The Ethereum Consensus Network has a Privacy Issue

Lucianna Kiffer
Research Assistant Professor • IMDEA Networks Institute

Lioba Heimbach
ETH Zurich

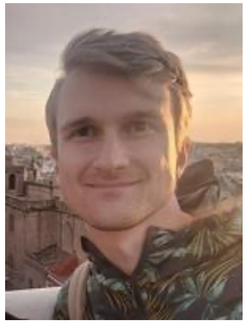

Yann Vonlathen
ETH Zurich

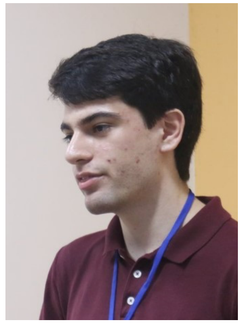

Juan Villacis
University of Bern

Prof. Roger Wattenhofer
ETH Zurich

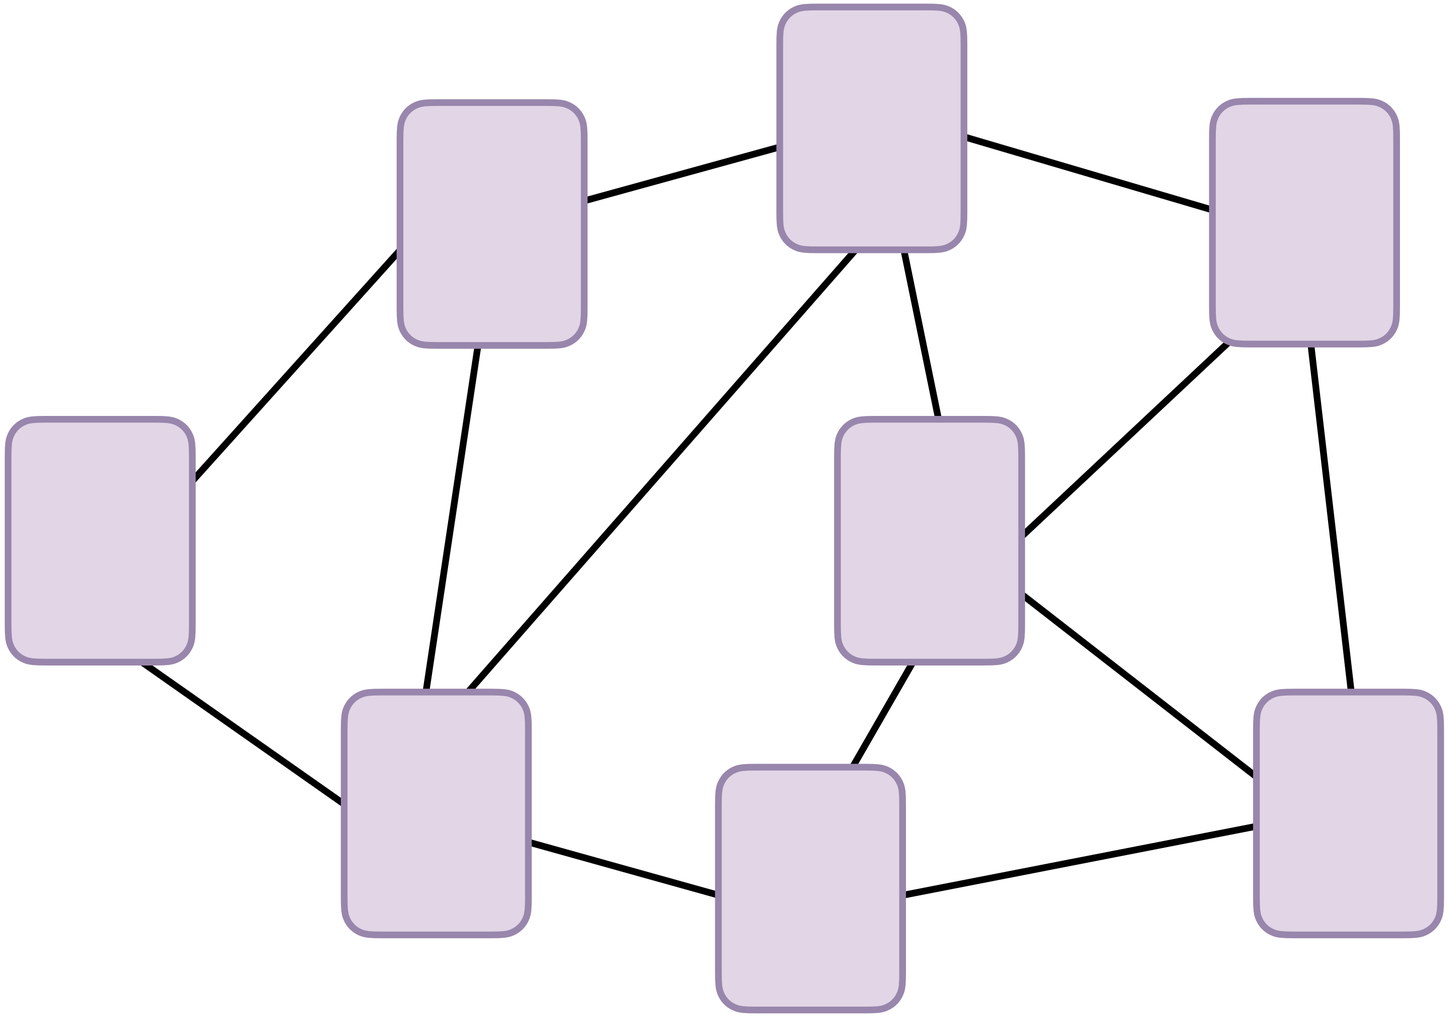

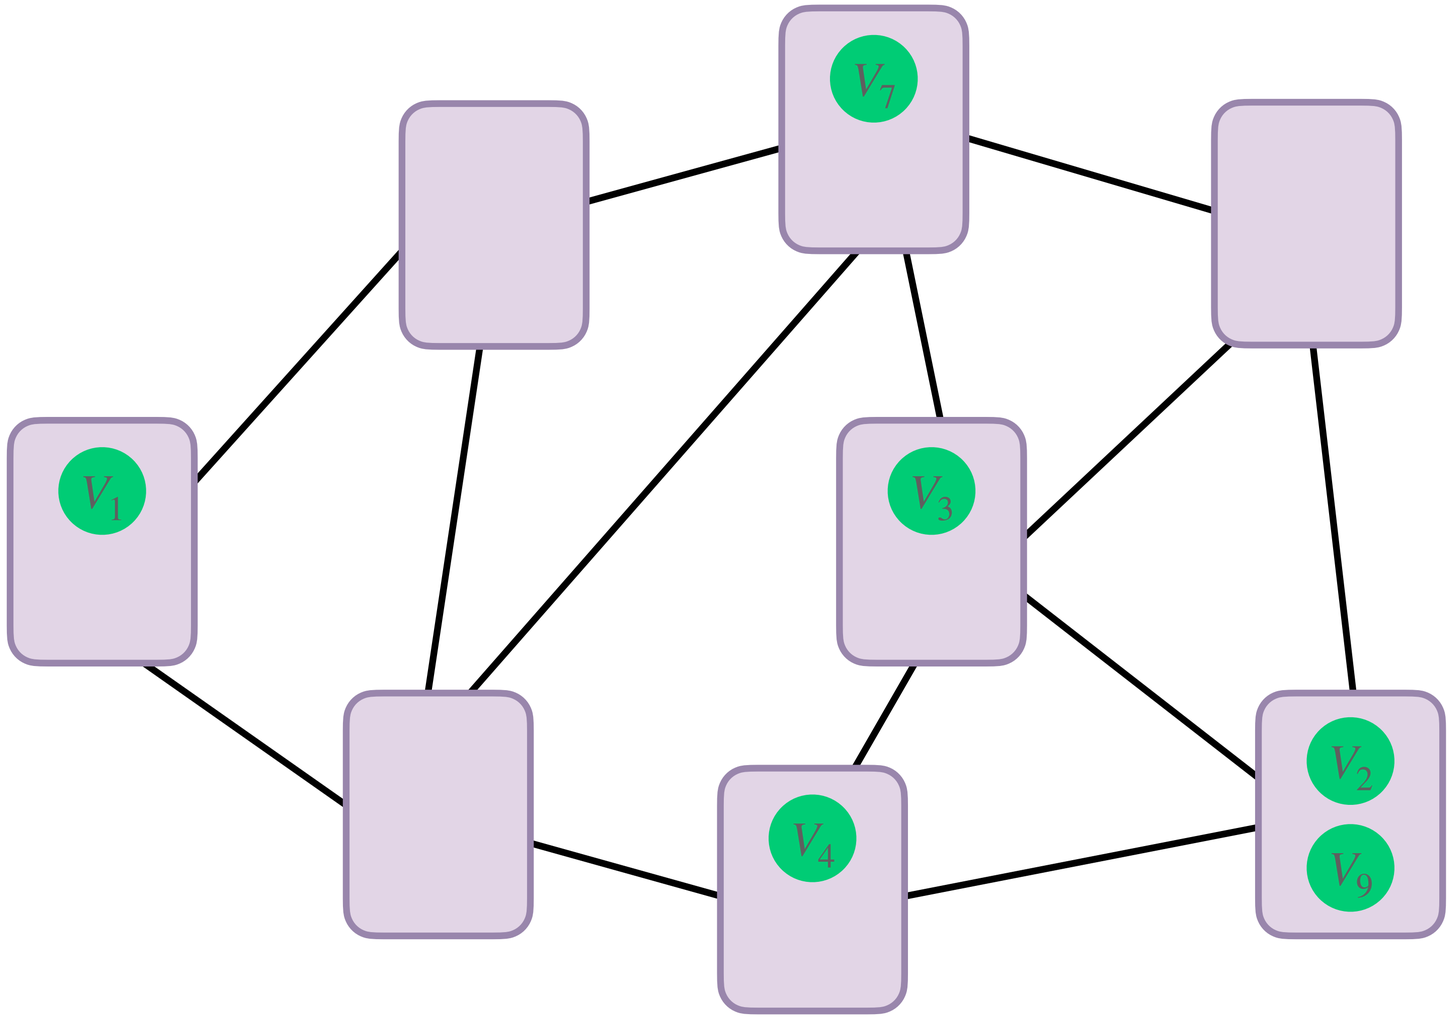
V_7
V_1
V_3
V_2
V_4
V_9

- Entities **in charge of consensus**
- Have unique Validator ID (per 32 ETH)

