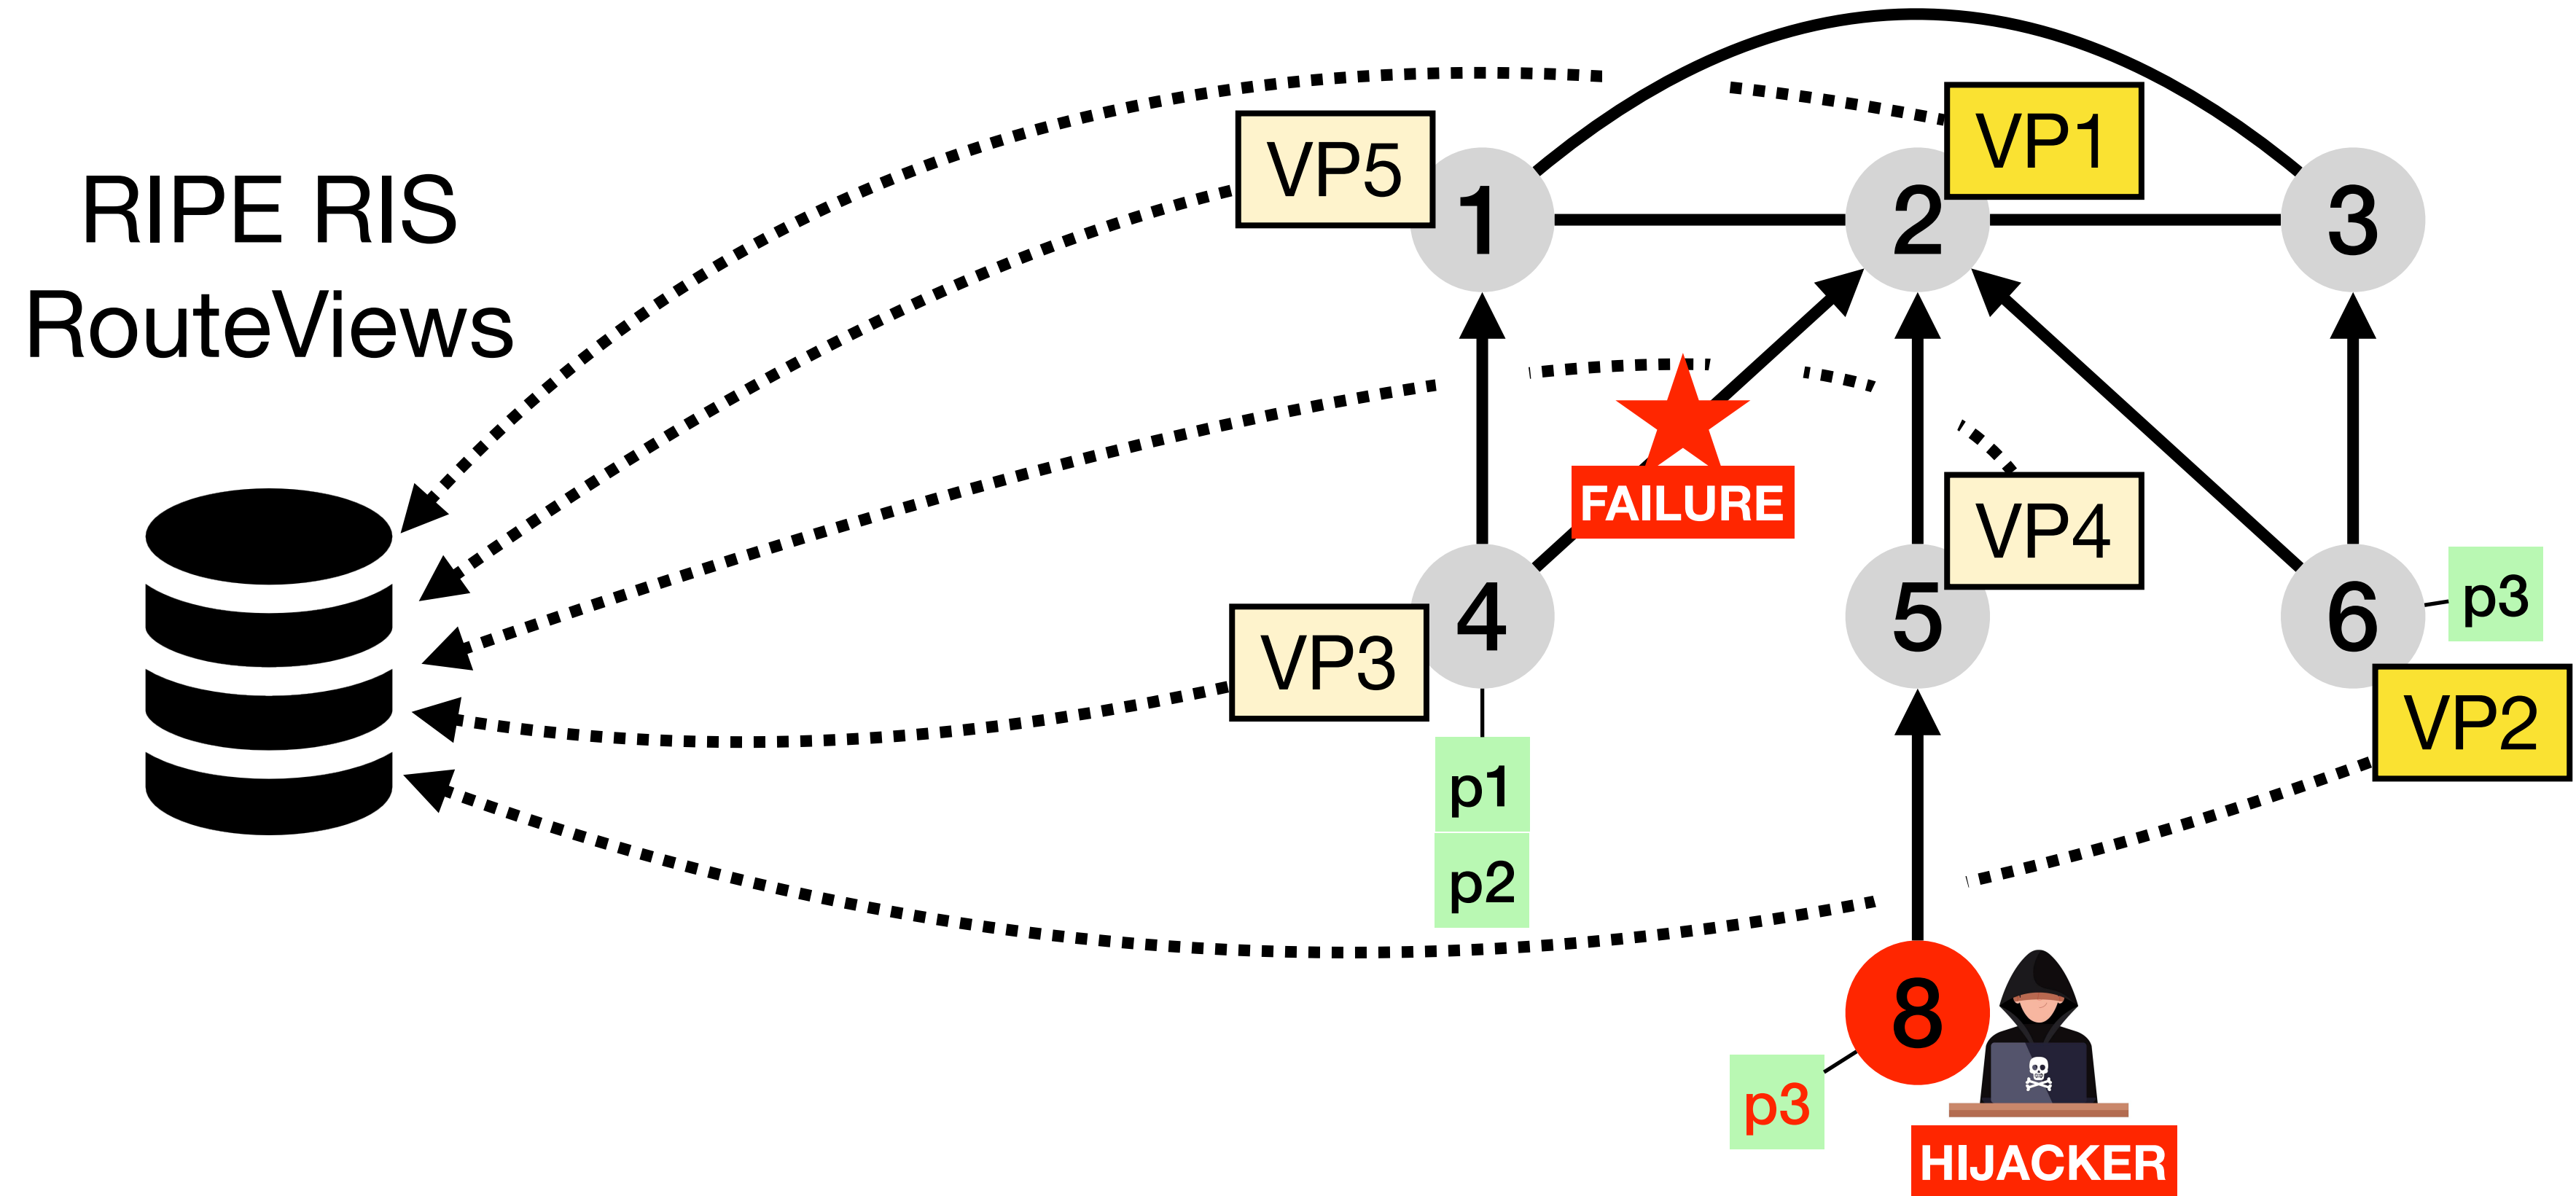
Overshoot: deploying as many VPs as possible



Overshoot: deploying as many VPs as possible To prevent missing important information

Collected routes

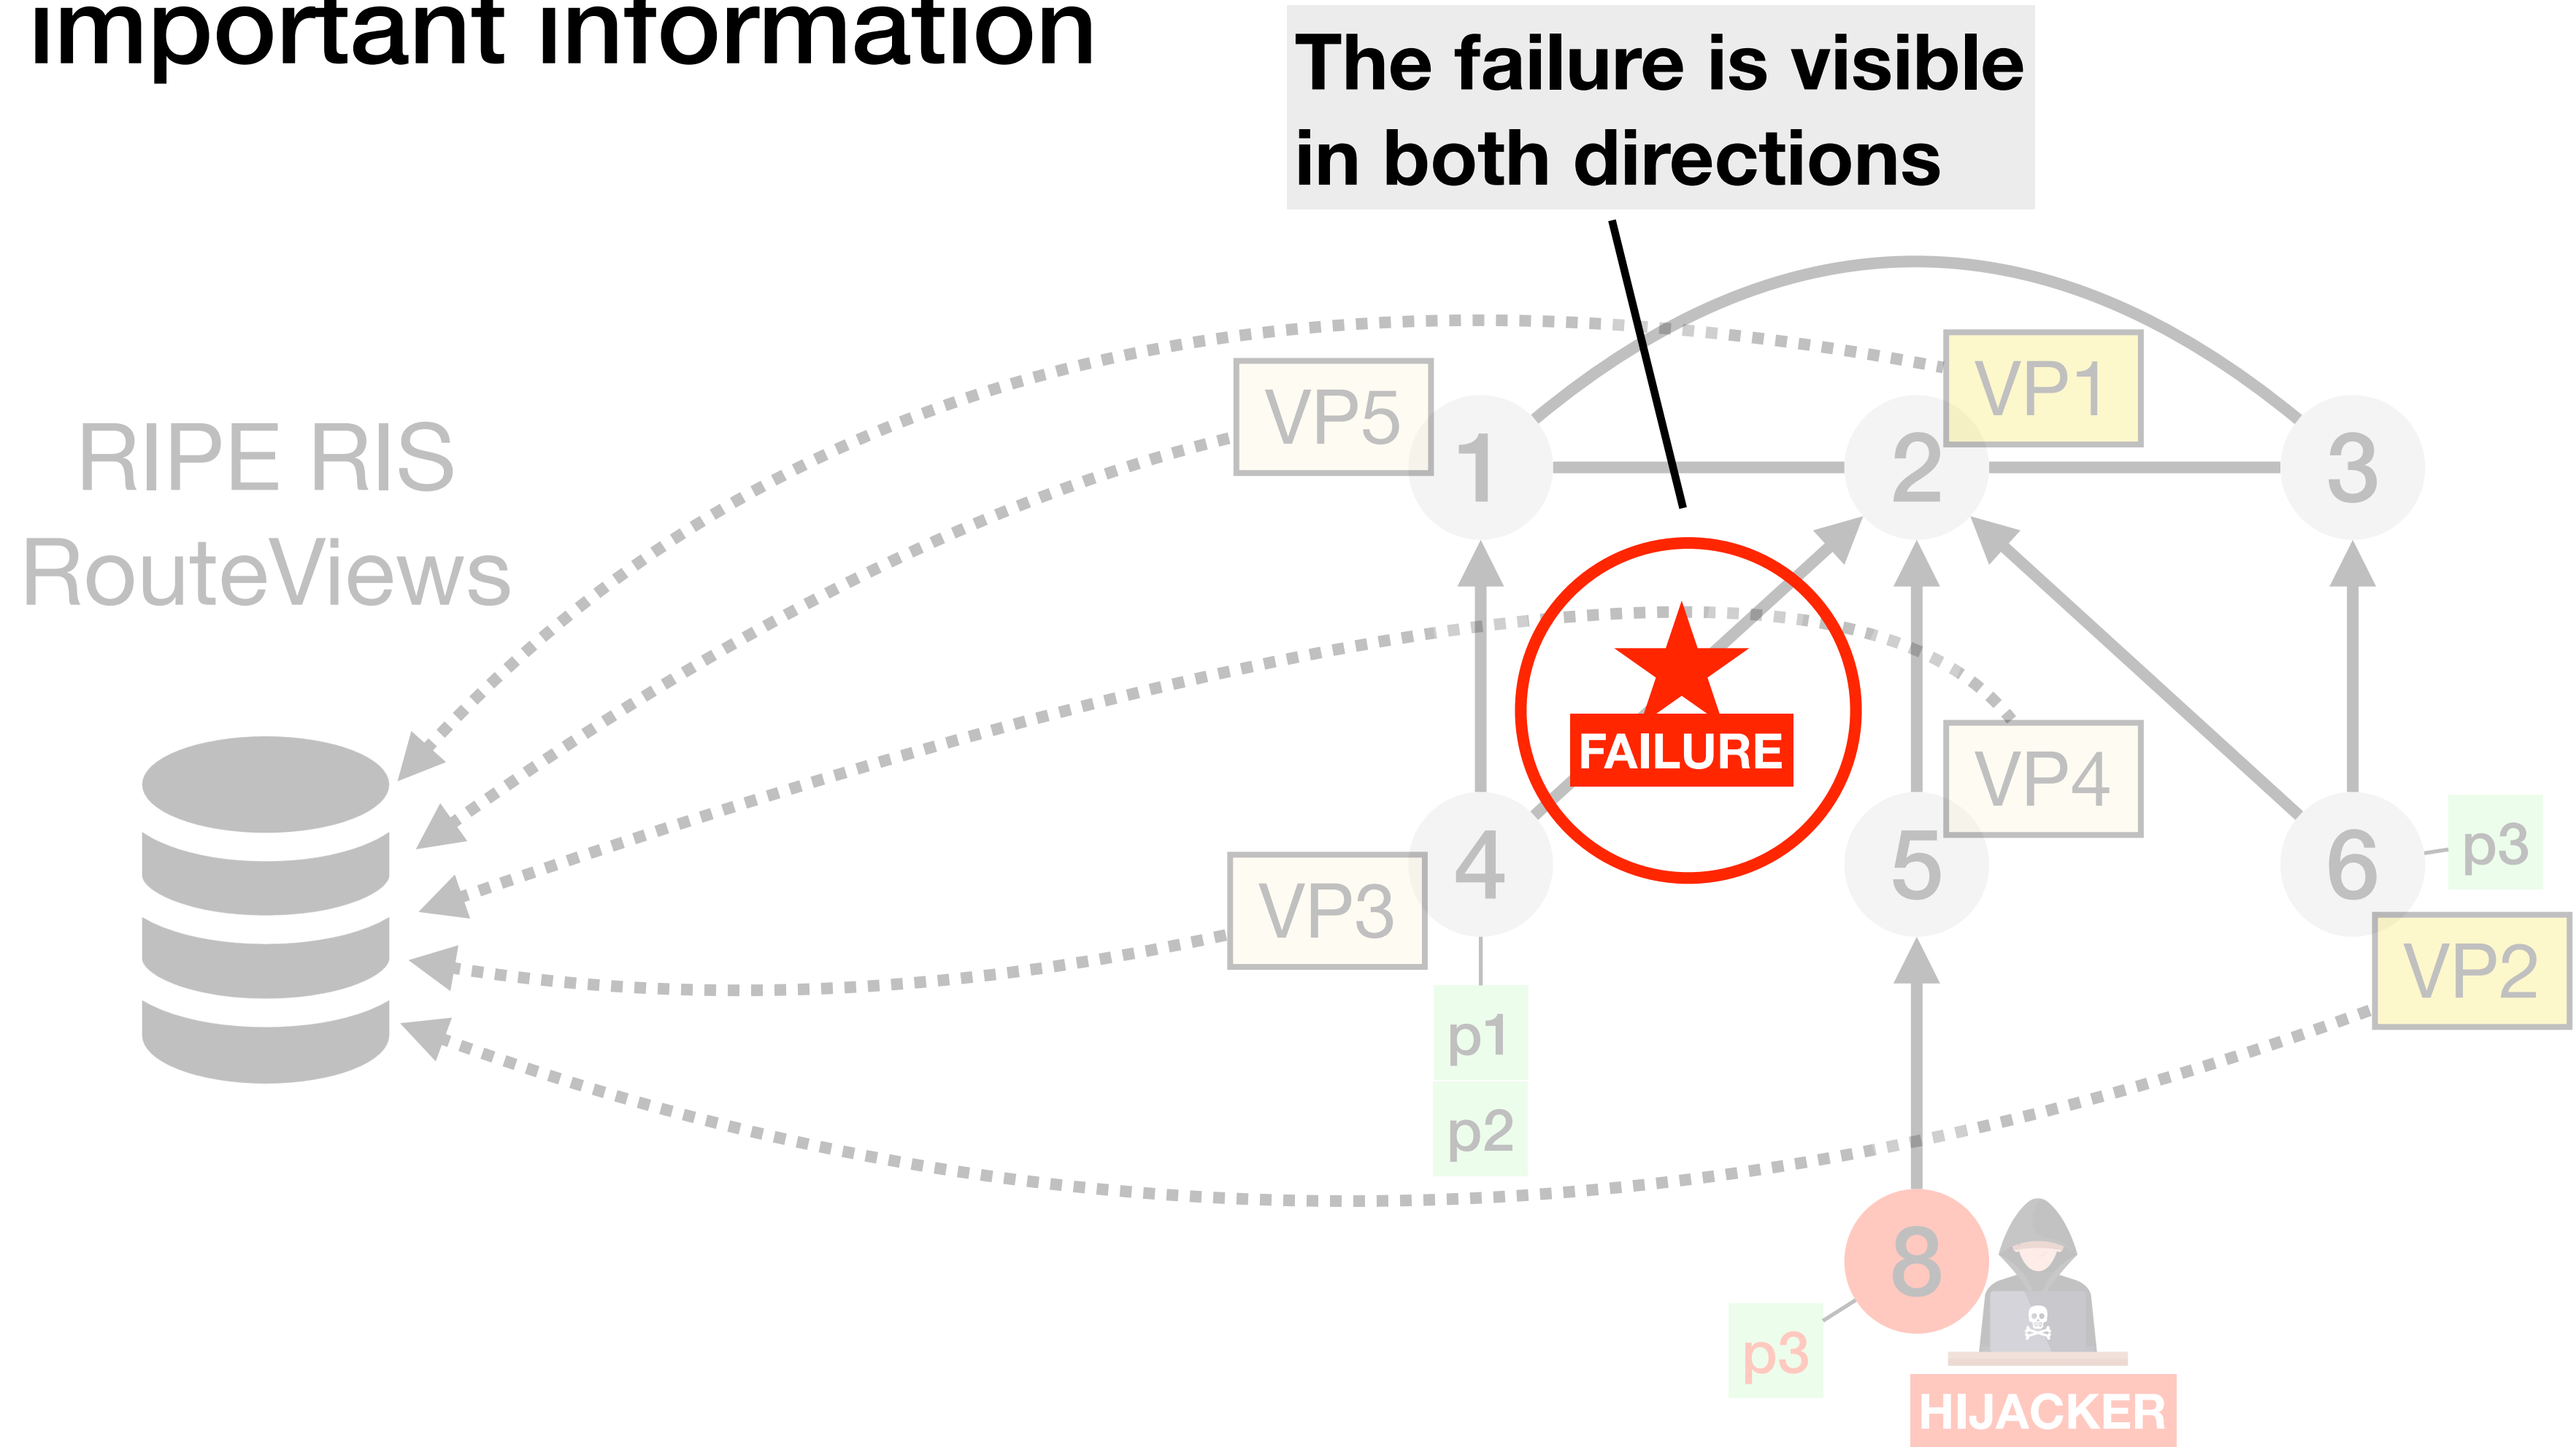
VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8



