

# The Next Generation of BGP Data Collection Platforms

**Thomas Alfroy**

University of Strasbourg

AIMS GMI Workshop

25 June 2024


*Joint work with:*

Thomas Holterbach

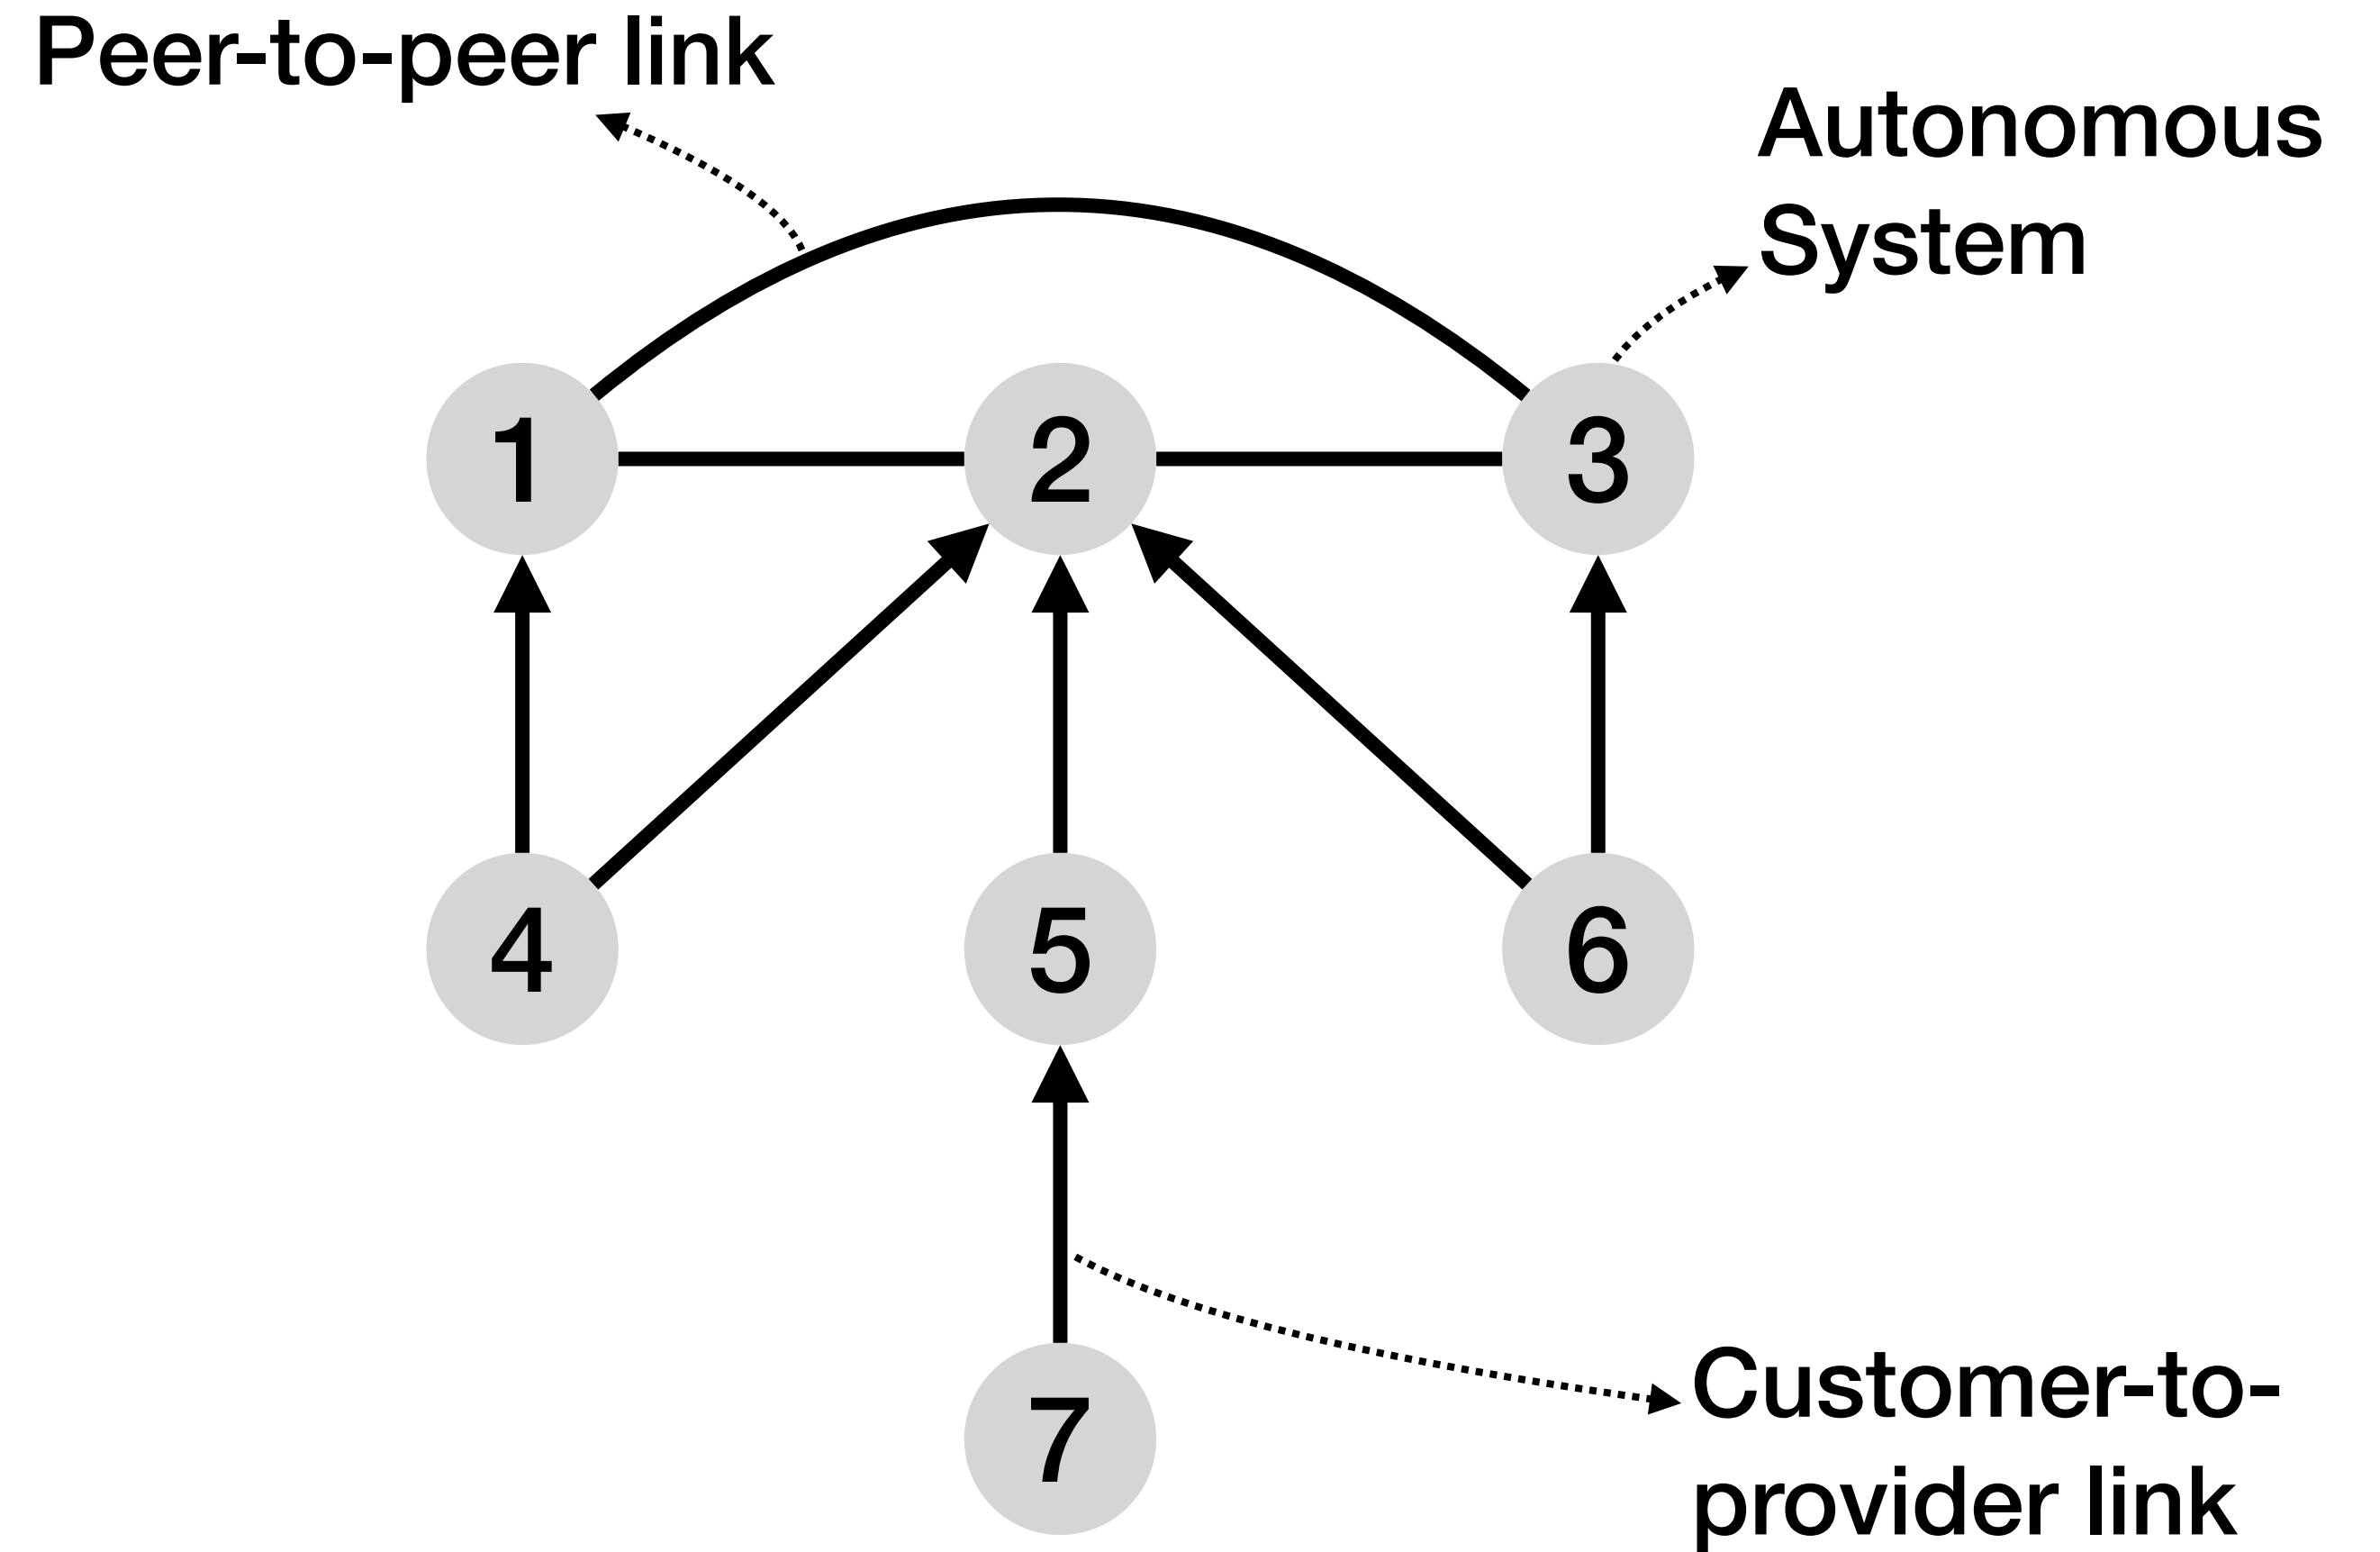
Thomas Krenc

KC Claffy

Cristel Pelsser

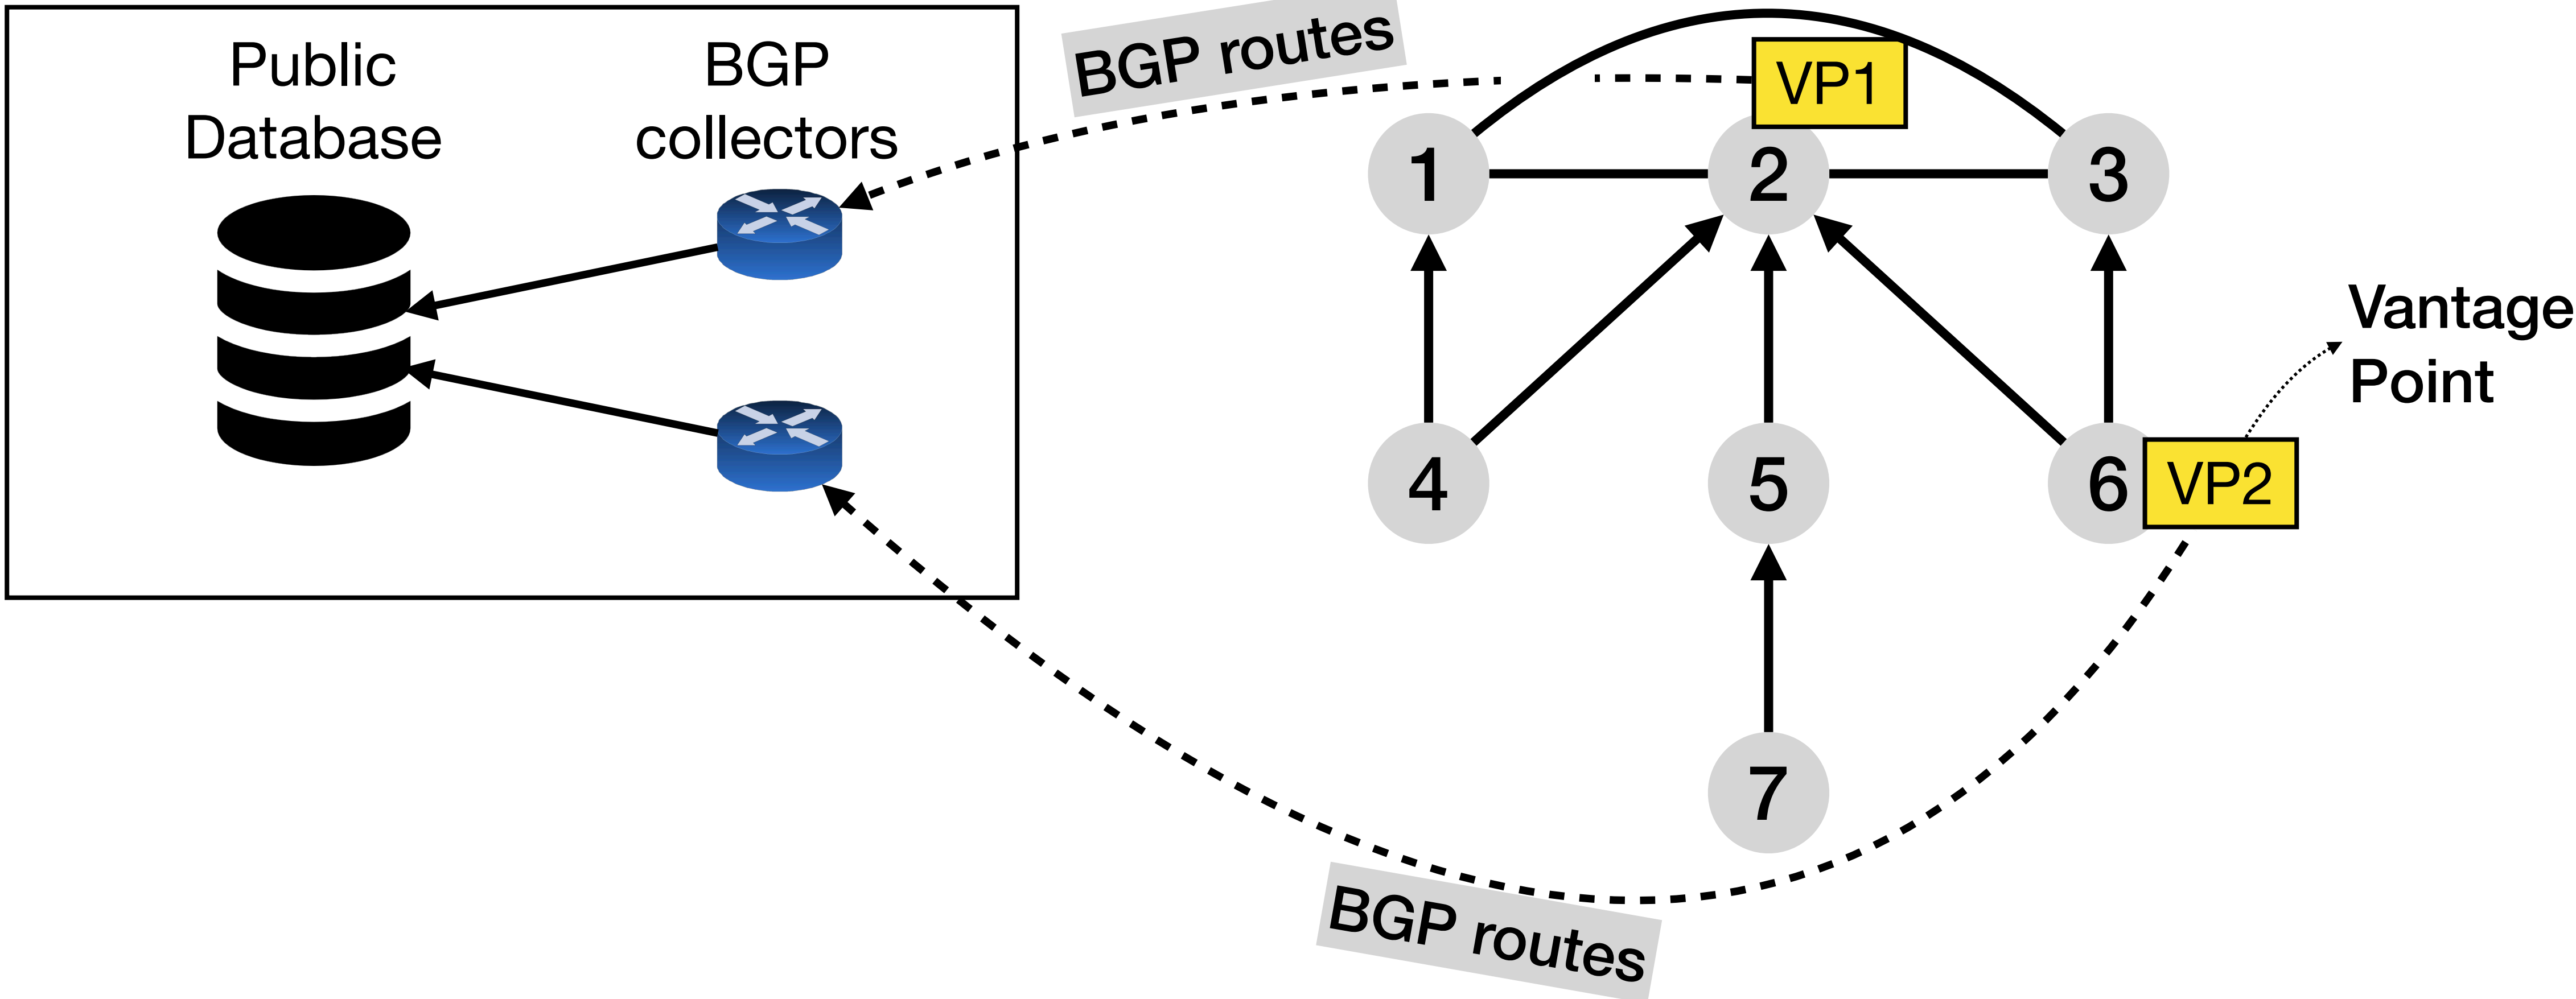


# Public BGP routes are collected by RIPE RIS and RouteViews



# 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)**

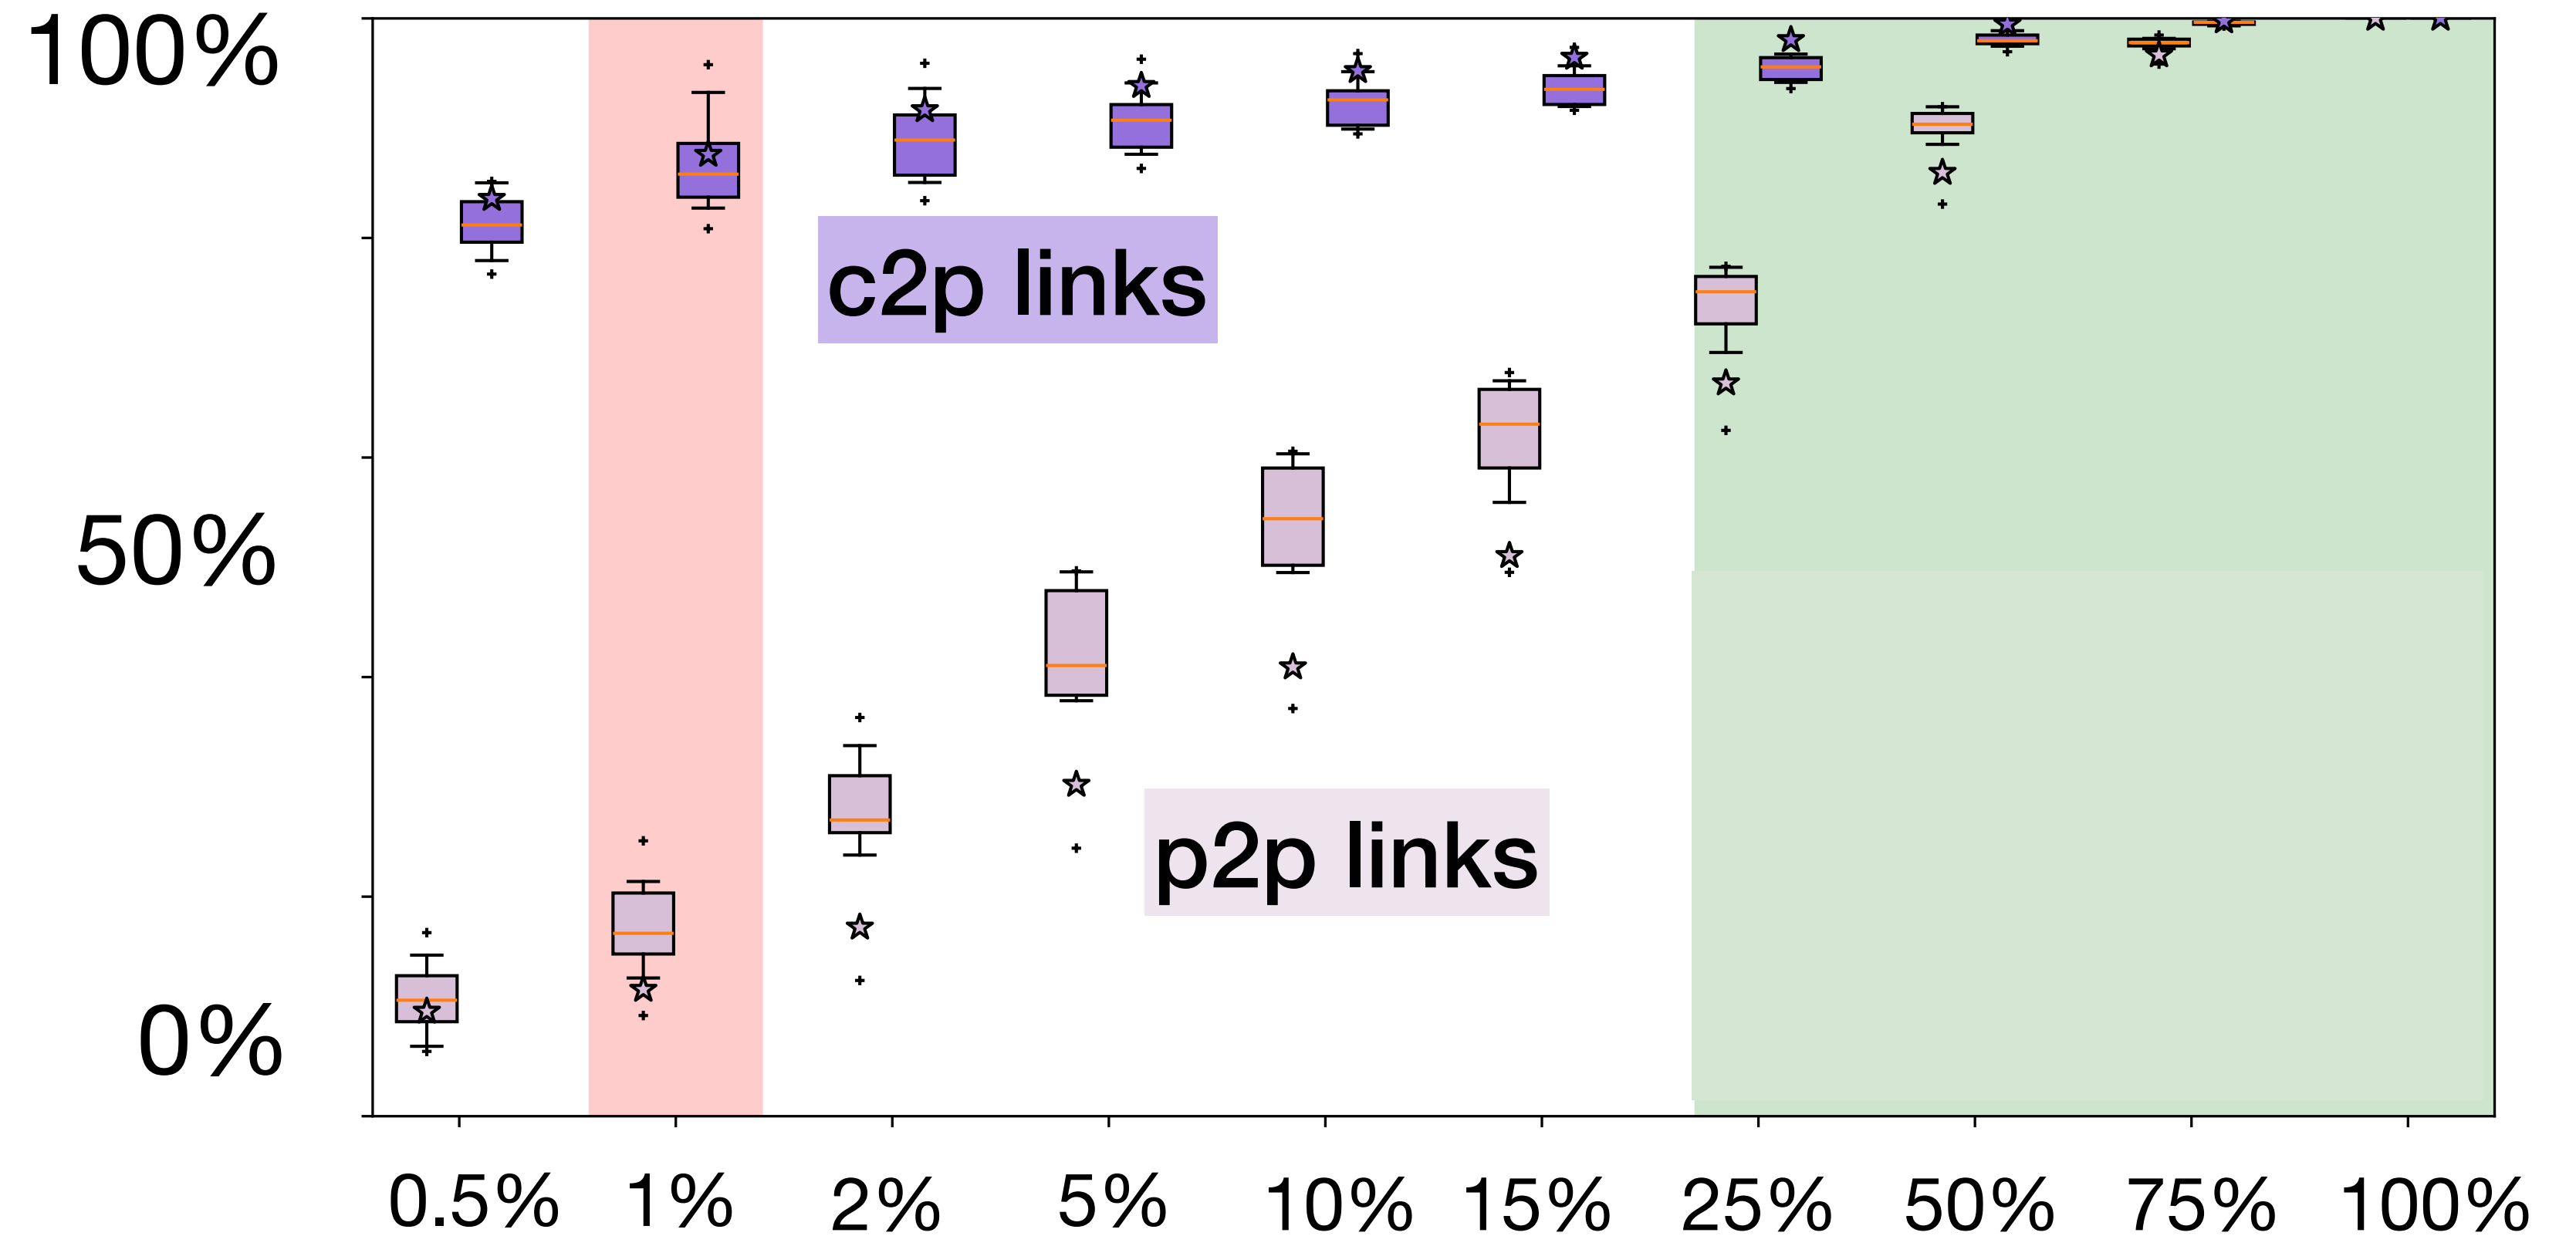
Results from 10 simulations  
on "mini" Internets with 6k ASes

# 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 6k 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

Percentage of forged-origin hijacks detected



Results from 10 simulations on “mini” Internets with 6k ASes

Percentage of ASes hosting a VP (custom scale)

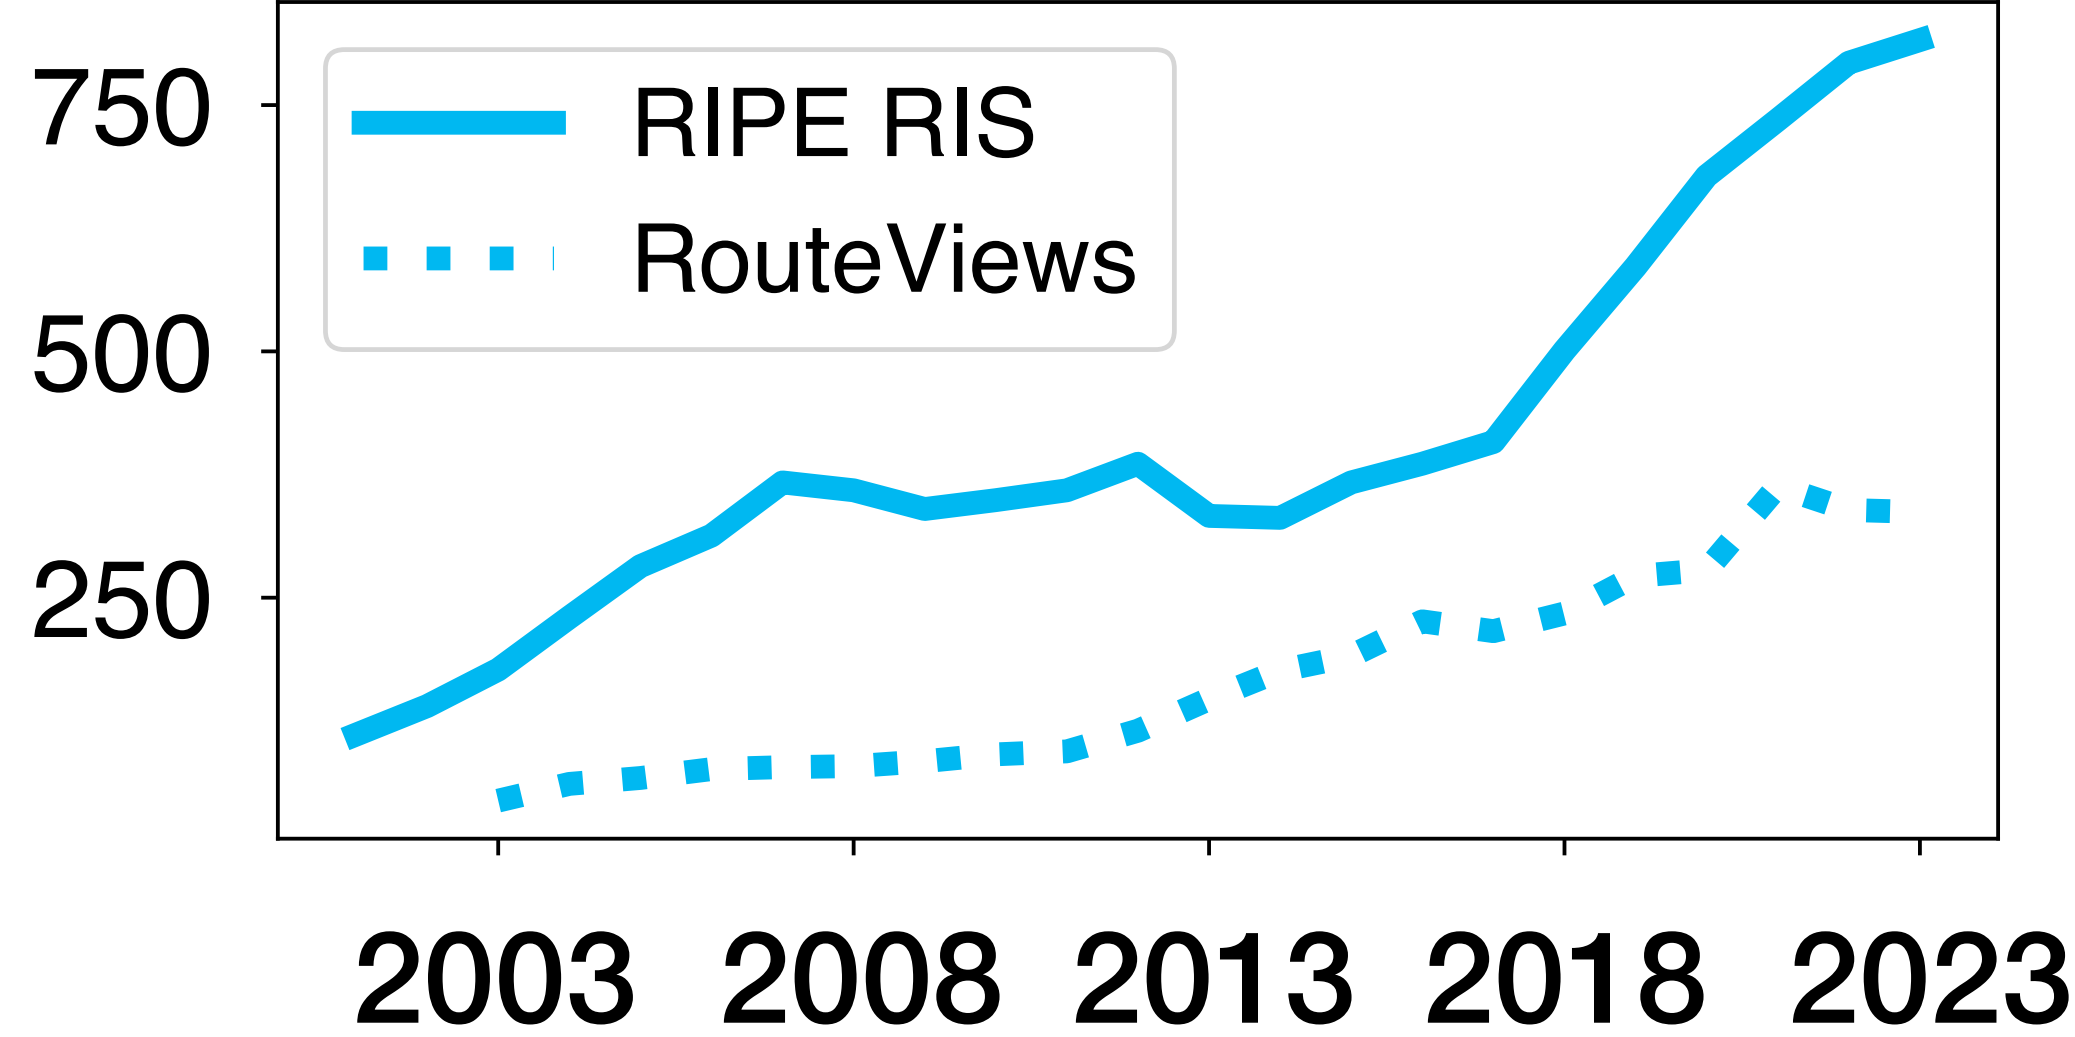
**Three observations motivate reevaluating  
how we collect BGP routes**