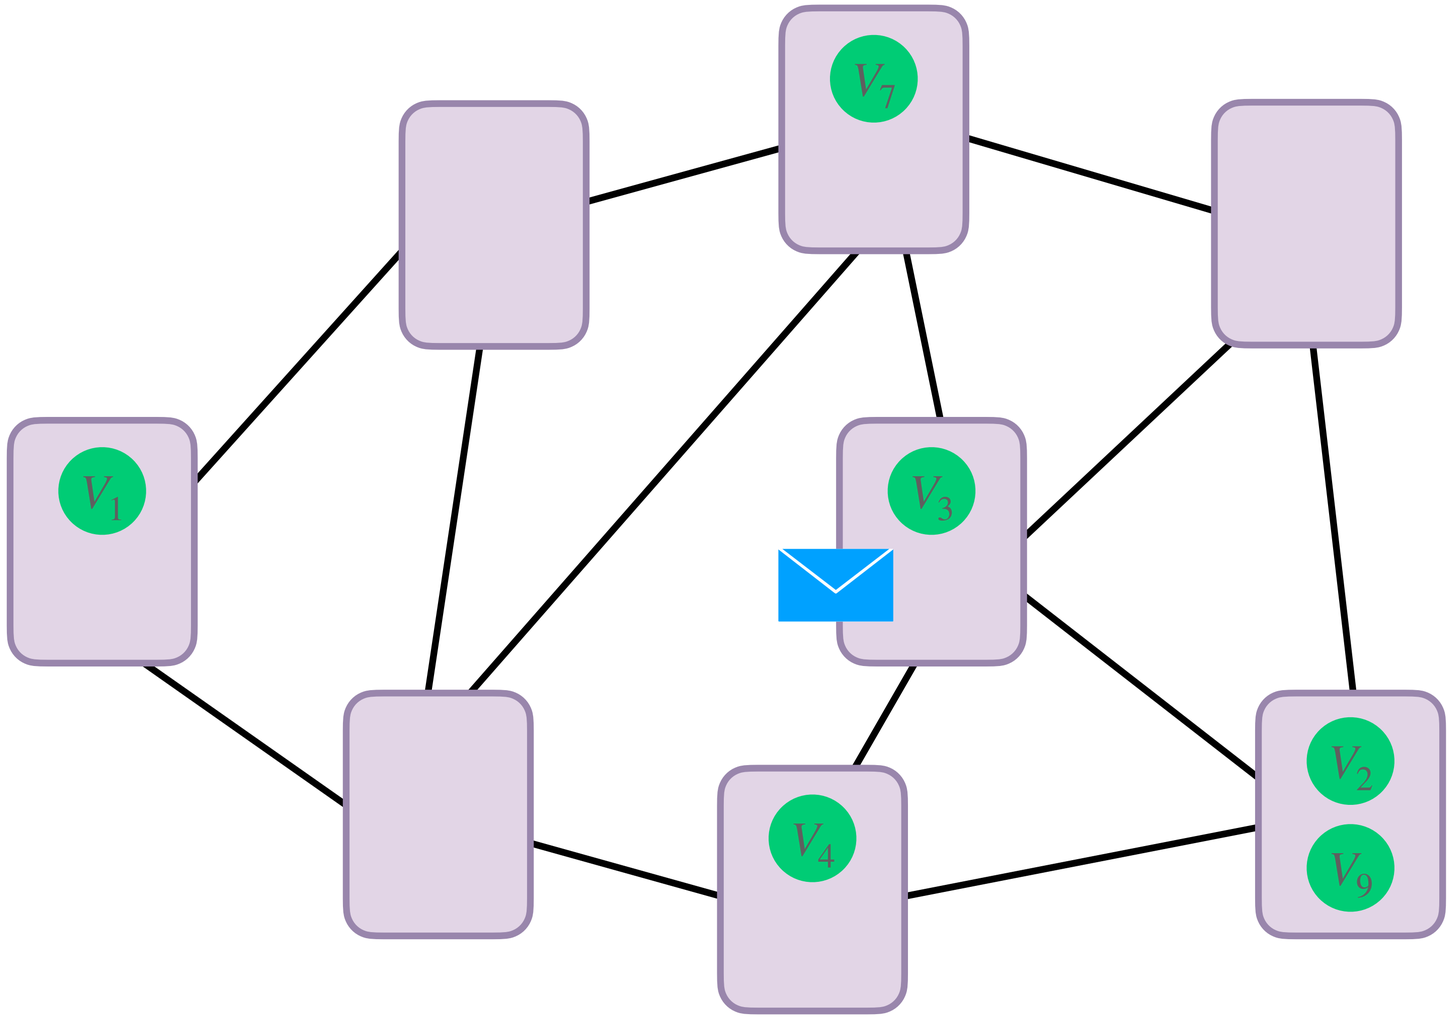
V_7
V_1
V_3
V_4
V_2
V_9

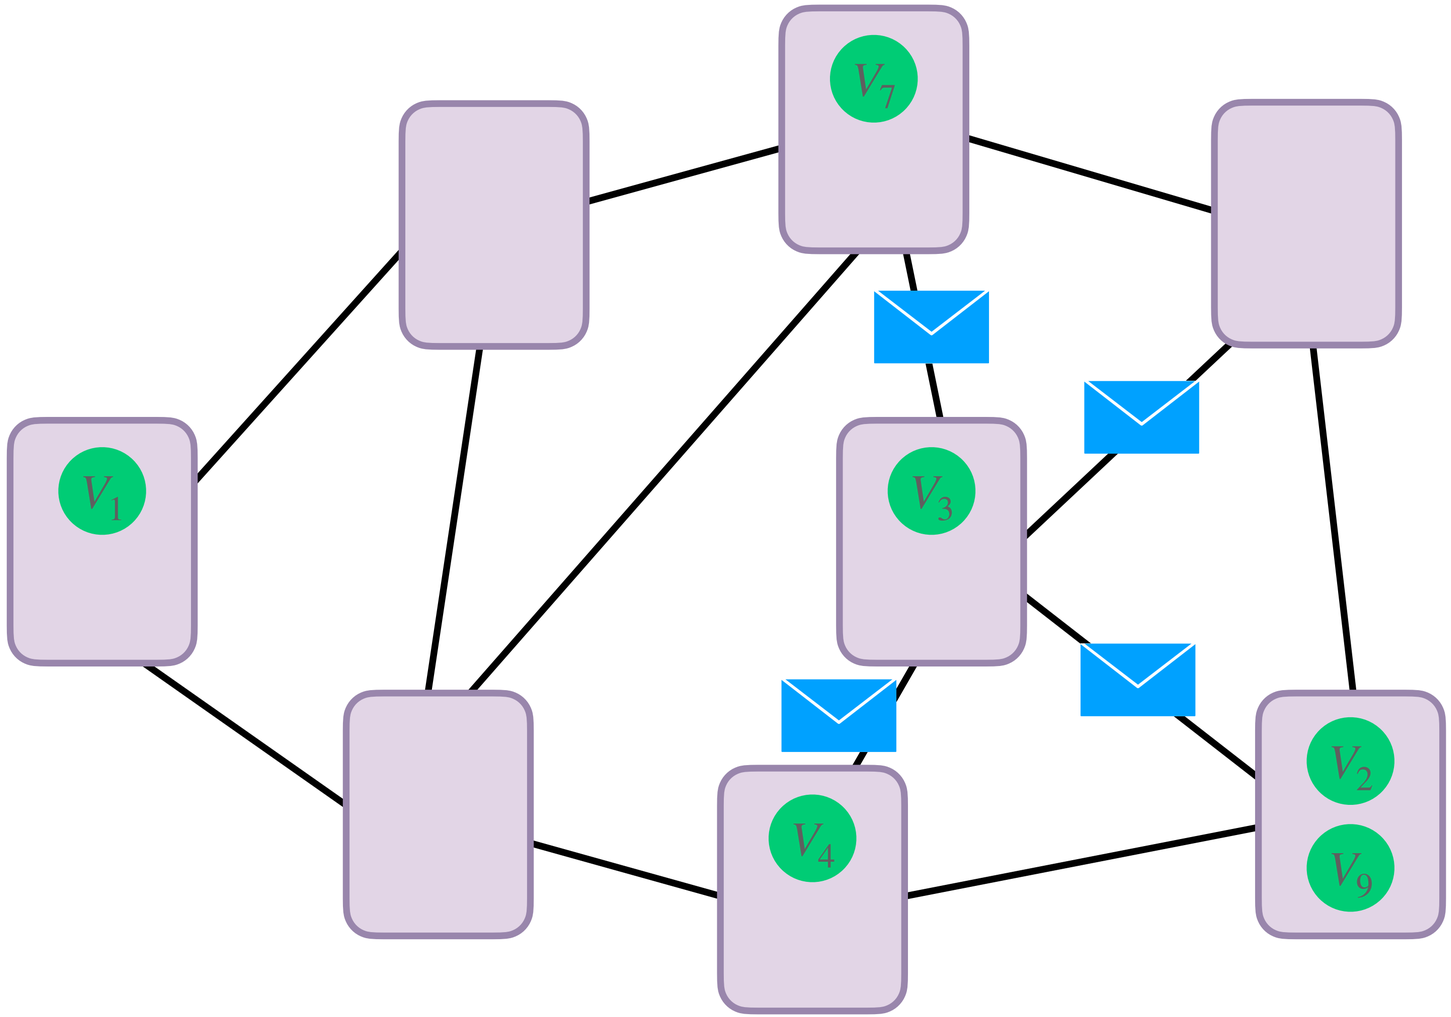