Overshoot: deploying as many VPs as possible To prevent missing important information

Collected routes

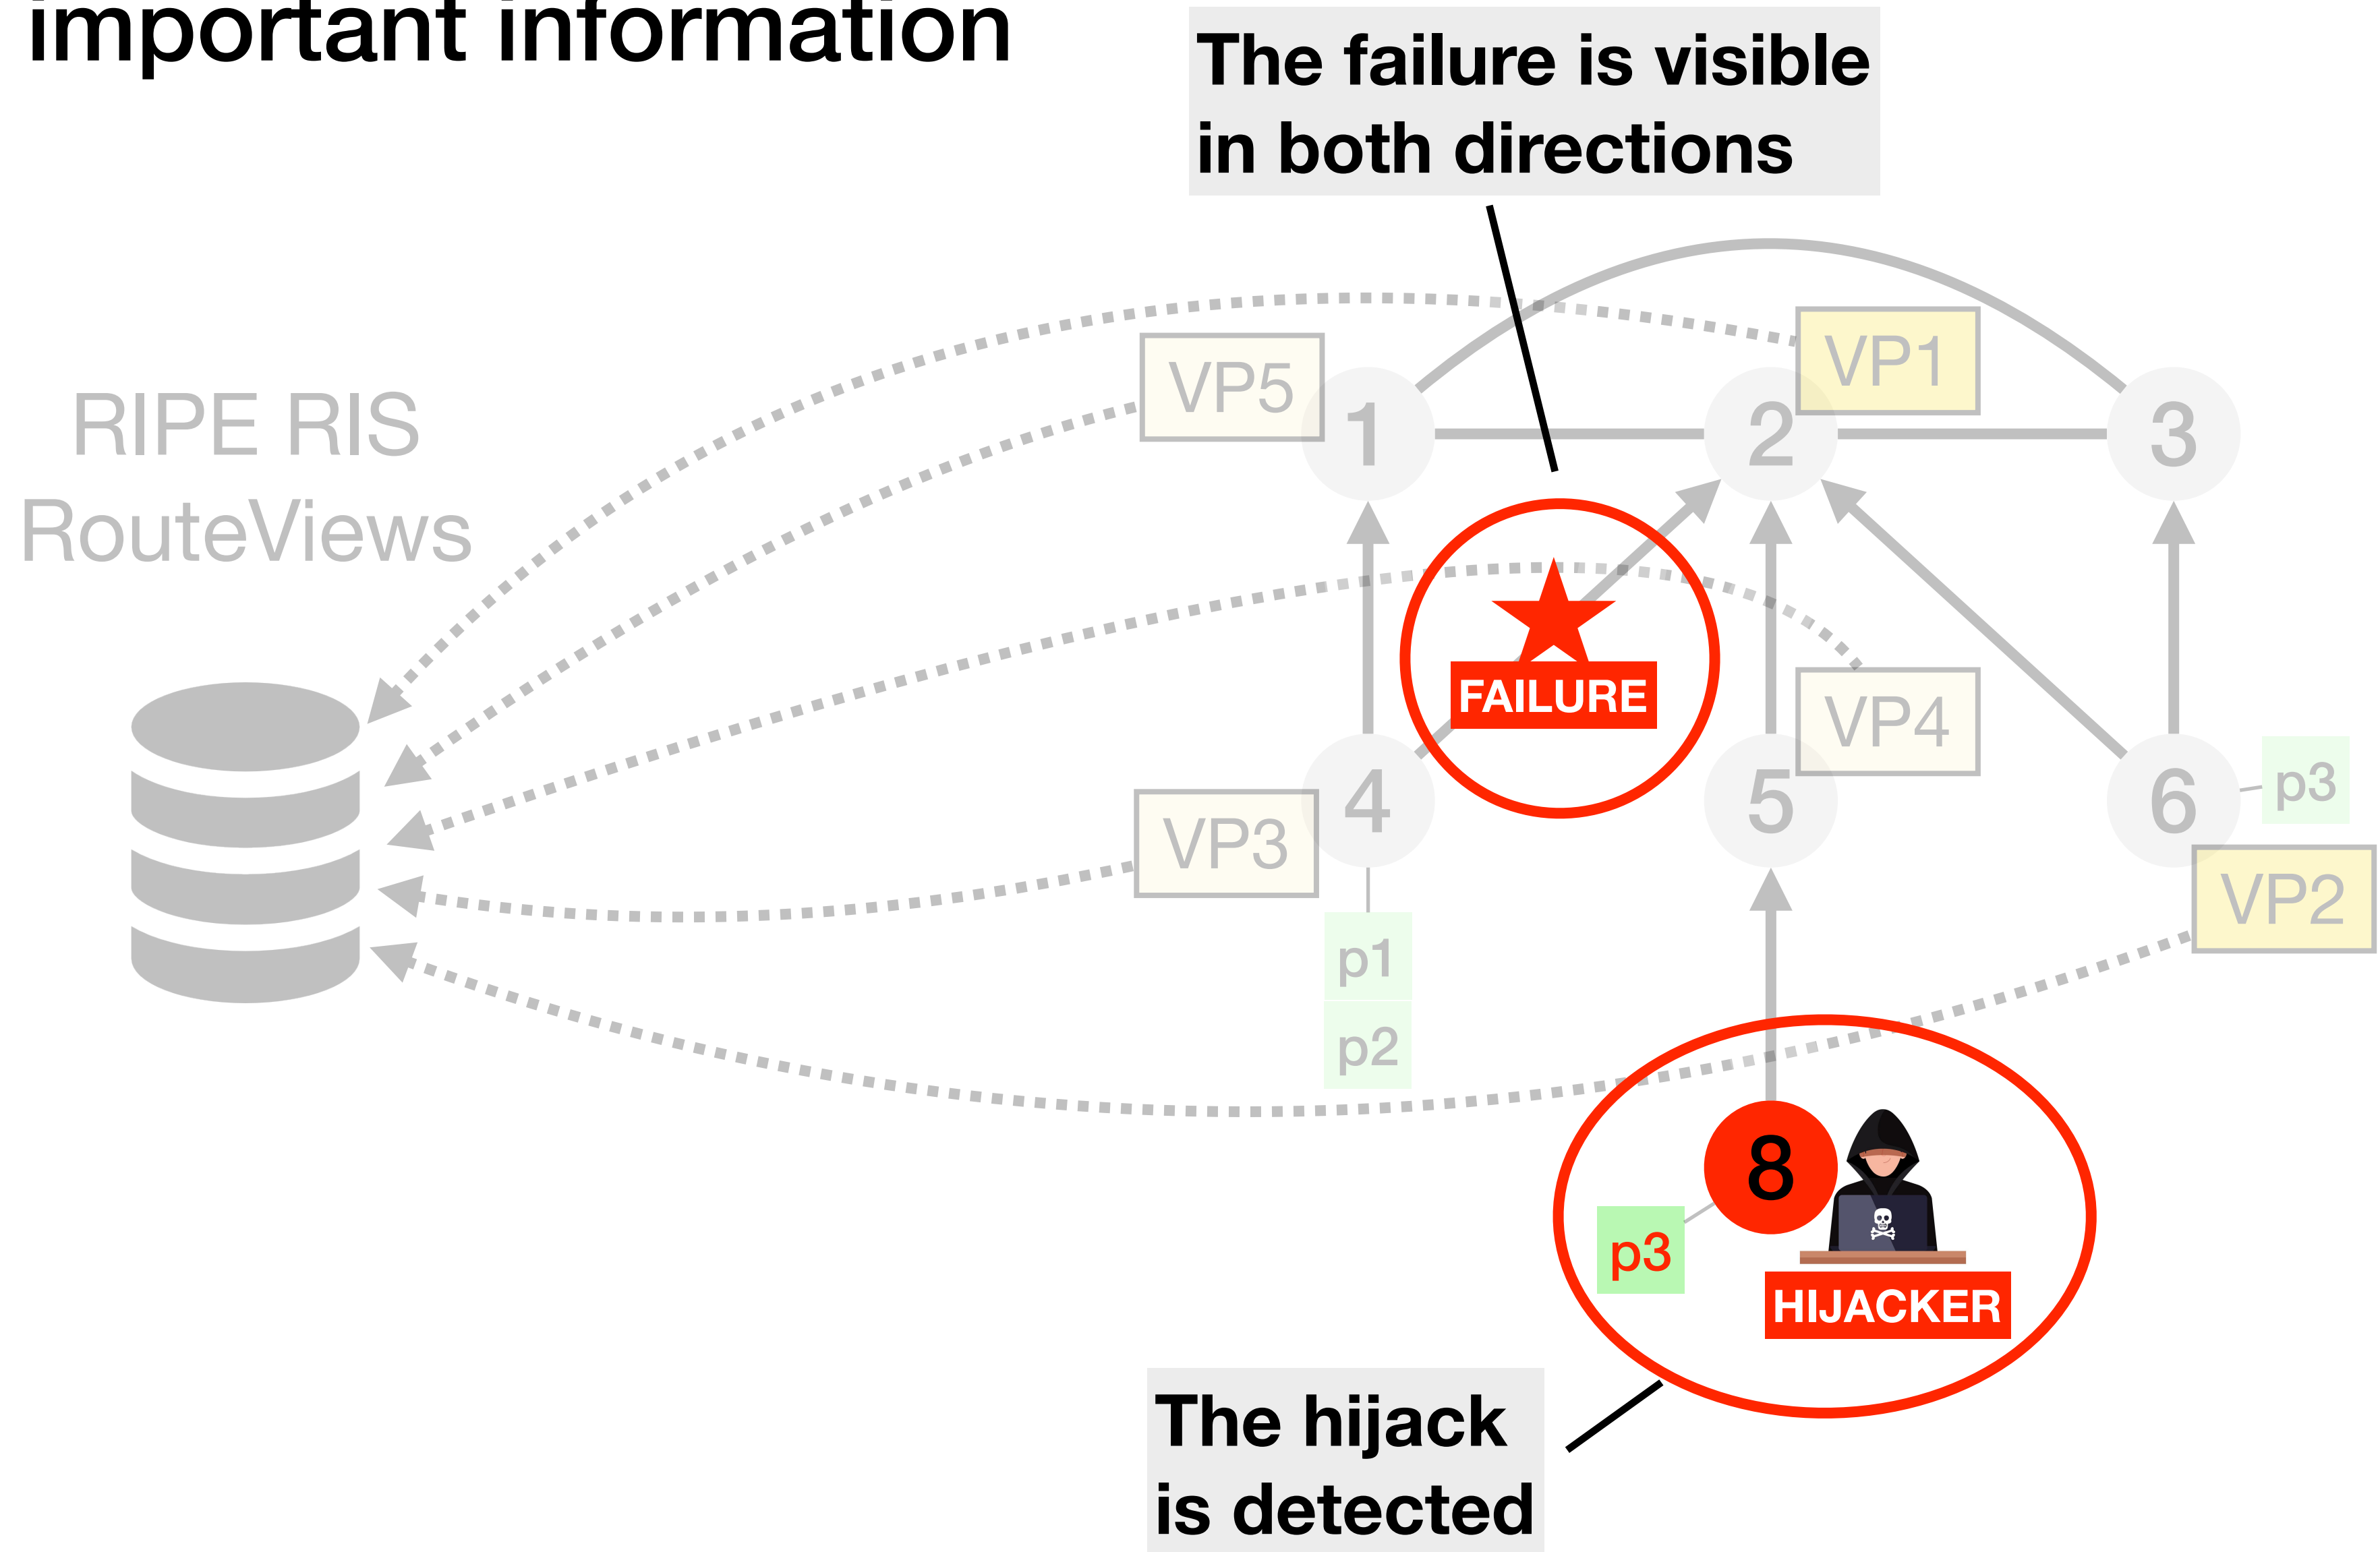
VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8