**Observation #1: RIPE RIS and RouteViews lack coverage**

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

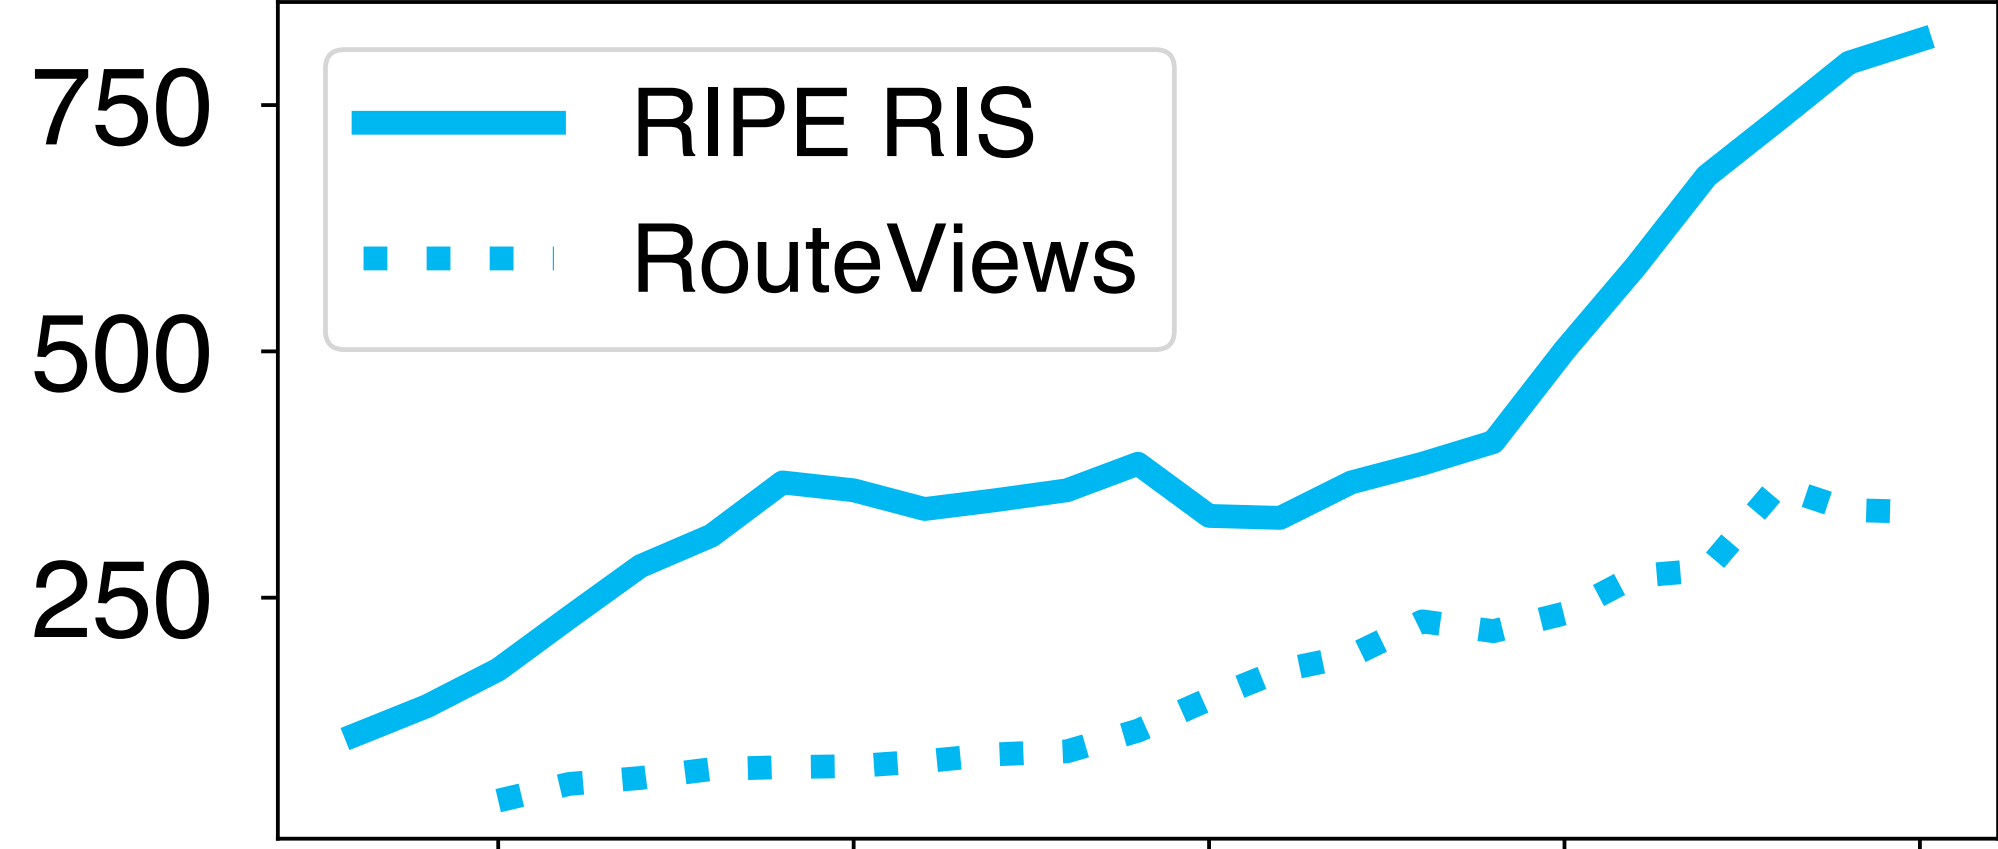
# Despite deploying new VPs, RIS and RouteViews' coverage is flat due the growing size of the Internet

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



Three 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

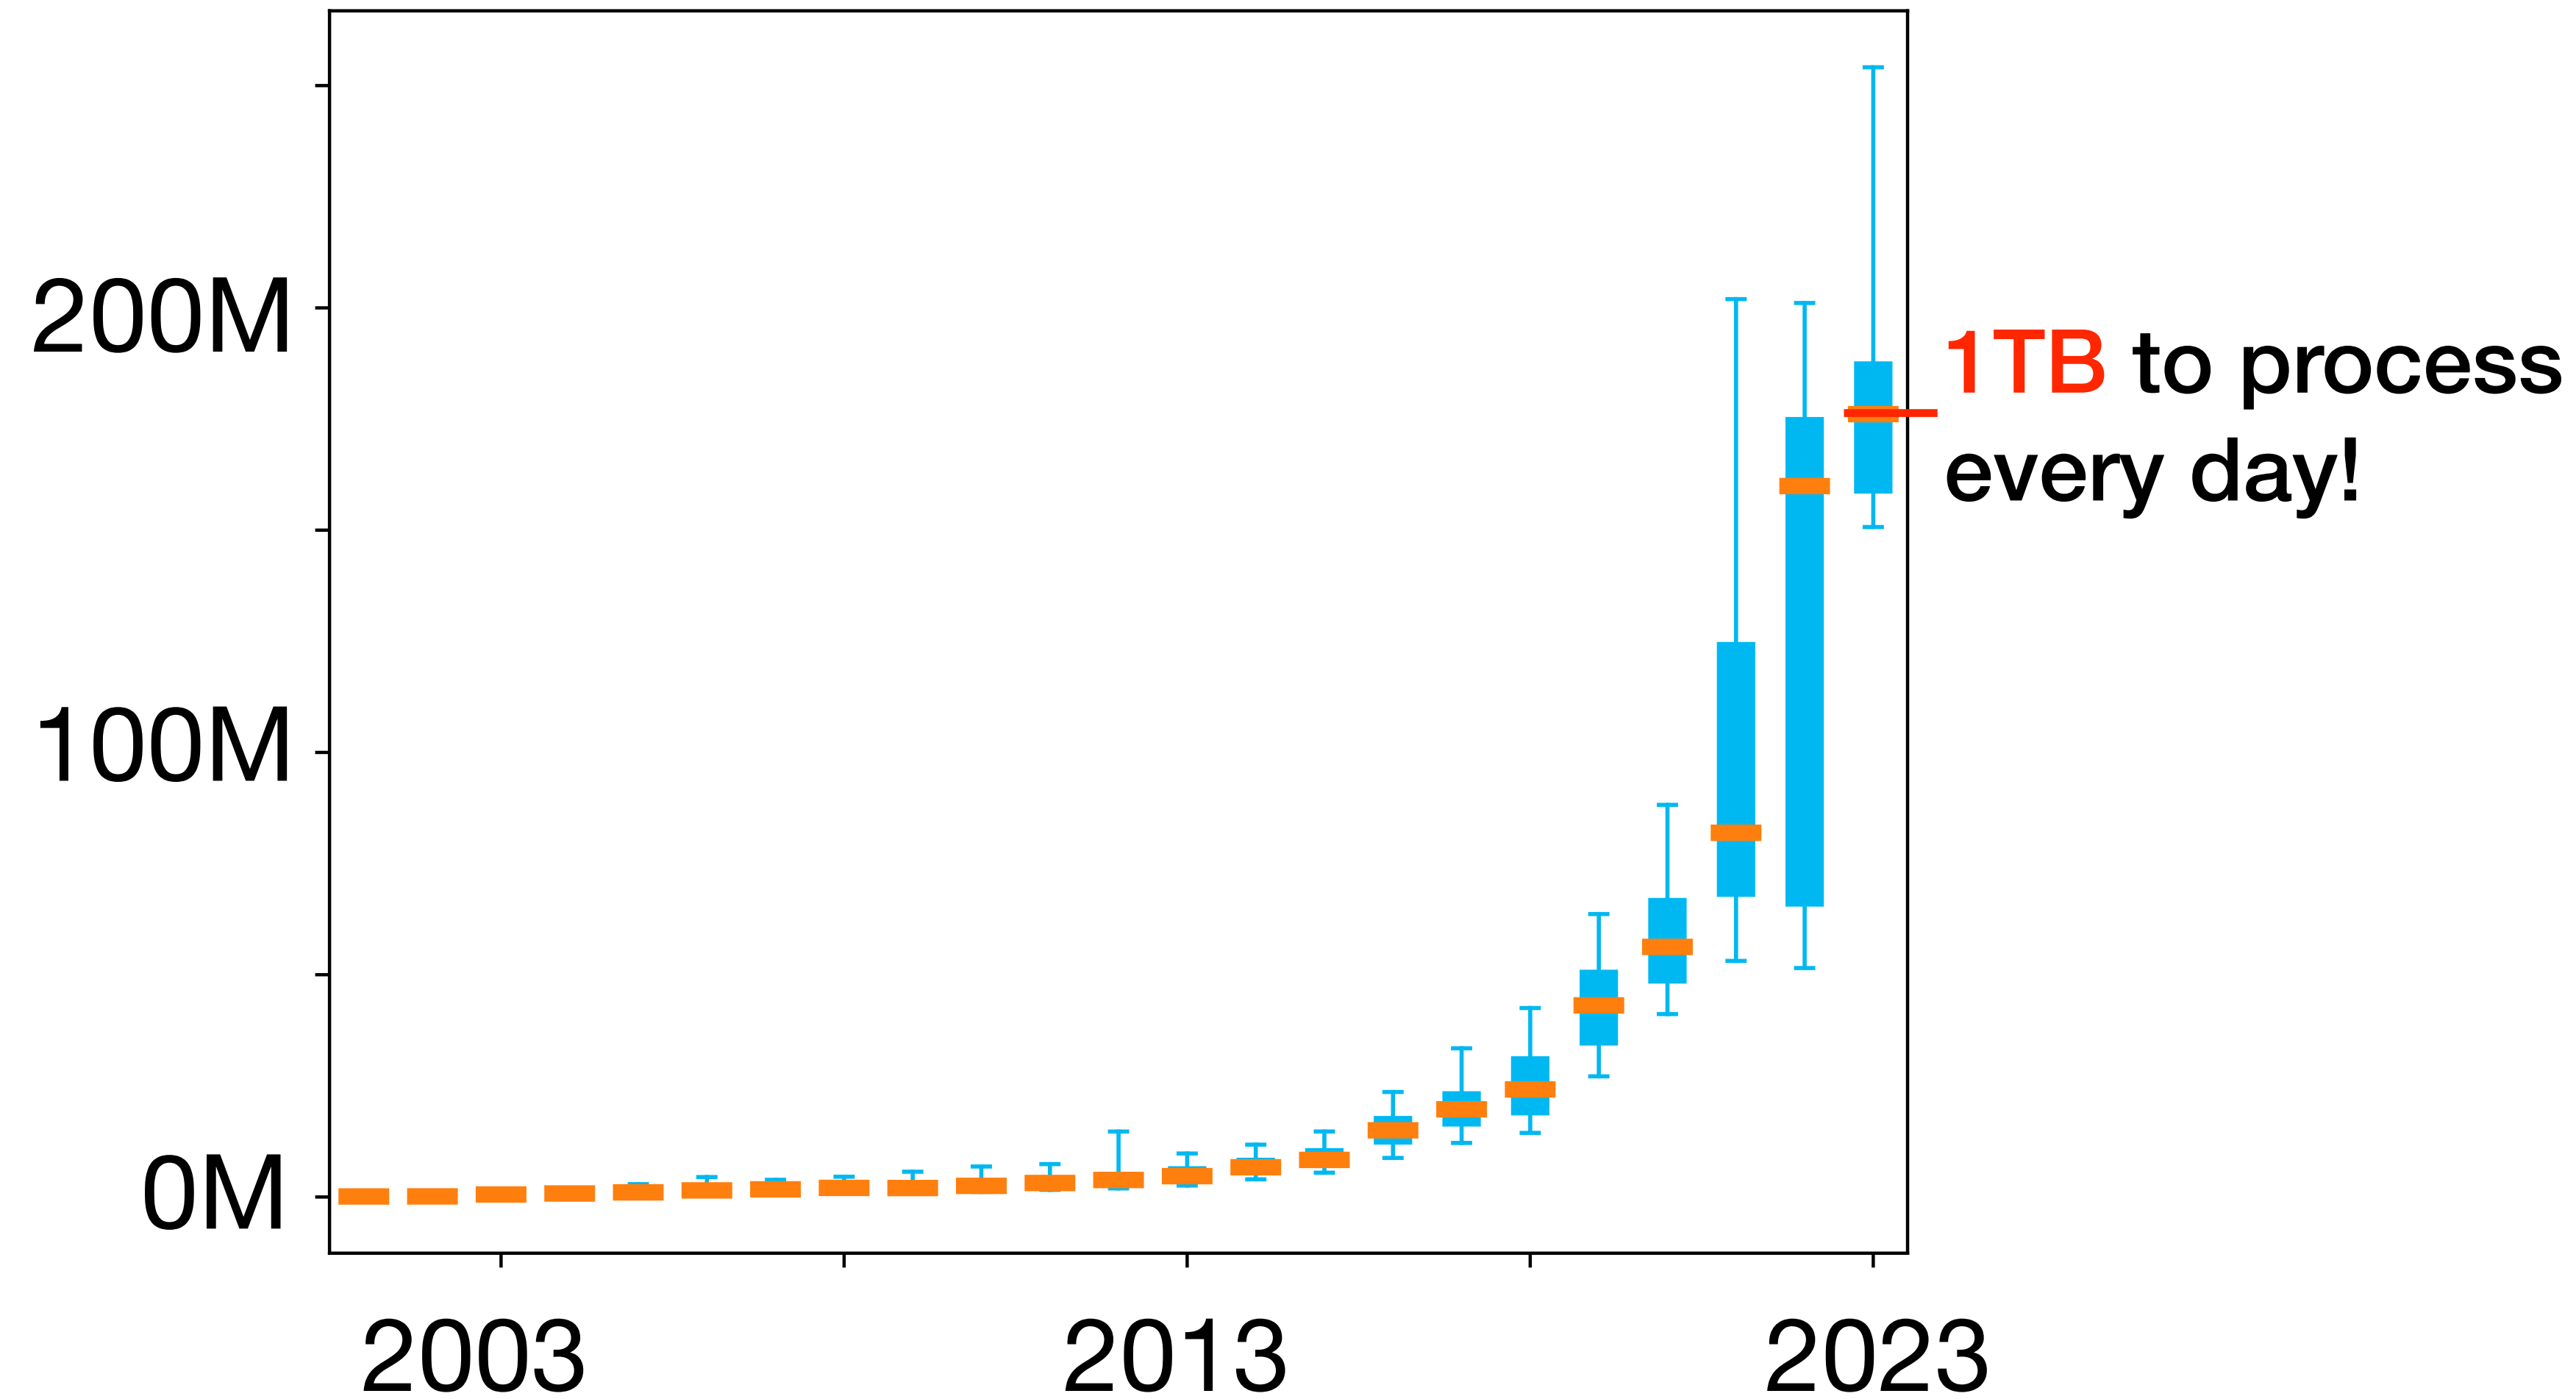
# of BGP routes collected per hour  
(*RIS + RouteViews*)