V_7
V_1
V_3
V_4
V_2
V_9

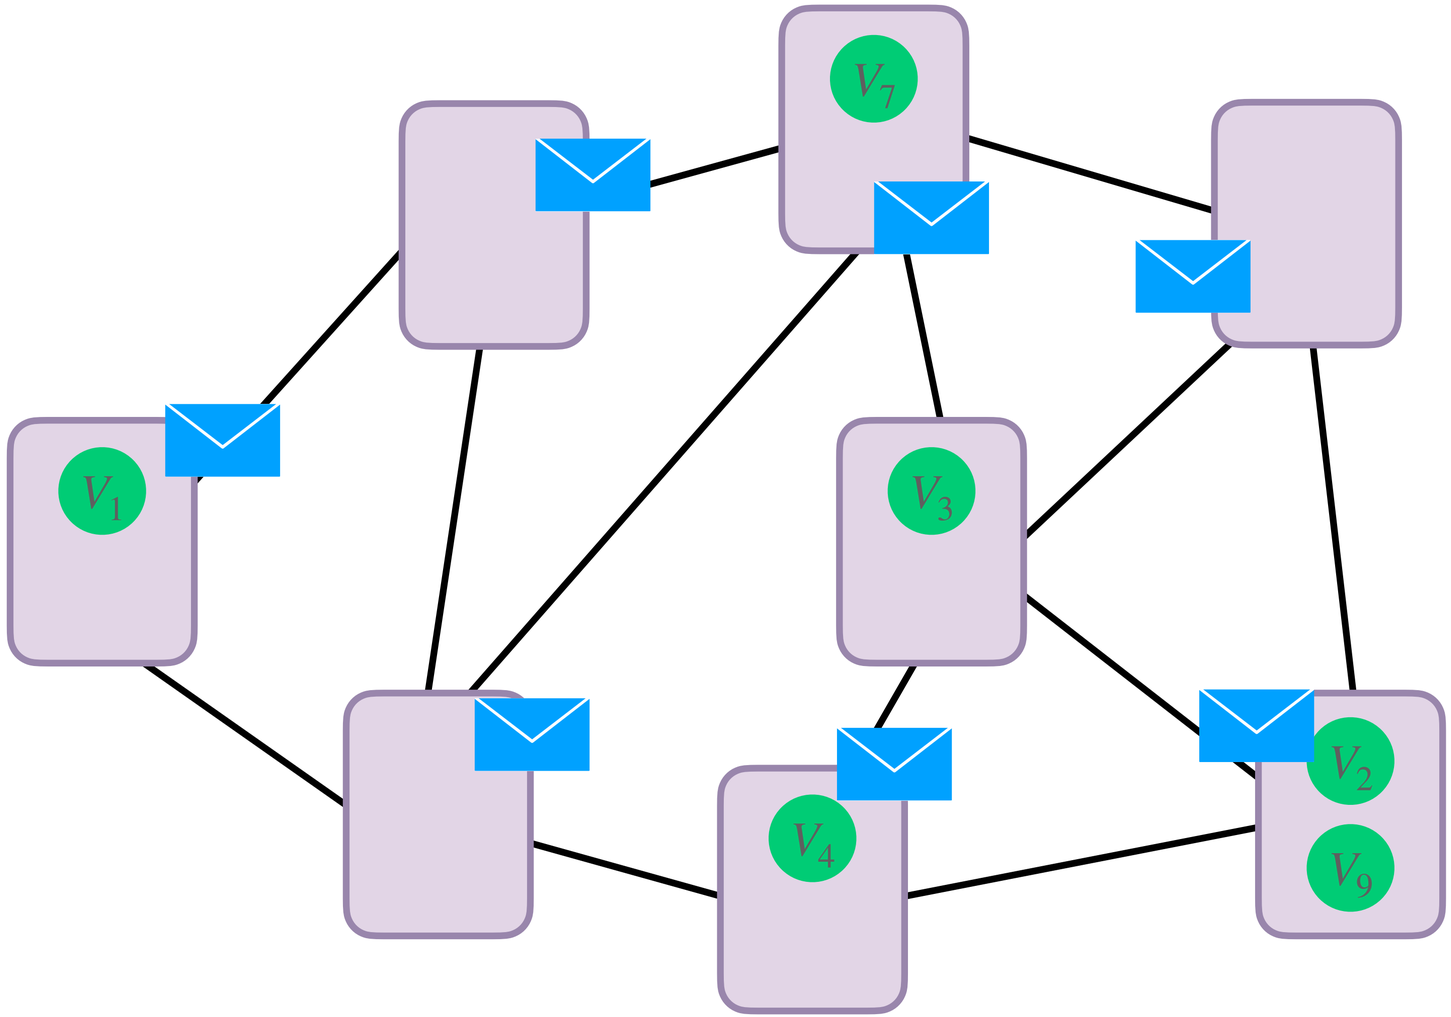
V_7
V_1