Overshoot: deploying as many VPs as possible To prevent missing important information

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8



The “overshoot-and-discard” data collection paradigm



Overshoot: We collect data from as many VPs as possible
To prevent missing important information

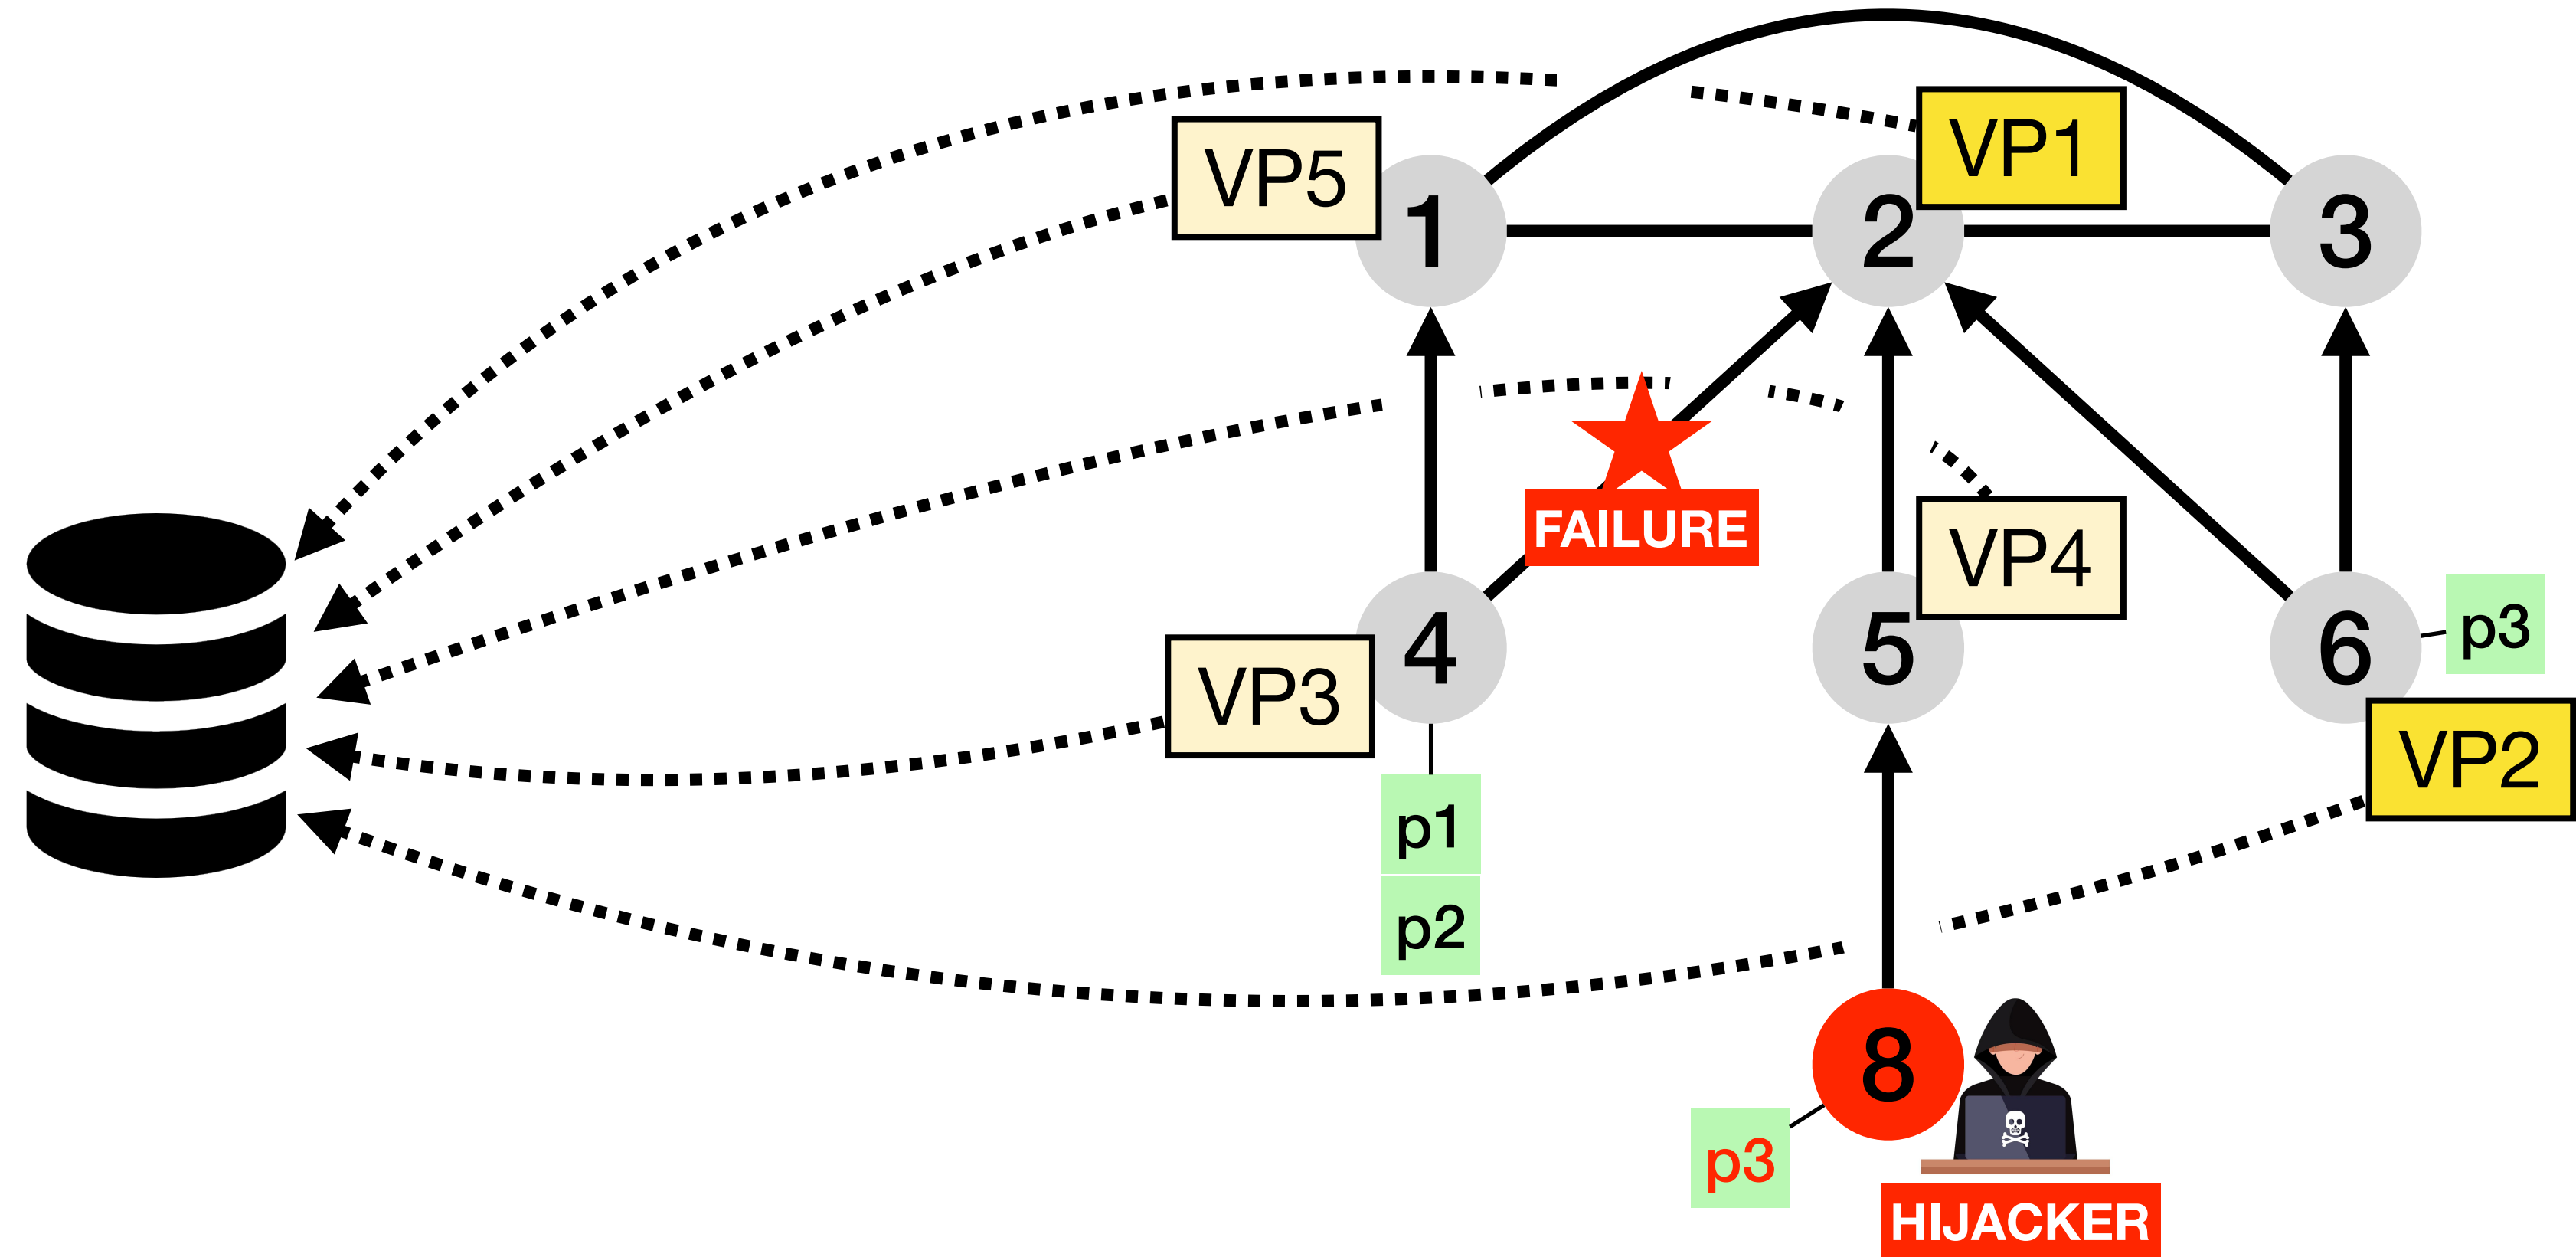


Discard: We filter out the redundant BGP routes
To reduce the volume of data collected

Discard: redundant BGP routes are discarded using filters

Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8



Discard: redundant BGP routes are discarded using filters

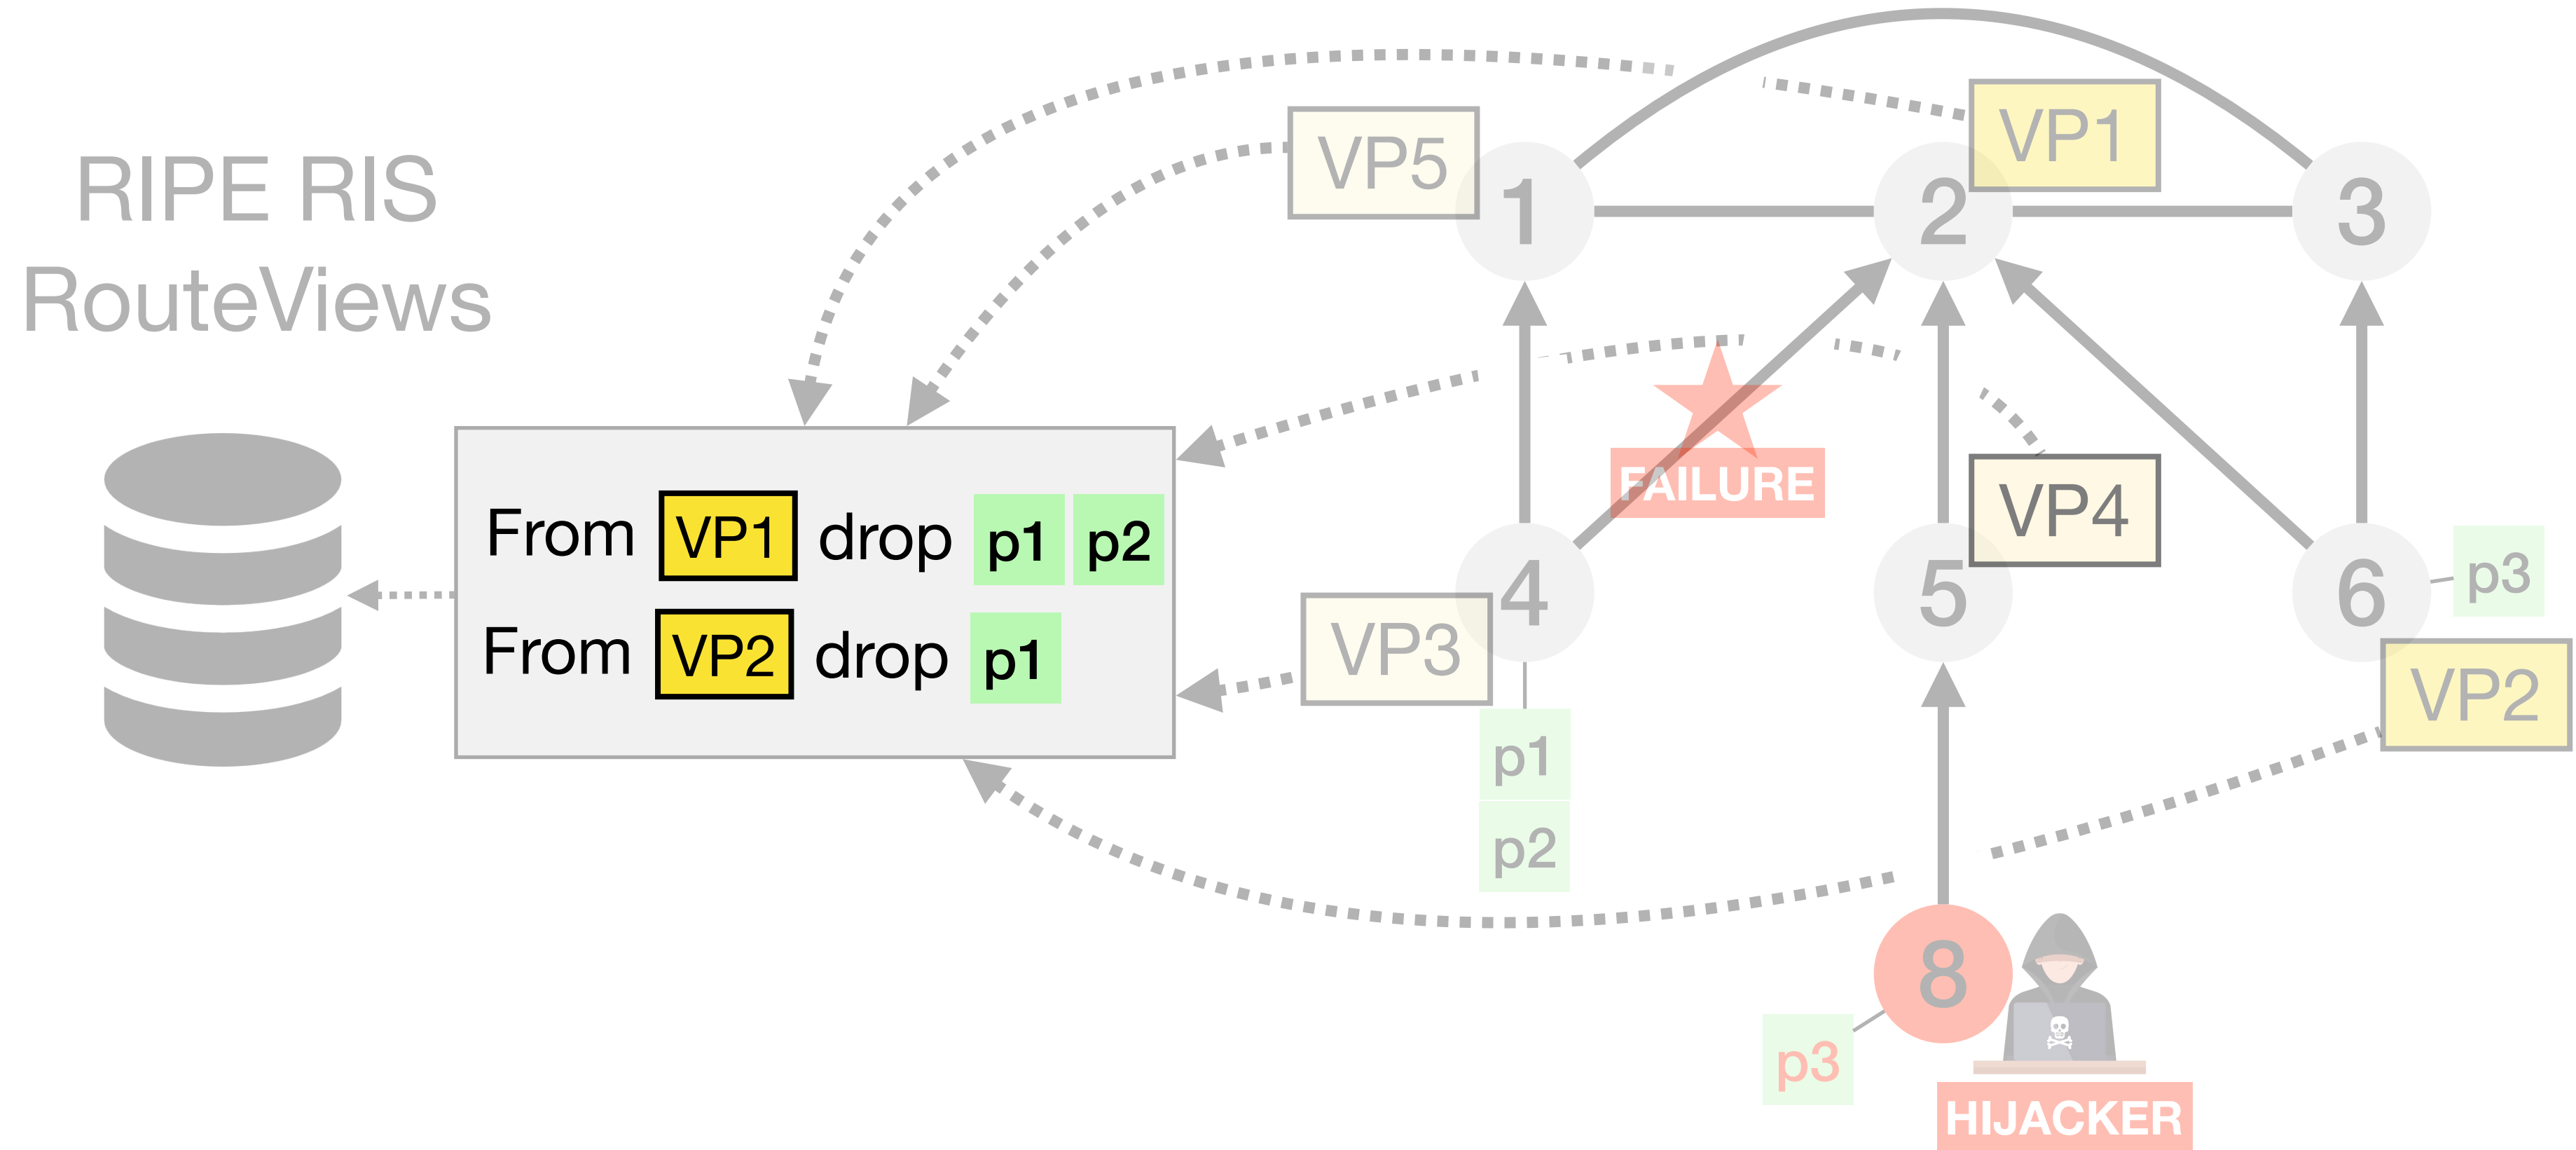
Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	3 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8