The number of routes collected increases **quadratically**

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



# The Next Generation of BGP Data Collection Platforms


AIMS GMI Workshop  
25 June 2024




# Outline

1. We observe that BGP routes are often redundant

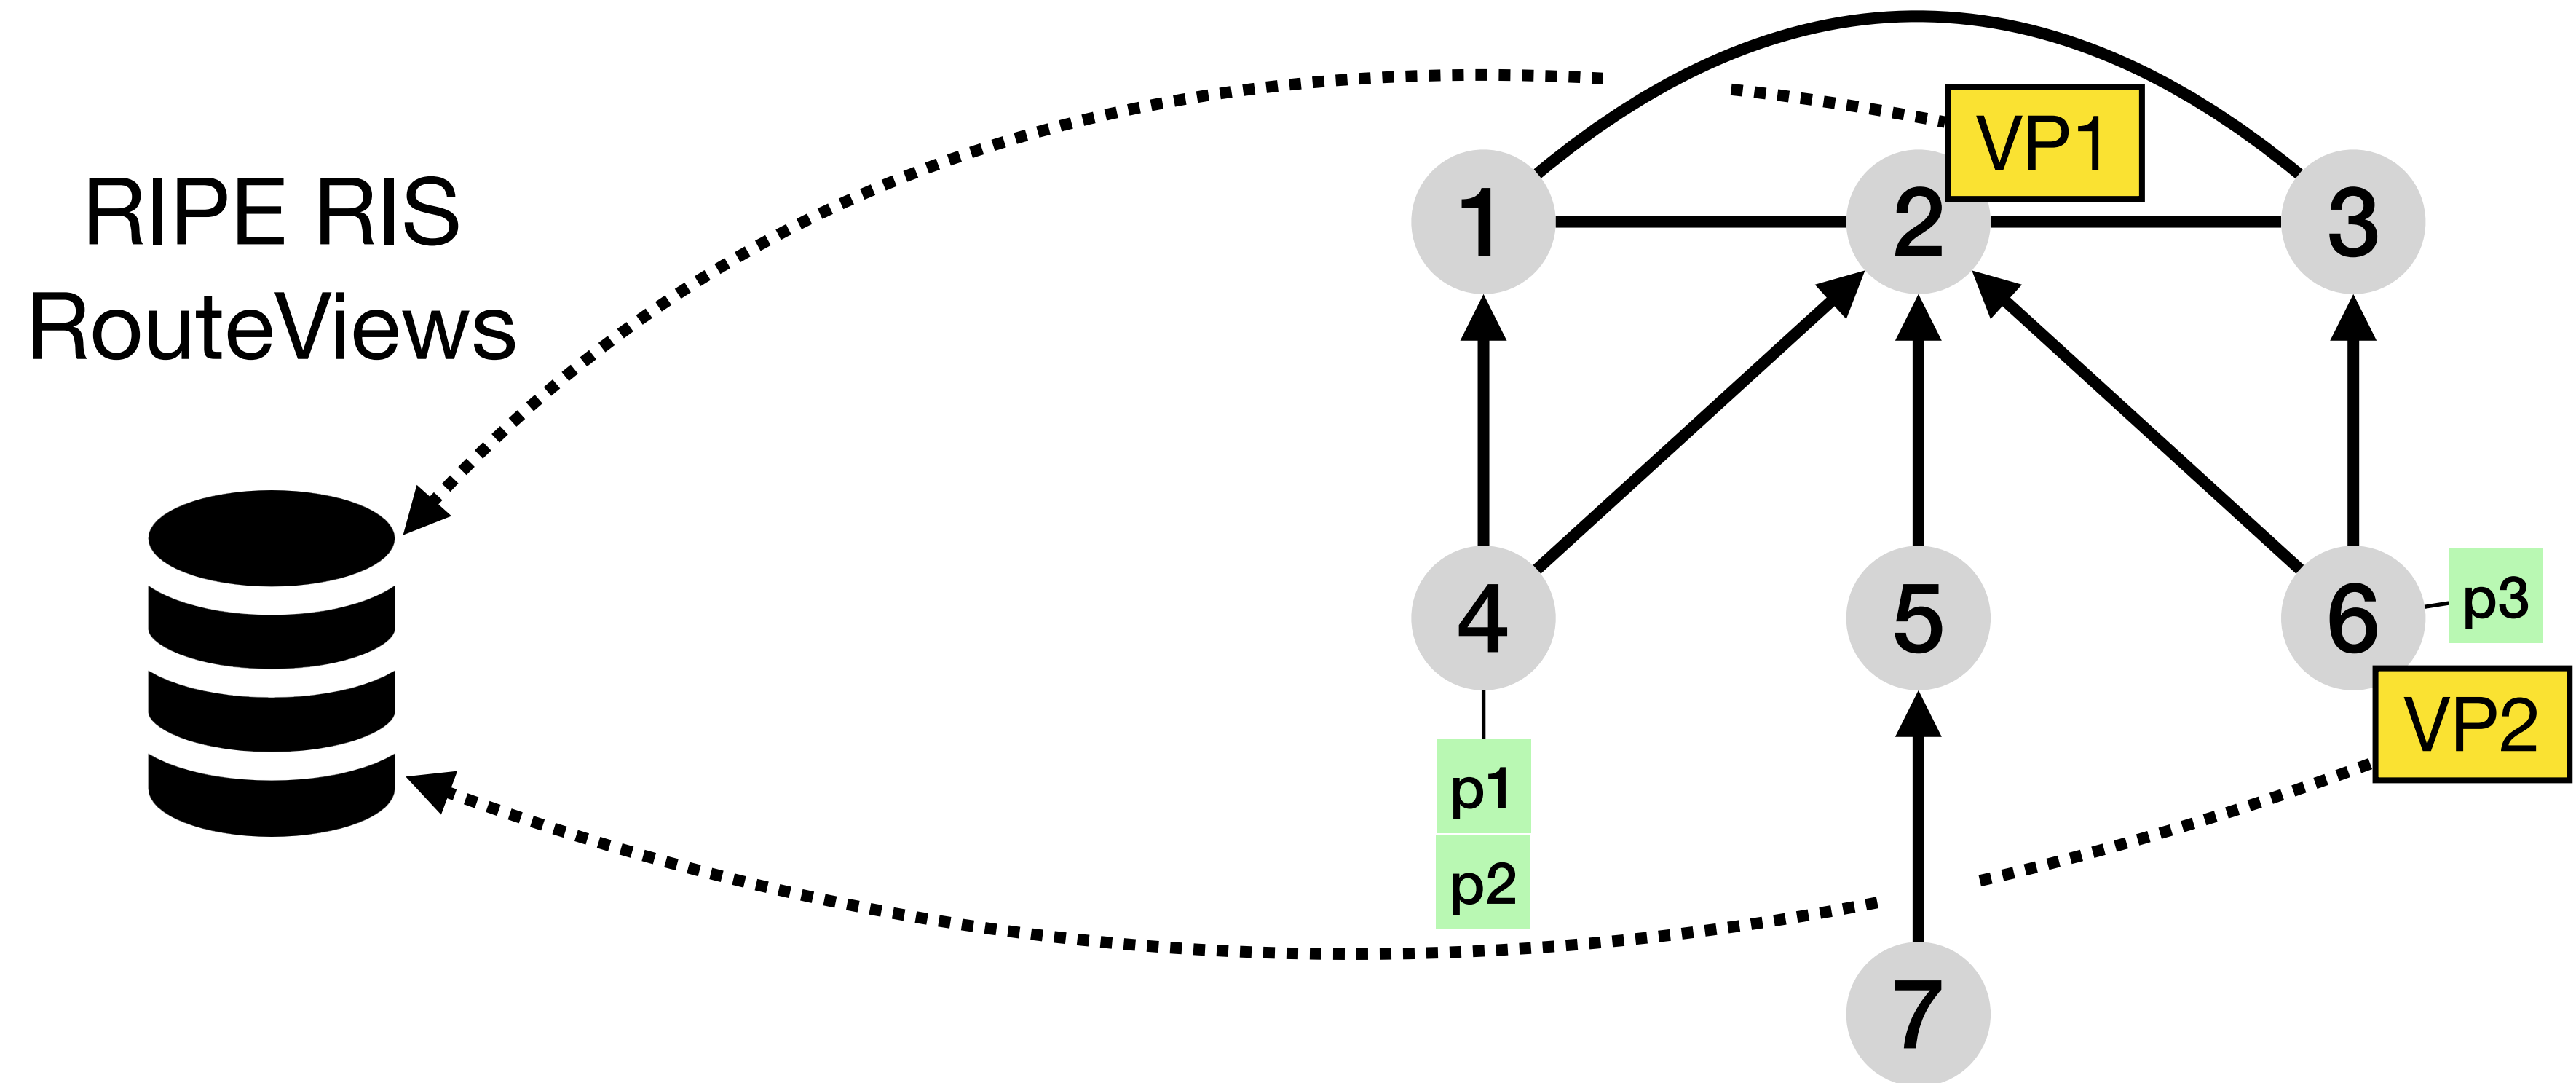
# BGP routes can be redundant



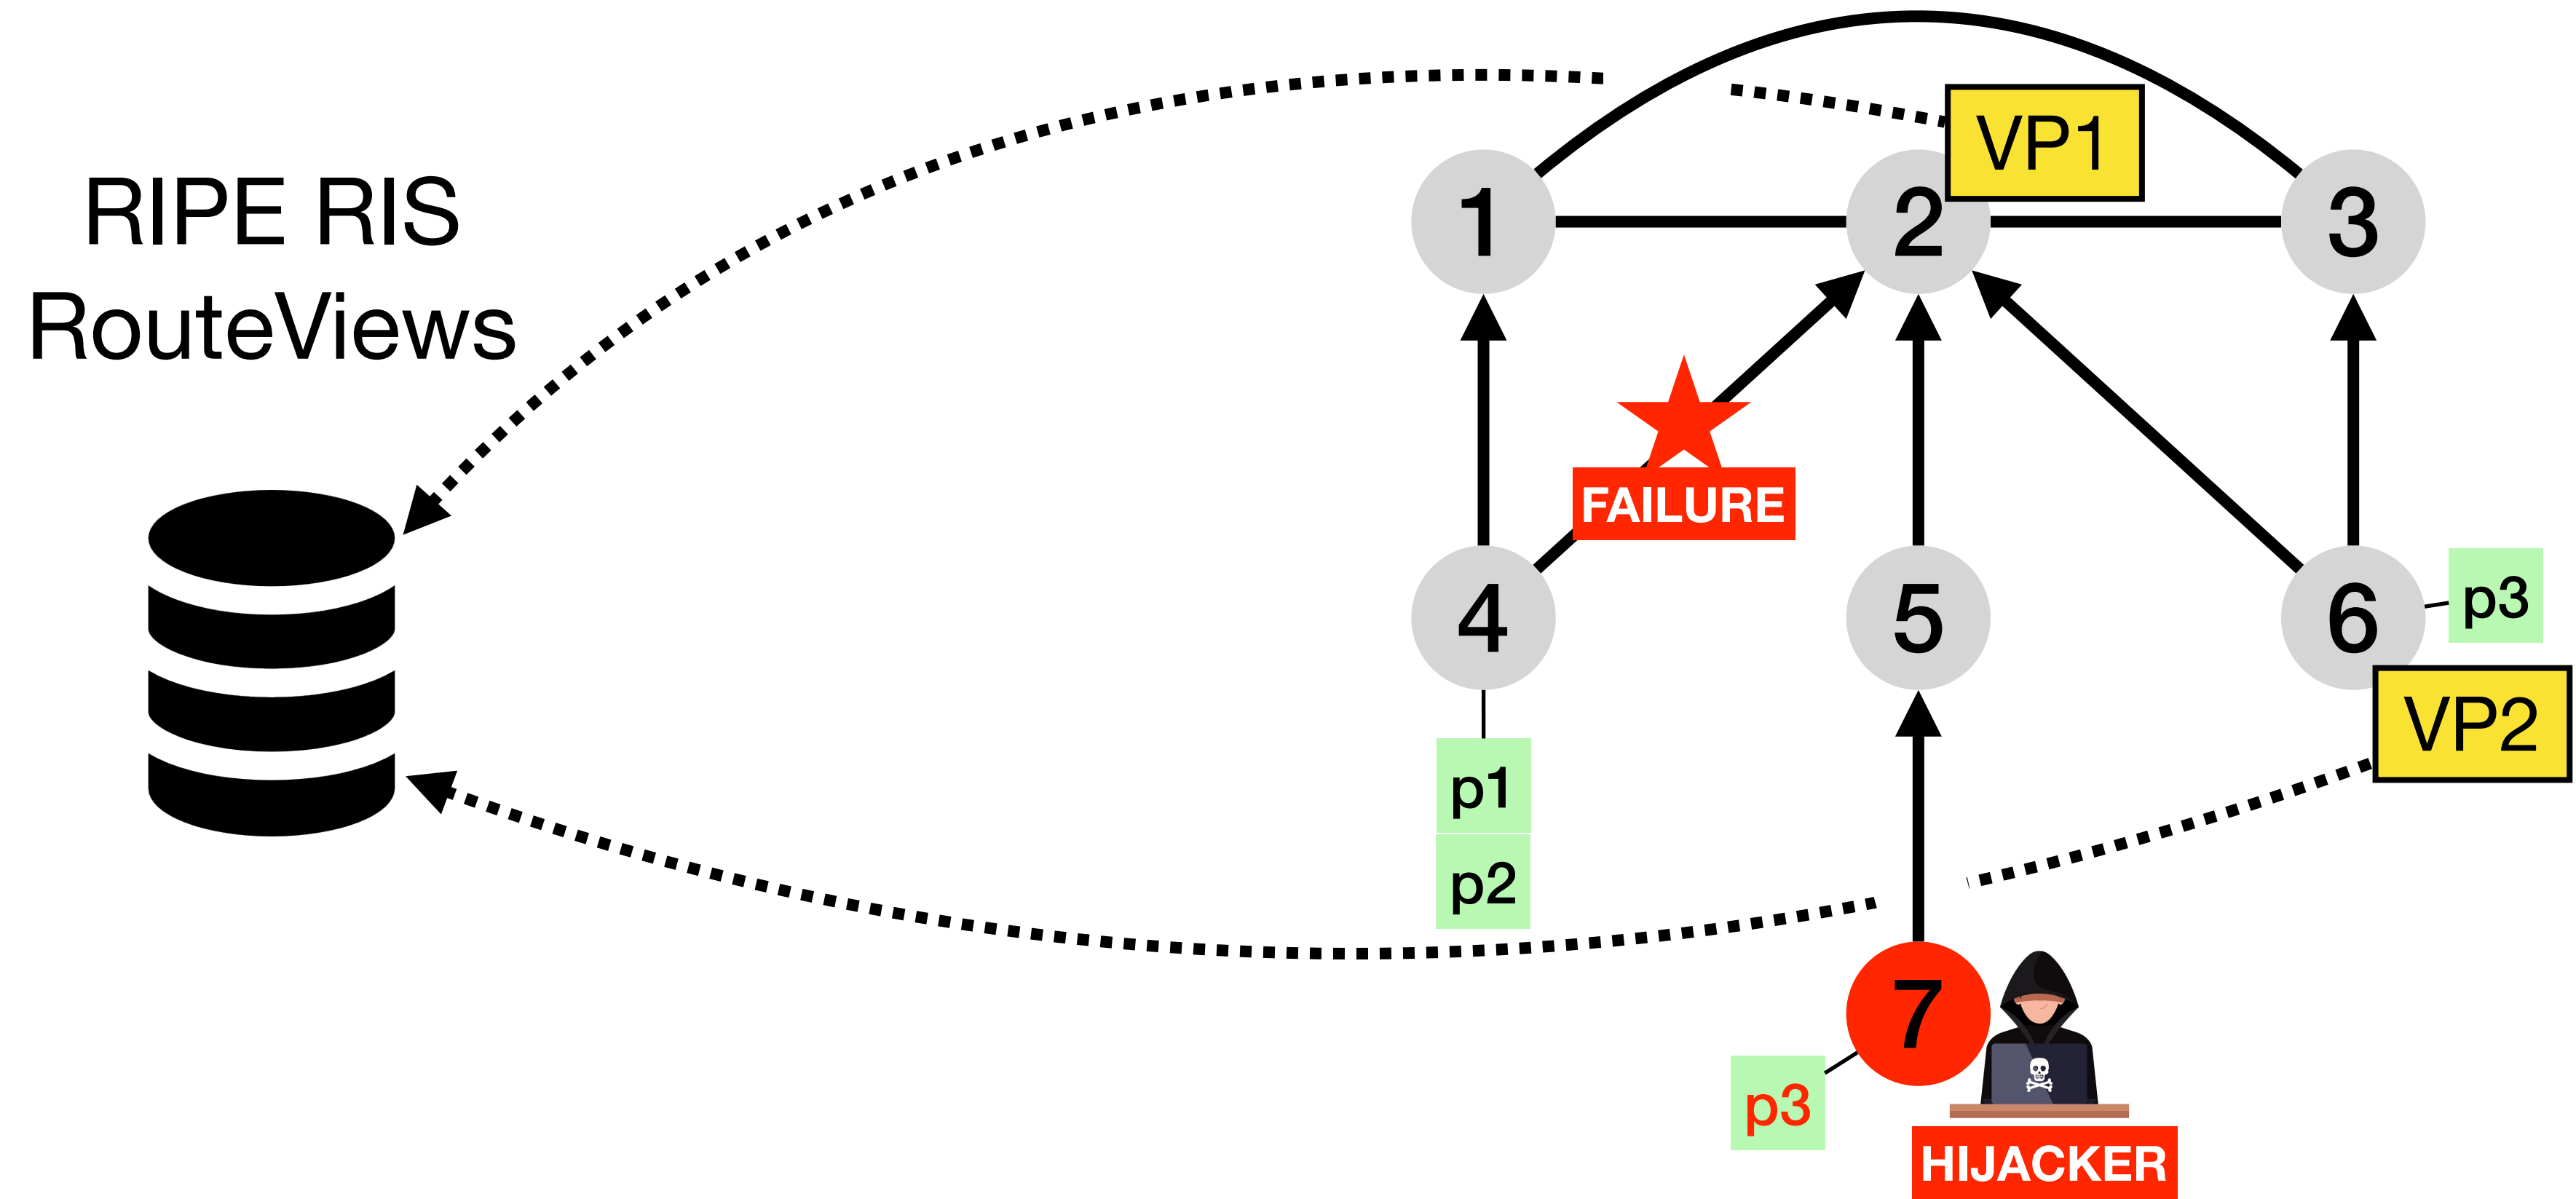
# BGP routes can be redundant



# BGP routes can be redundant



# BGP routes can be redundant





# BGP routes can be redundant

Collected routes

VP	prefix	AS path

RIPE RIS  
RouteViews



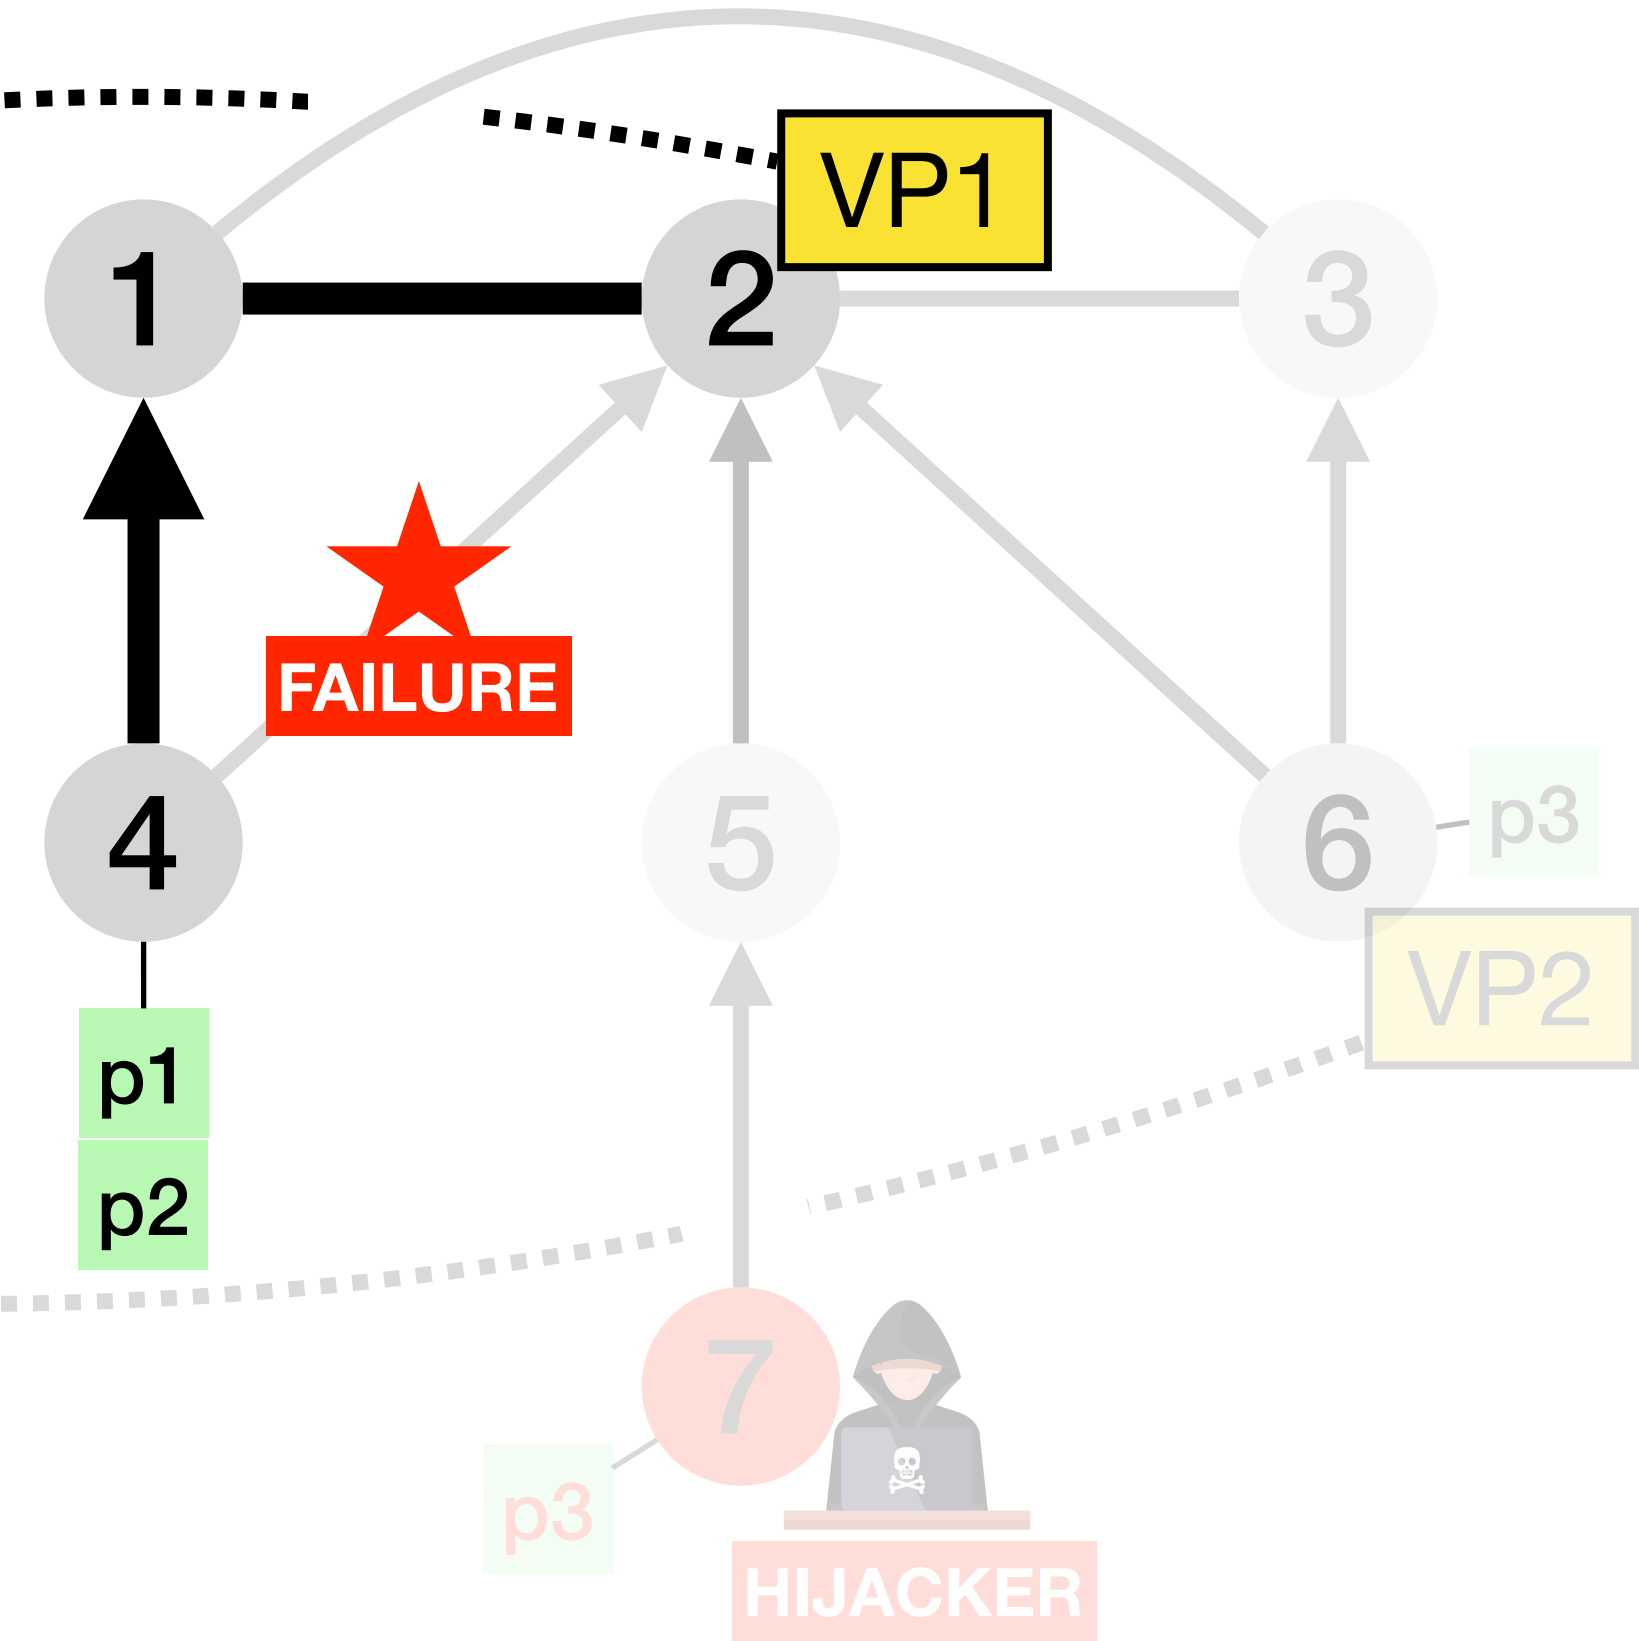



# BGP routes can be redundant

Collected routes

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

RIPE RIS  
RouteViews

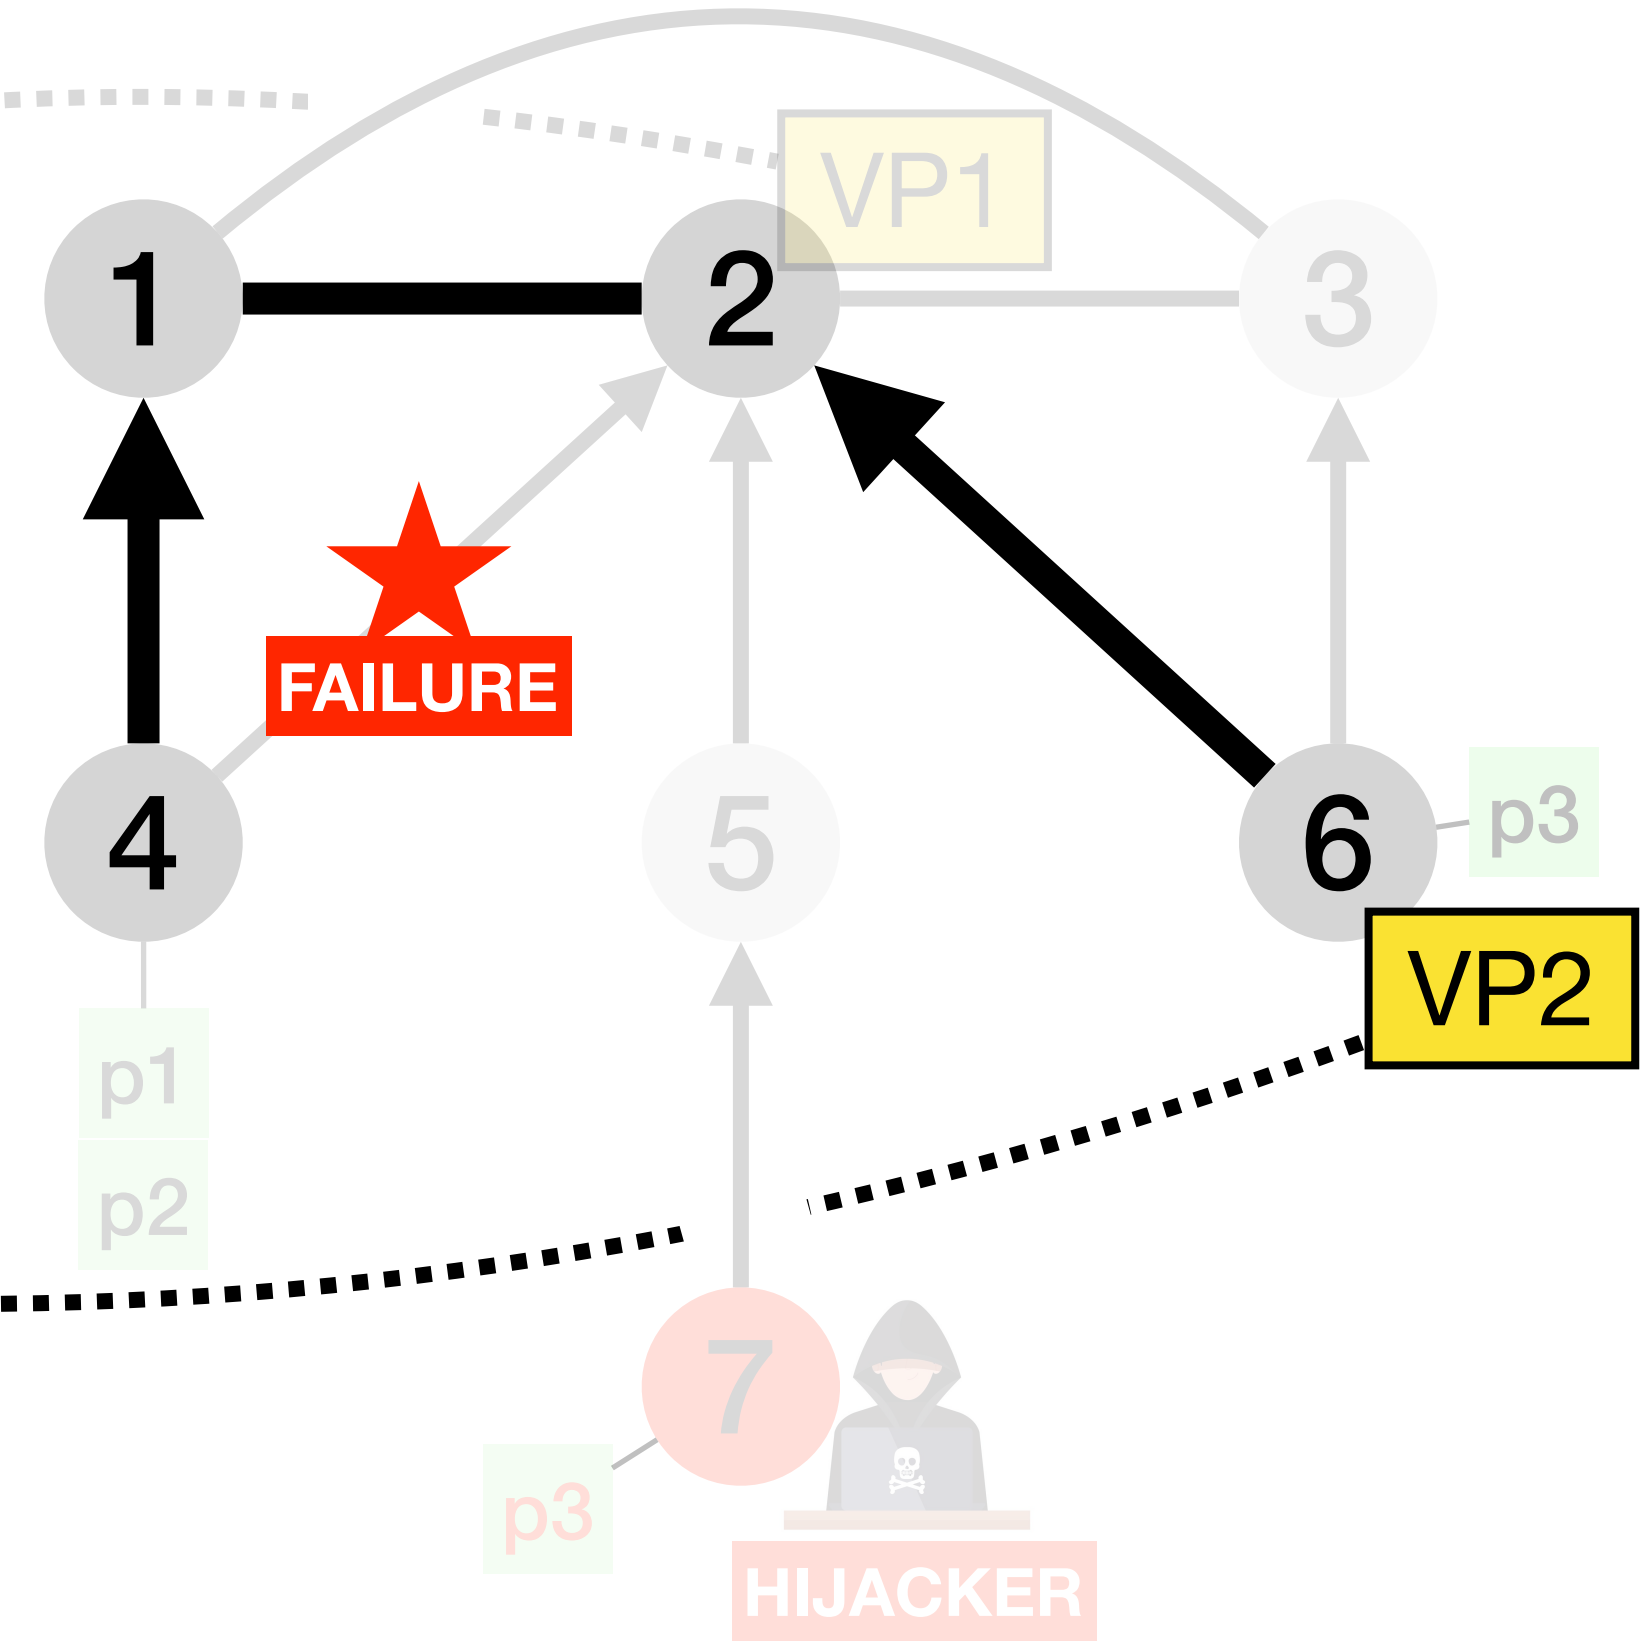



# BGP routes can be redundant

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

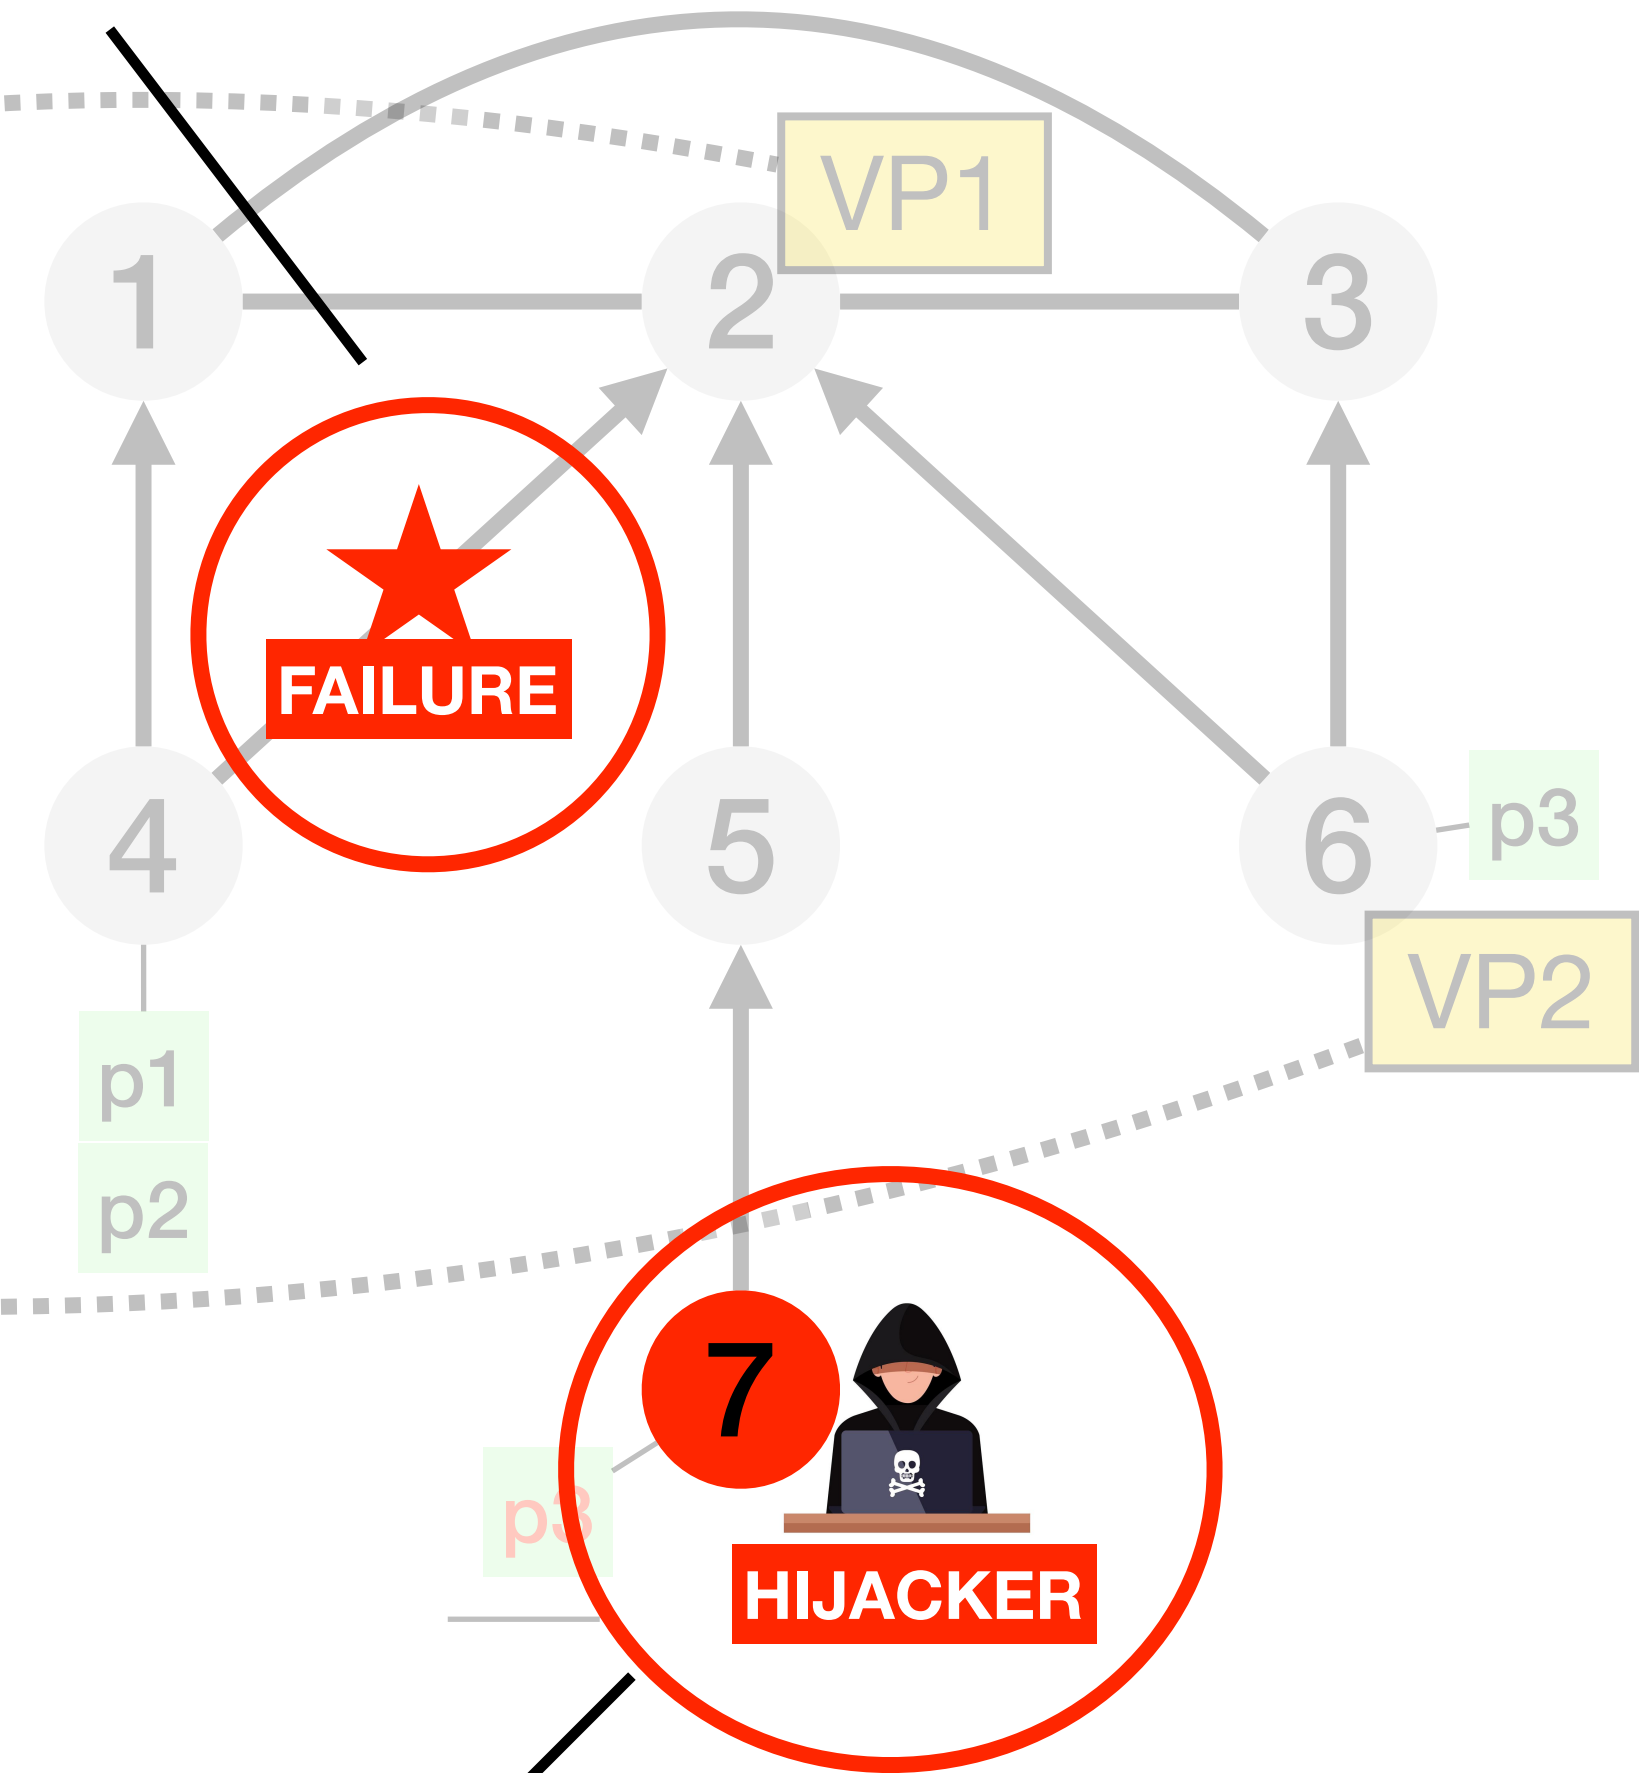
The failure is visible  
in one direction only

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**

# Large Hadron Collider (LHC) generates billions of collisions



# Large Hadron Collider (LHC) generates billions of collisions

Only 0.0006% of generated collisions are actually relevant



# Large Hadron Collider (LHC) generates billions of collisions

Only 0.0006% of generated collisions are actually relevant

They rely on custom **hardware** and **algorithms** to discard uninteresting data prior processing





# Large Hadron Collider (LHC) generates billions of collisions

Only 0.0006% of generated collisions are actually relevant


They rely on custom **hardware** and **algorithms** to discard uninteresting data prior processing



They are using an "**overshoot-and-discard**" collection strategy

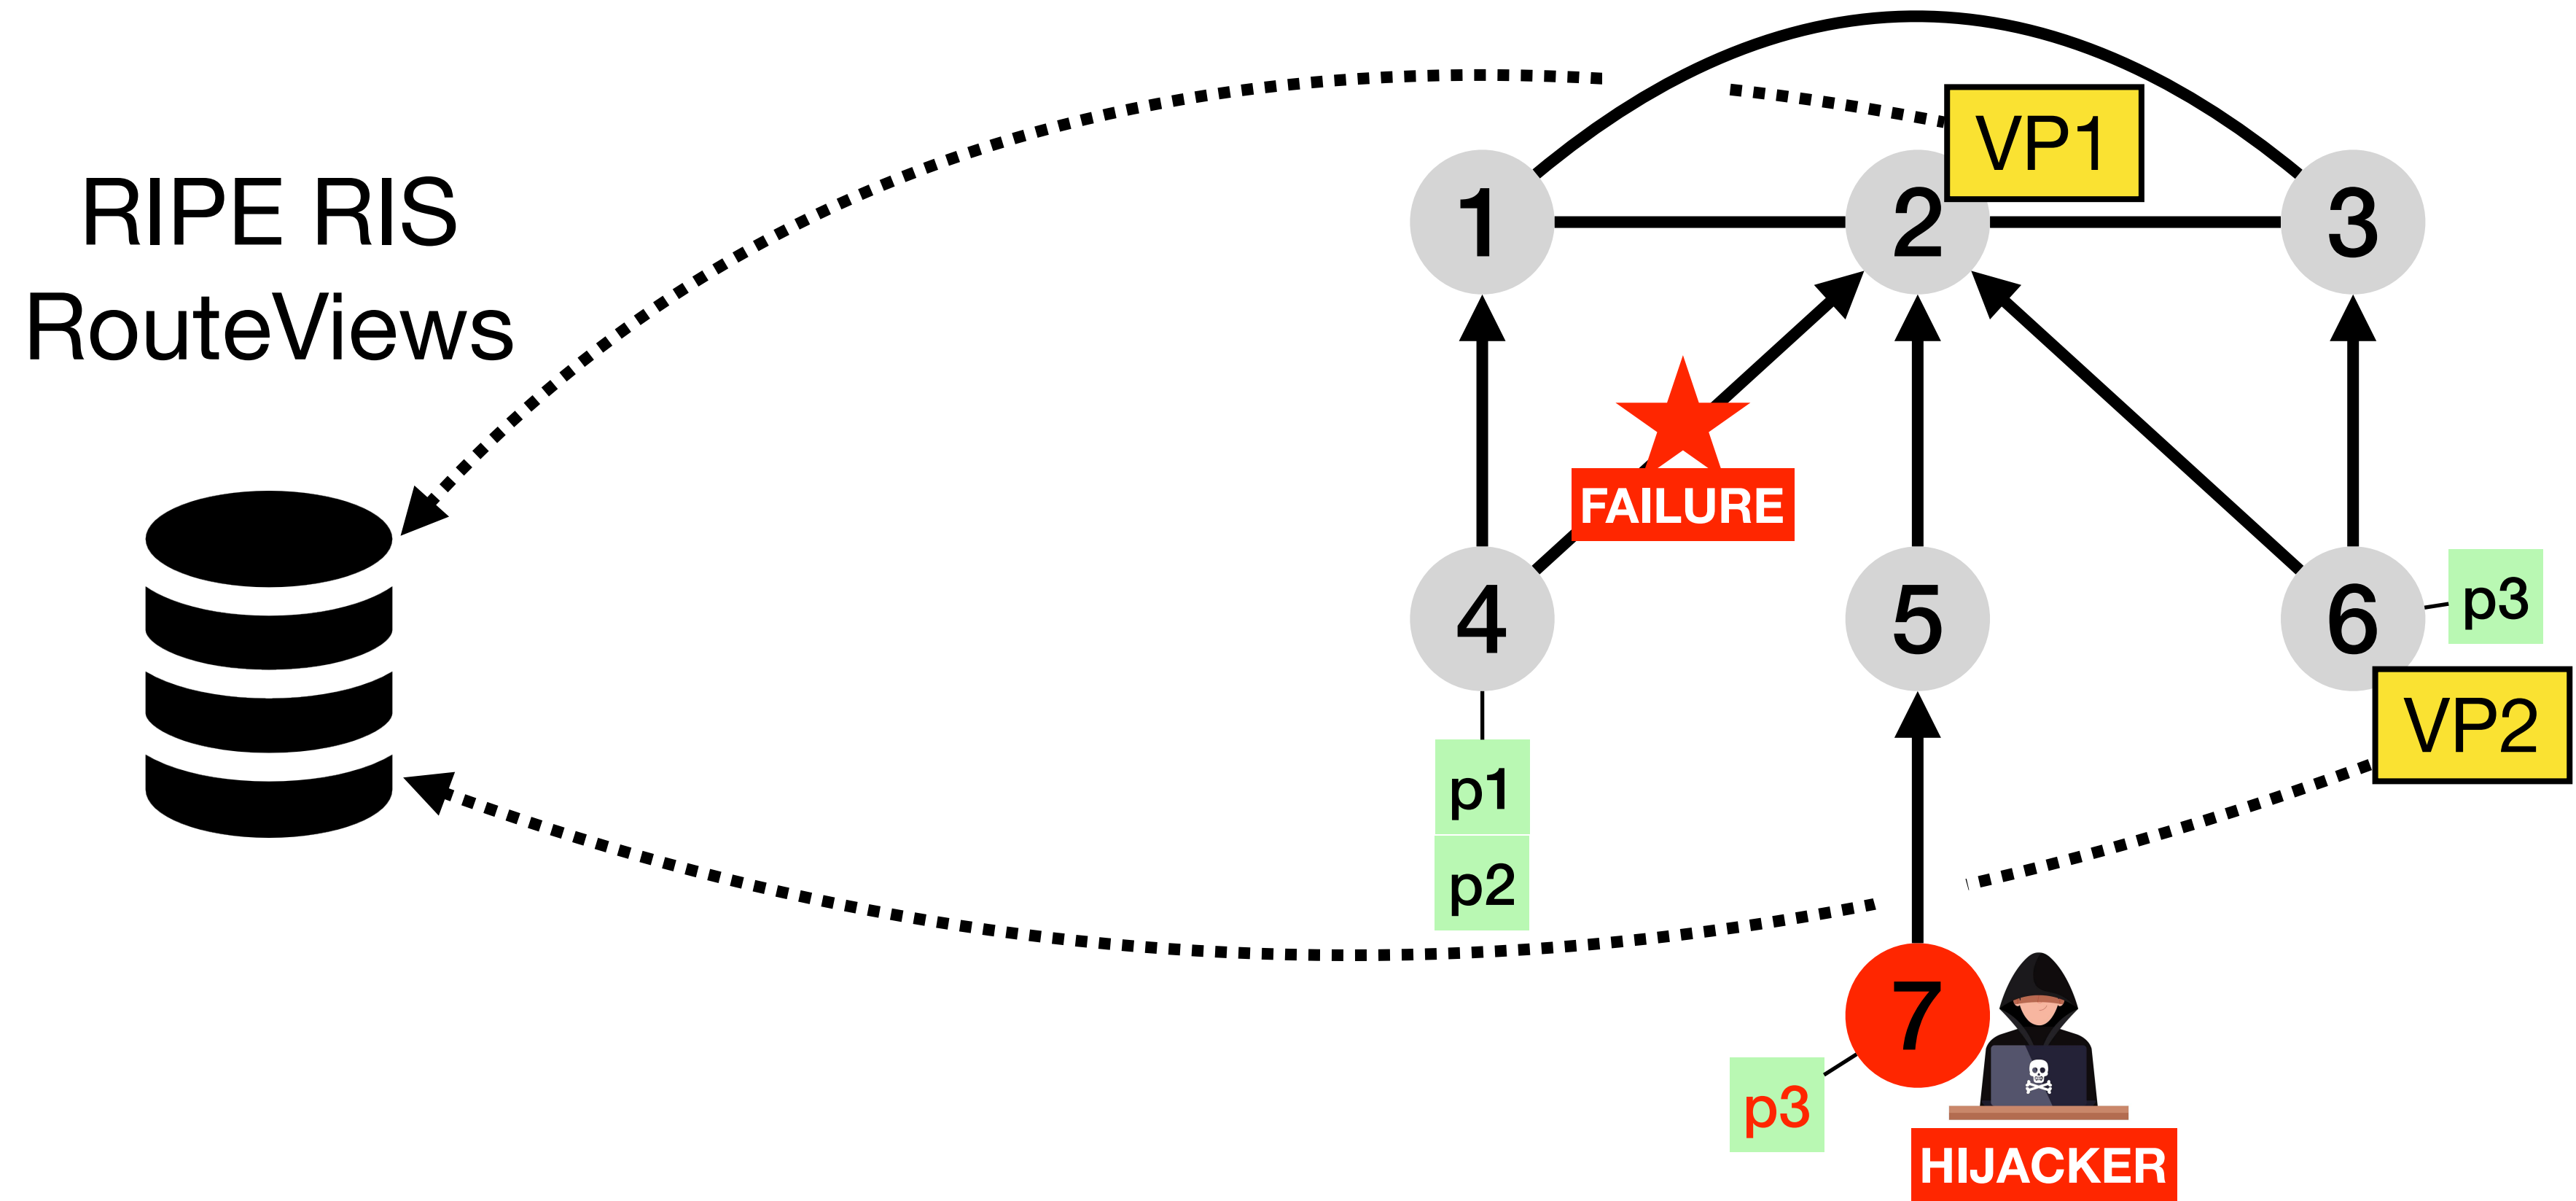
The “overshoot-and-discard” data collection paradigm can be adapted to BGP data collection