V_3
V_2
V_4
V_9

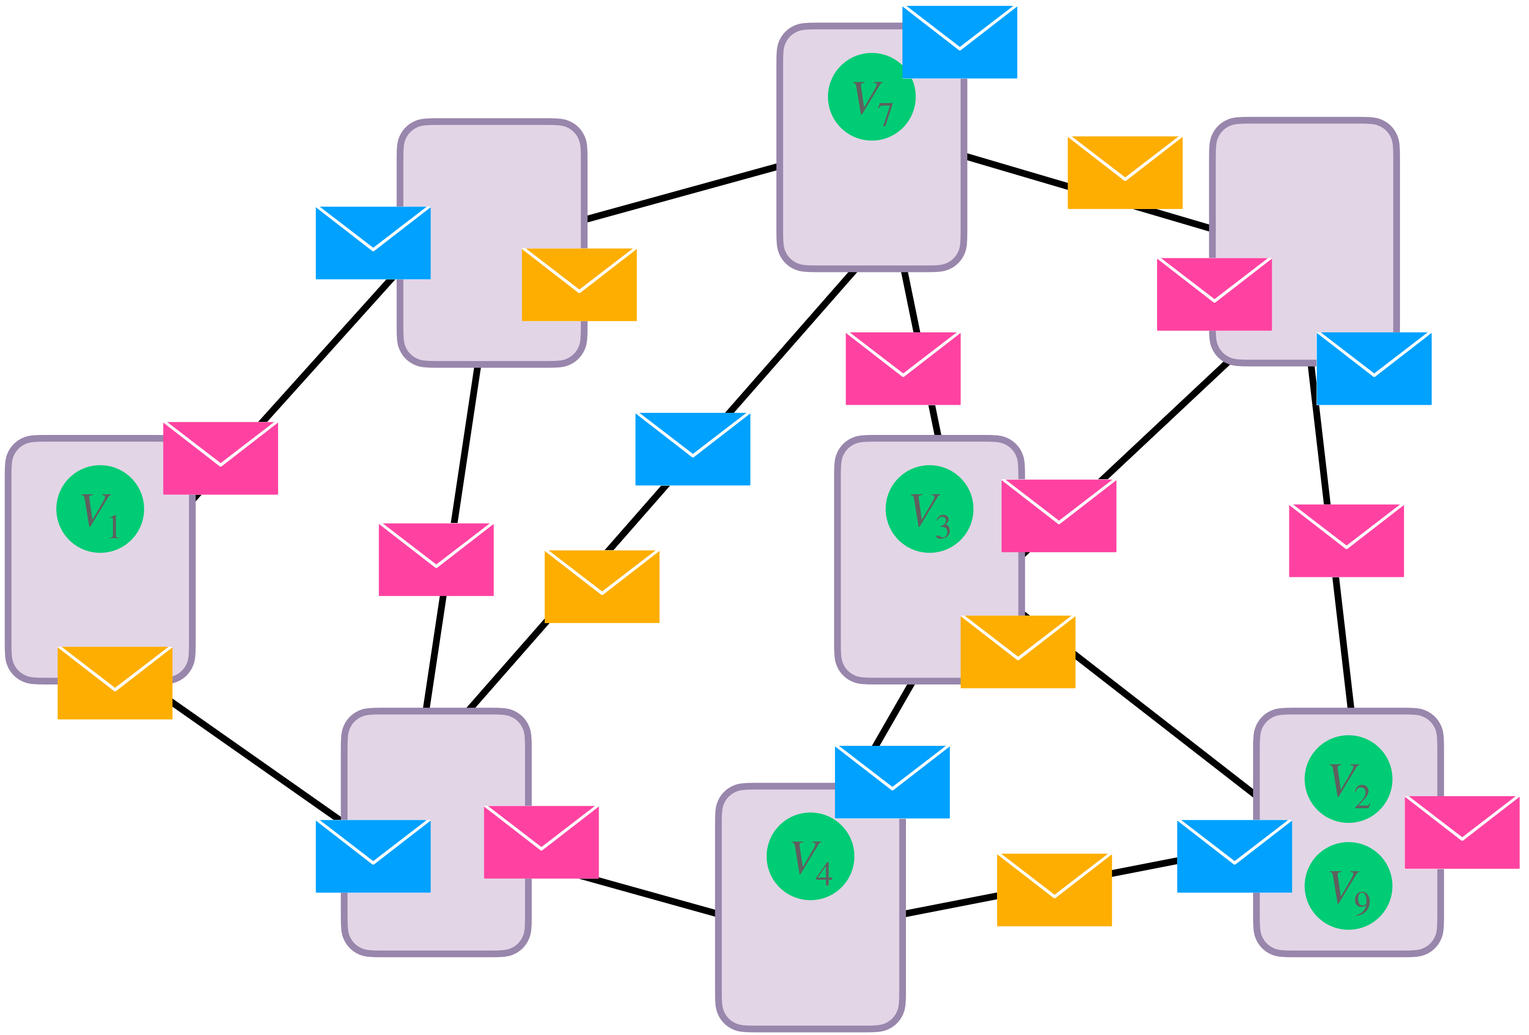
V_7
V_1
V_3
V_2
V_9
V_4