RIPE RIS
RouteViews



From VP1 drop p1 p2
From VP2 drop p1



Discard: redundant BGP routes are discarded using filters

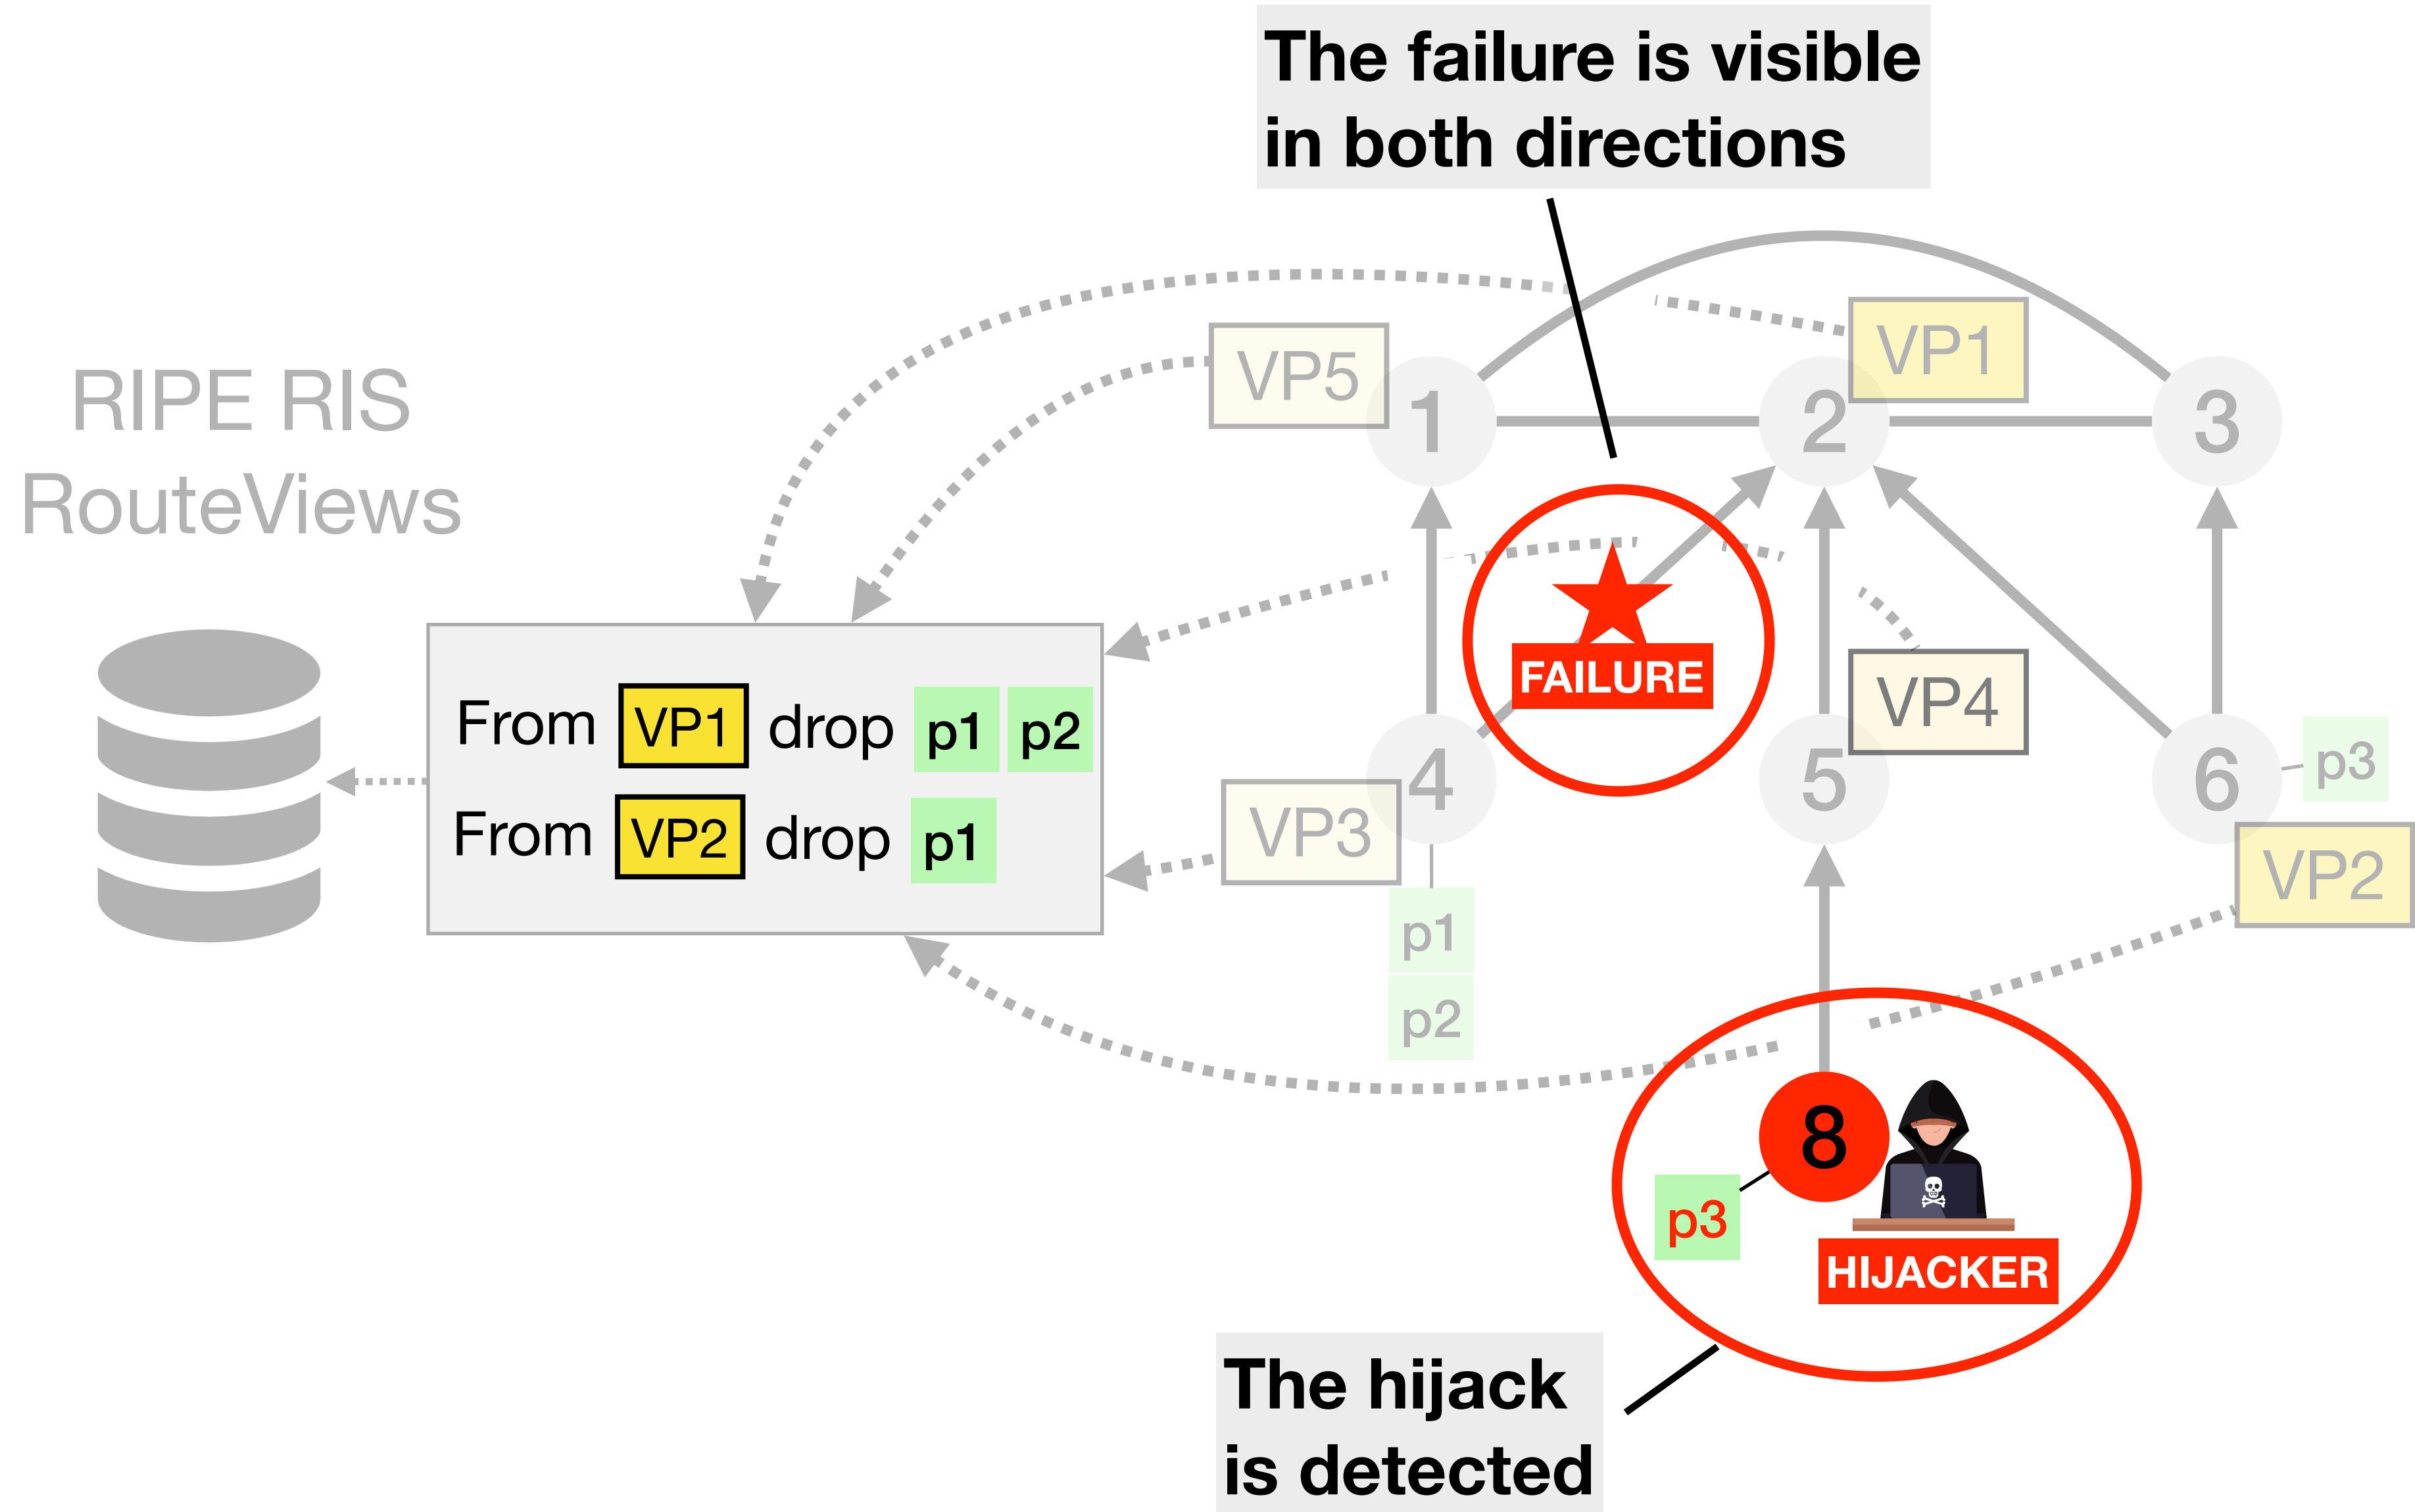
Collected routes

VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	3 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 8