# The “overshoot-and-discard” data collection paradigm can be adapted to BGP data collection

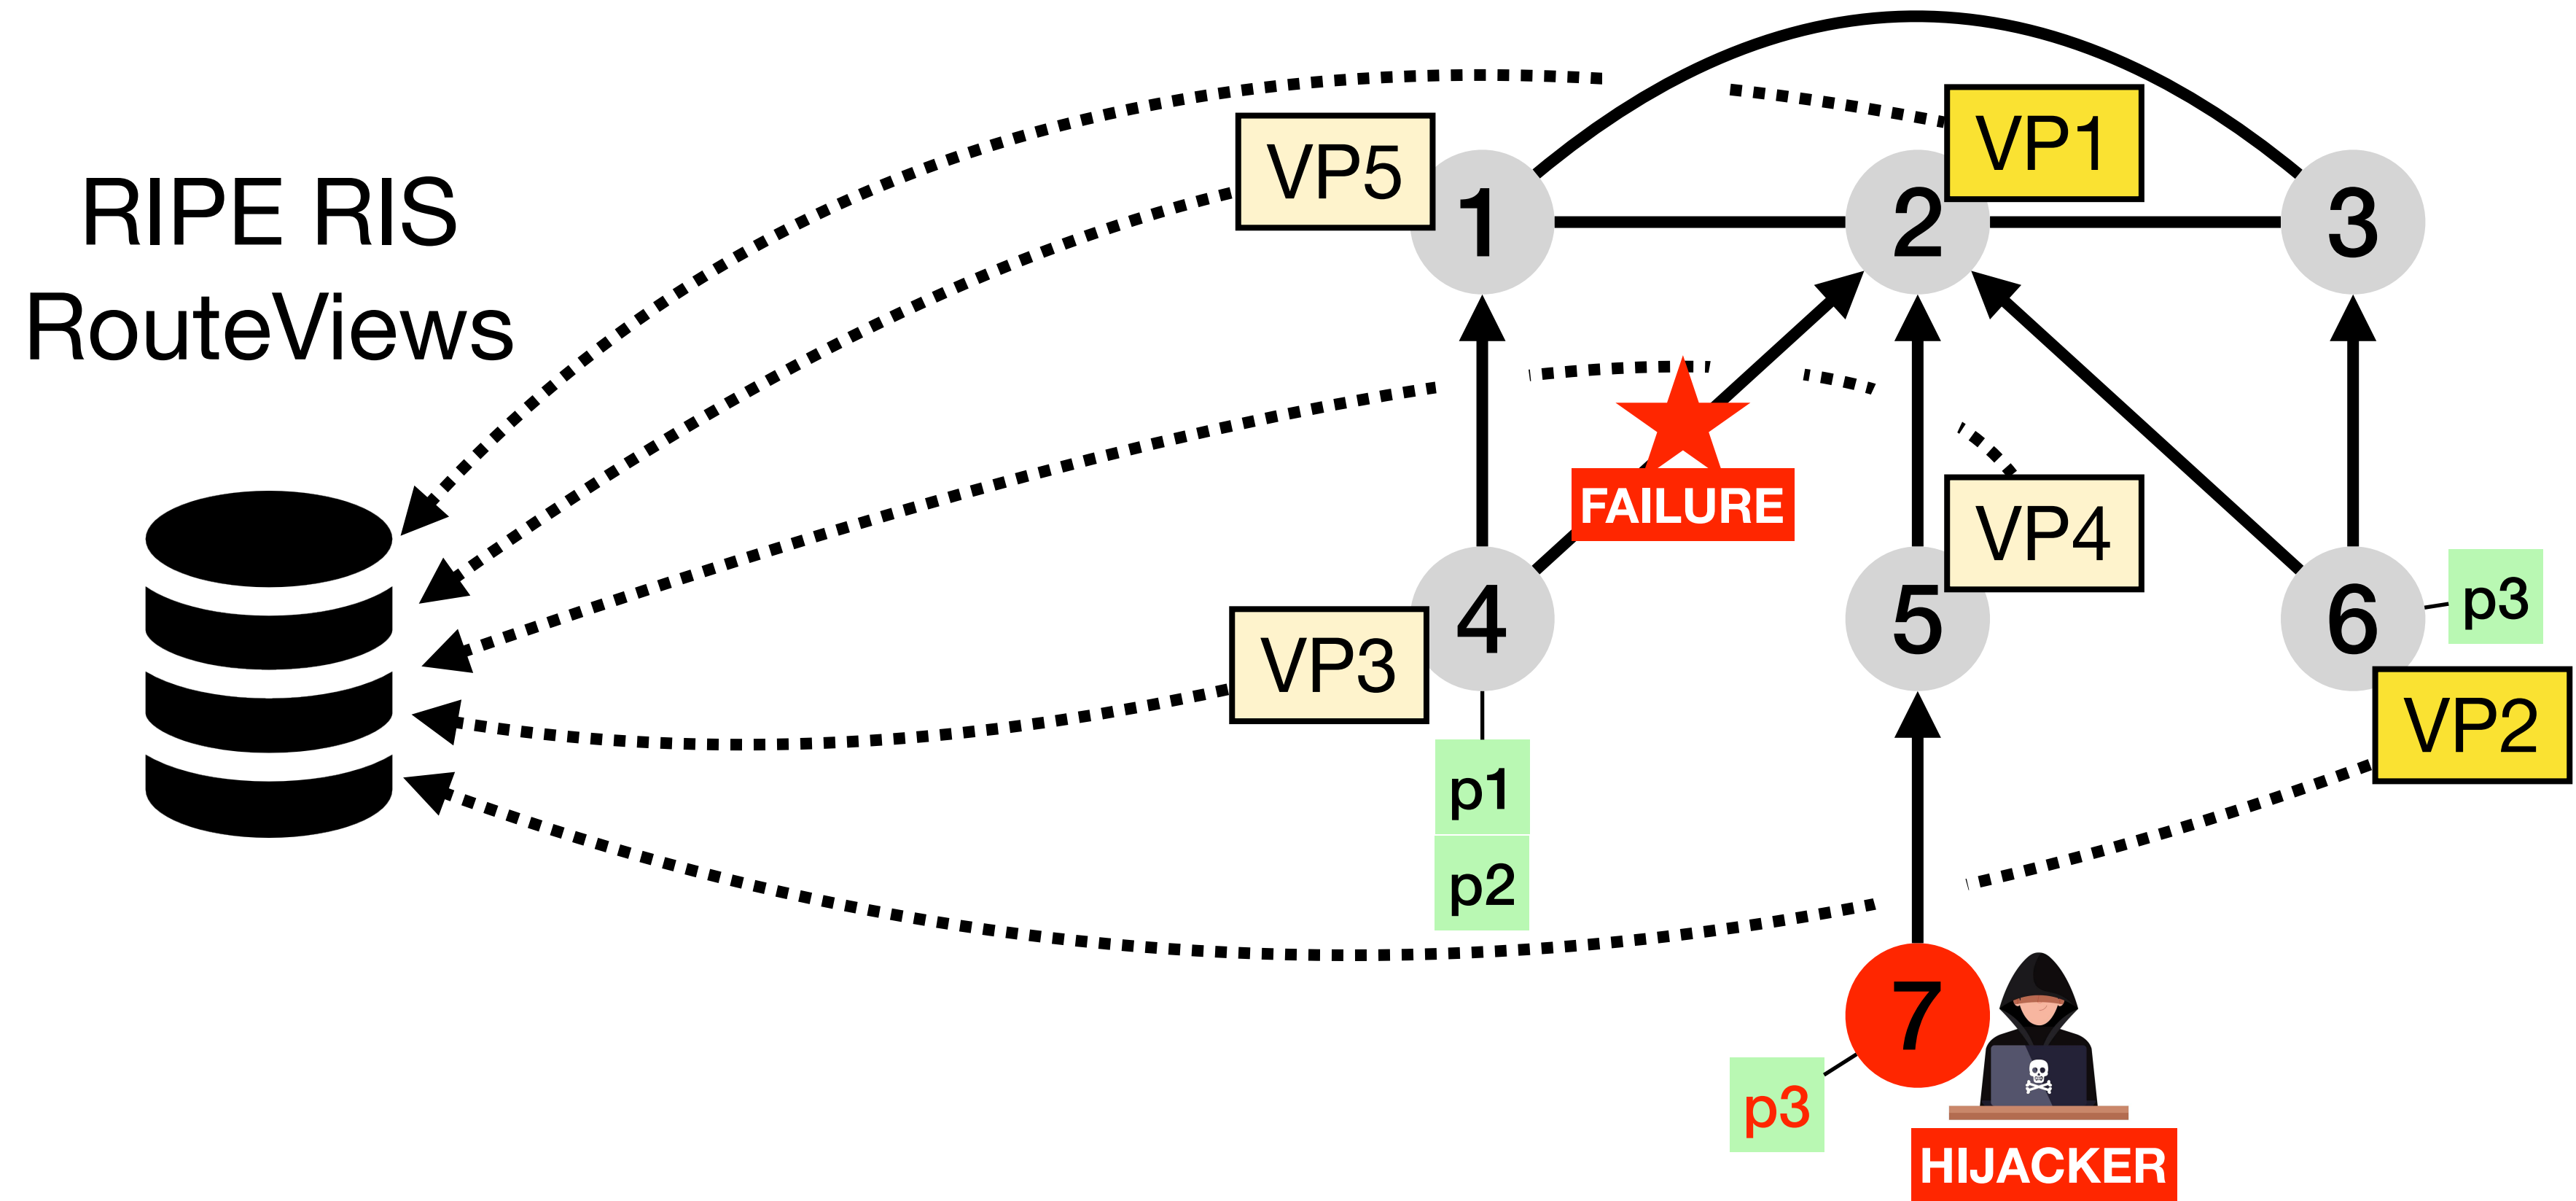


Overshoot: We collect data from as many VPs as possible  
*To prevent missing important information*

# Overshoot: deploying as many VPs as possible



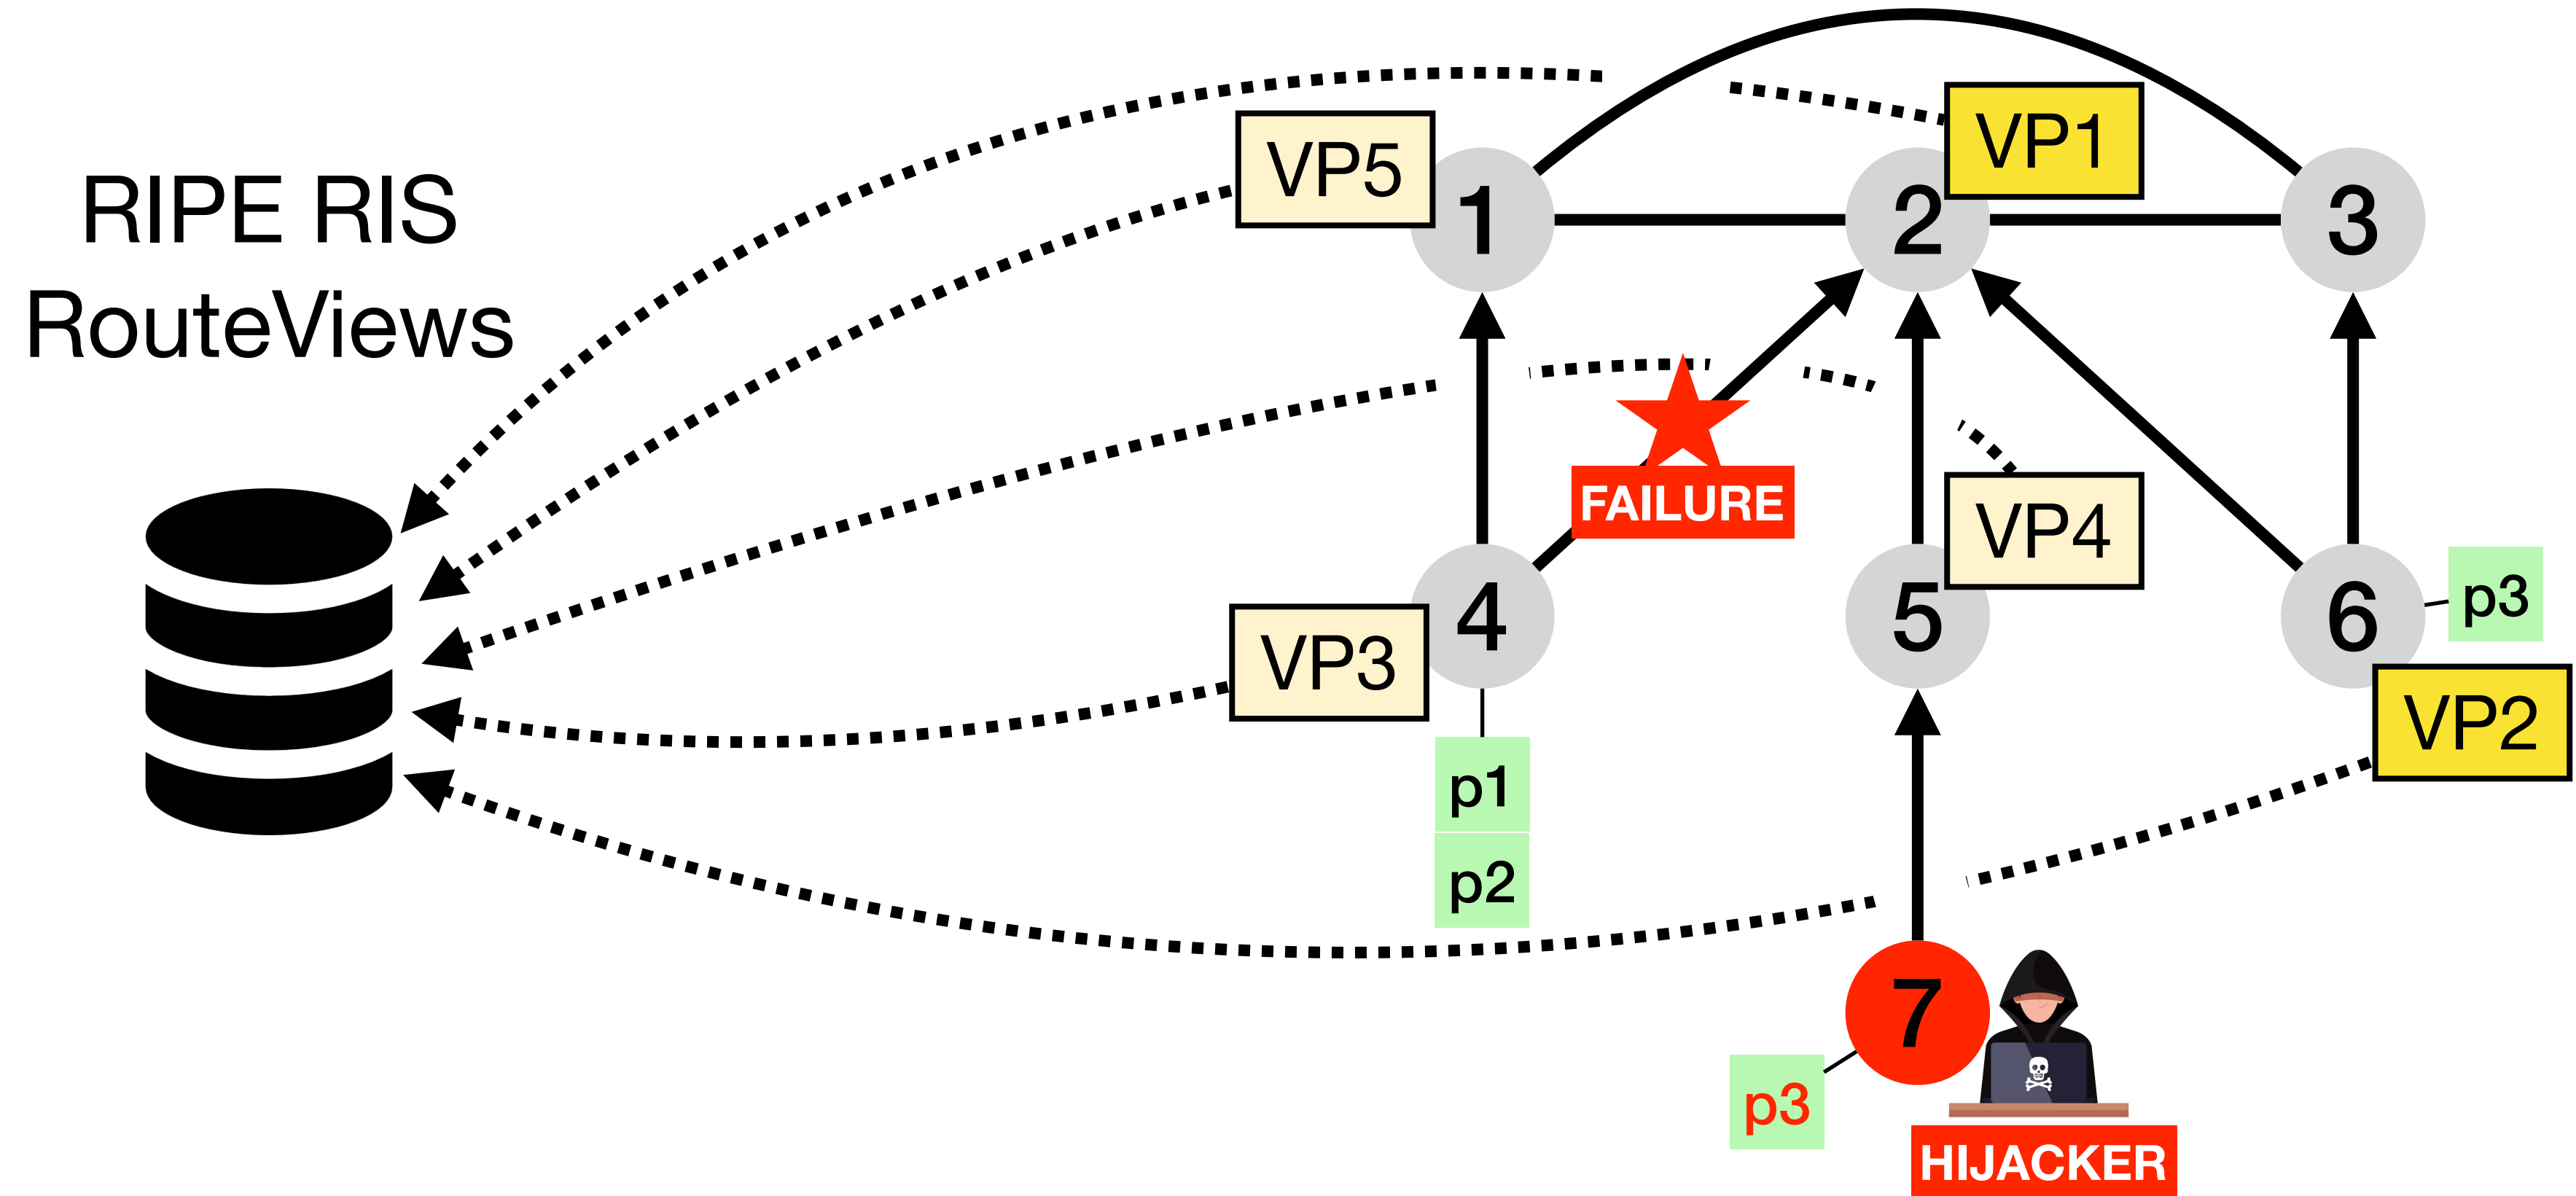
# Overshoot: deploying as many VPs as possible



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

Collected routes

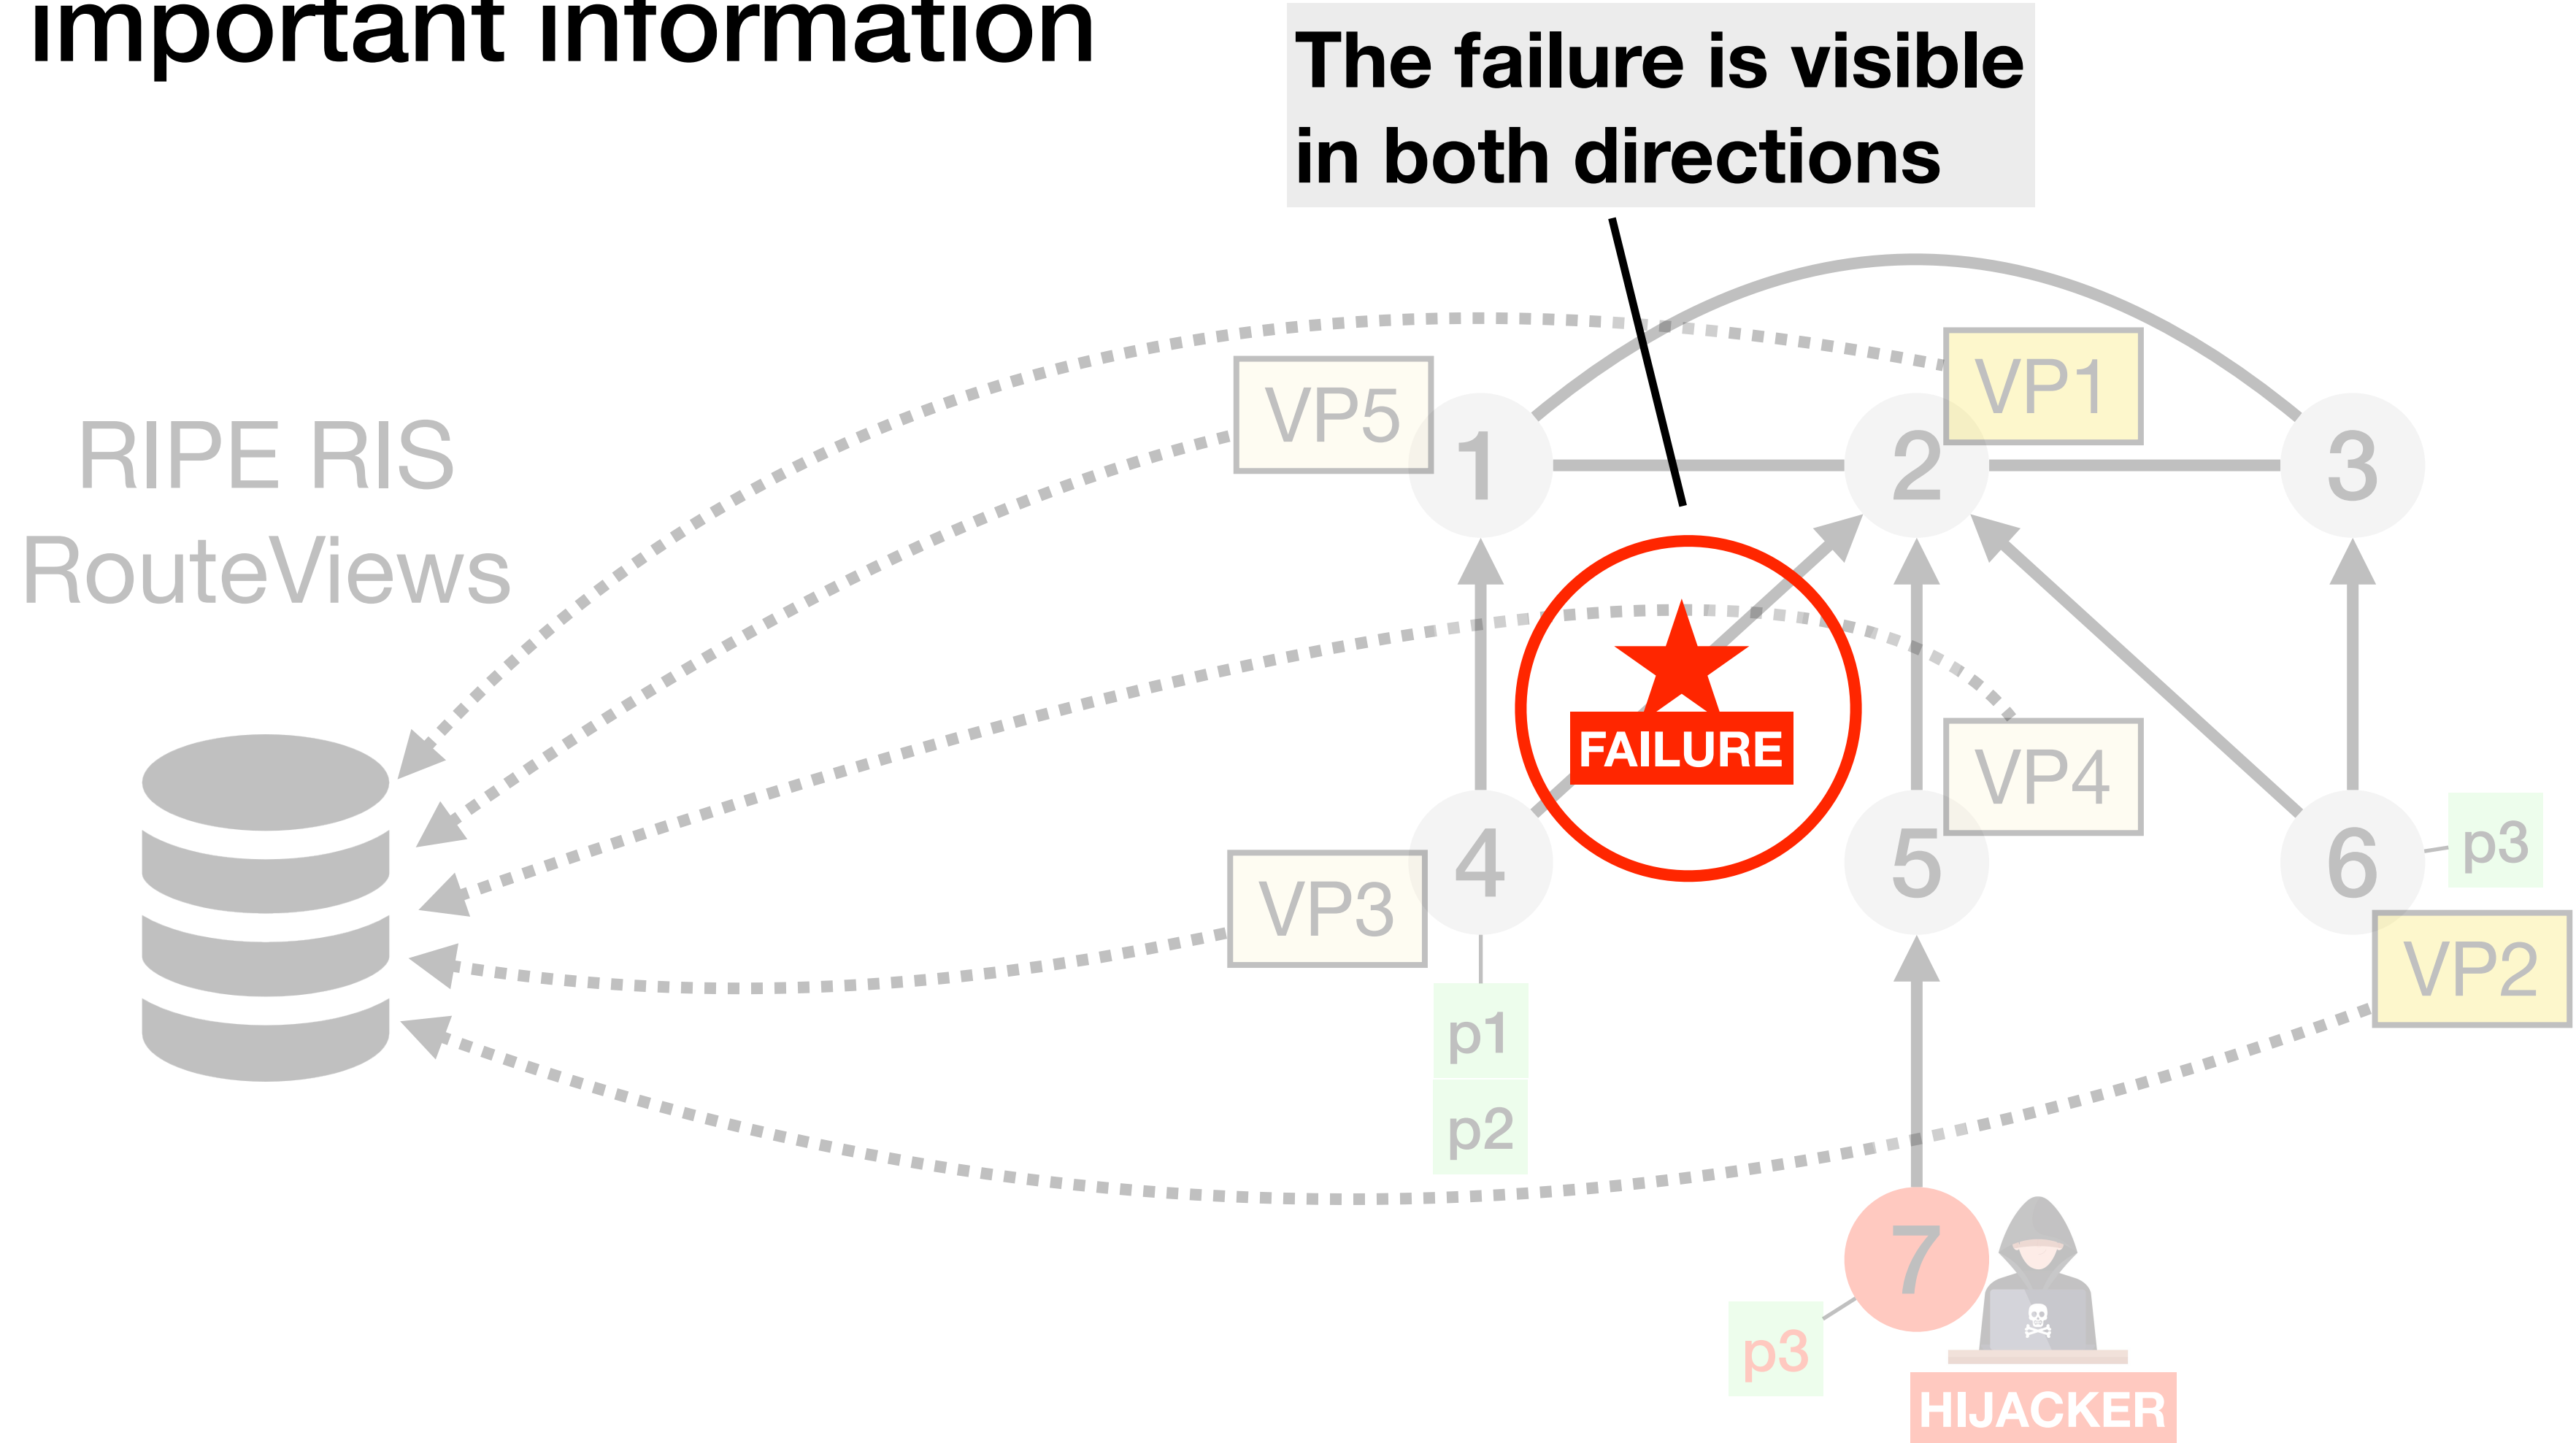
VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 2 1 4
VP3	p3	4 1 2 6
VP4	p3	5 7