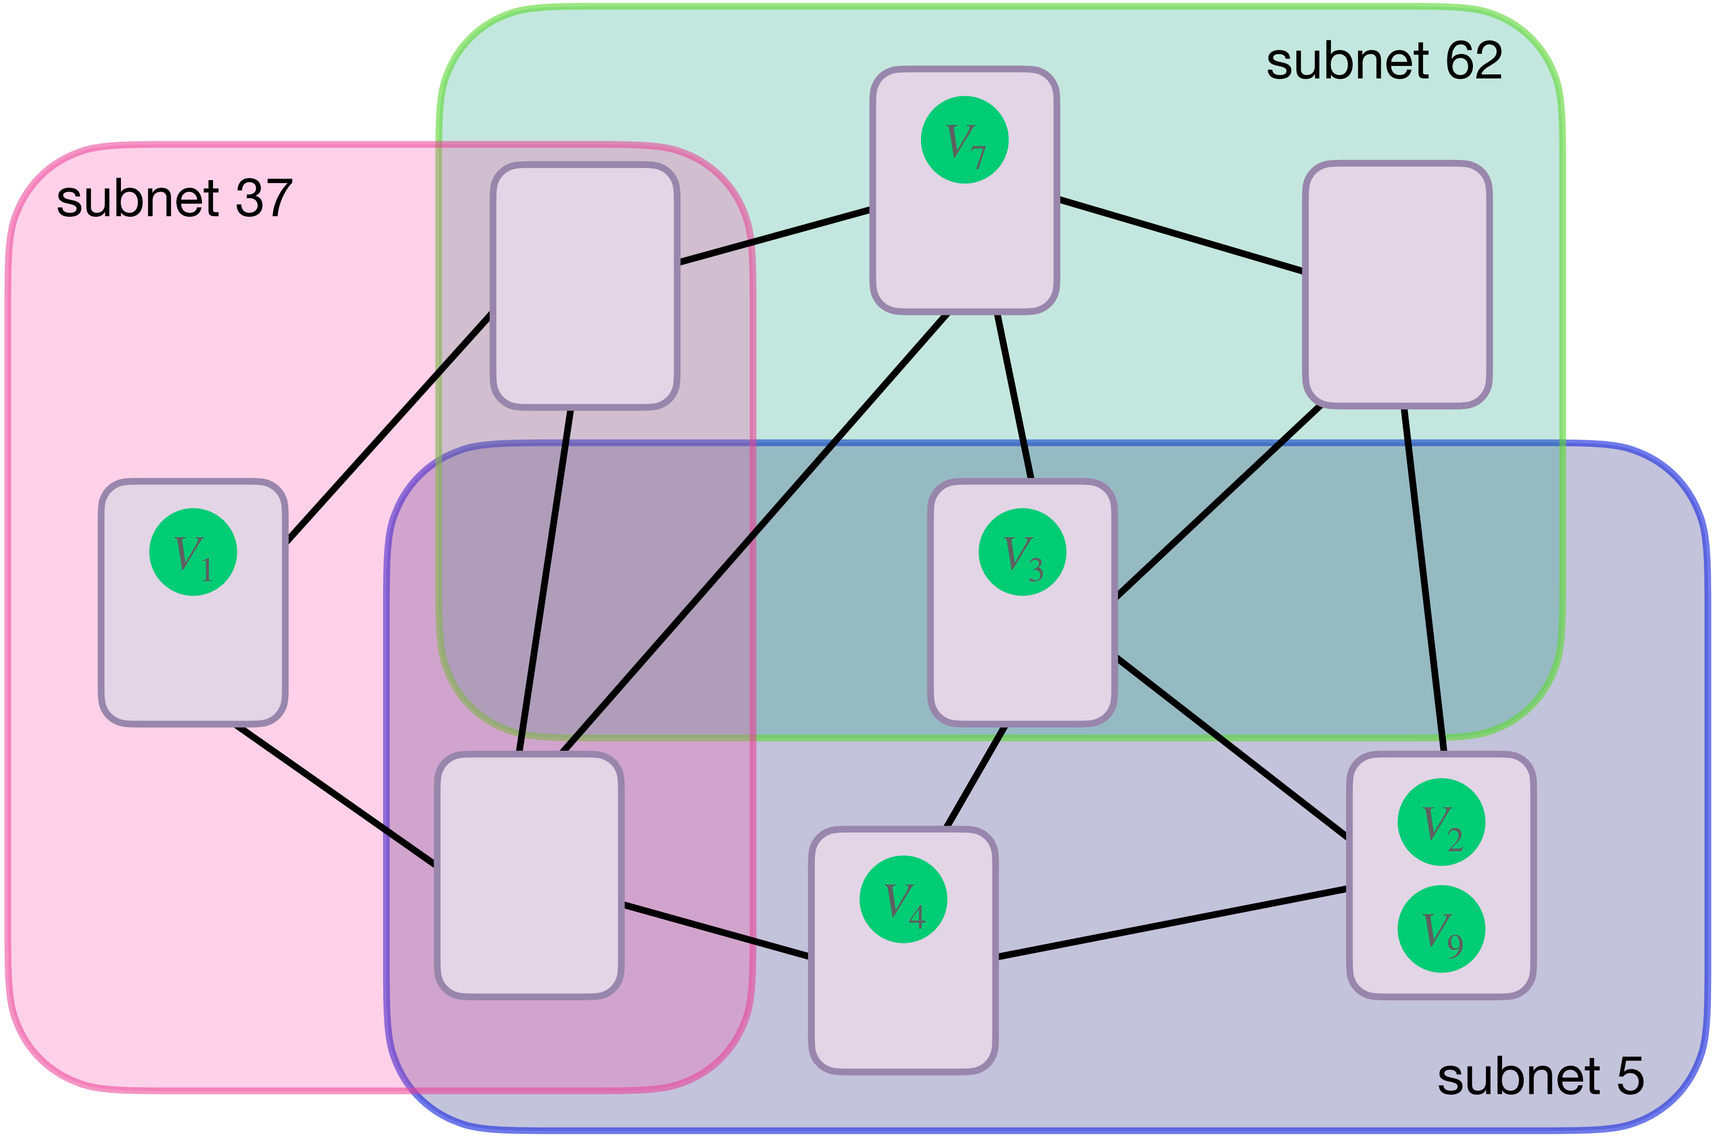
subnet 62
subnet 37
V_7
V_1
V_3
V_4
V_2
V_9
subnet 5