RIPE RIS
RouteViews



From **VP1** drop **p1** **p2**
From **VP2** drop **p1**




The failure is visible
in both directions

The hijack
is detected


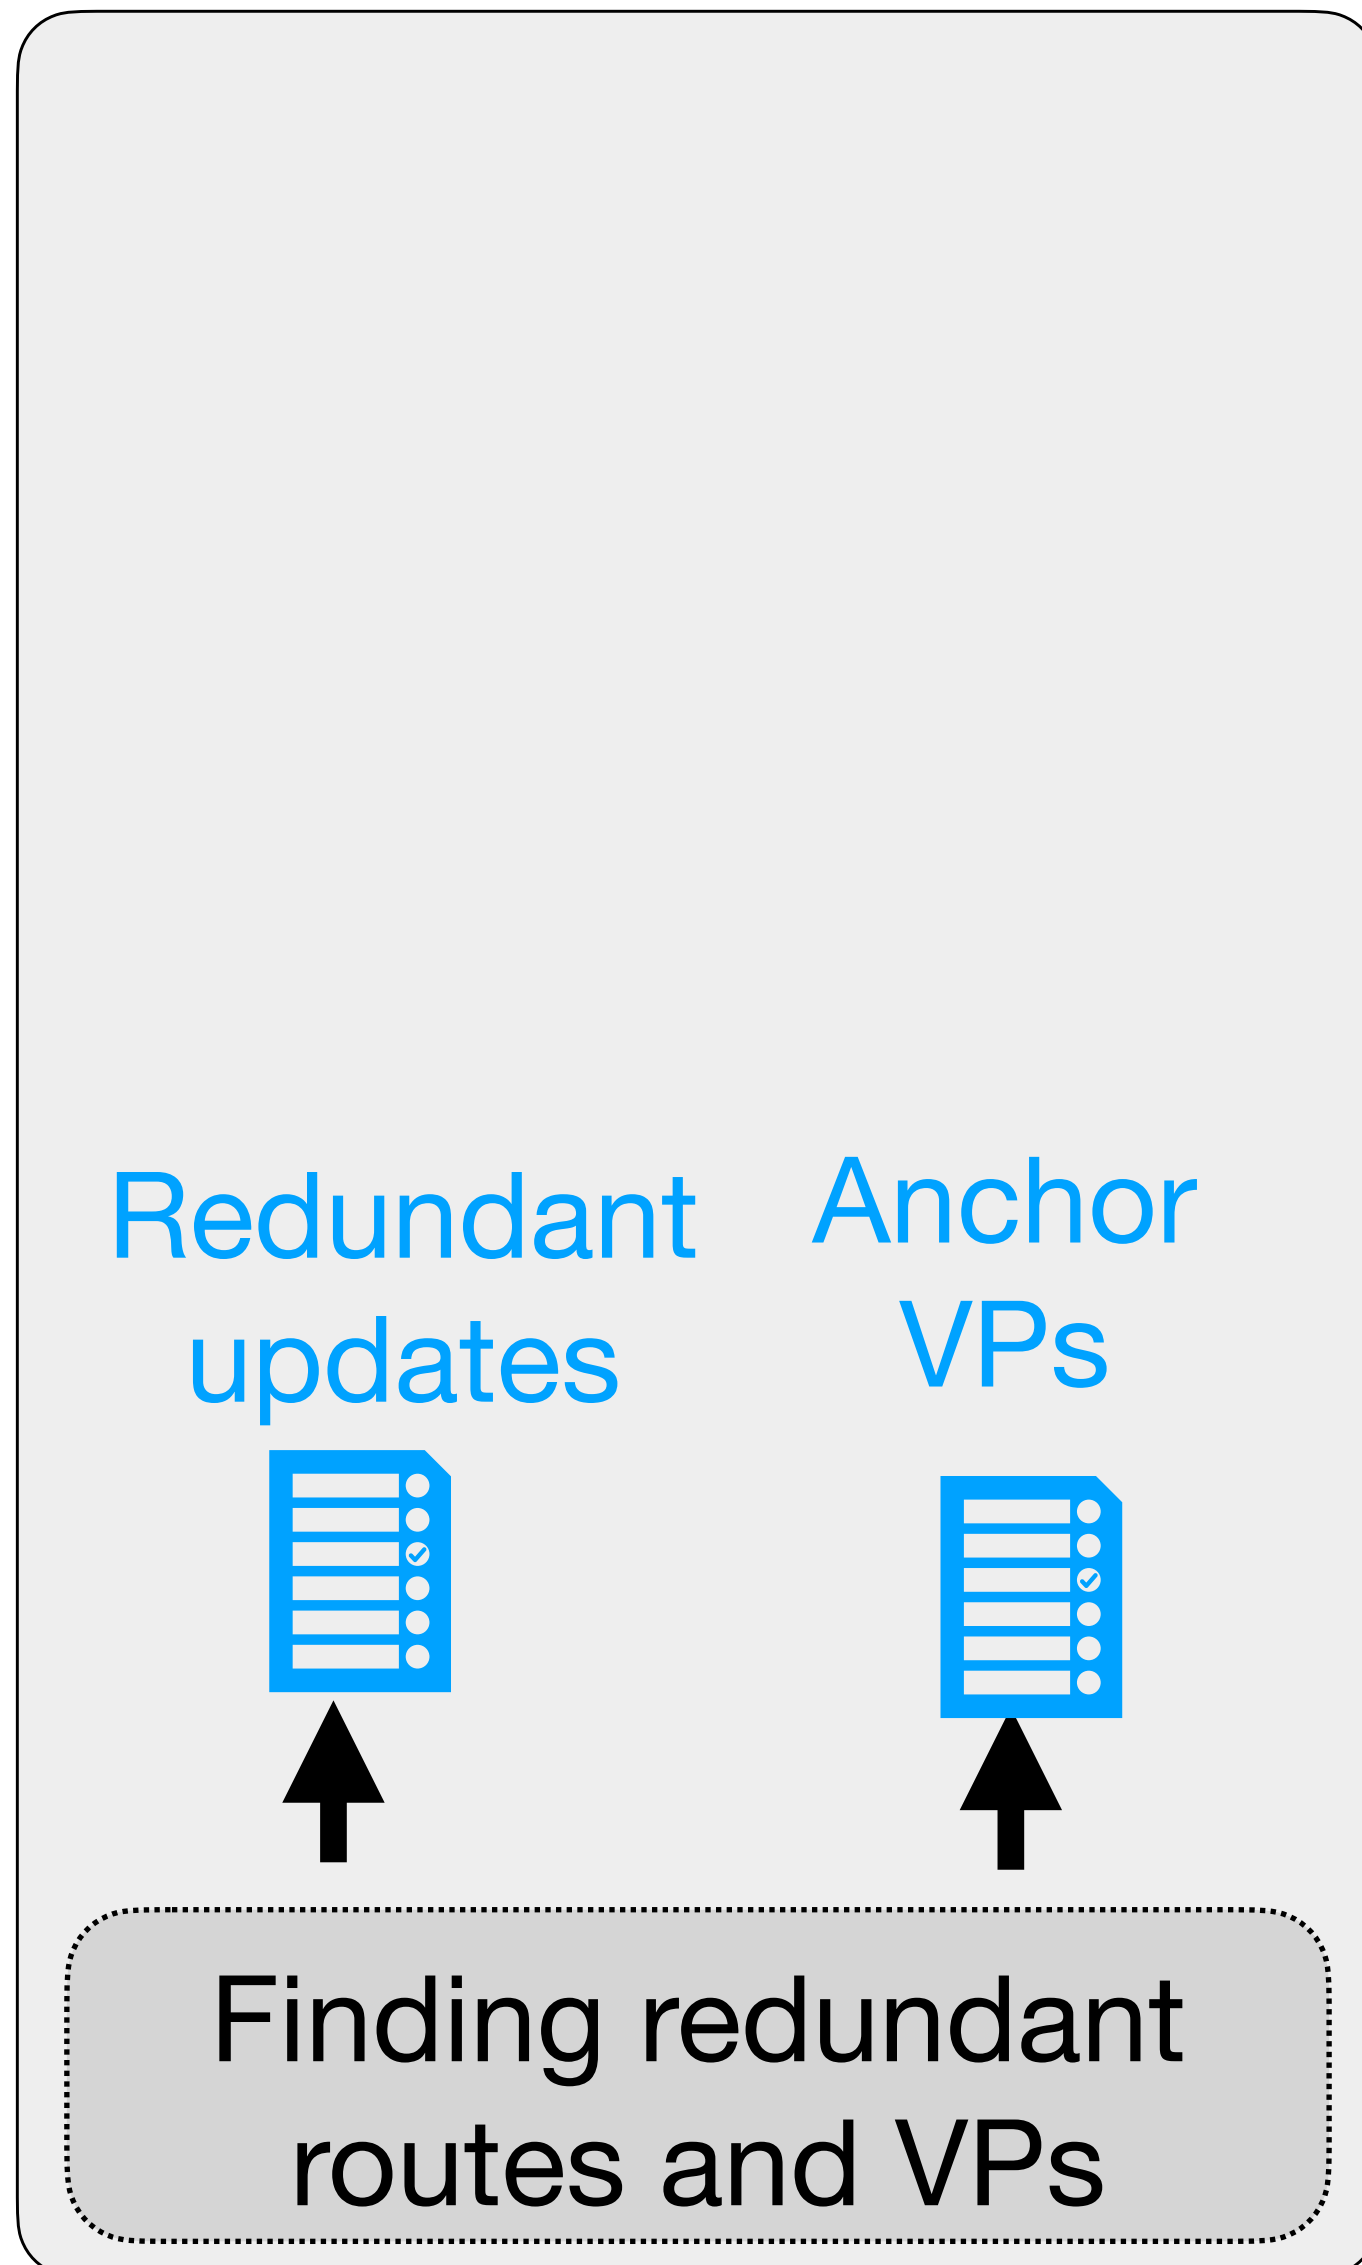
Outline

1. We observe that BGP routes are often redundant
2. Redundant BGP routes enable an overshoot-and-discard collection scheme
3. ***GILL***: A BGP data collection platform that discards redundant routes to scales to tens of thousands of VPs

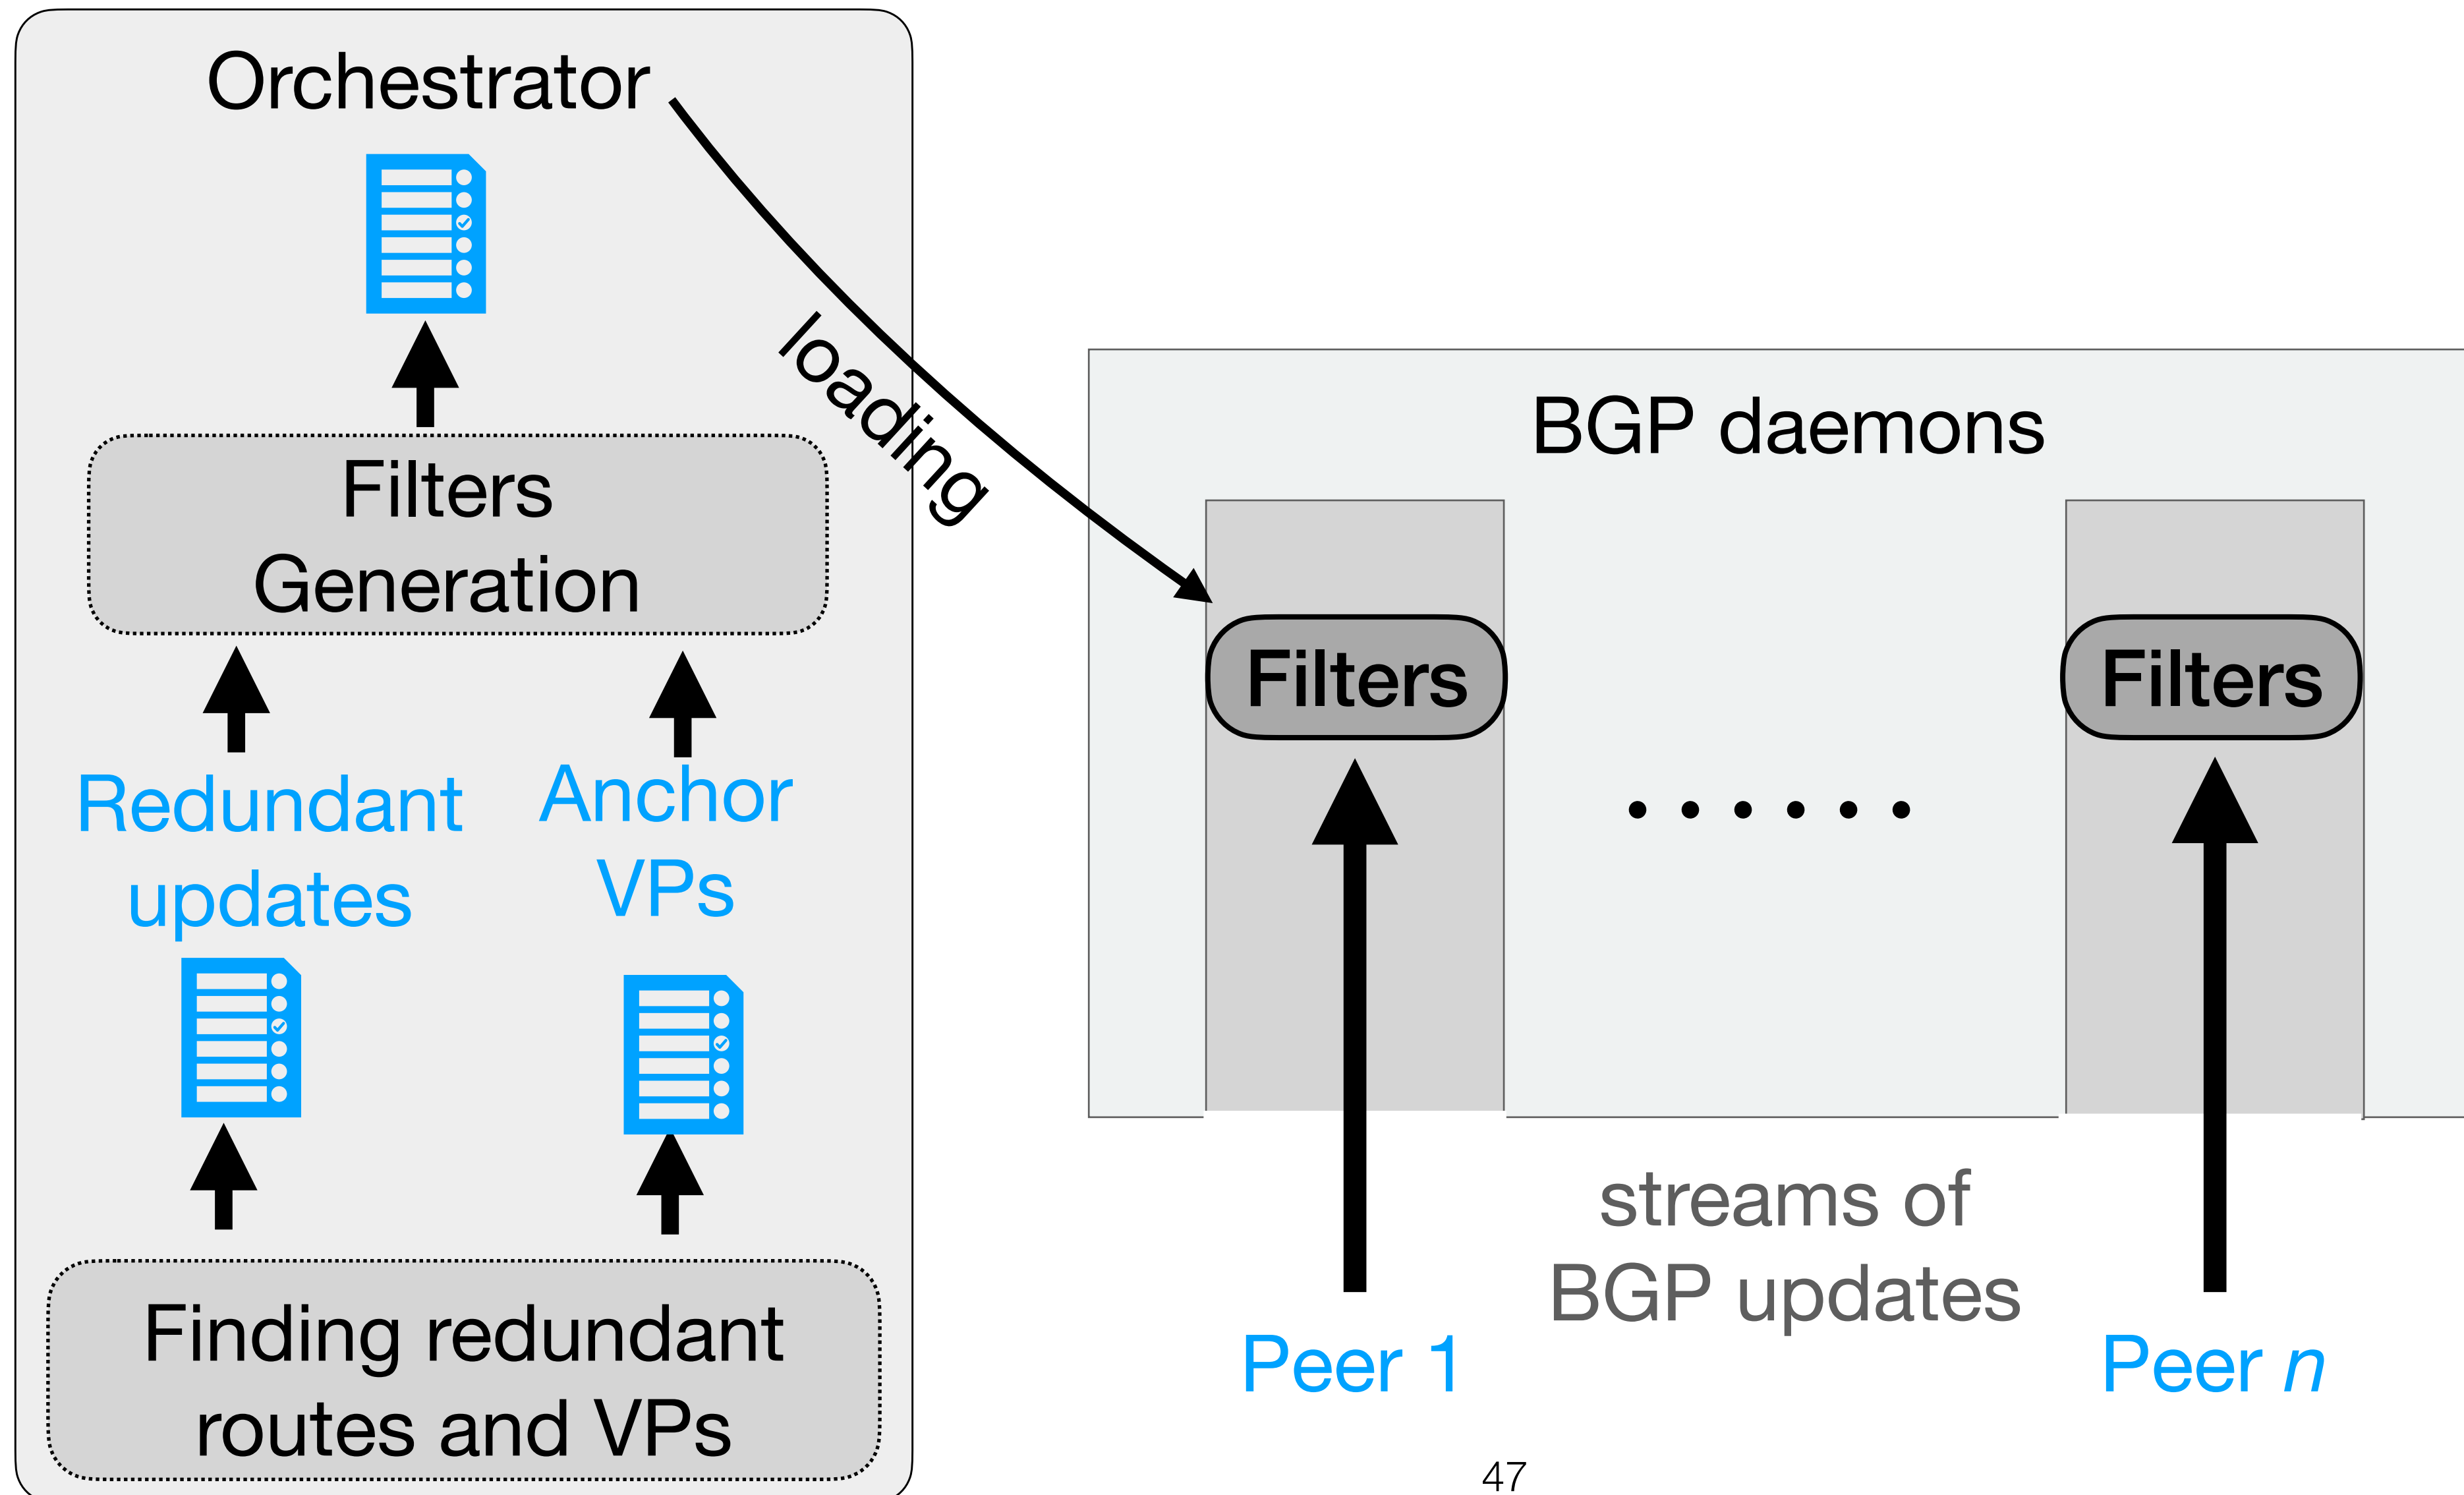
GILL uses BGP daemons written in C and optimised to collect BGP routes