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

## Collected routes

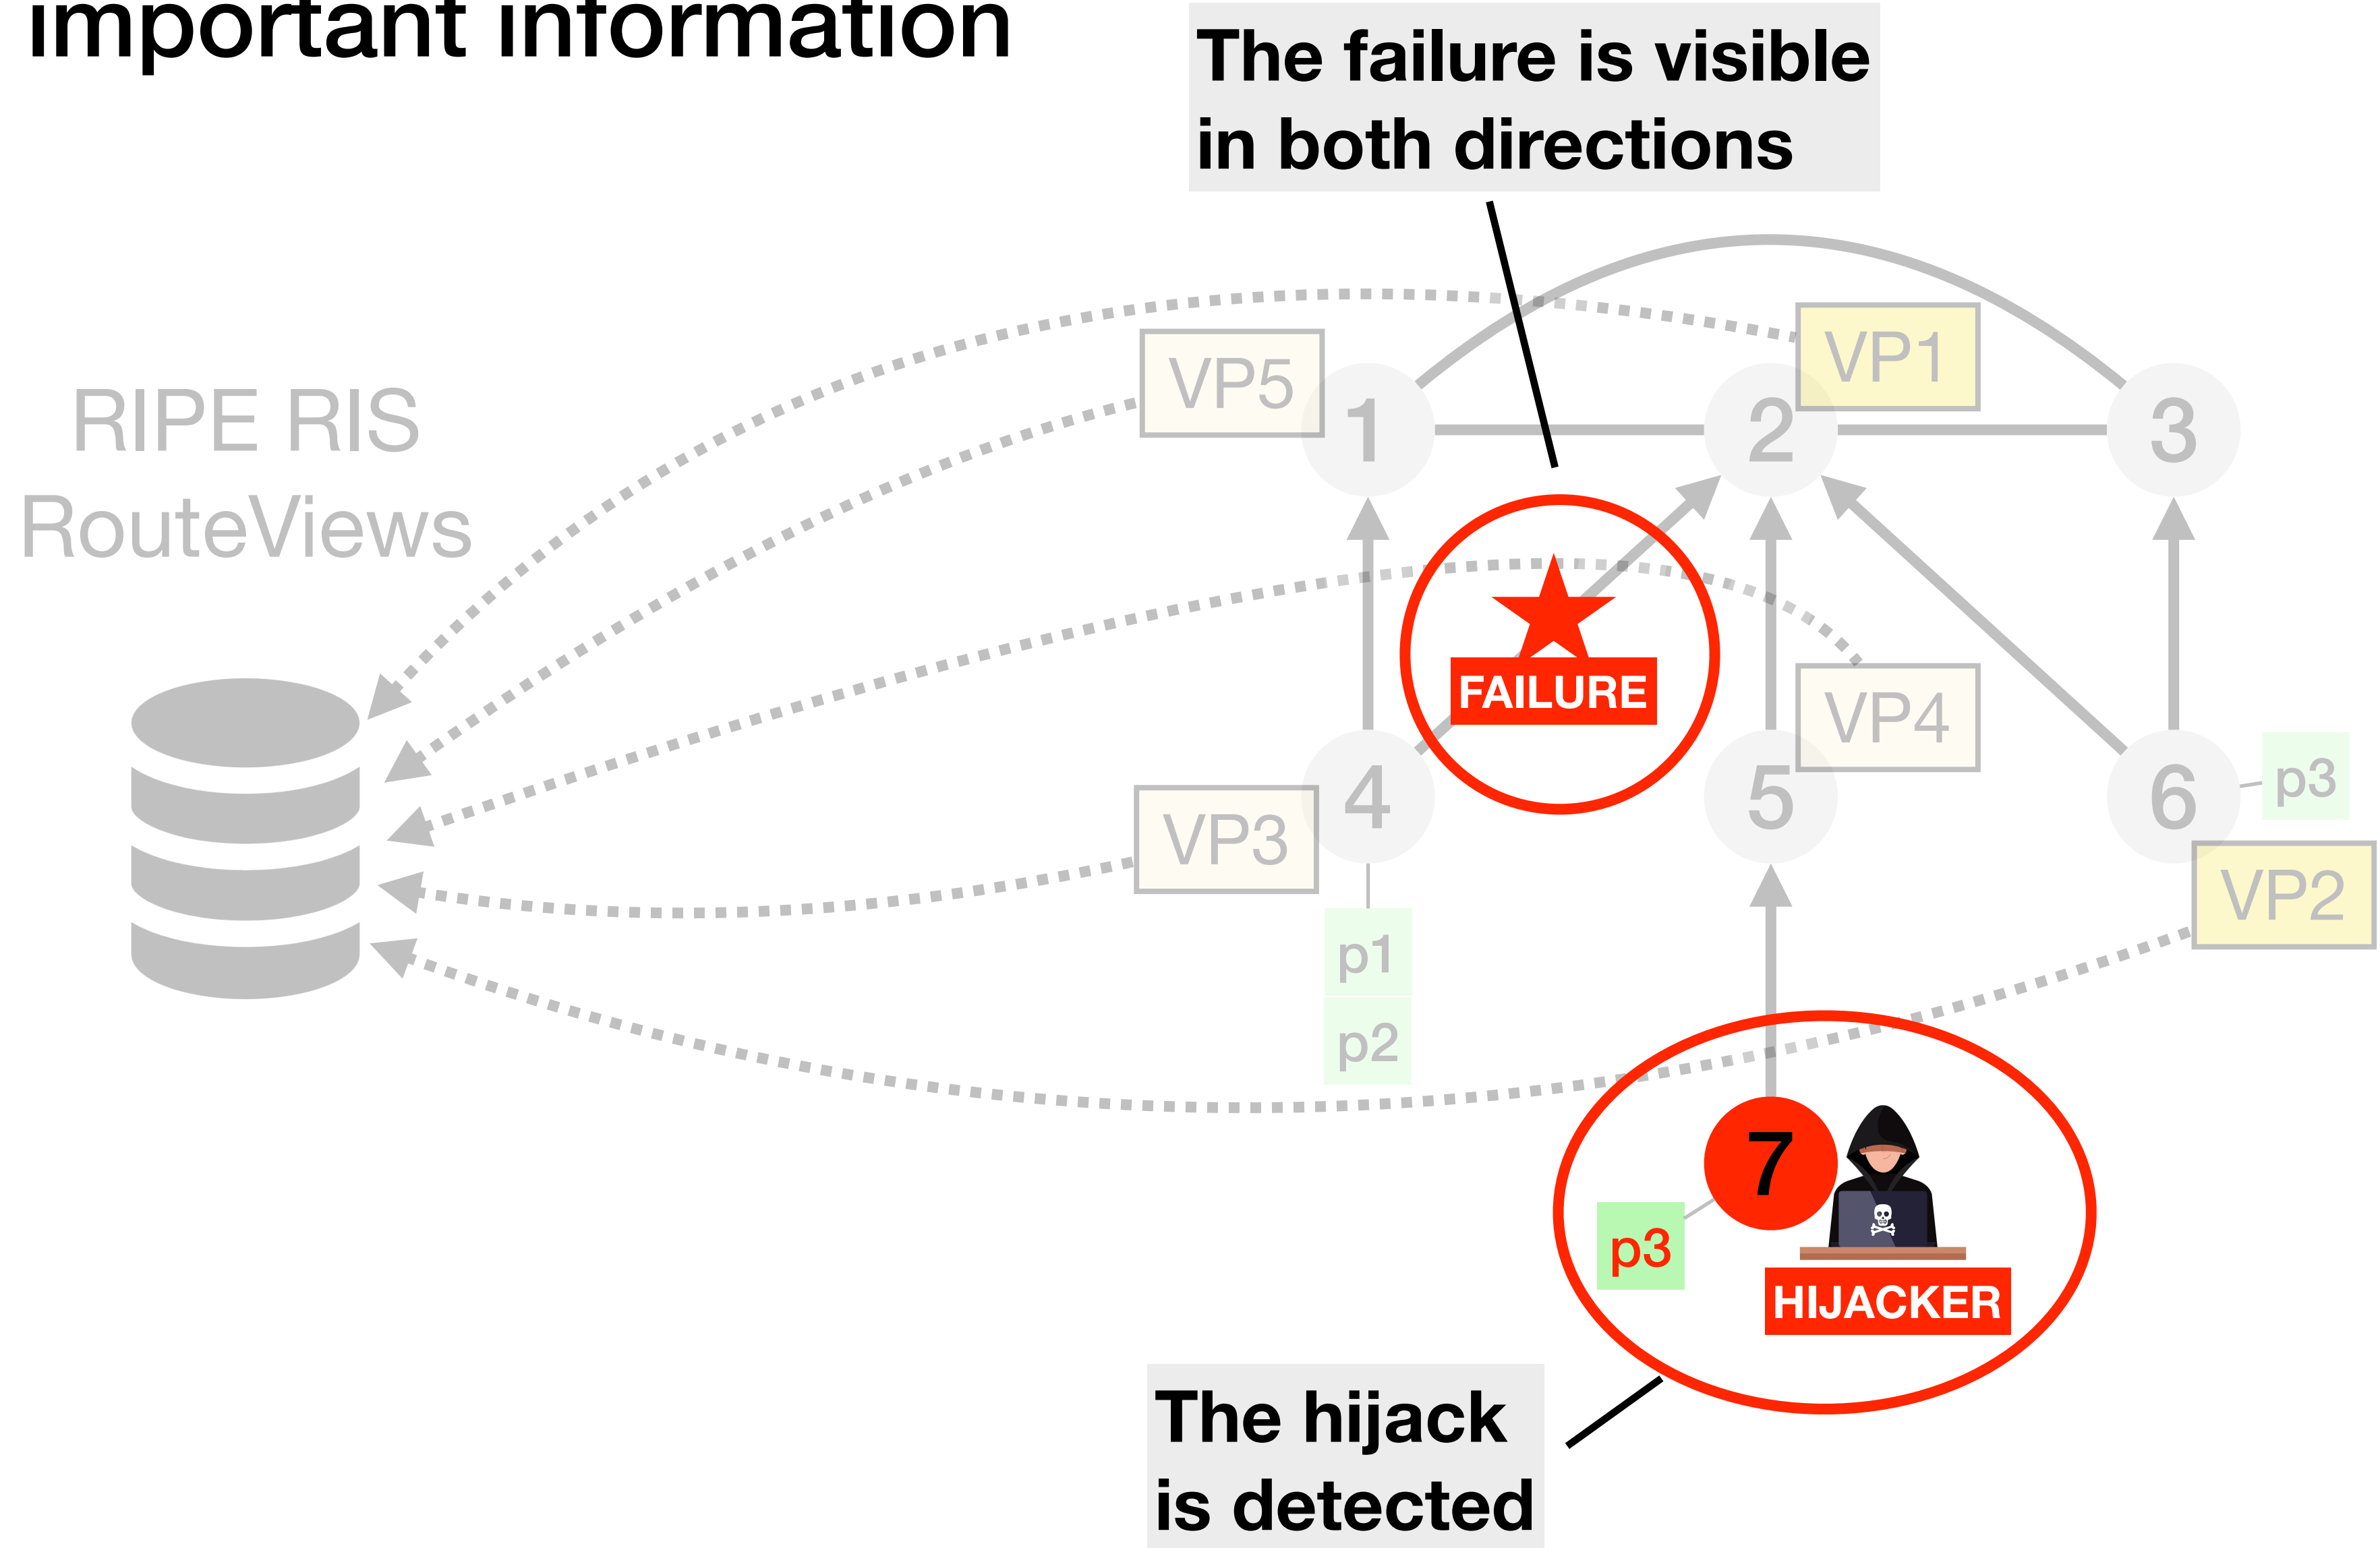
VP	prefix	AS path
VP1	p1	2 1 4
VP1	p2	2 1 4
VP2	p1	6 2 1 4
VP2	p2	6 <b>2 1 4</b>
VP3	p3	<b>4 1 2 6</b>
VP4	p3	5 7



# 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 <b>2 1 4</b>
VP3	p3	<b>4 1 2 6</b>
VP4	<b>p3</b>	<b>5 7</b>





# The “overshoot-and-discard” data collection paradigm can be adapted to BGP data collection



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

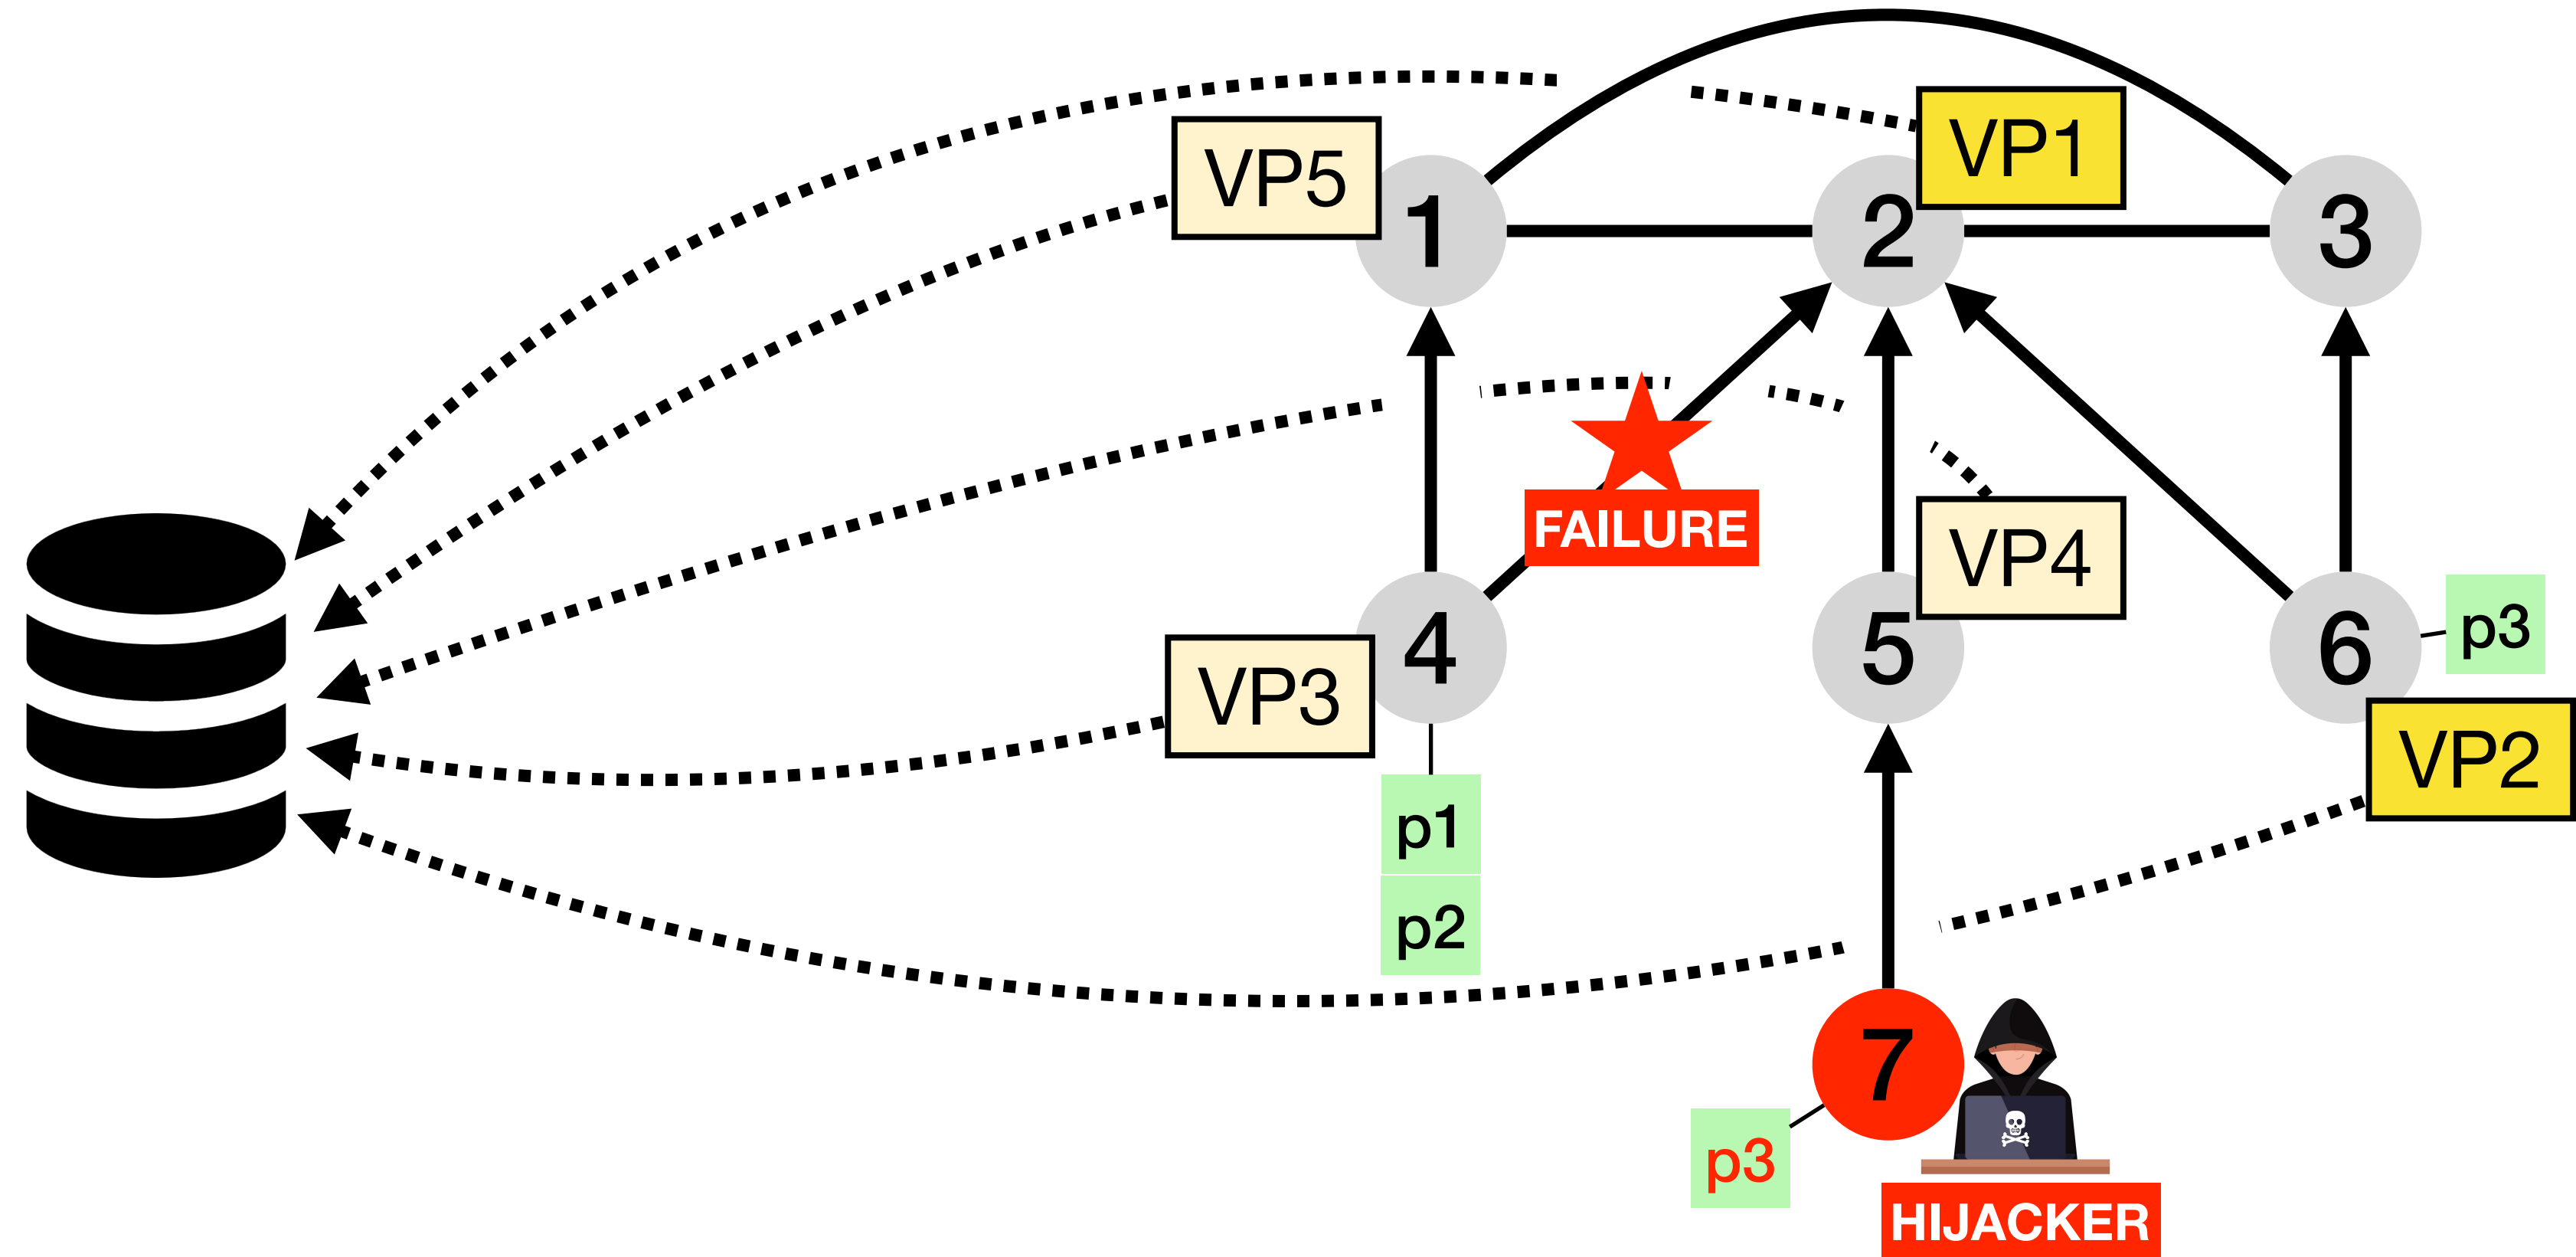


Discard: We filter out the redundant BGP routes  
*To reduce the volume of data collected*

# 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 <b>2 1 4</b>
VP3	p3	<b>4 1 2 6</b>
VP4	<b>p3</b>	<b>5 7</b>



# Discard: redundant BGP routes are discarded using filters

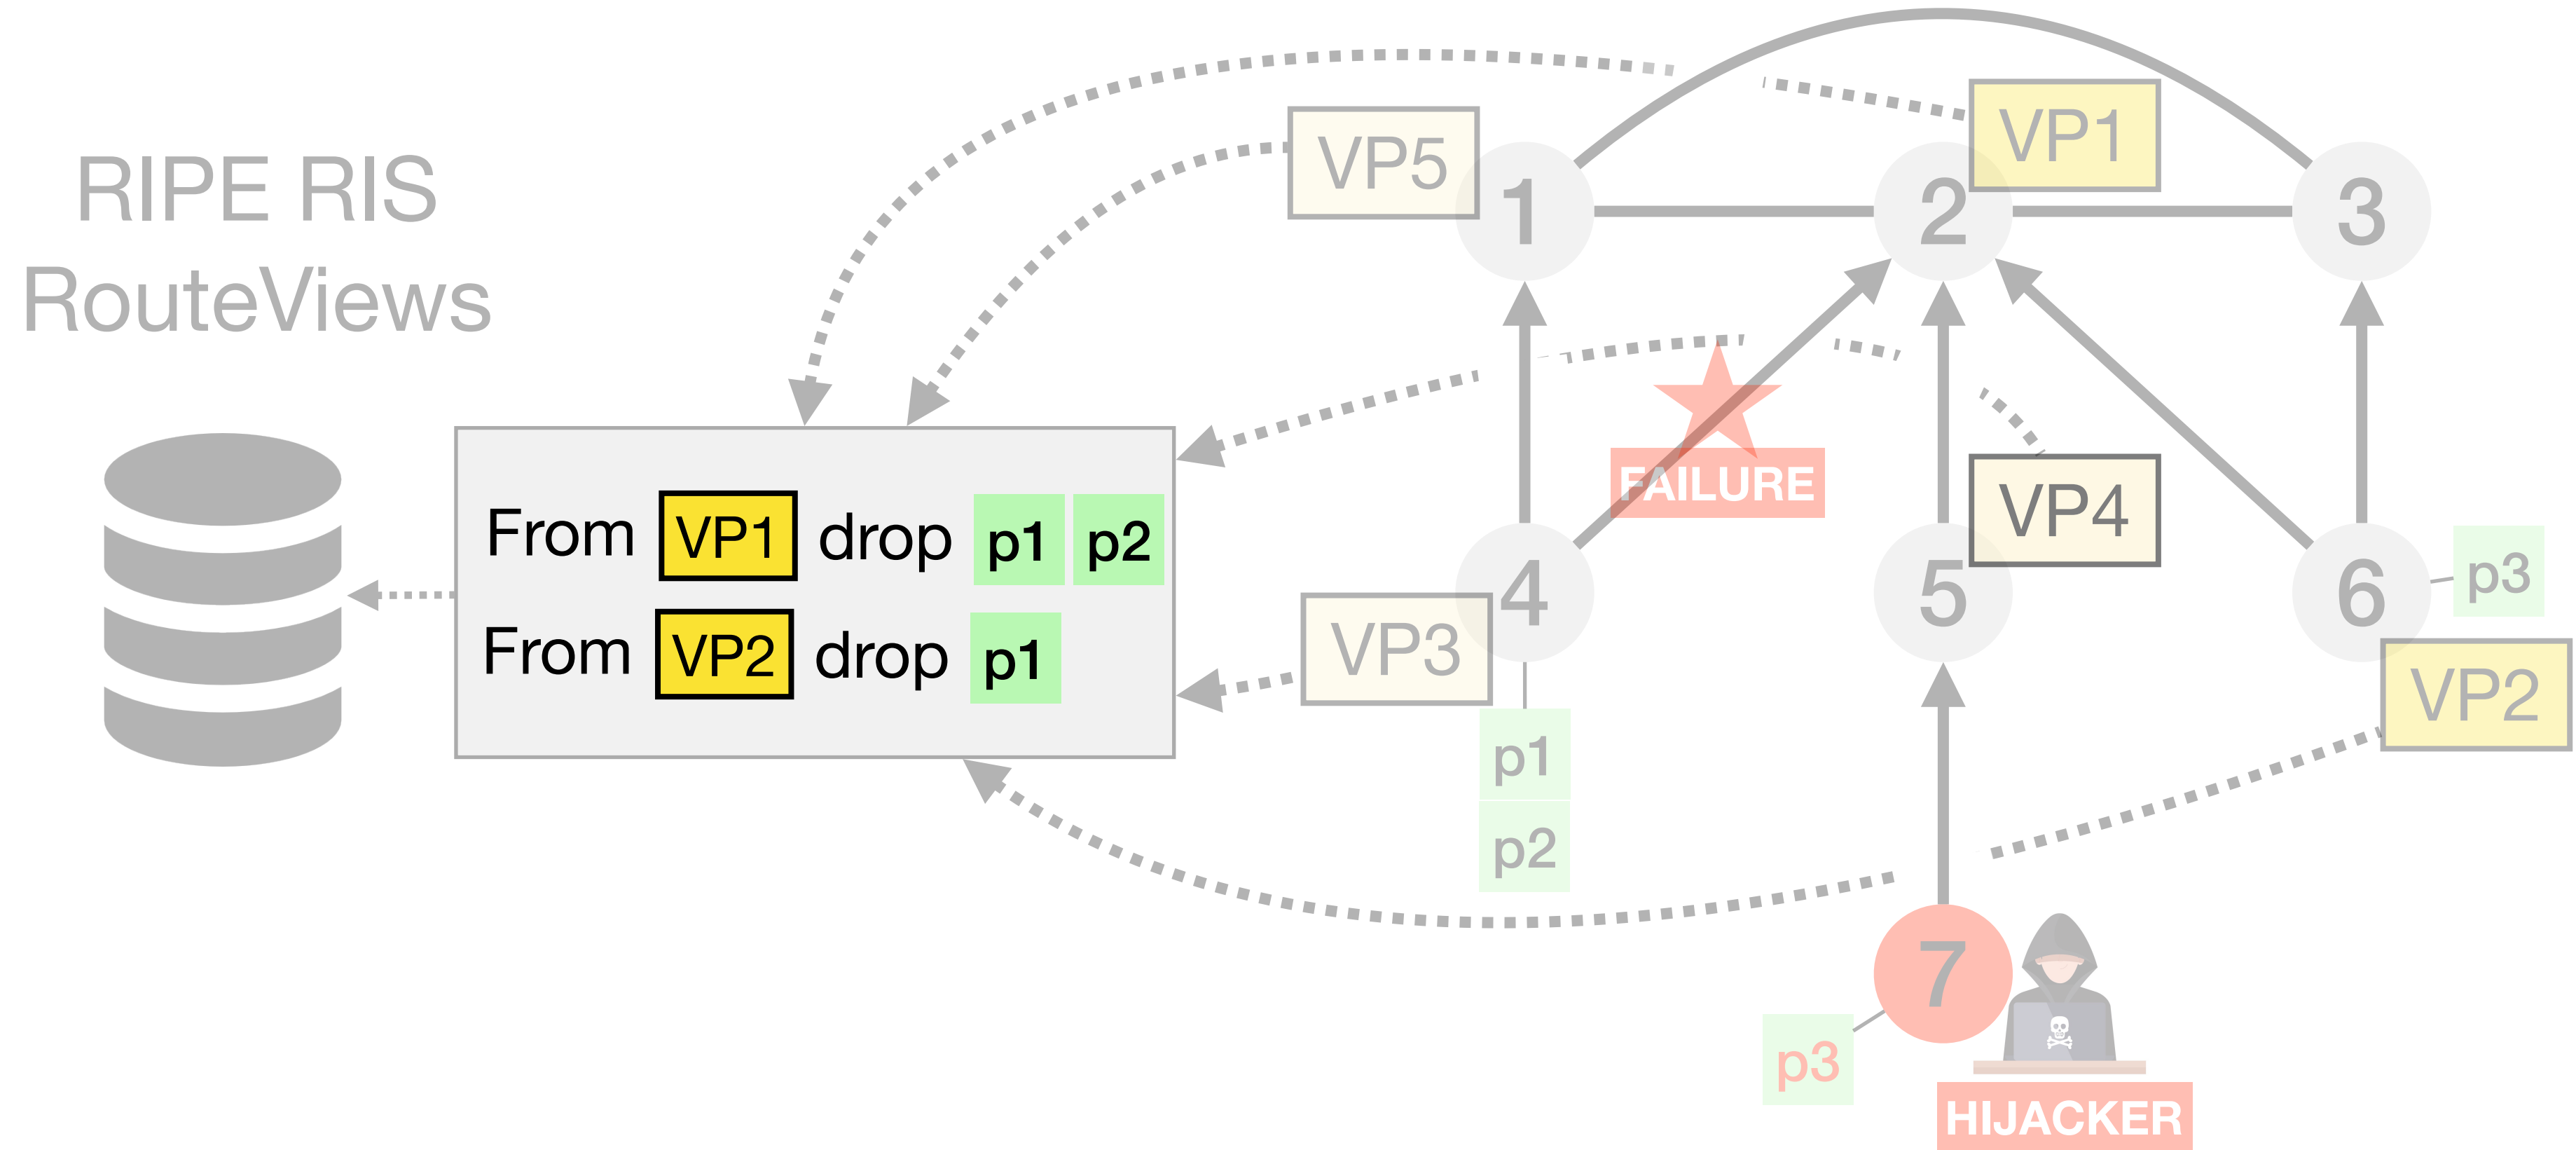
Collected routes

VP	prefix	AS path
<del>VP1</del>	<del>p1</del>	<del>2 1 4</del>
<del>VP1</del>	<del>p2</del>	<del>2 1 4</del>
<del>VP2</del>	<del>p1</del>	<del>3 2 1 4</del>
VP2	p2	6 <b>2 1 4</b>
VP3	p3	<b>4 1 2 6</b>
VP4	<b>p3</b>	<b>5 7</b>

RIPE RIS  
RouteViews



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



# Discard: redundant BGP routes are discarded using filters

Collected routes

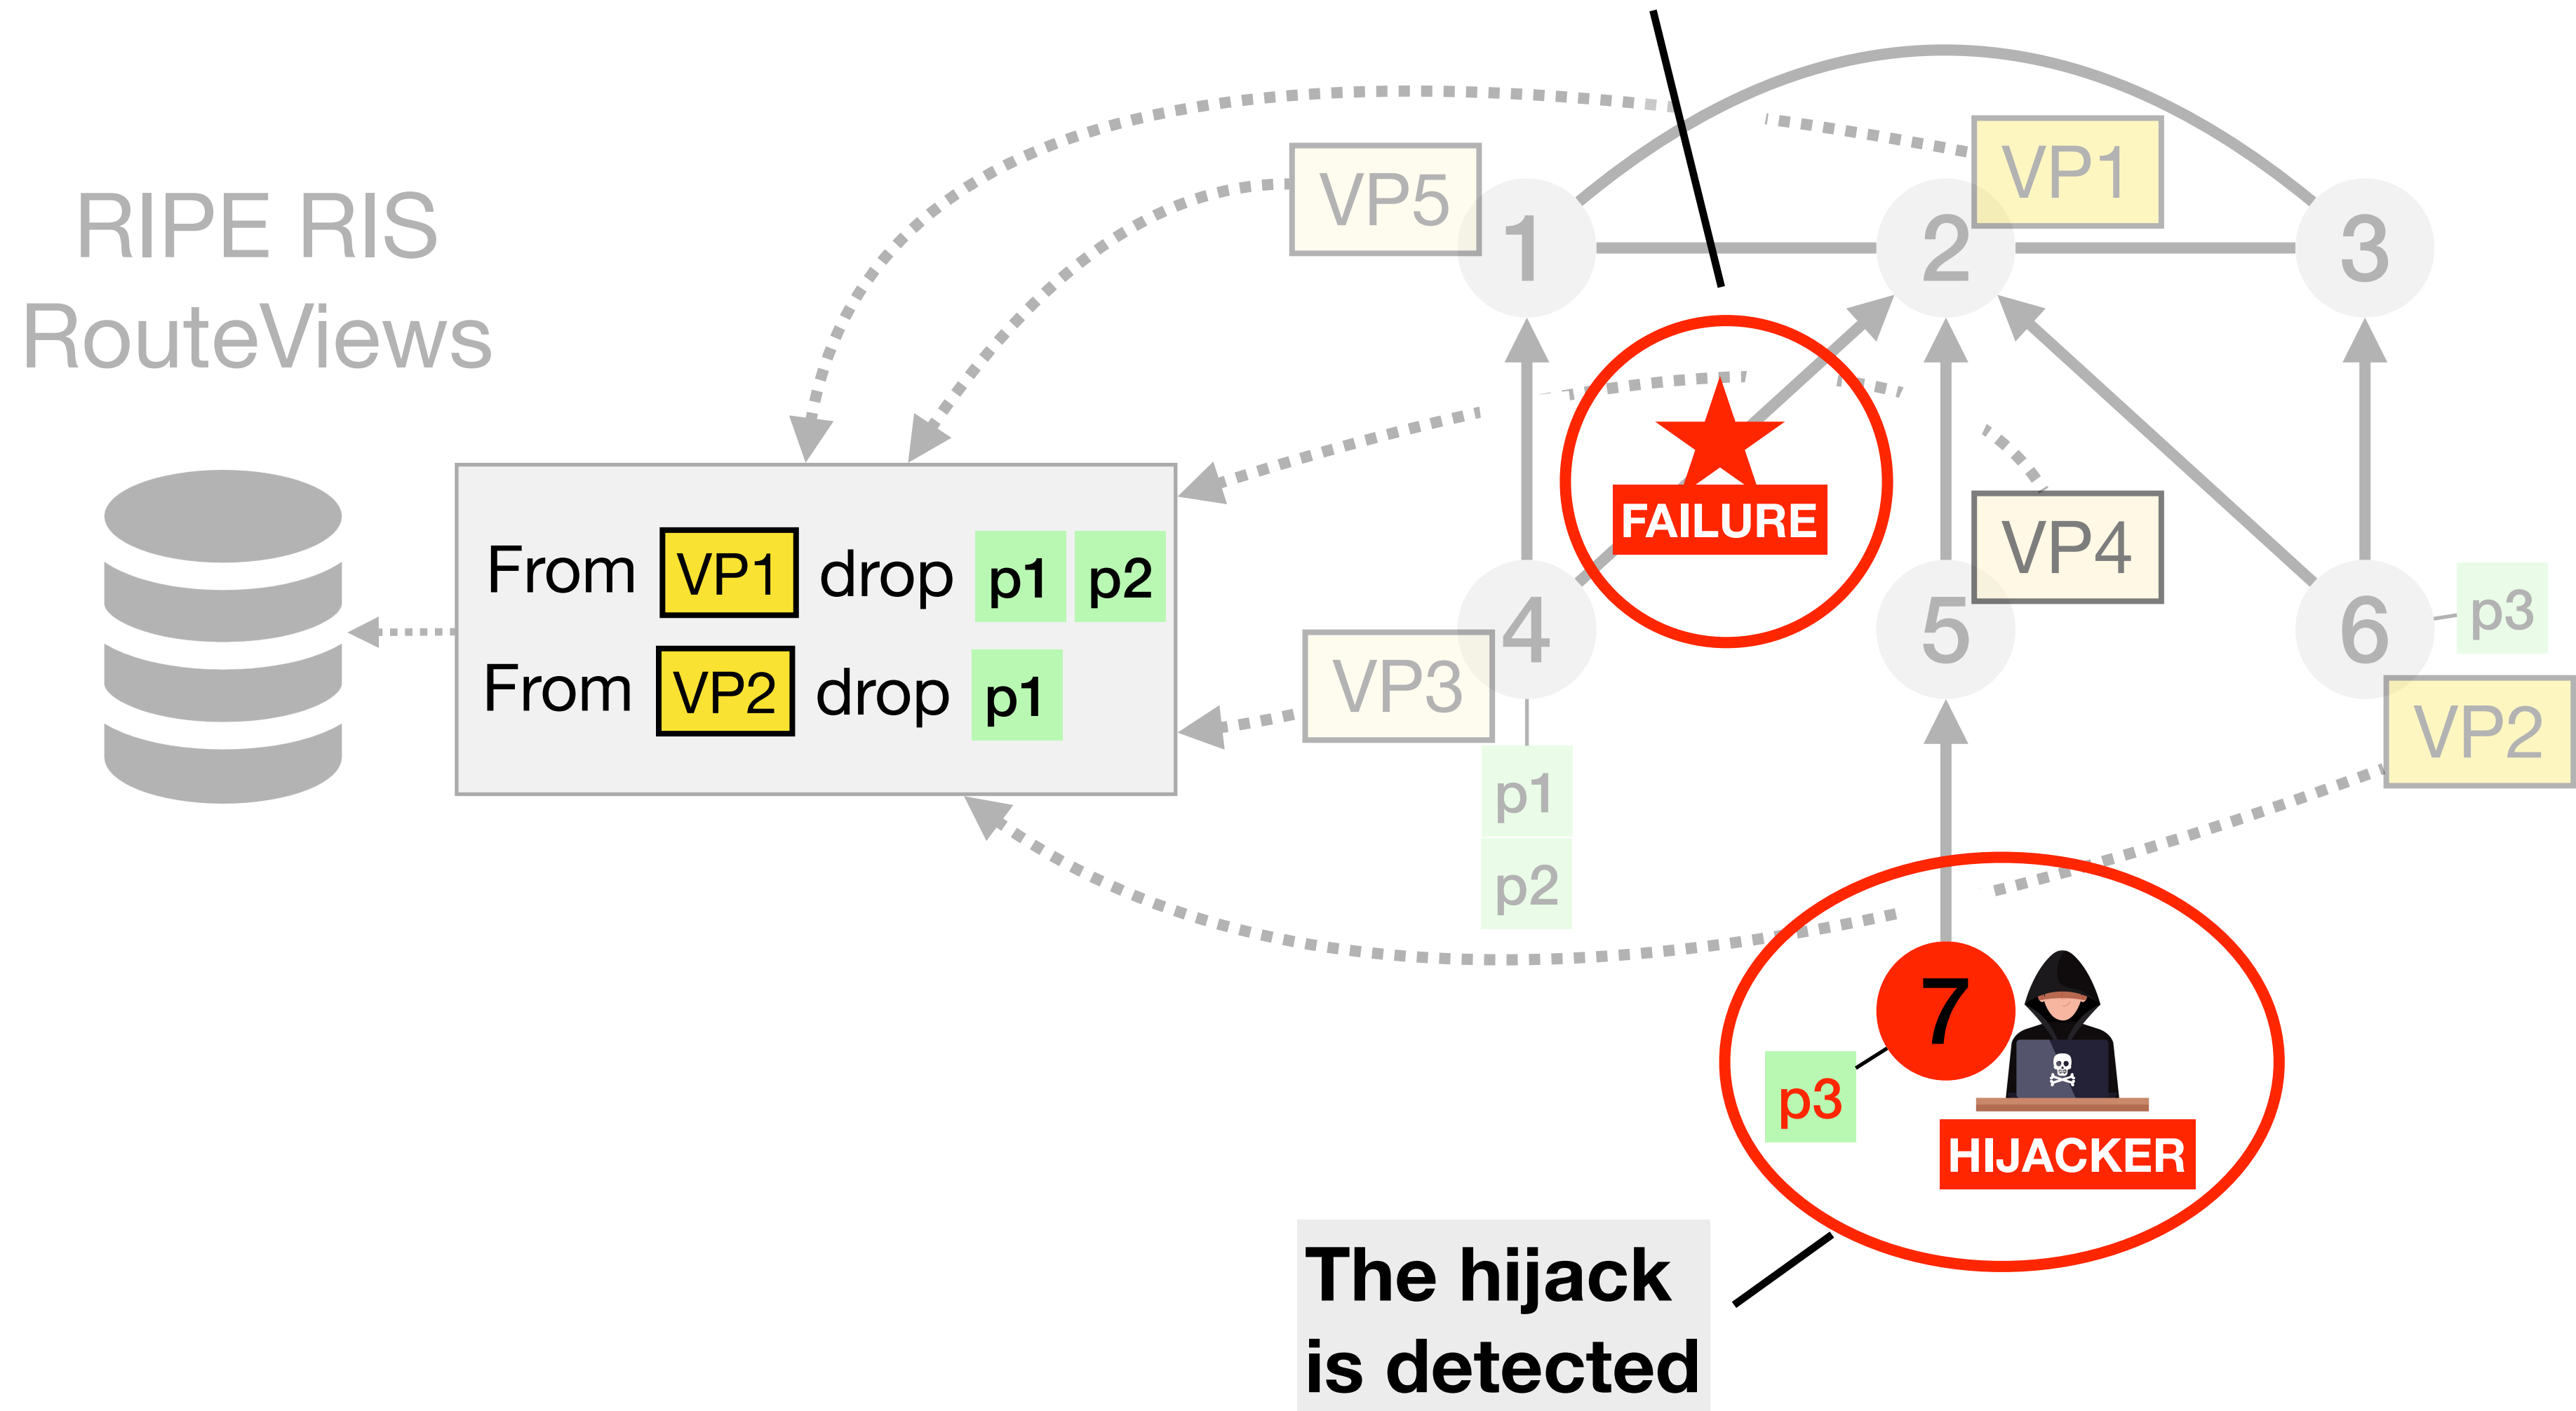
VP	prefix	AS path
<del>VP1</del>	<del>p1</del>	<del>2 1 4</del>
<del>VP1</del>	<del>p2</del>	<del>2 1 4</del>
<del>VP2</del>	<del>p1</del>	<del>3 2 1 4</del>
VP2	p2	6 <b>2 1 4</b>
VP3	p3	<b>4 1 2 6</b>
VP4	<b>p3</b>	<b>5 7</b>