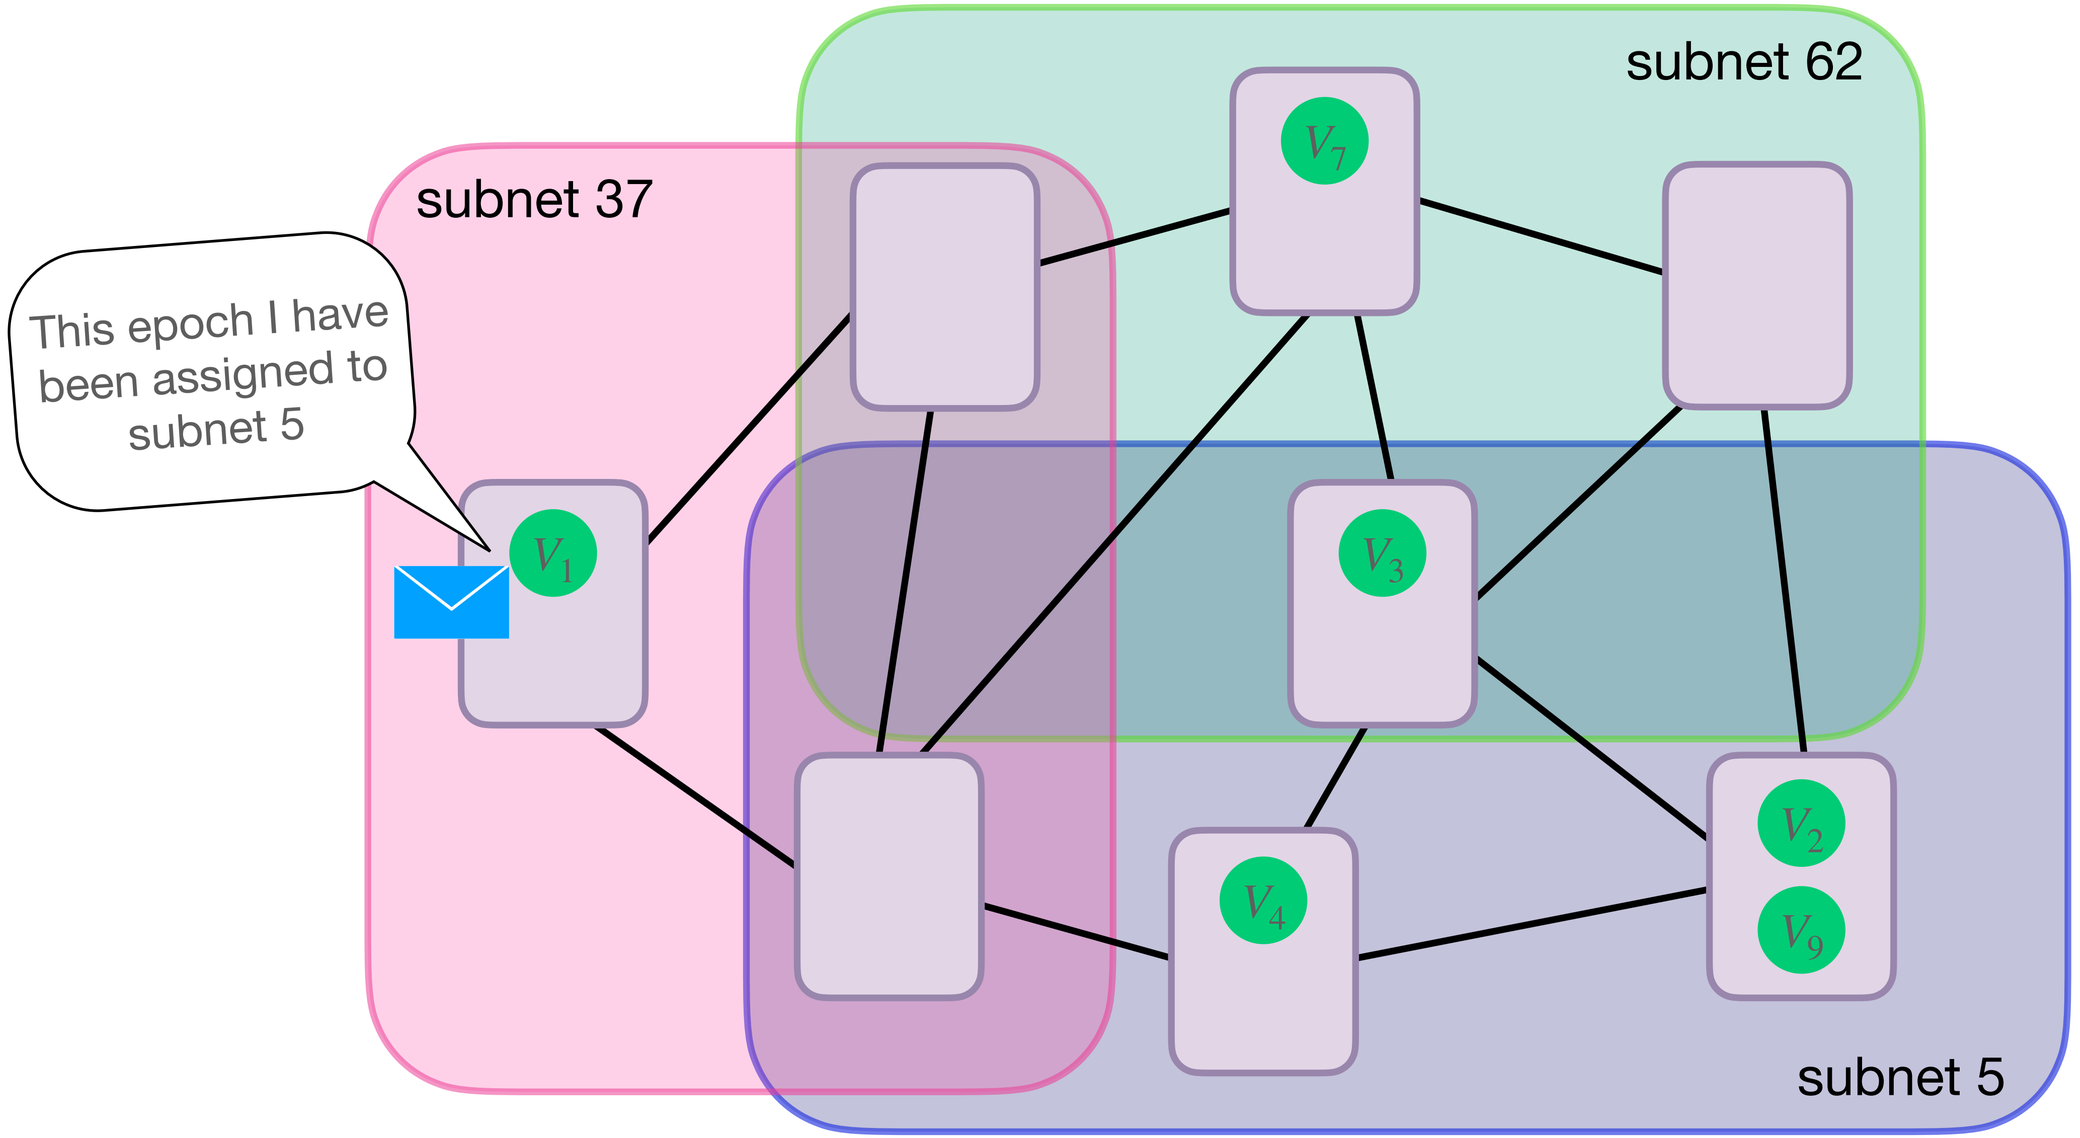
subnet 62
subnet 37
This epoch I have been assigned to subnet 5
V_7
V_1
V_3
V_4
V_2
V_9
subnet 5