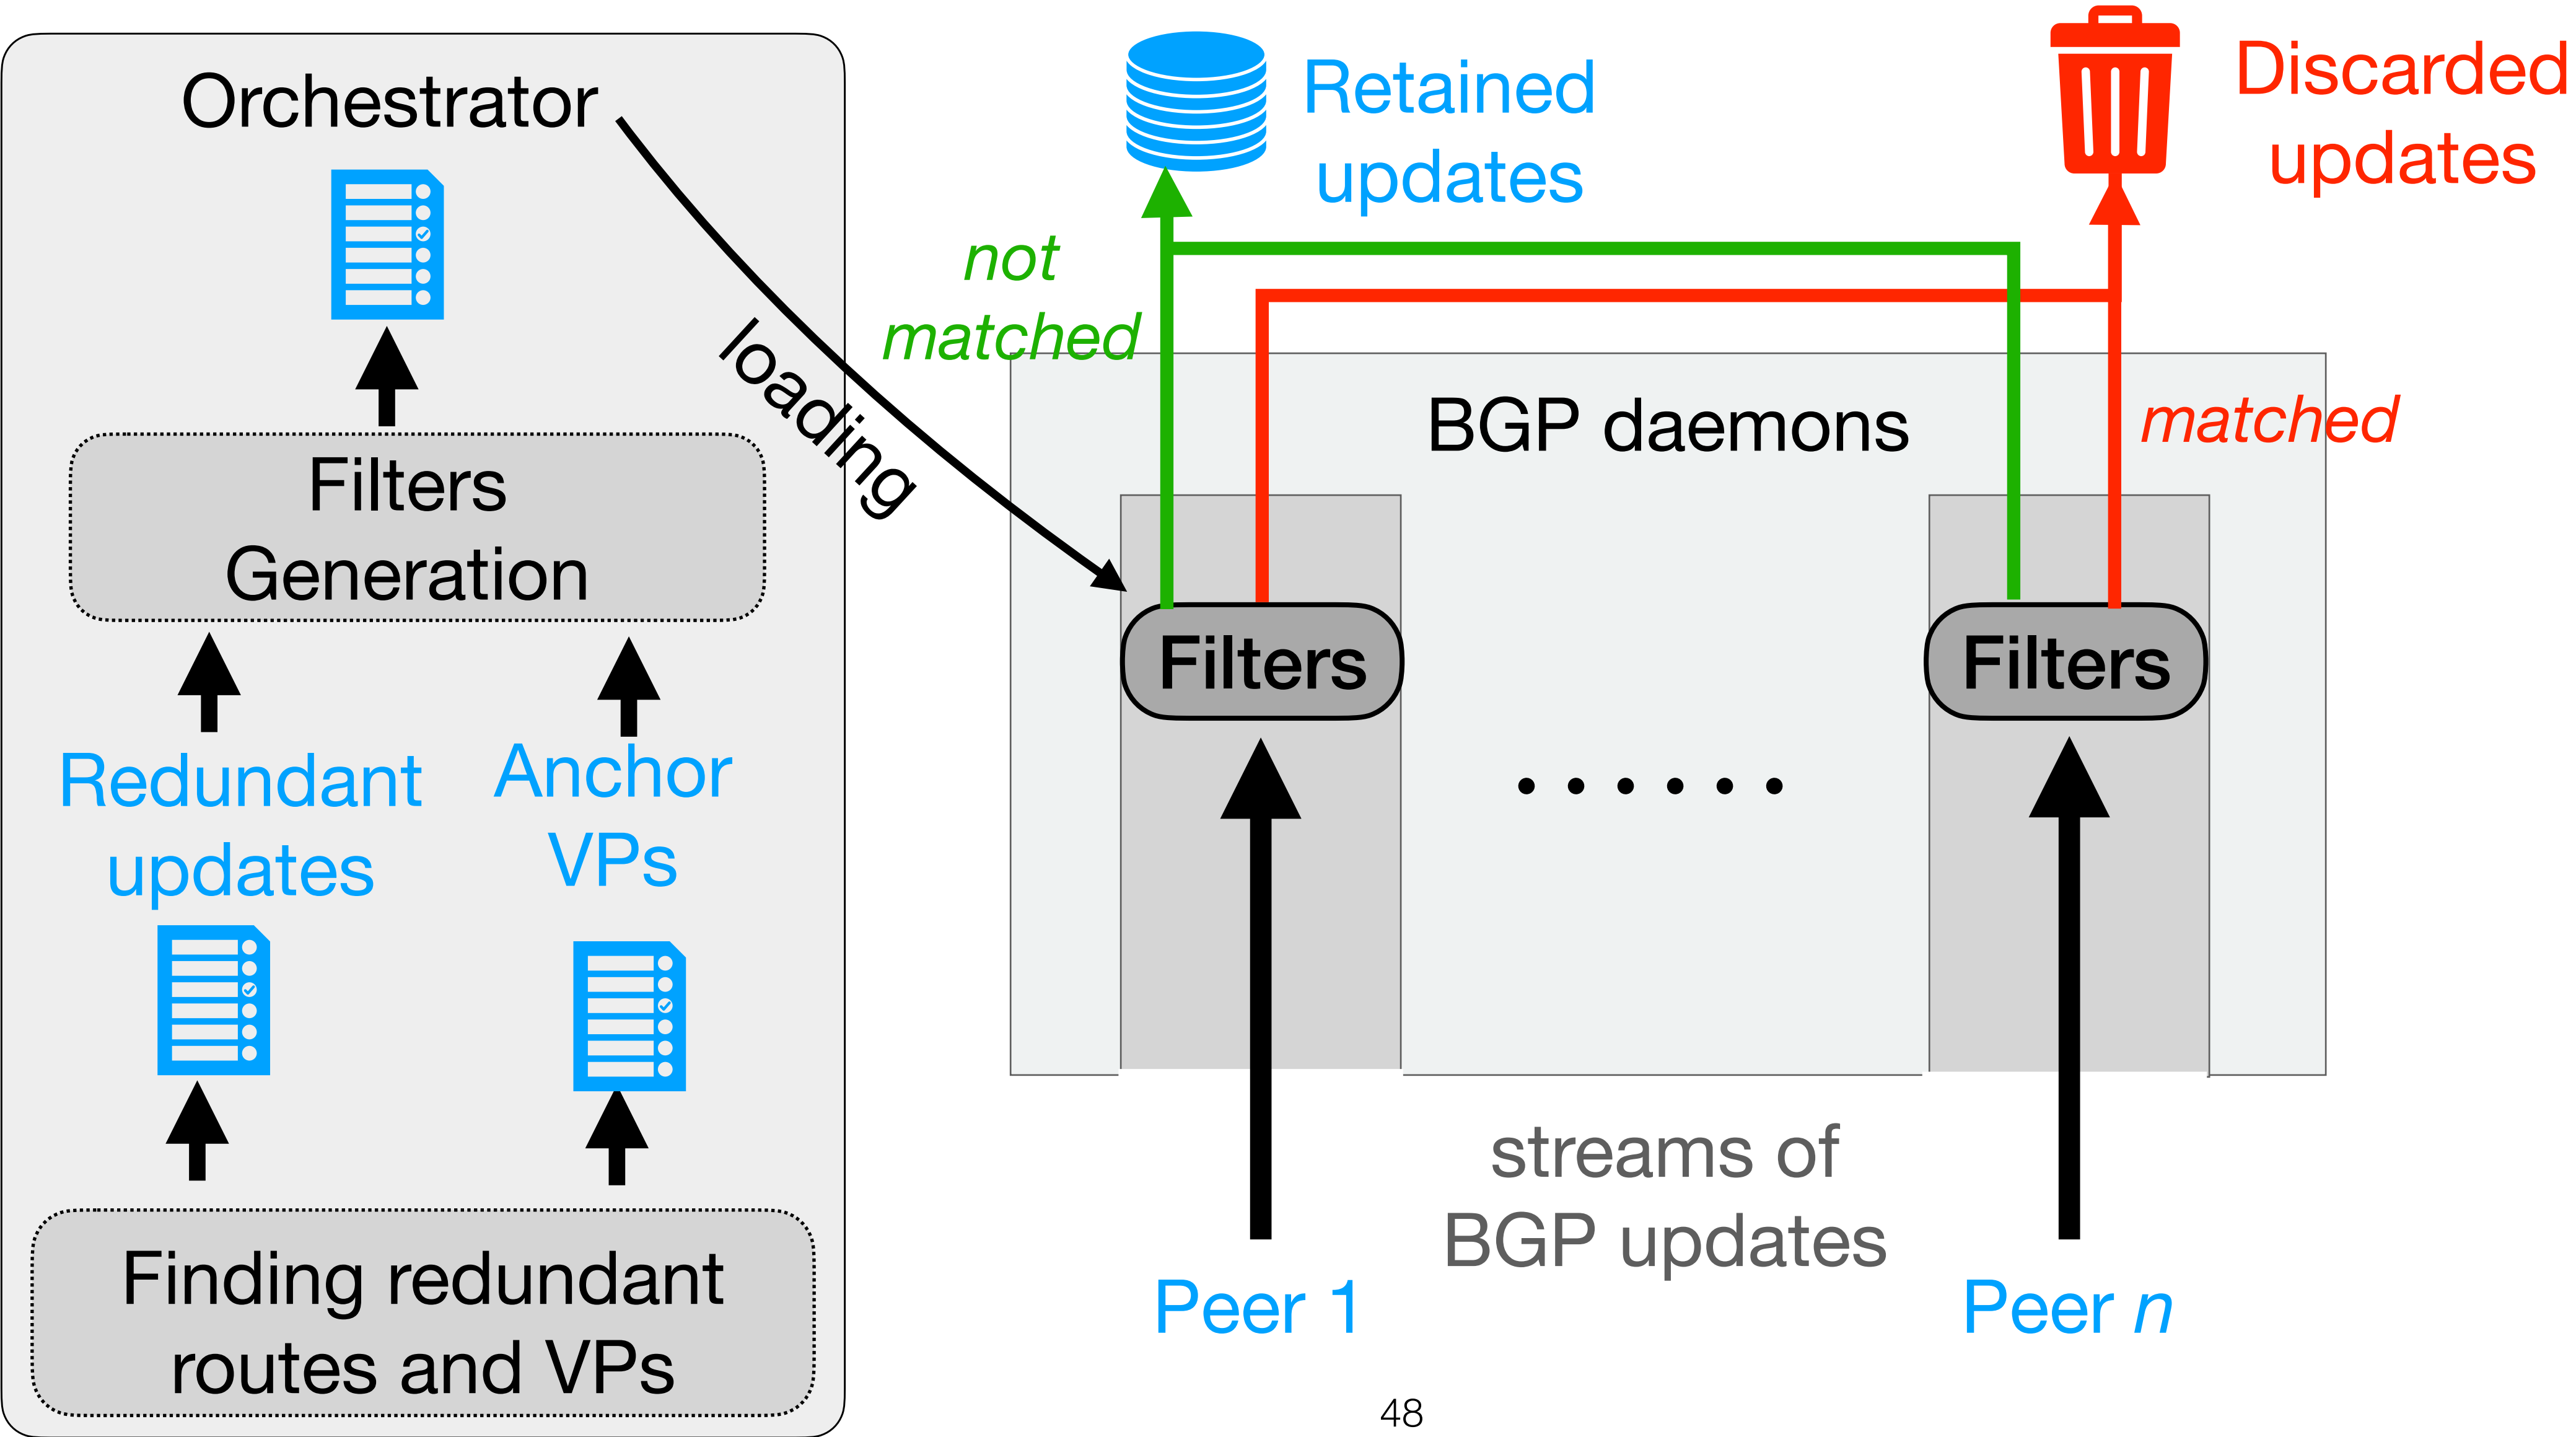
GILL finds redundant updates and anchors VP