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 system that measures redundancy between BGP routes and generates filters that discard redundant routes

***GILL*** selects the updates to retain  
using a new metric called the **Reconstitution Power (RP)**



***GILL*** selects the updates to retain  
using a new metric called the **Reconstitution Power (RP)**

Original set  
of updates



***GILL*** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**

Original set of updates





Sampled update








***GILL*** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**


Original set of updates



Sampled update





Reconstituted Updates






***GILL*** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**




Original set of updates



Sampled update



Reconstituted Updates

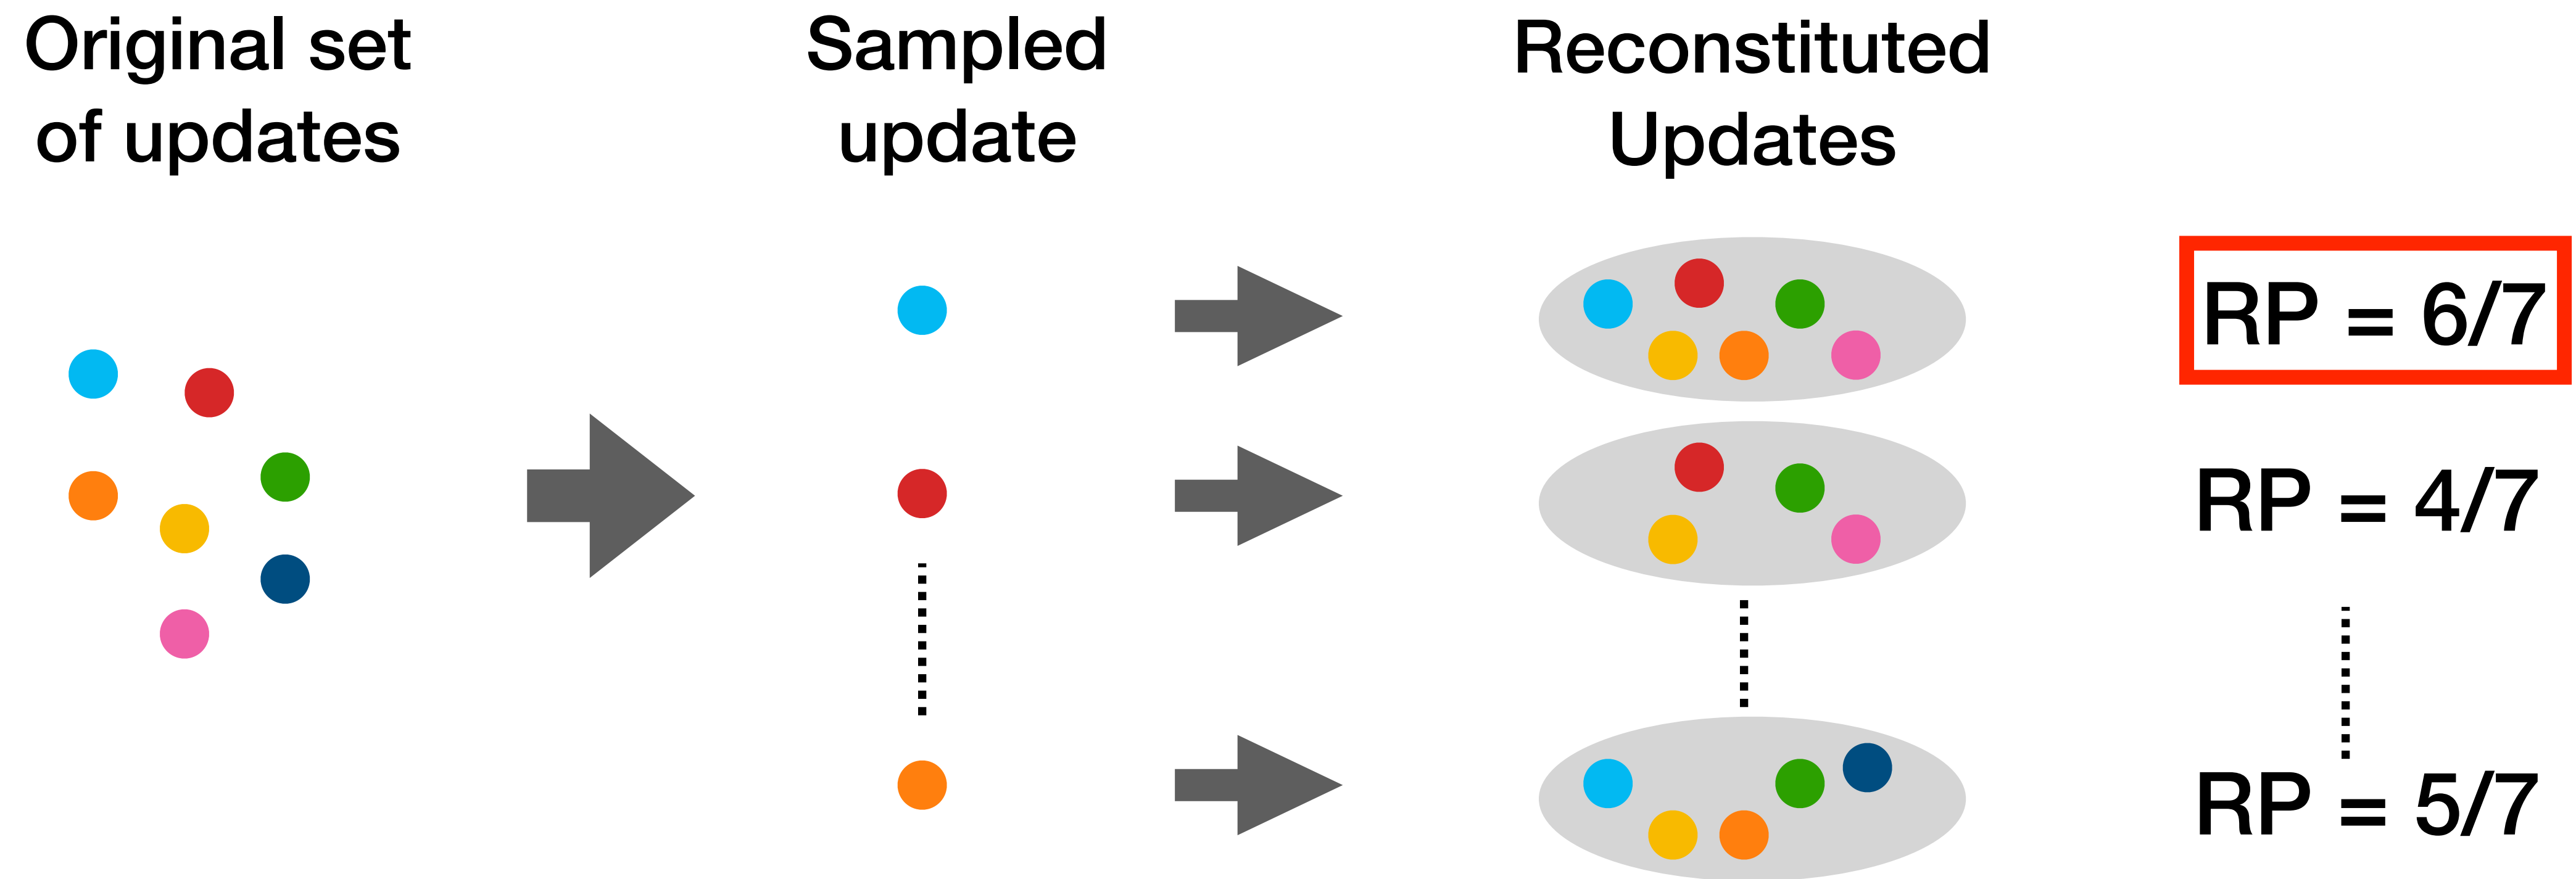


$$RP = 6/7$$

$$RP = 4/7$$

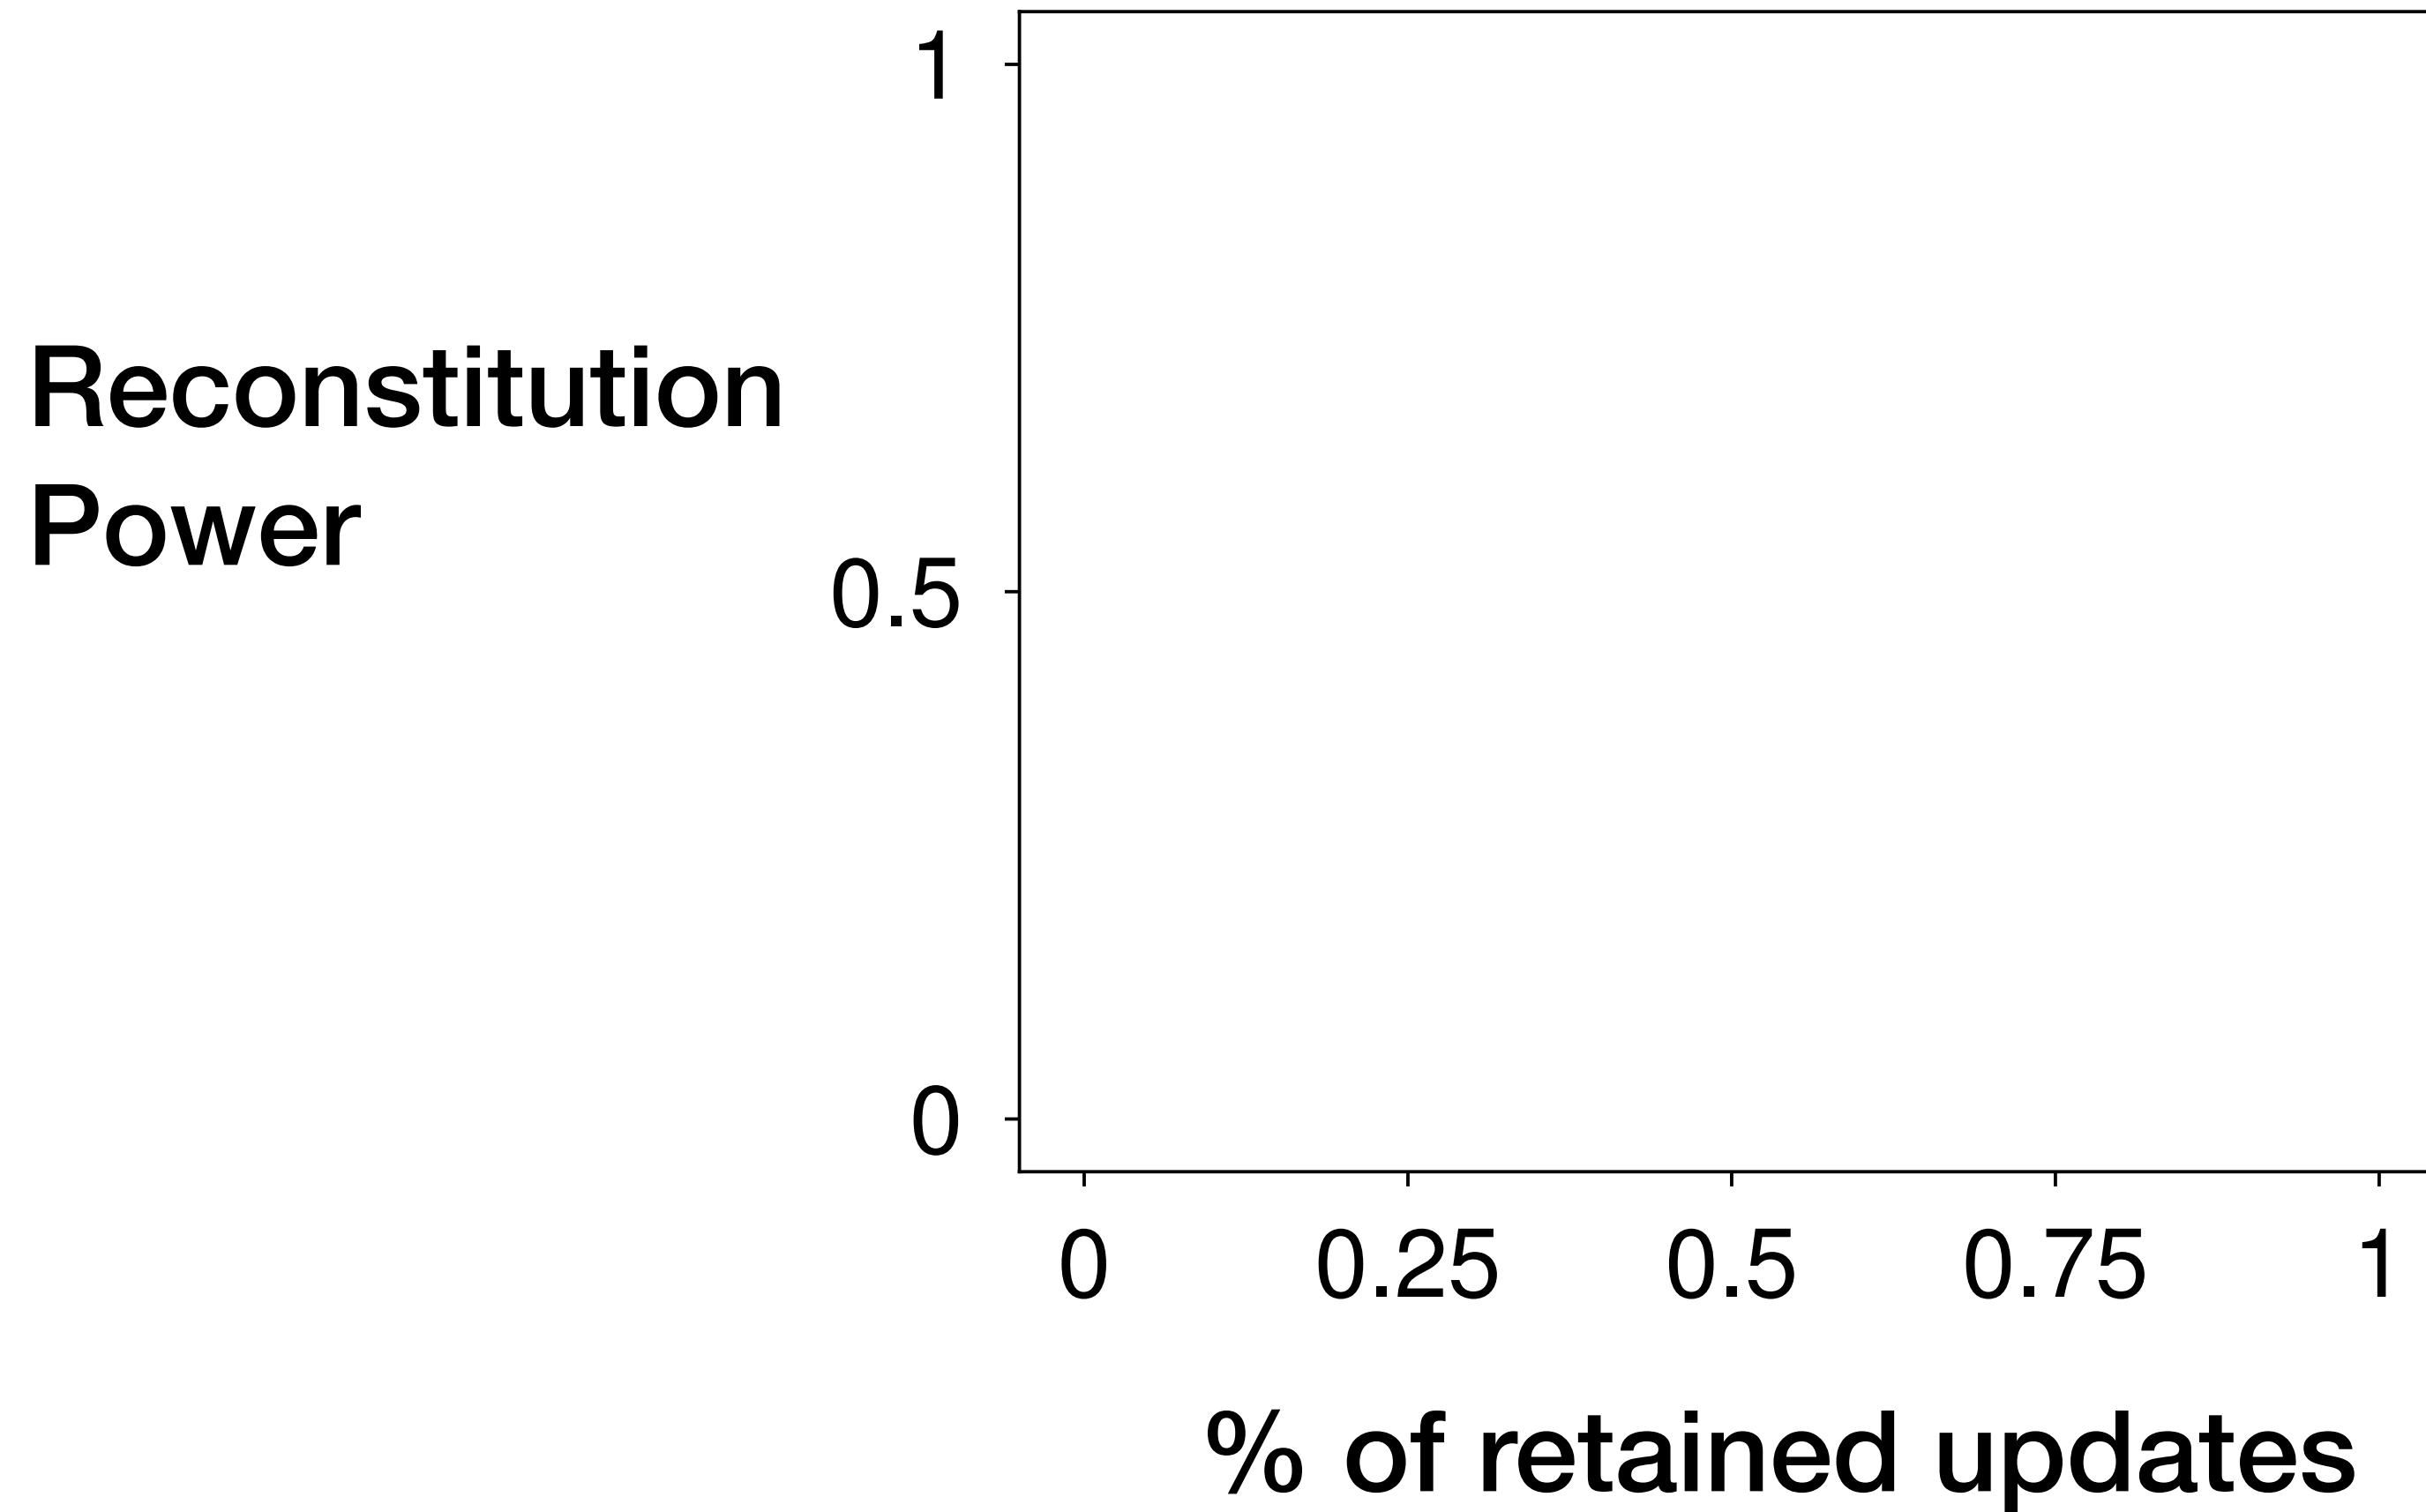
$$RP = 5/7$$

**GILL** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**




**GILL** iteratively selects the updates with the **higher** Reconstitution Power

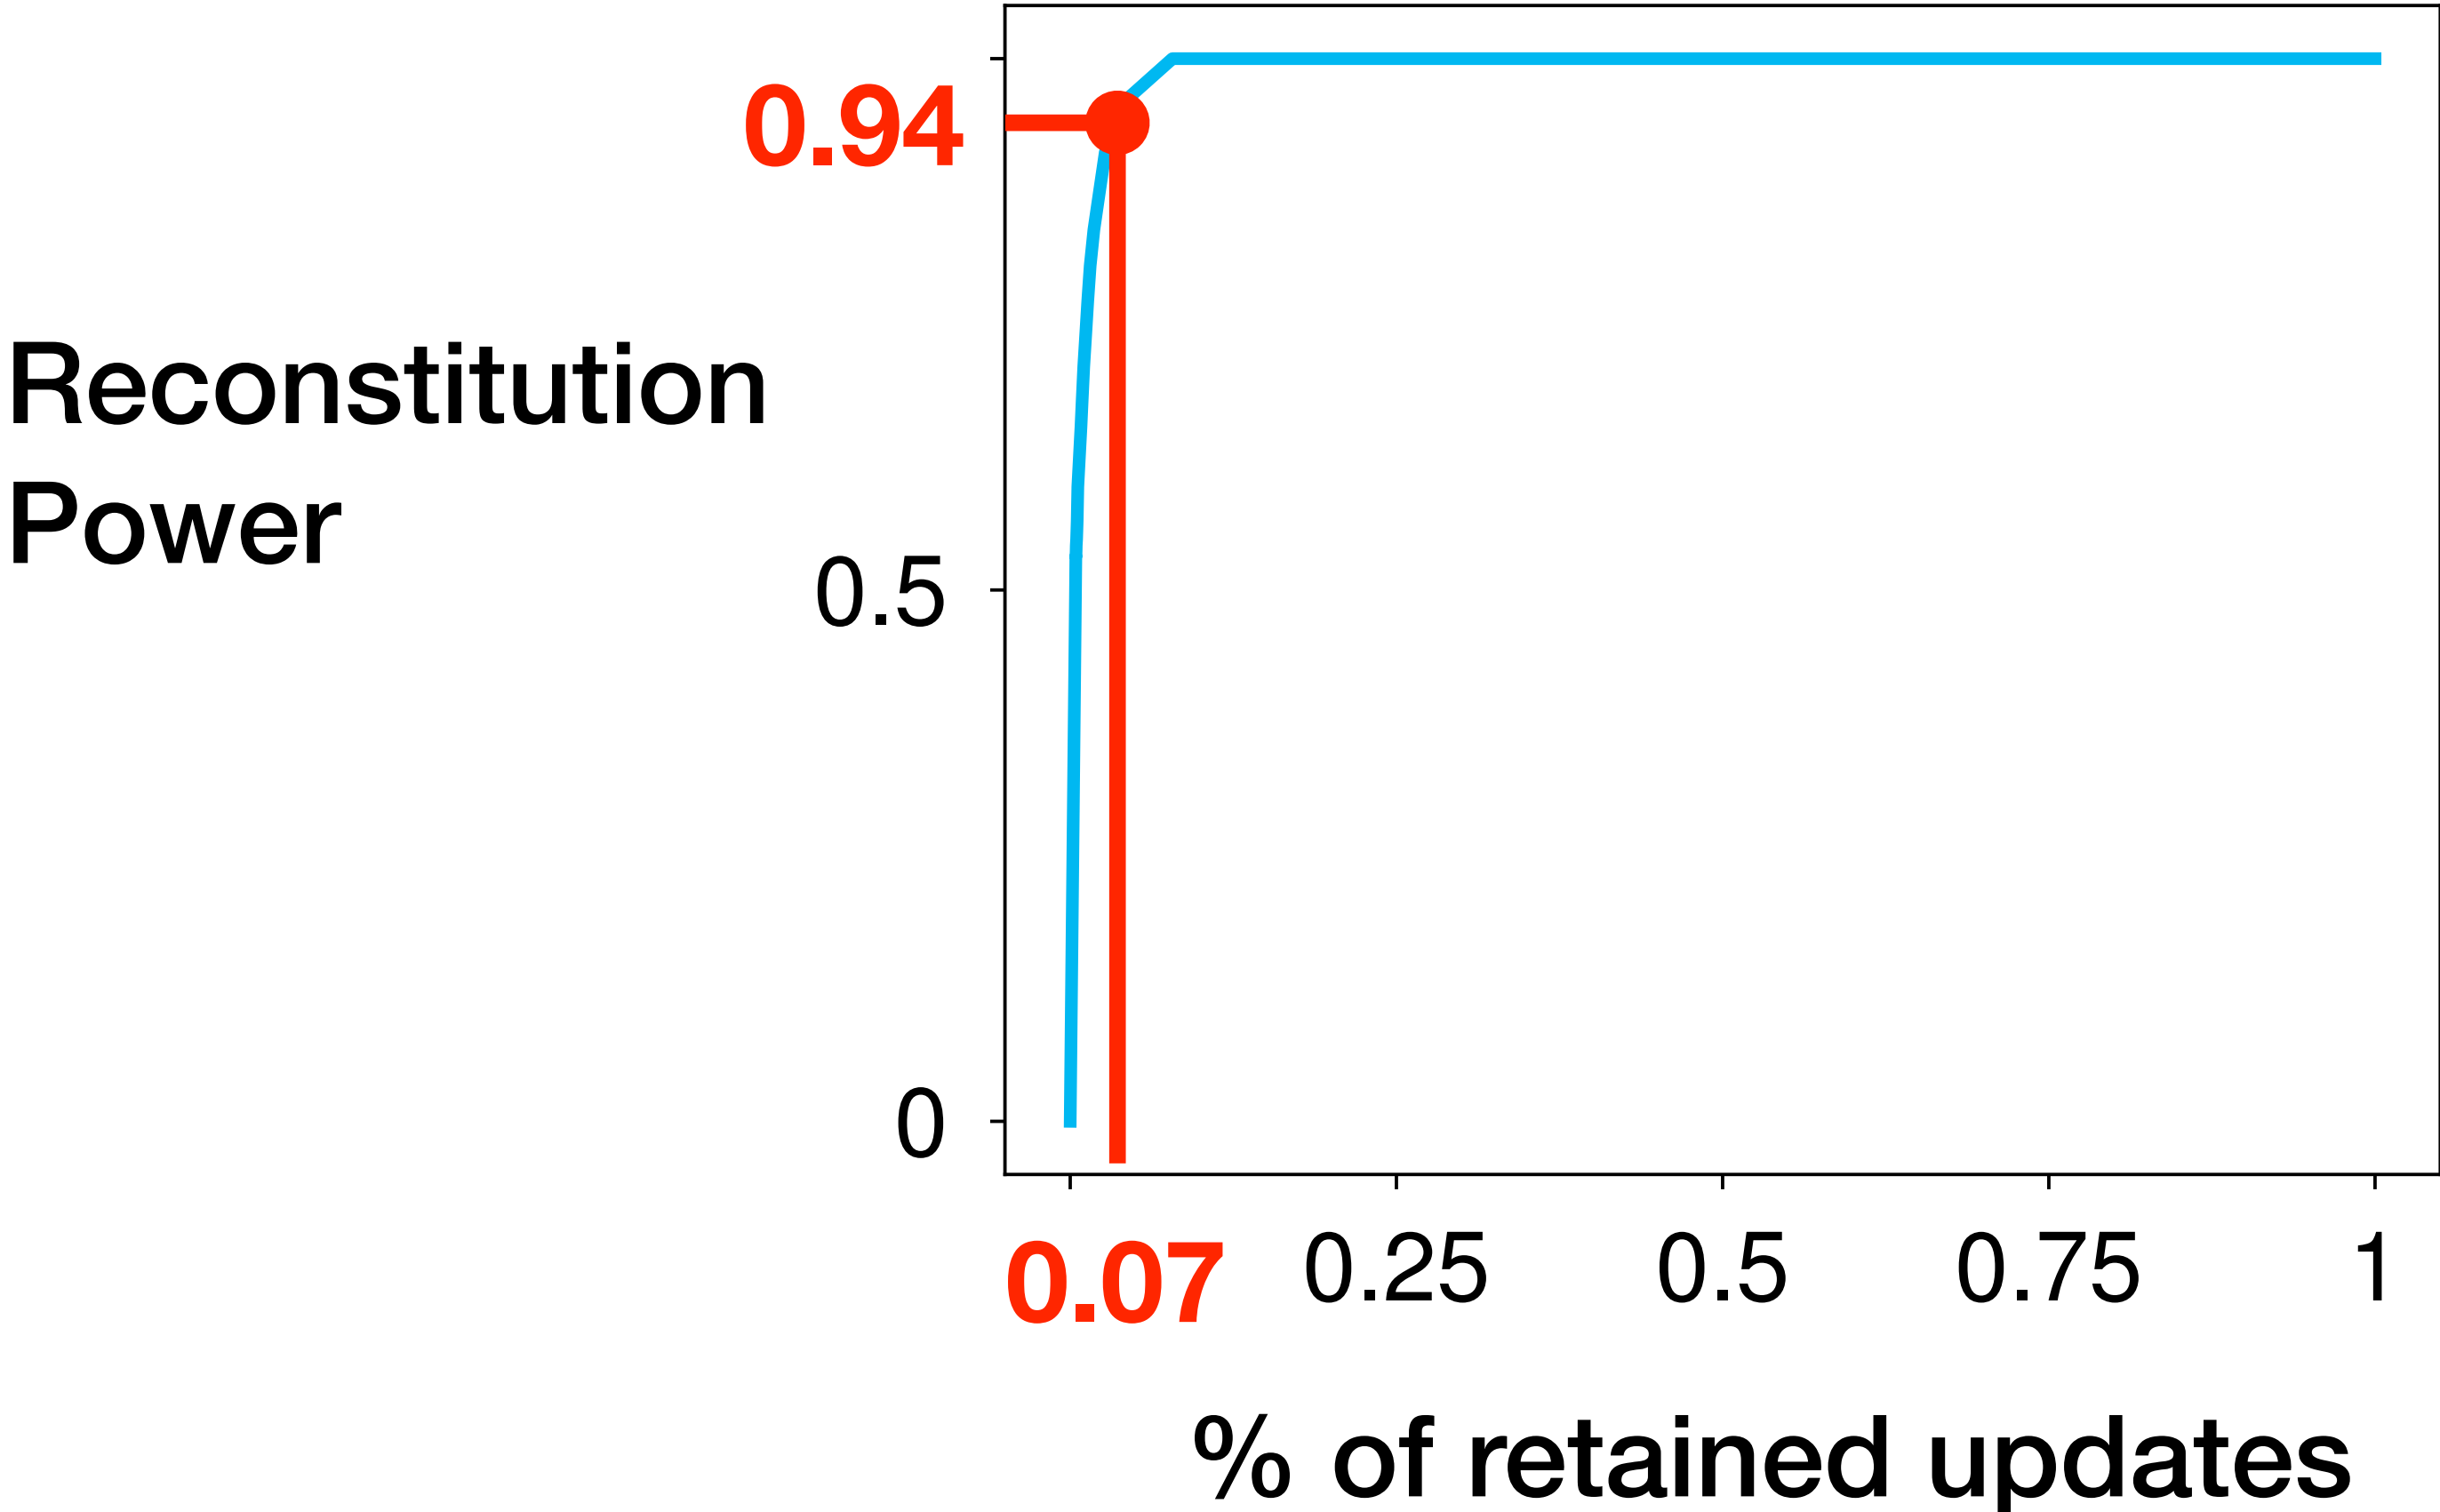
***GILL*** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**



***GILL*** selects the updates to retain using a new metric called the **Reconstitution Power (RP)**



*GILL* selects the updates to retain using a new metric called the **Reconstitution Power (RP)**




We can reconstitute **94%** of the updates from **7%** of them

***GILL* builds filters that discriminate  
retained updates from redundant updates**

***GILL*** builds filters that discriminate  
retained updates from redundant updates



Original set  
of updates







# *GILL* builds filters that discriminate retained updates from redundant updates

Original set of updates





Retained updates




Redundant updates

# *GILL* builds filters that discriminate retained updates from redundant updates


Original set of updates



Retained updates





**F I L T E R S**




Redundant updates

# *GILL* builds filters that discriminate retained updates from redundant updates

Original set of updates




Retained updates



**F I L T E R S**



Redundant updates




How much information do we lose?

# *GILL* builds filters that discriminate retained updates from redundant updates

Original set of updates




Retained updates



**FILTERS**

Redundant updates



How much information do we lose?  
**Not much!**

# Outline

1. We observe that BGP routes are often redundant

2. Redundant BGP routes enable an overshoot-and-discard collection scheme

3. *GILL*: a system that measures redundancy between BGP routes and generates filters that discard redundant routes

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	2%	0%	X
<i>GILL</i>			

Results from simulations  
on “mini” Internets with 6k 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	2%	0%	X	20%	37%	73%
<i>GILL</i>						

Results from simulations  
on “mini” Internets with 6k 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)
<b>Current approach</b>	2%	0%	X ↑ Same number of updates	20%	37%	73%
<b><i>GILL</i></b>	50%	96%	X ↓			

Results from simulations on “mini” Internets with 6k 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)
<b>Current approach</b>	2%	0%	X ↑ Same number of updates	20%	37%	73%
<b>GILL</b>	50%	96%	X ↓	61%	80%	82%

Results from simulations on “mini” Internets with 6k ASes

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

<https://bgproutes.io/>



## 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



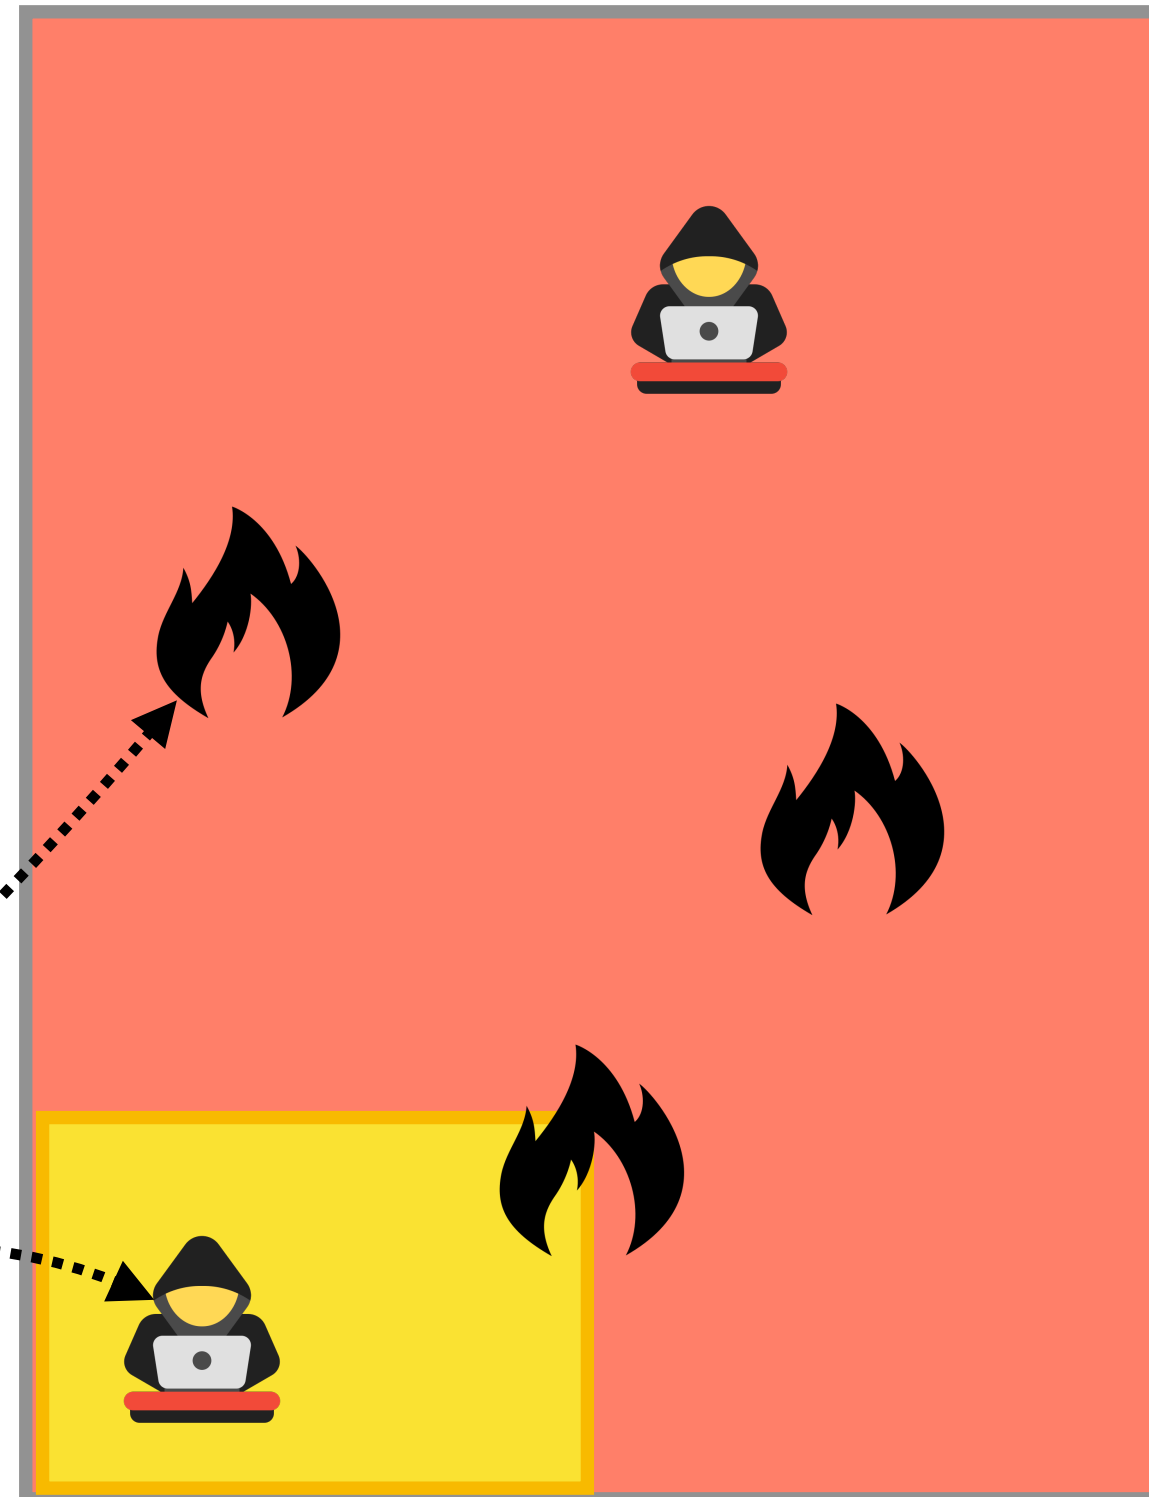
Collected BGP routes

Collected but discarded BGP routes

### Today

### Naive approach

### GILL



Useful bits of data



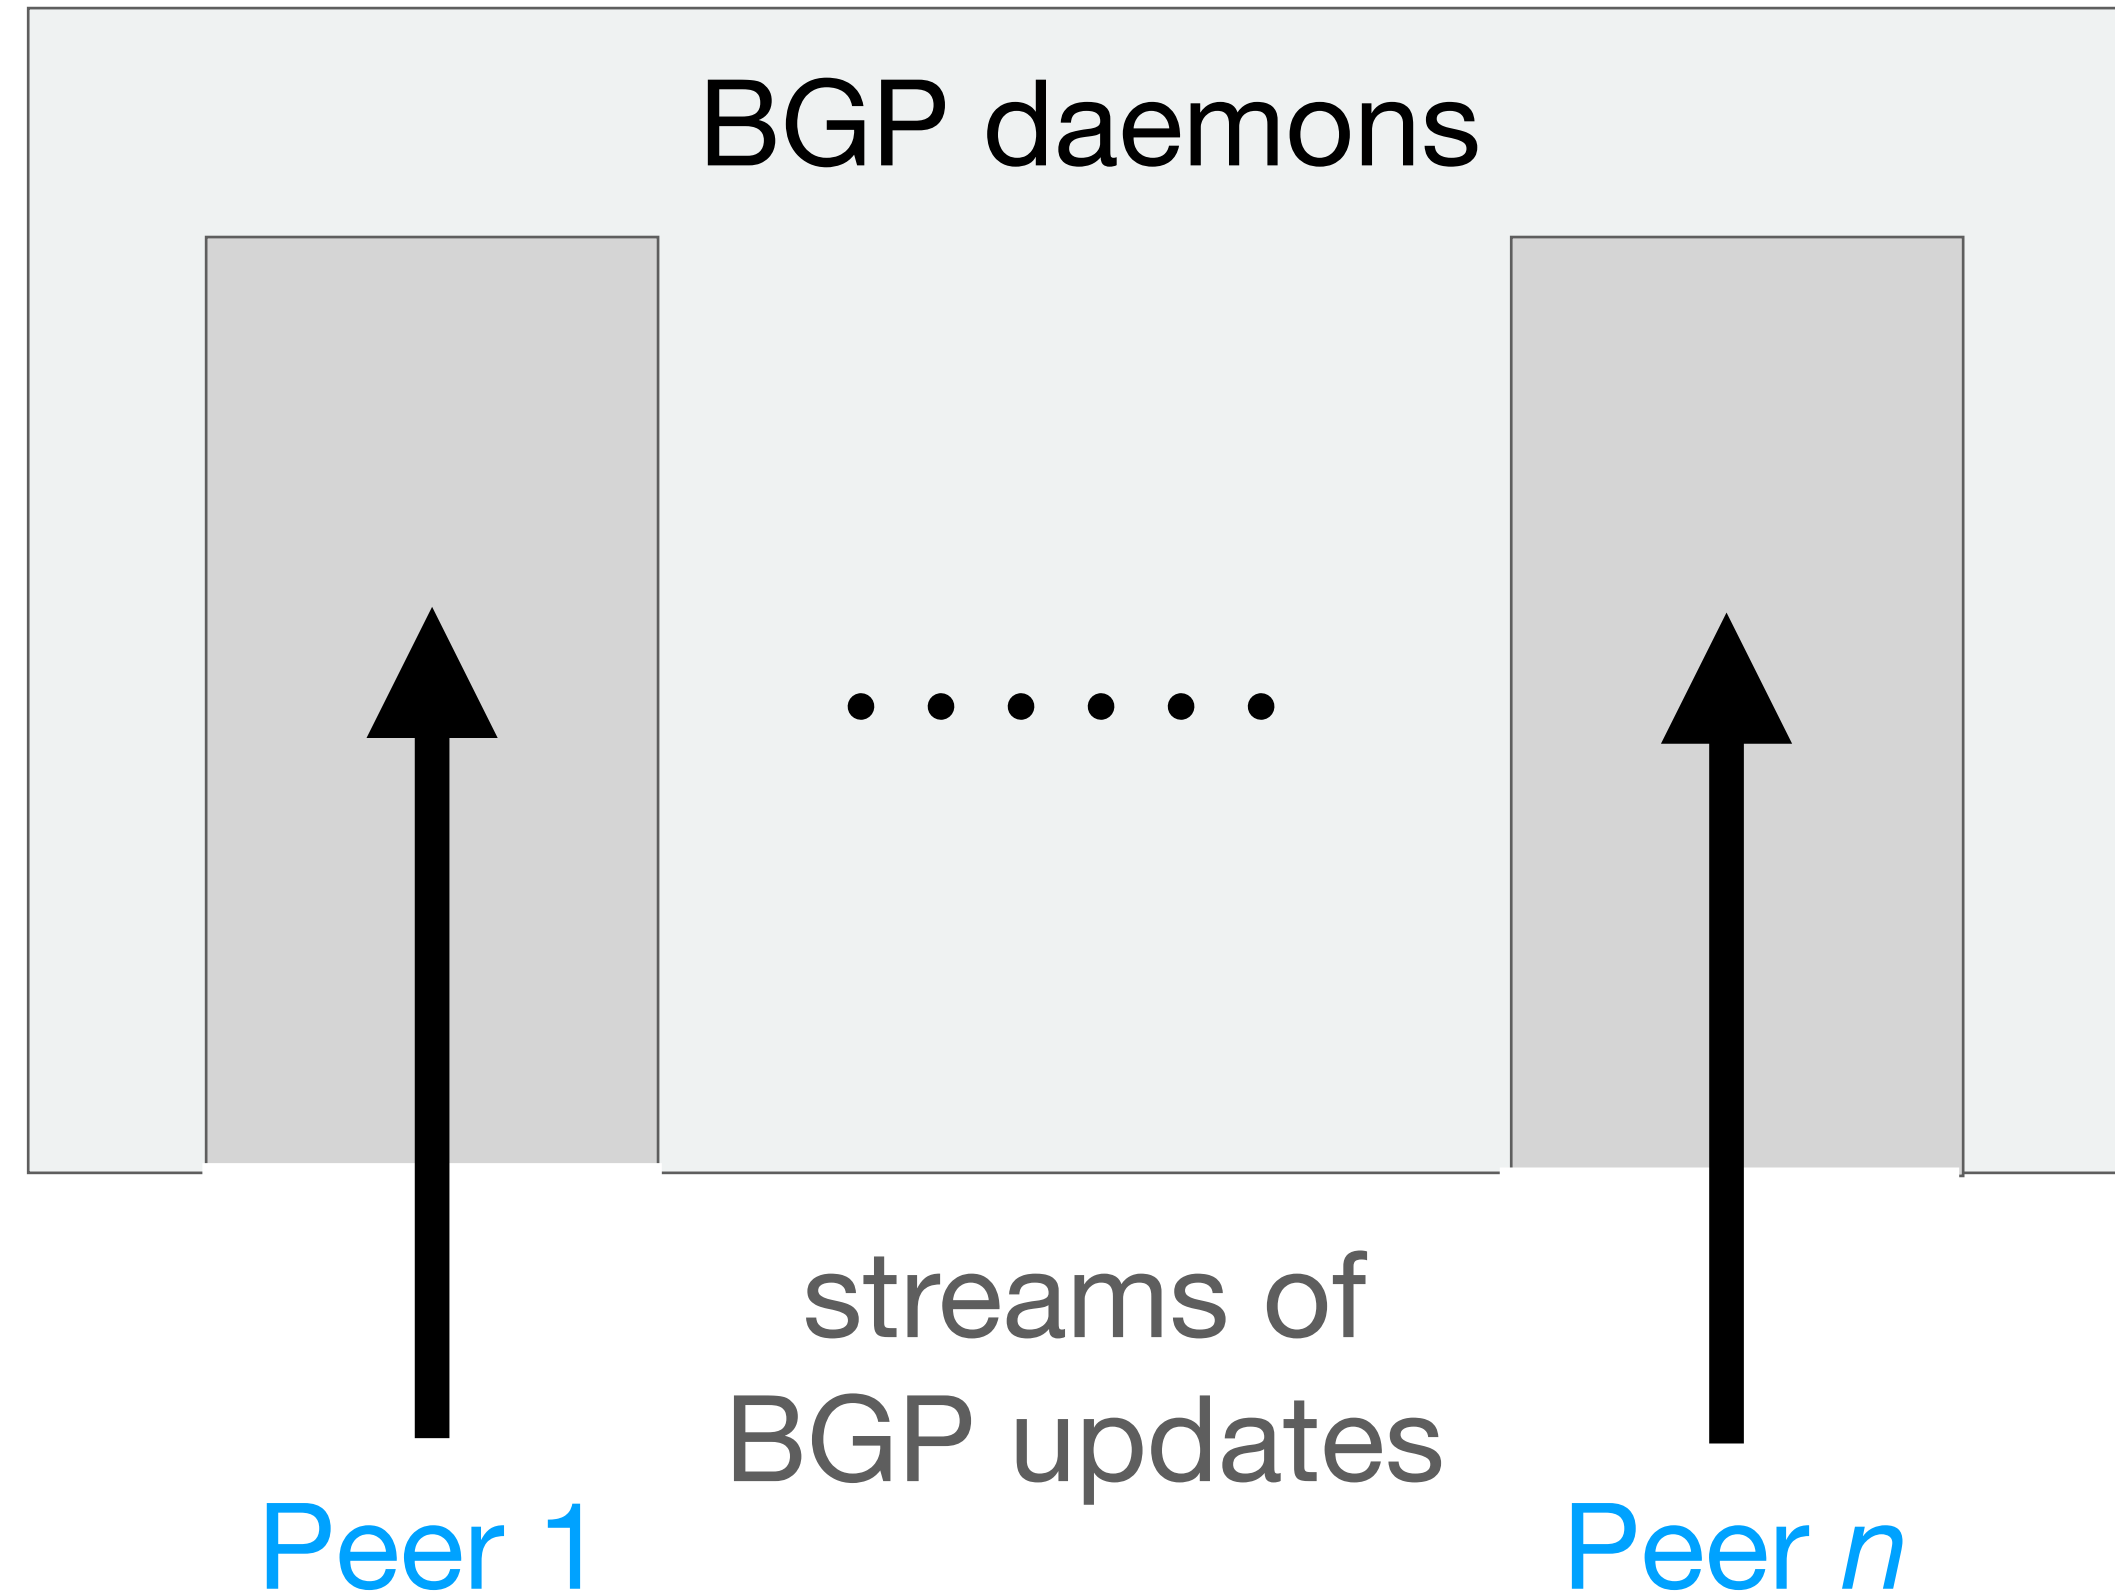
Useful bits are missed




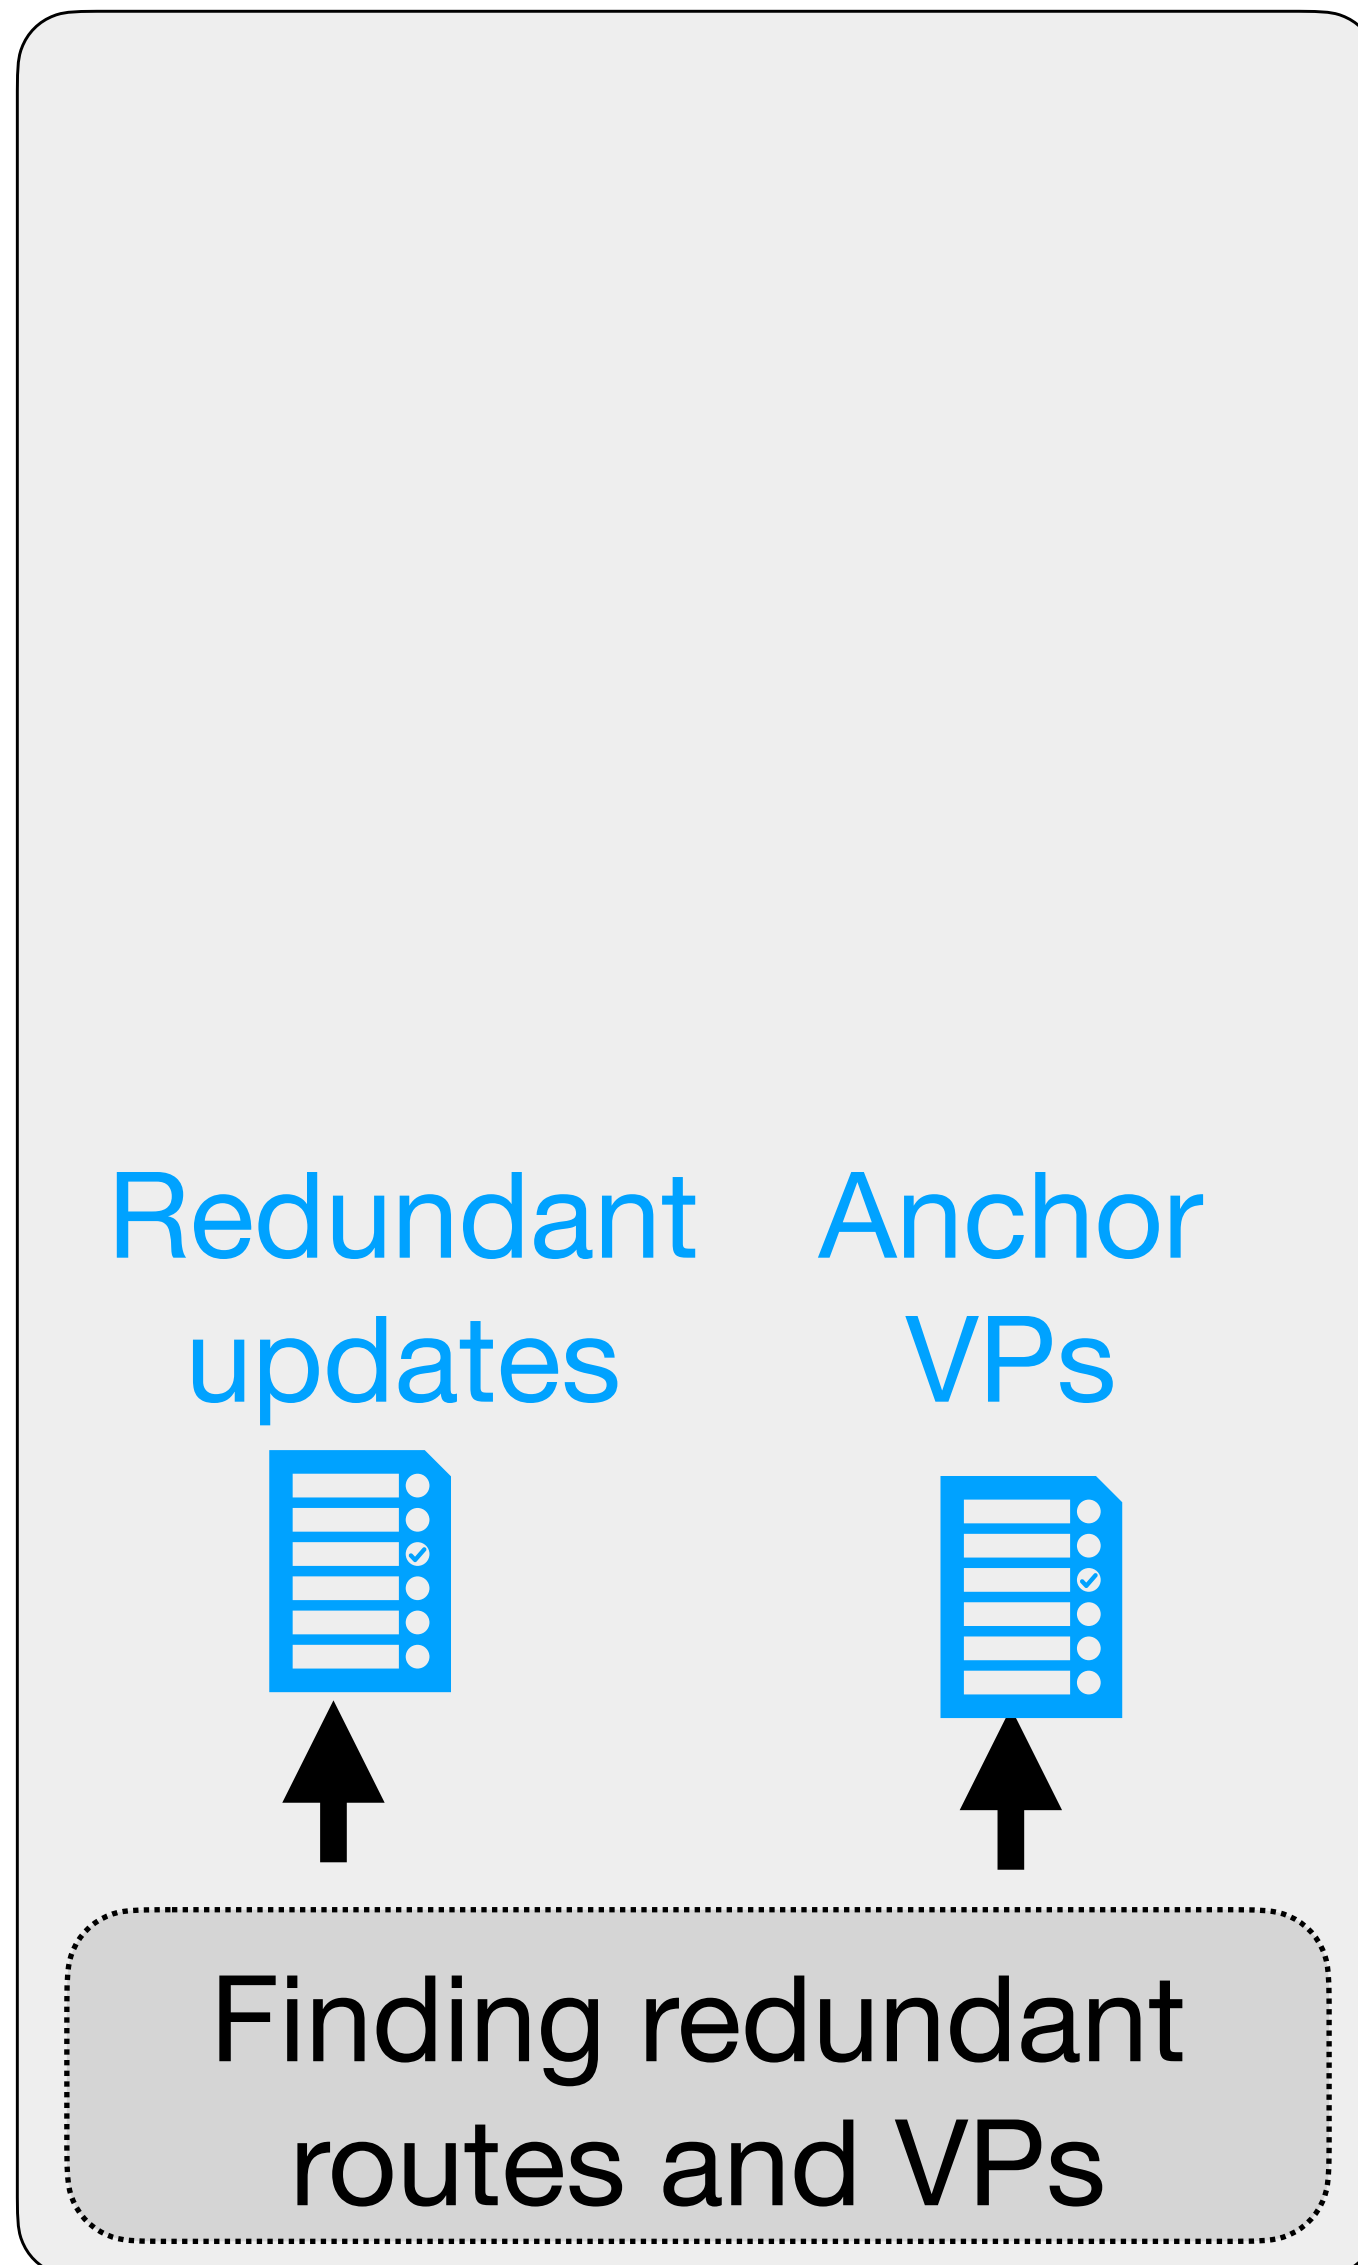
Data management problems



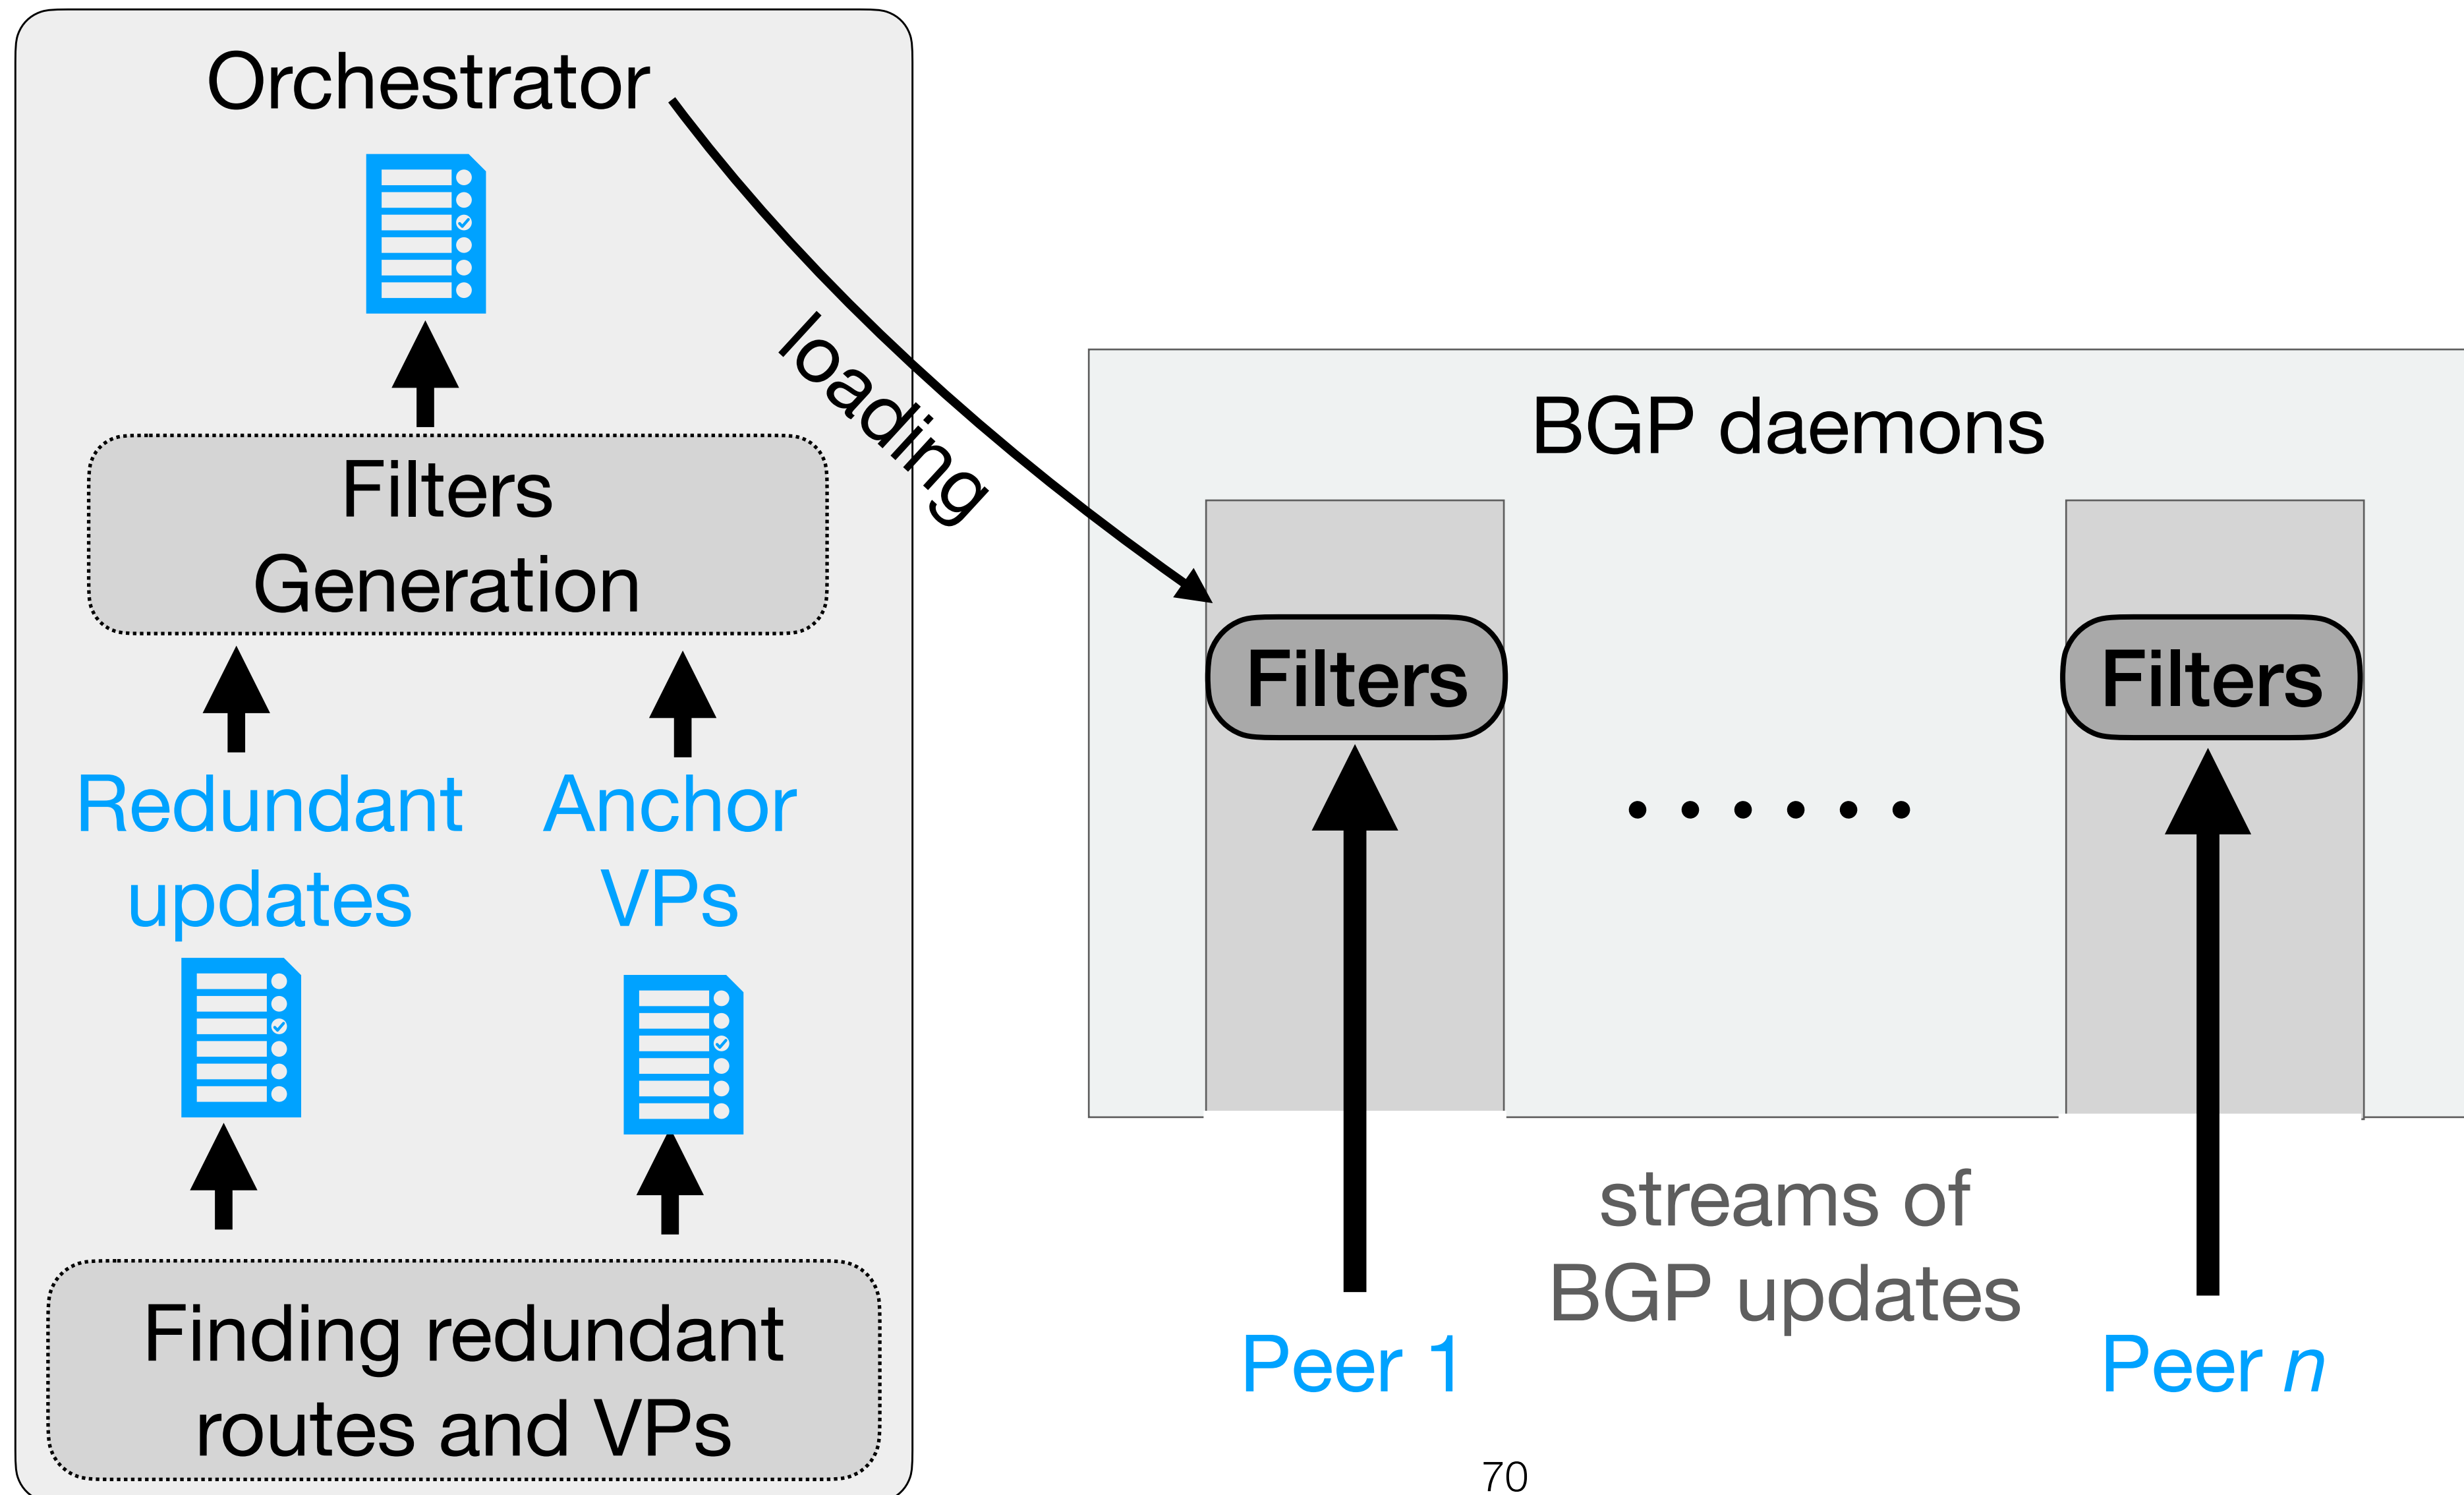
***GILL*** uses BGP daemons written in C and optimized to collect BGP routes