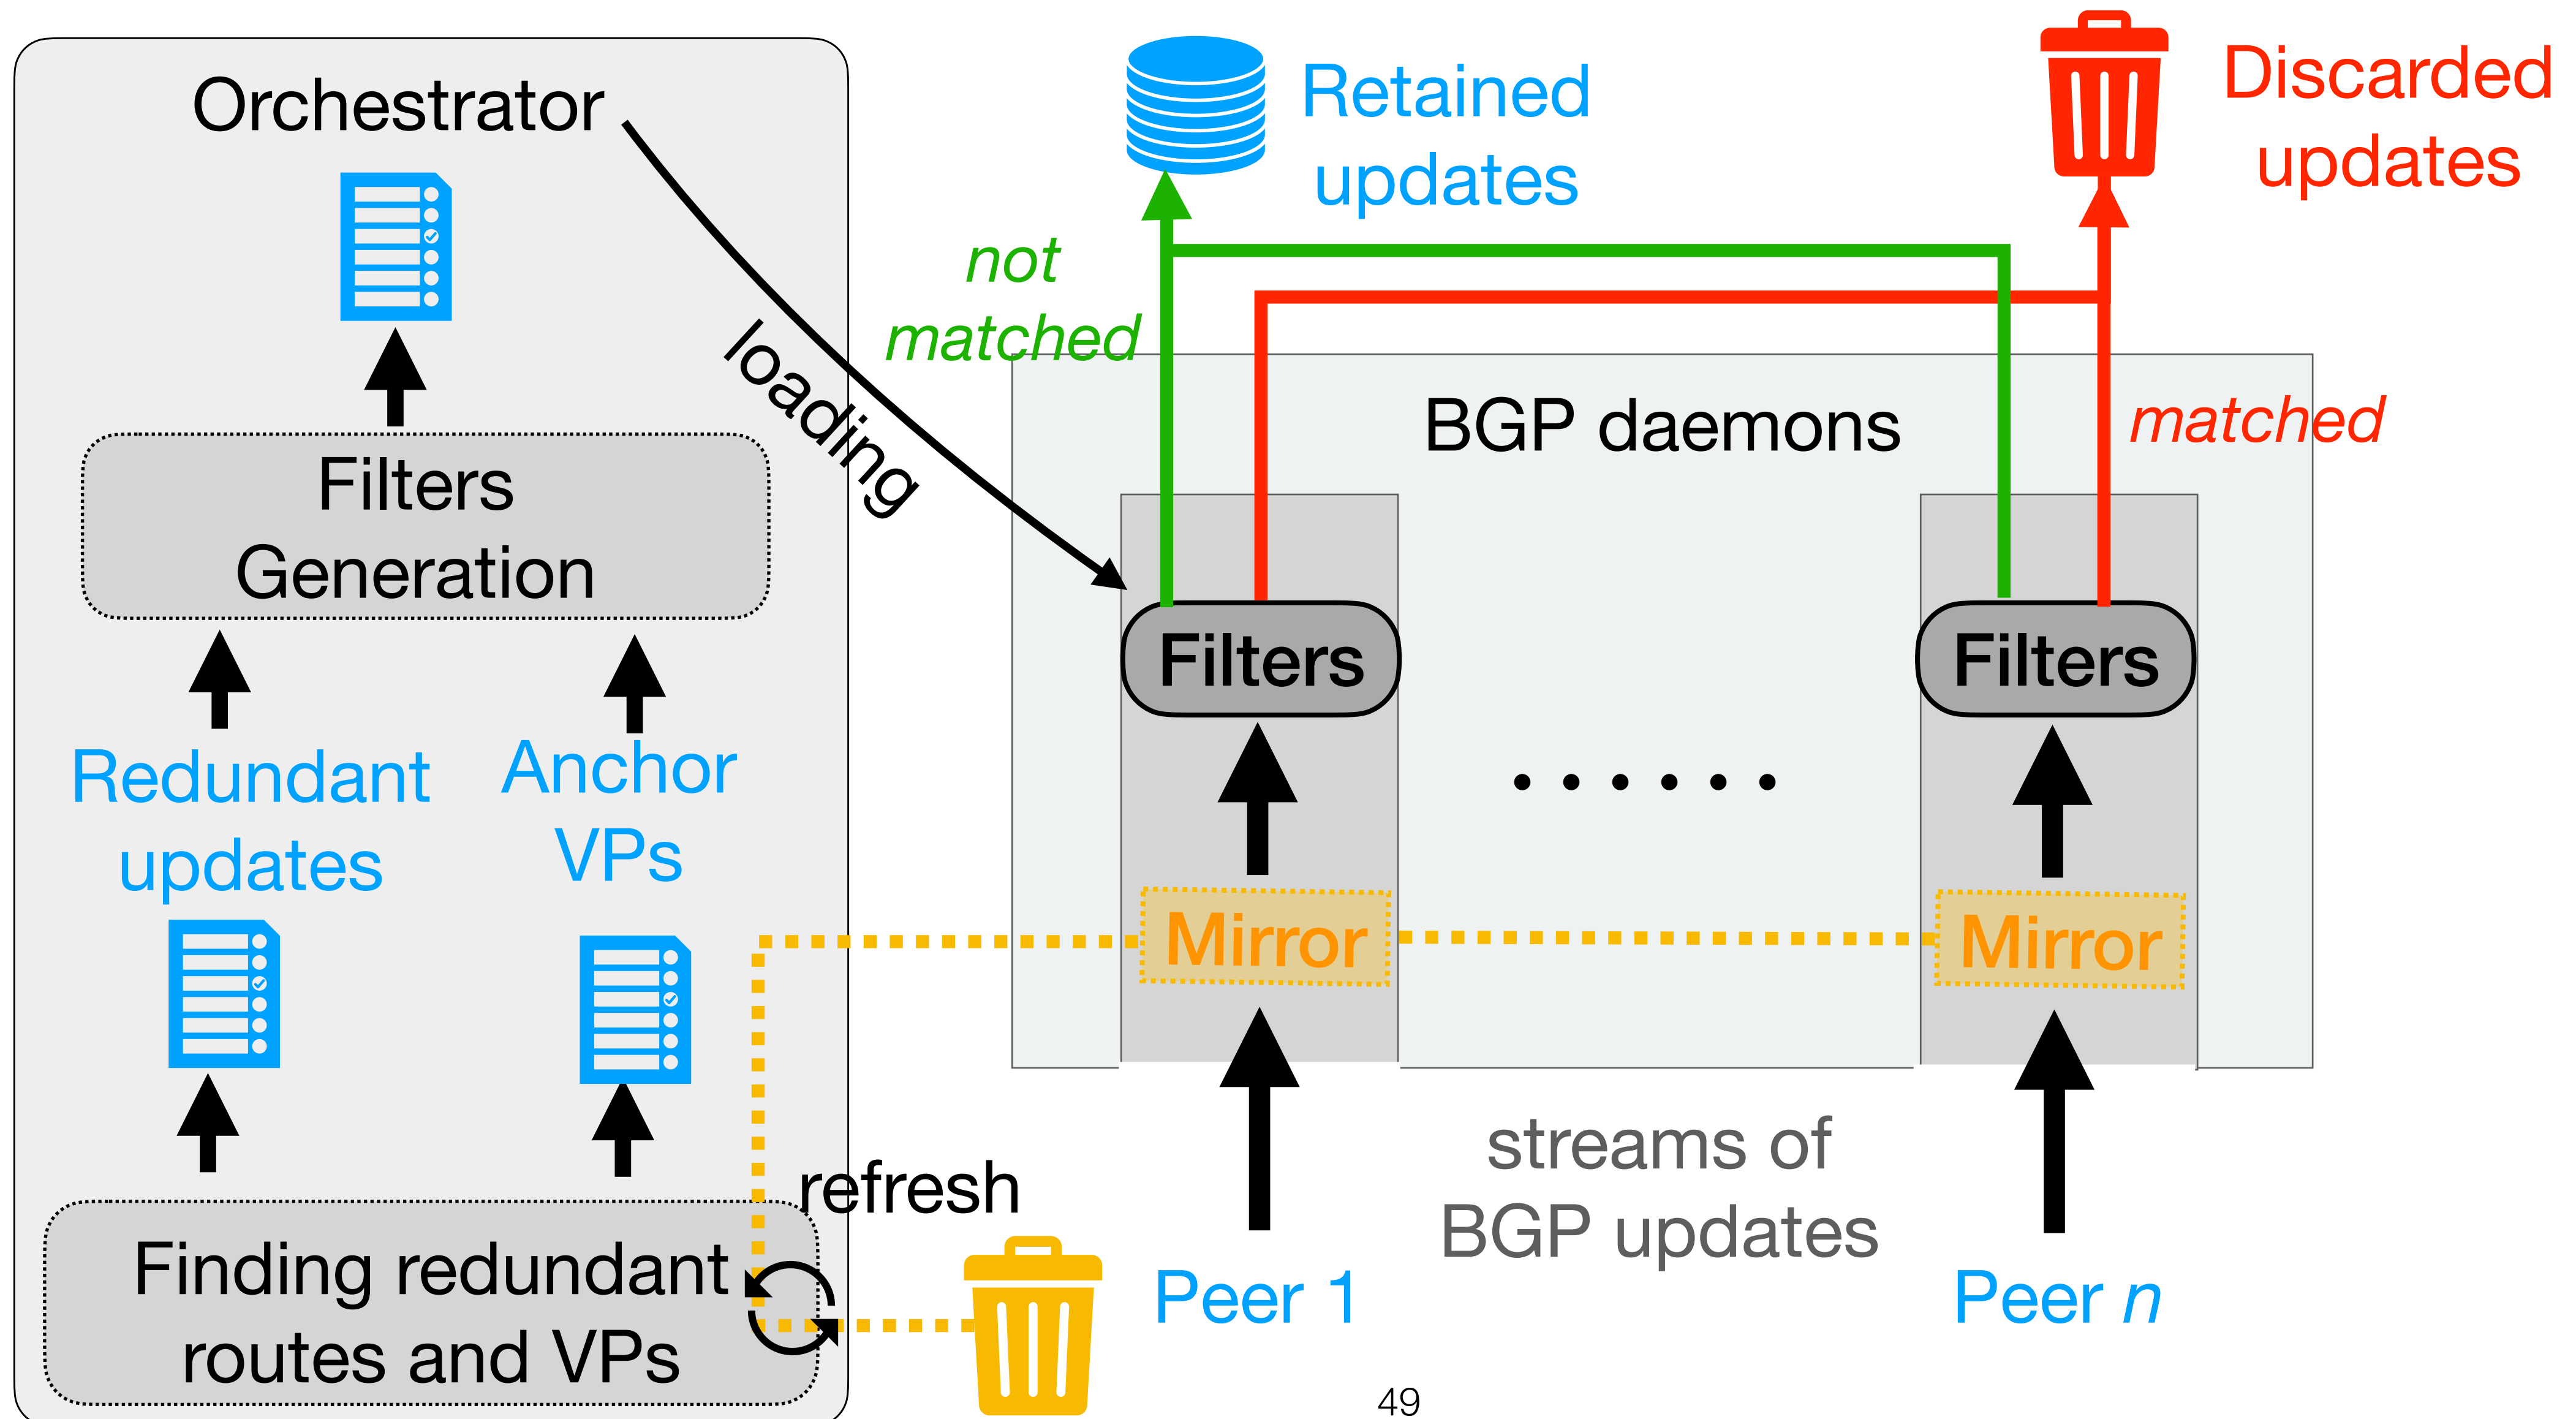
GILL finds redundant updates and anchors VP



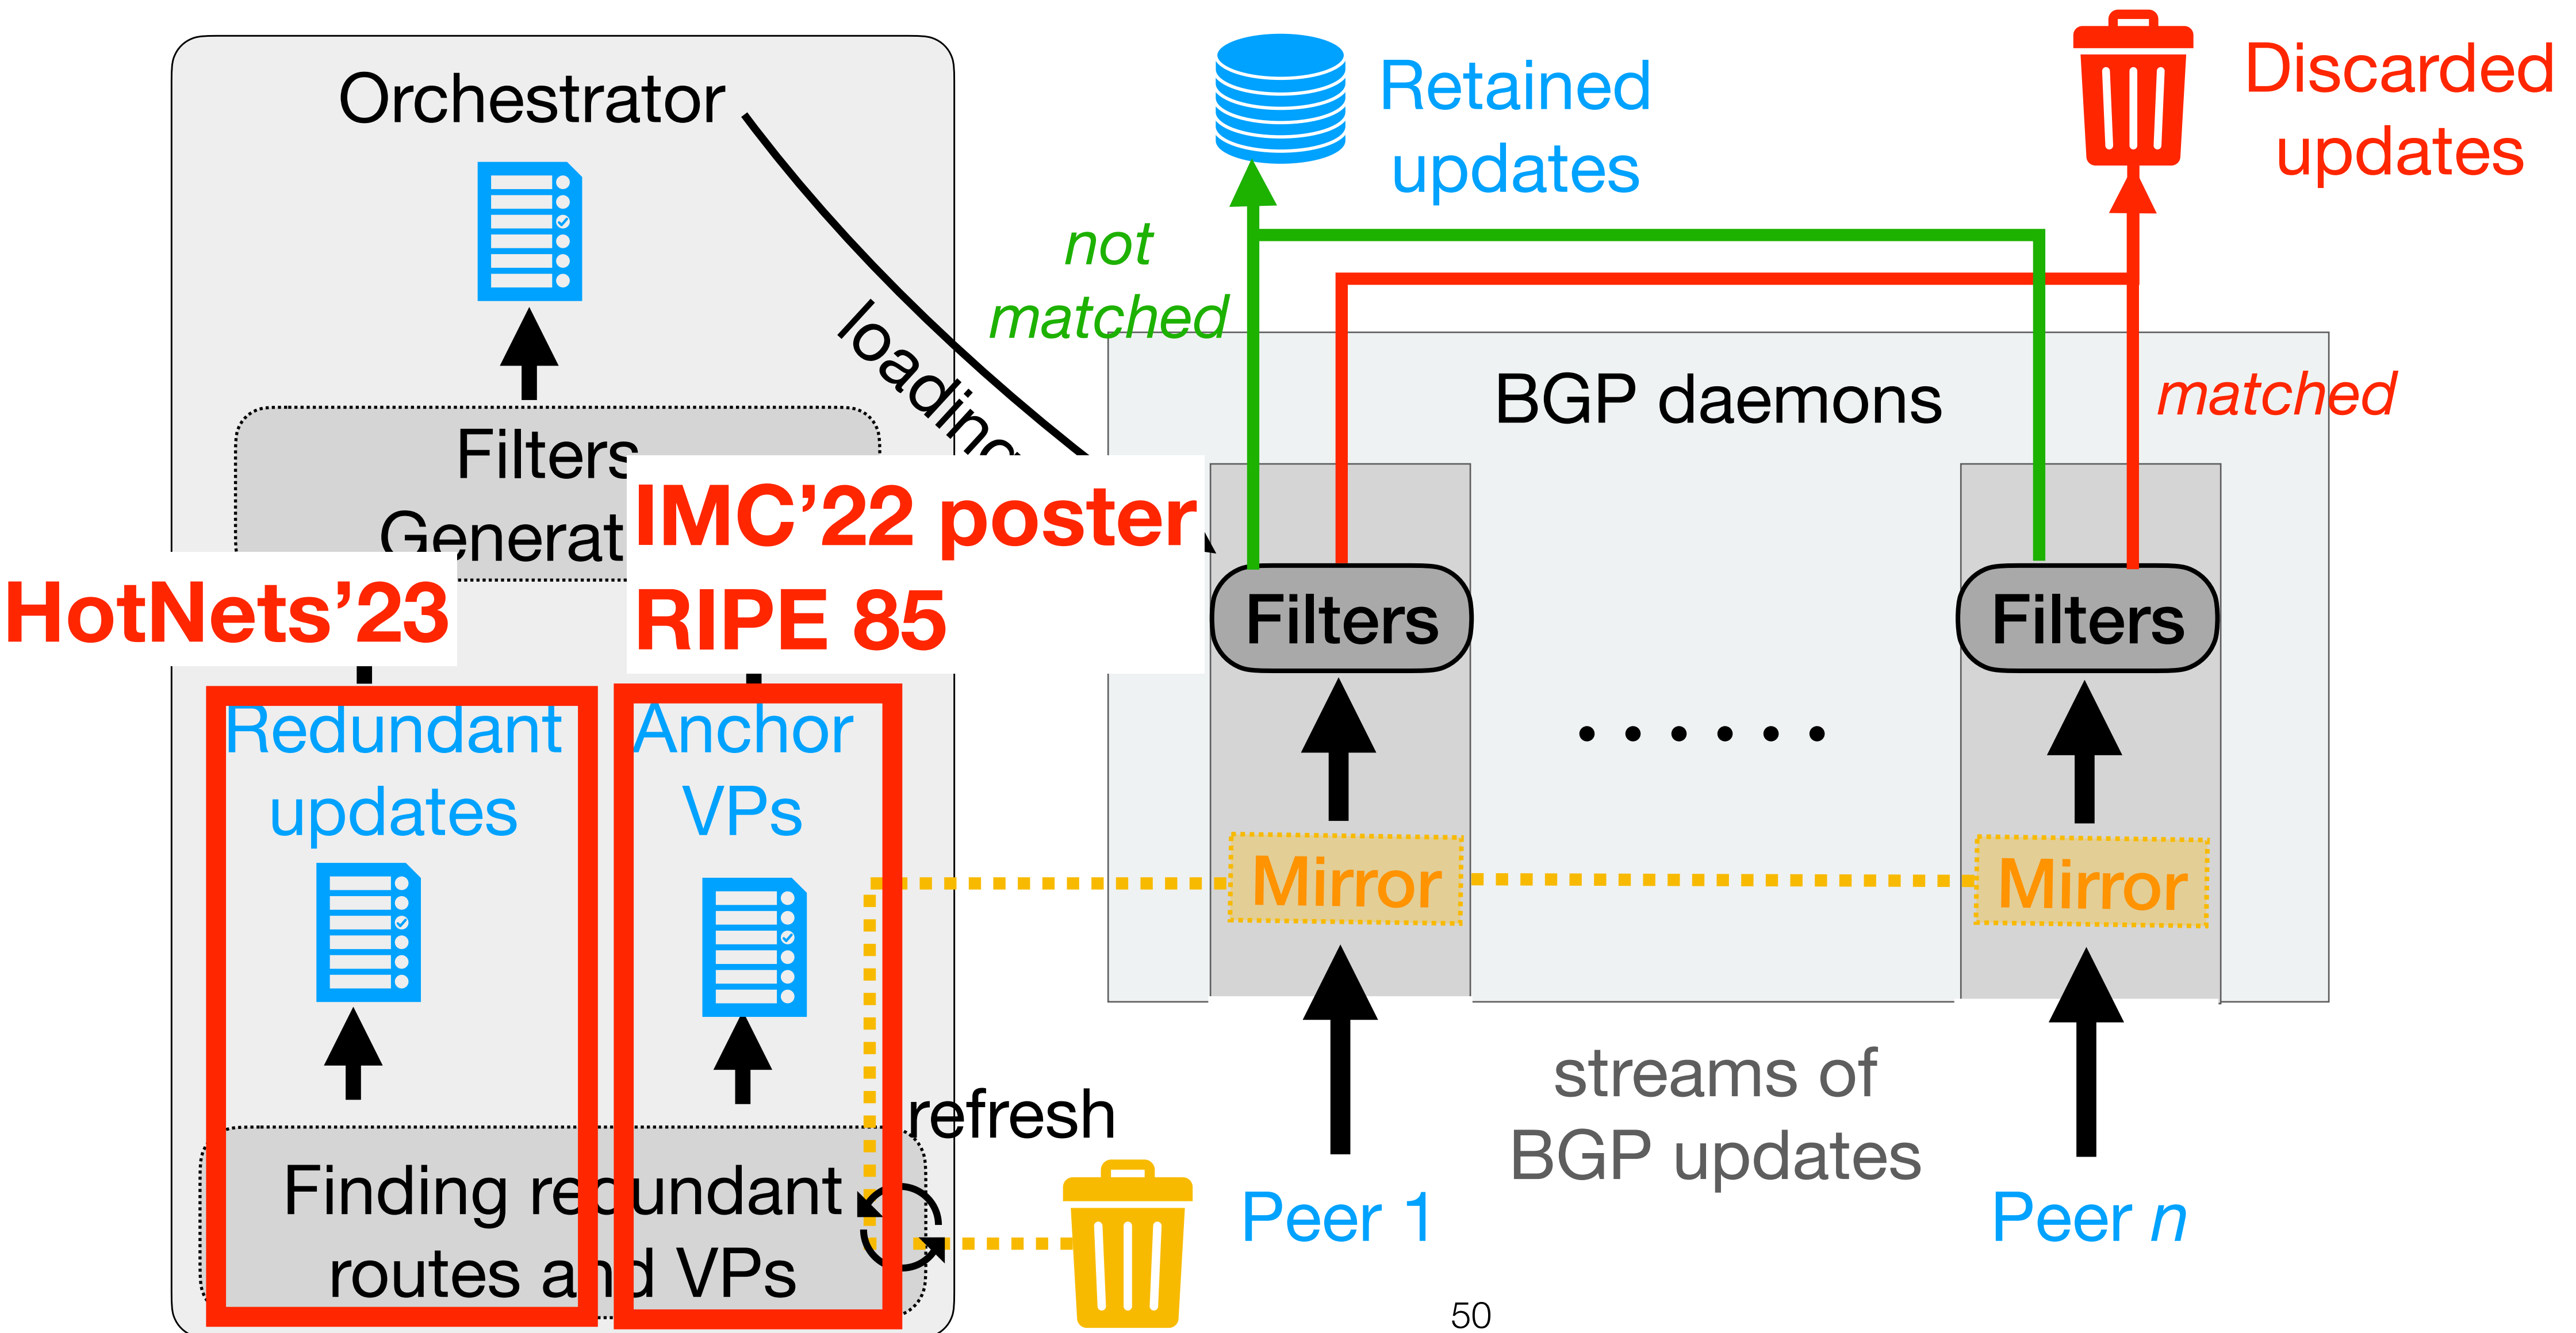
GILL computes filters, loads them into the BGP daemons and discards the filtered routes



GILL updates filters over time using an out-of-band mirroring scheme



GILL updates filters over time using an out-of-band mirroring scheme



Gill finds redundant BGP data
without optimising a particular objective

**Key Intuition: A set of BGP updates is redundant if it can
probabilistically be reconstituted from another set of updates**

Gill finds redundant BGP data
without optimising a particular objective

**Key Intuition: A set of BGP updates is redundant if it can
probabilistically be reconstituted from another set of updates**

See our HotNets'23 paper

Outline

1. We observe that BGP routes are often redundant
2. Redundant BGP routes enable an overshoot-and-discard collection scheme
3. *GILL*: A BGP data collection platform that discards redundant routes to scales to tens of thousands of VPs
4. ***GILL*'s long-term impact is significant for various objectives**

GILL's long-term impact

Platform's settings

	coverage	% of discarded BGP updates	# of stored BGP updates
Current approach	1%	0%	X
<i>GILL</i>			

Results from simulations
on “mini” Internets with 1k ASes

GILL's long-term impact

Platform's settings

Use cases

	coverage	% of discarded BGP updates	# of stored BGP updates	Topology mapping (p2p links)	Failure localisation (p2p links)	Hijacks detected (Type-1)
Current approach	1%	0%	X	20%	37%	73%
<i>GILL</i>						

Results from simulations
on “mini” Internets with 1k ASes

GILL's long-term impact

Platform's settings

Use cases

	coverage	% of discarded BGP updates	# of stored BGP updates	Topology mapping (p2p links)	Failure localisation (p2p links)	Hijacks detected (Type-1)
Current approach	1%	0%	X ↑ Same order of magnitude	20%	37%	73%
<i>GILL</i>	50%	96%	X ↓			

Results from simulations on “mini” Internets with 1k ASes

GILL's long-term impact

Platform's settings

Use cases

	coverage	% of discarded BGP updates	# of stored BGP updates	Topology mapping (p2p links)	Failure localisation (p2p links)	Hijacks detected (Type-1)
Current approach	1%	0%	X ↑ Same order of magnitude	20%	37%	73%
GILL	50%	96%	X ↓	61%	80%	82%

Results from simulations on “mini” Internets with 1k ASes

A prototype of *GILL* is already up and running!

<https://bgproutes.quest>



Expanding BGP Data Horizons


BGP routes collected from operational routers are extremely valuable to monitor and study Internet routing. However, BGP data collection platforms as currently architected face fundamental challenges that threaten their long-term sustainability: their data comes with enormous redundancy and yet dangerous visibility gaps.

GILL is a new BGP routes collection platform that can collect routes from at least an order of magnitude more routers compared to existing platforms while limiting the increase in human effort and data volume.

GILL's key principle is an *overshoot-and-discard* collection scheme: Any AS can easily peer with GILL and export their routes. However, GILL only stores and makes available to users the nonredundant routes.

Coverage matters but is challenging

RIPE RIS and RouteViews, the two main BGP routes collection platforms, peer with routers from an increasing number of ASes (1500 in 2023).



Uncollected BGP routes

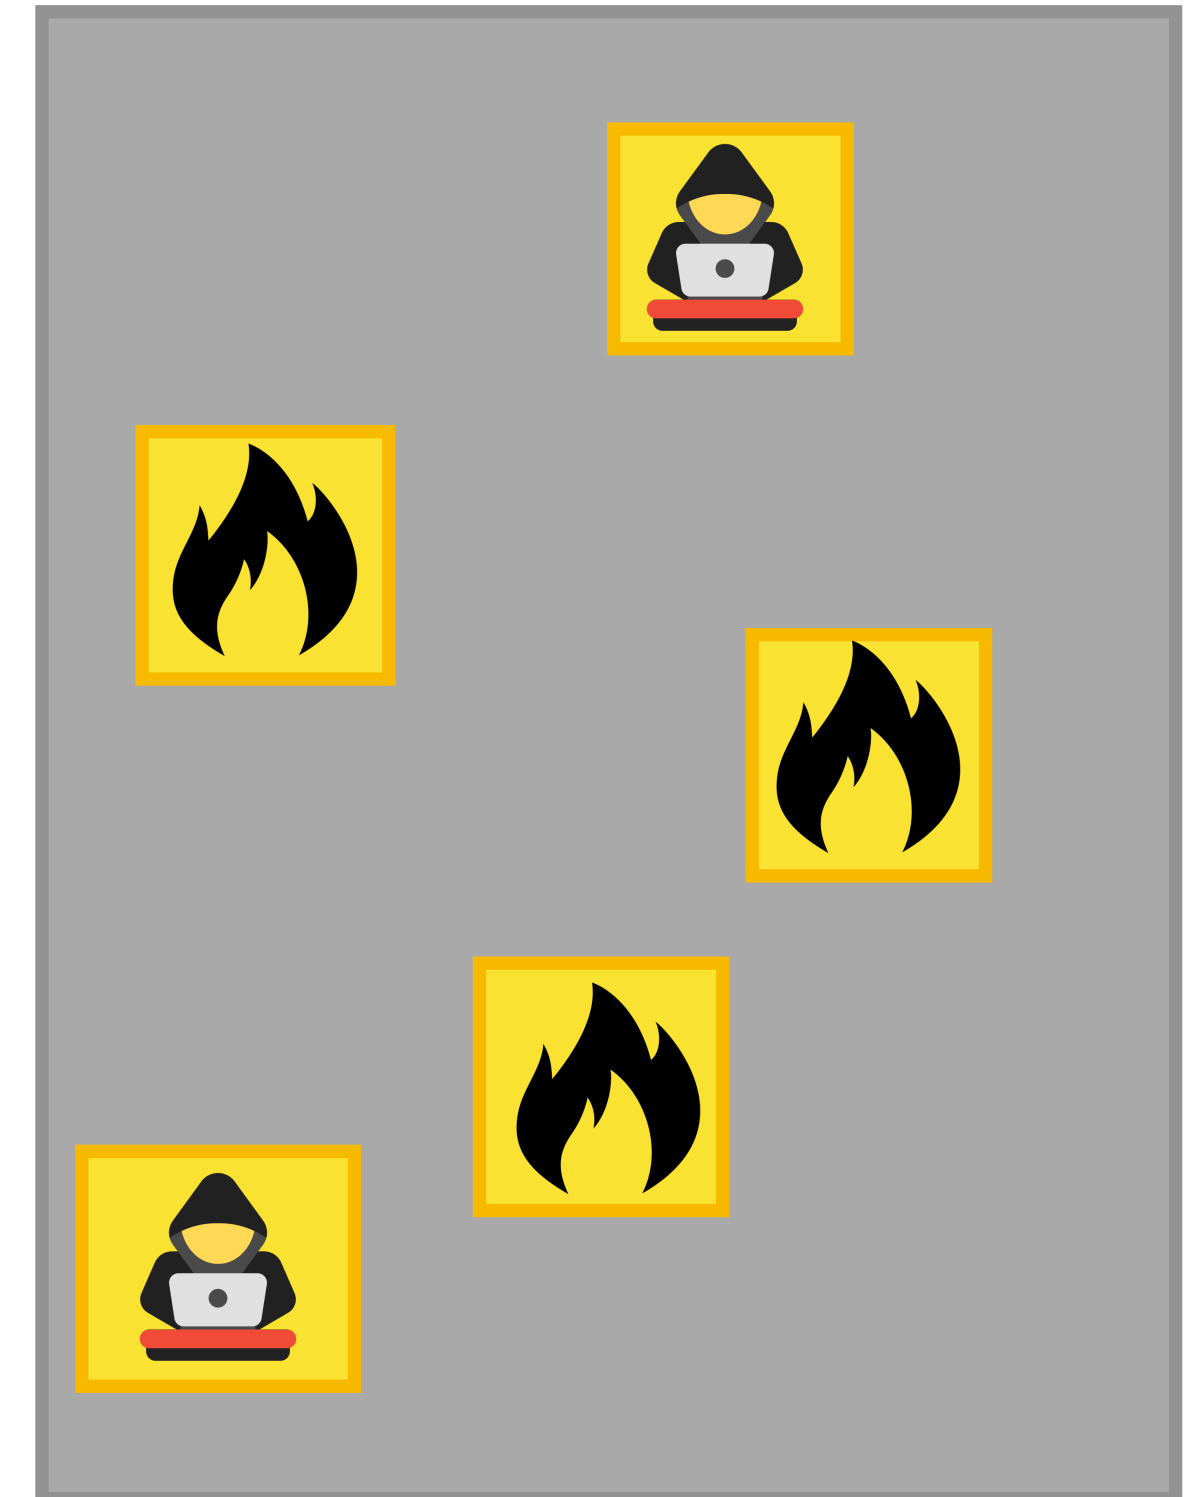


Collected BGP routes

Collected but discarded BGP routes

Today

Naive approach

GILL



Useful bits of data

 Useful bits are missed

 Data management problems