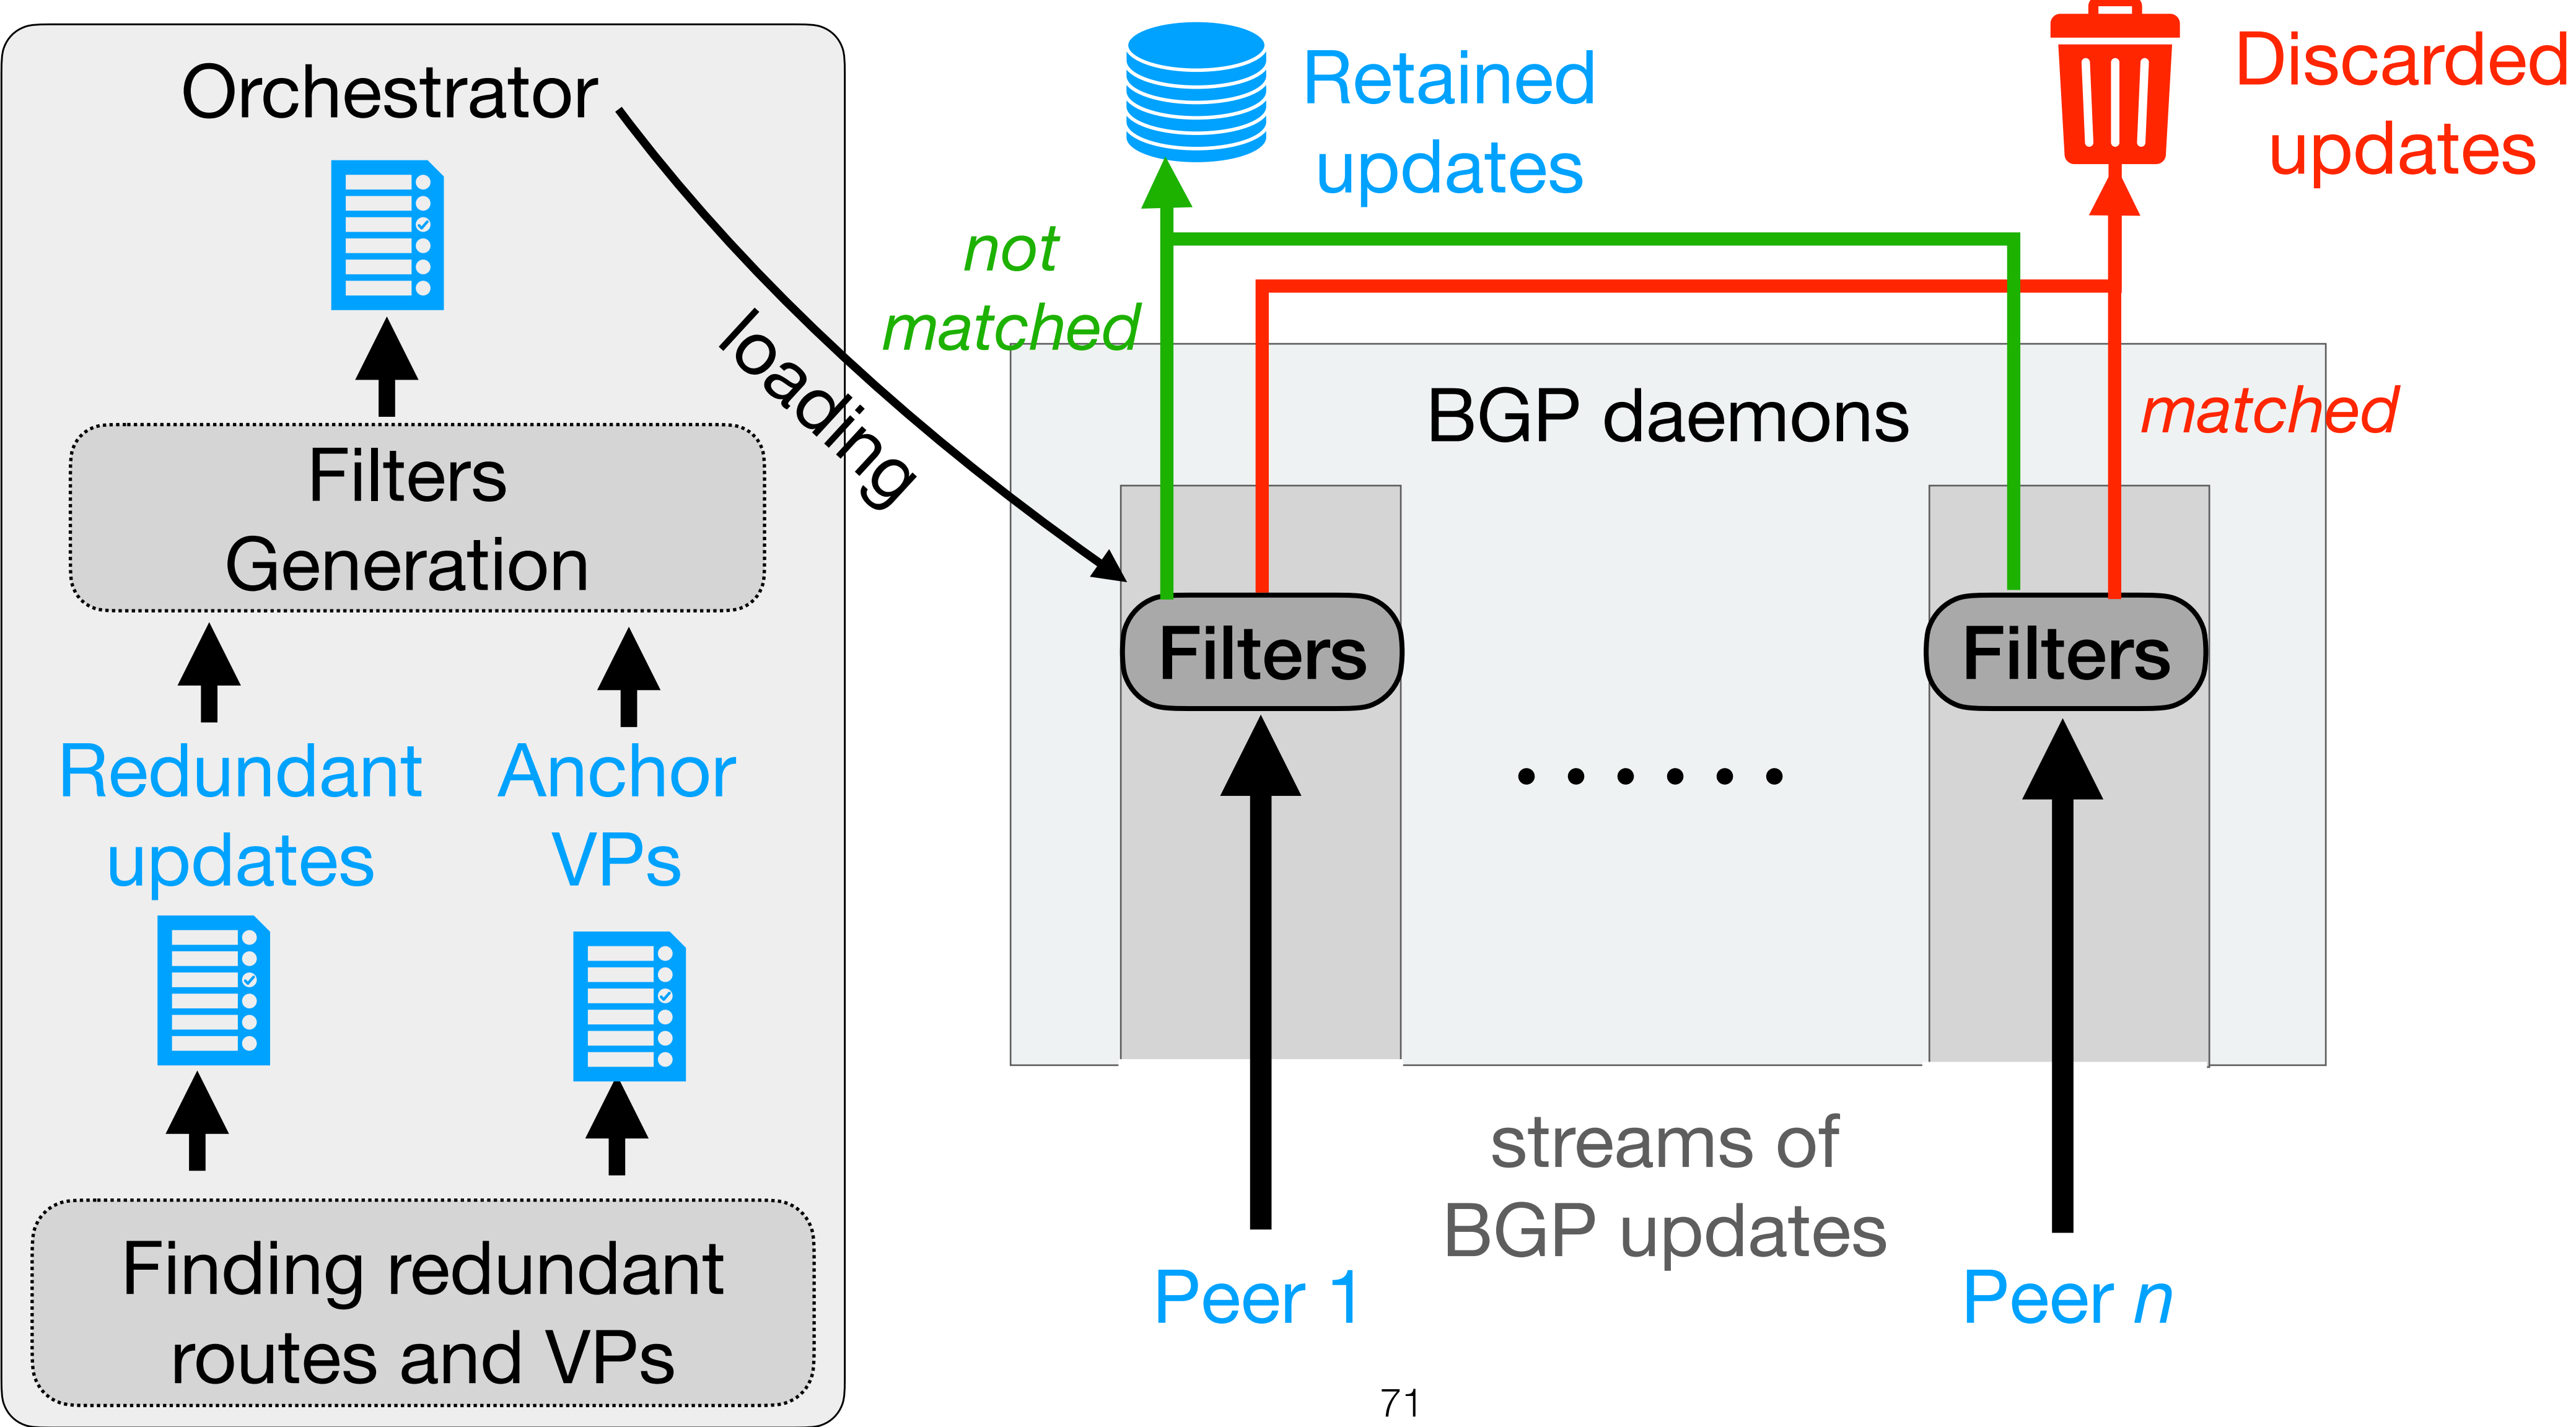
# *GILL* finds redundant updates and anchors VP



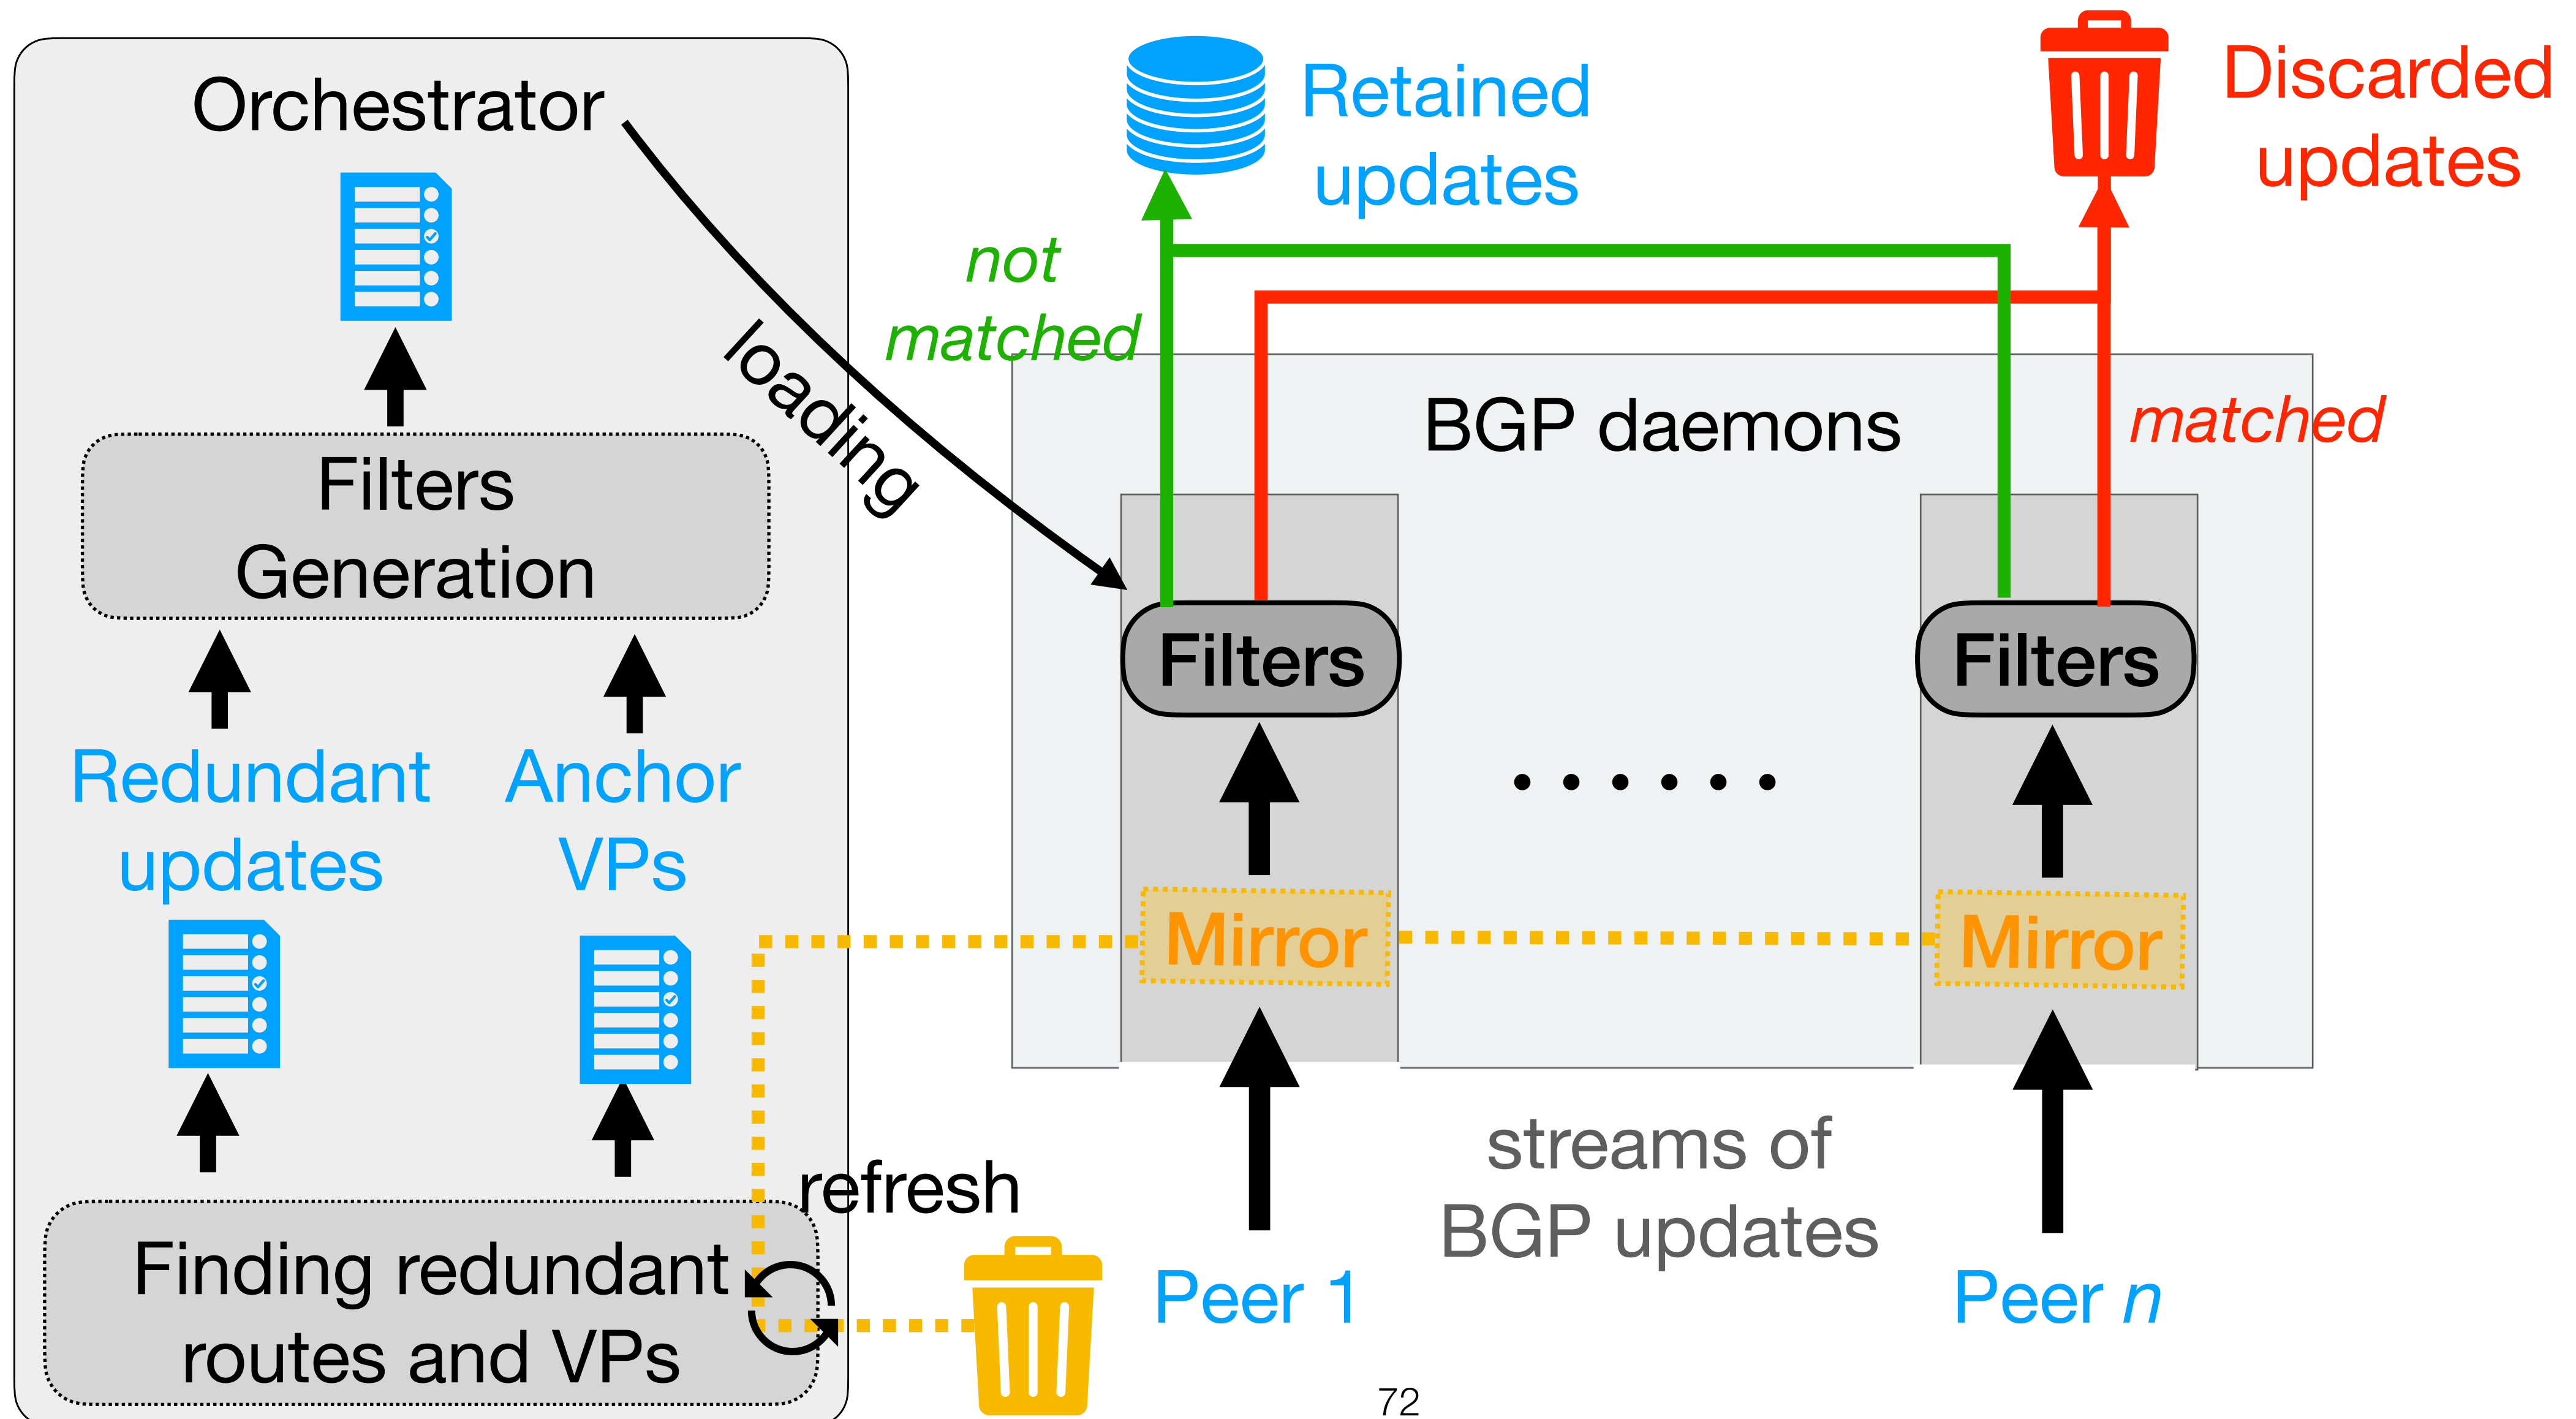
# GILL finds redundant updates and anchors VP



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

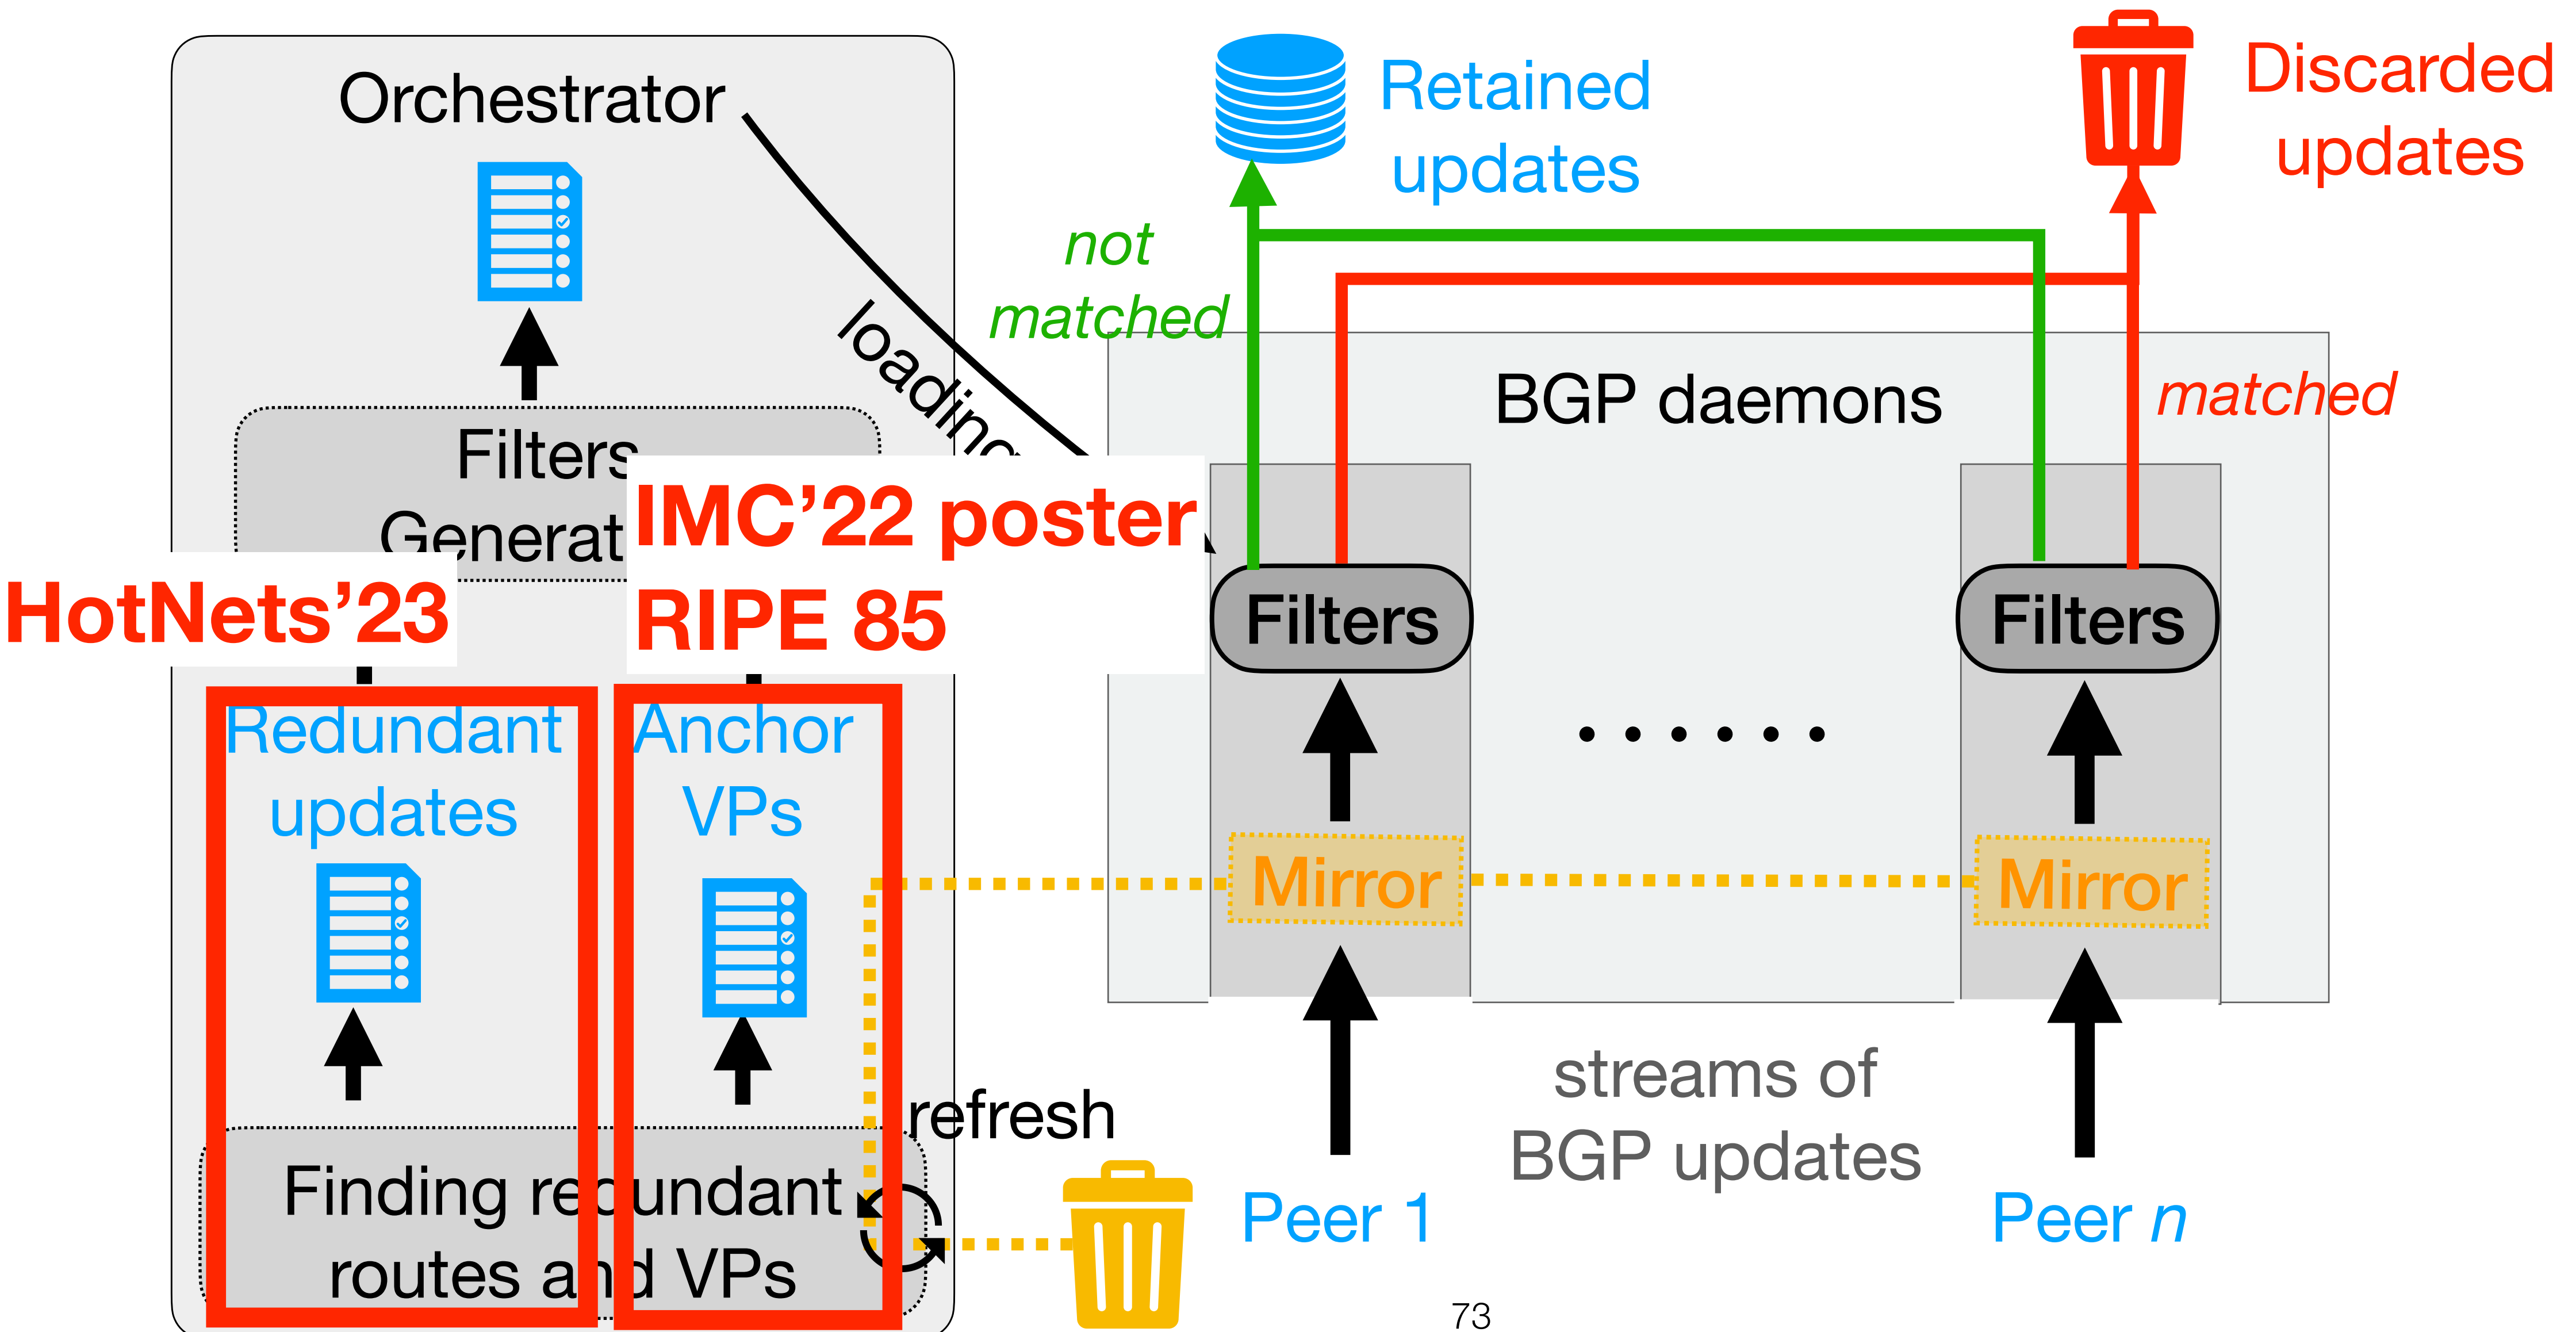


# GILL updates filters over time using an out-of-band mirroring scheme





# 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**


# There is a high level of redundancy in BGP data




Time + Prefix



Time + Prefix +  
AS path



Time + Prefix +  
AS path + Comm.



# Naive baselines fail to assess redundancy in BGP data