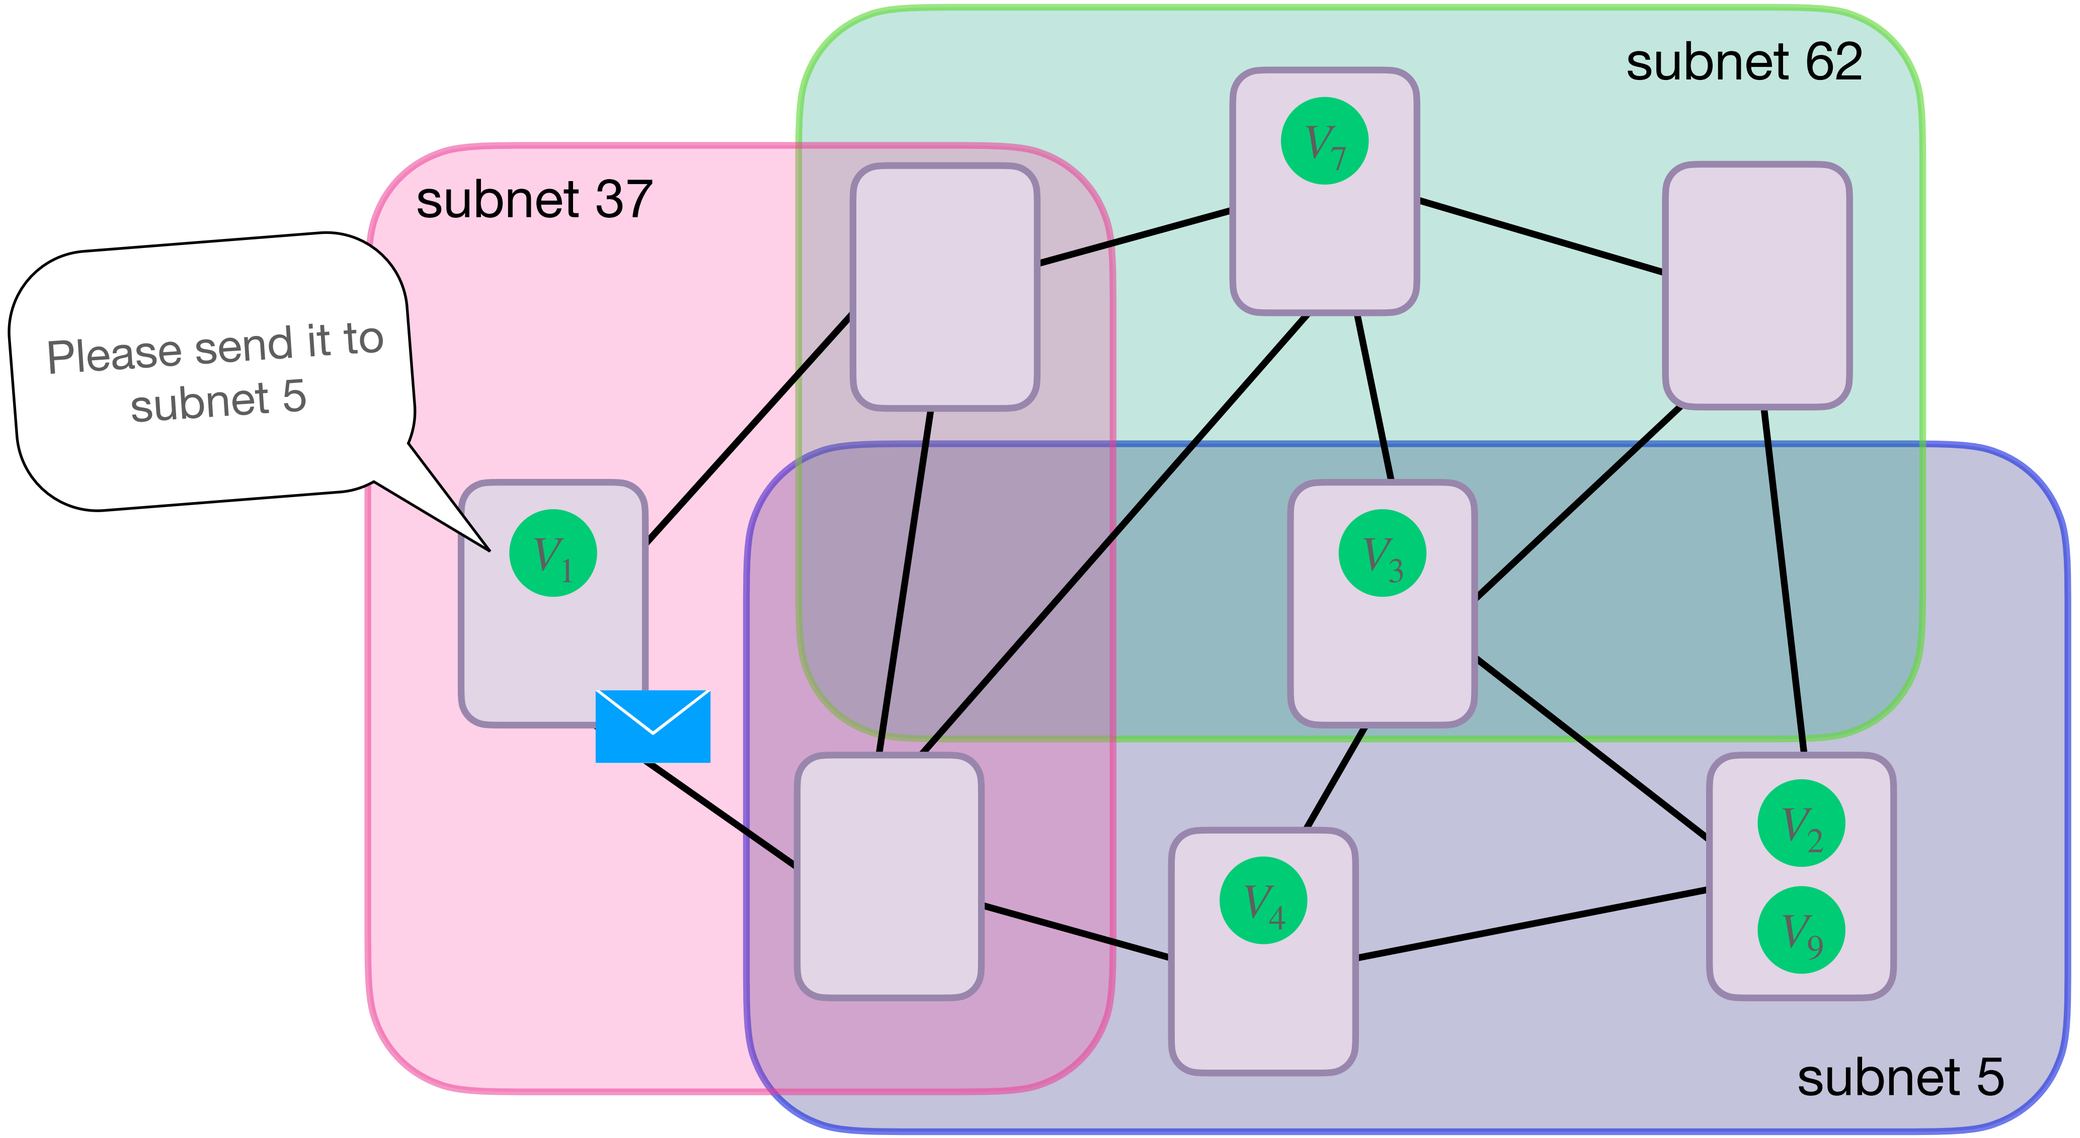
subnet 62
subnet 37
Please send it to subnet 5
V_7
V_1
V_3
V_4
V_2
V_9
subnet 5

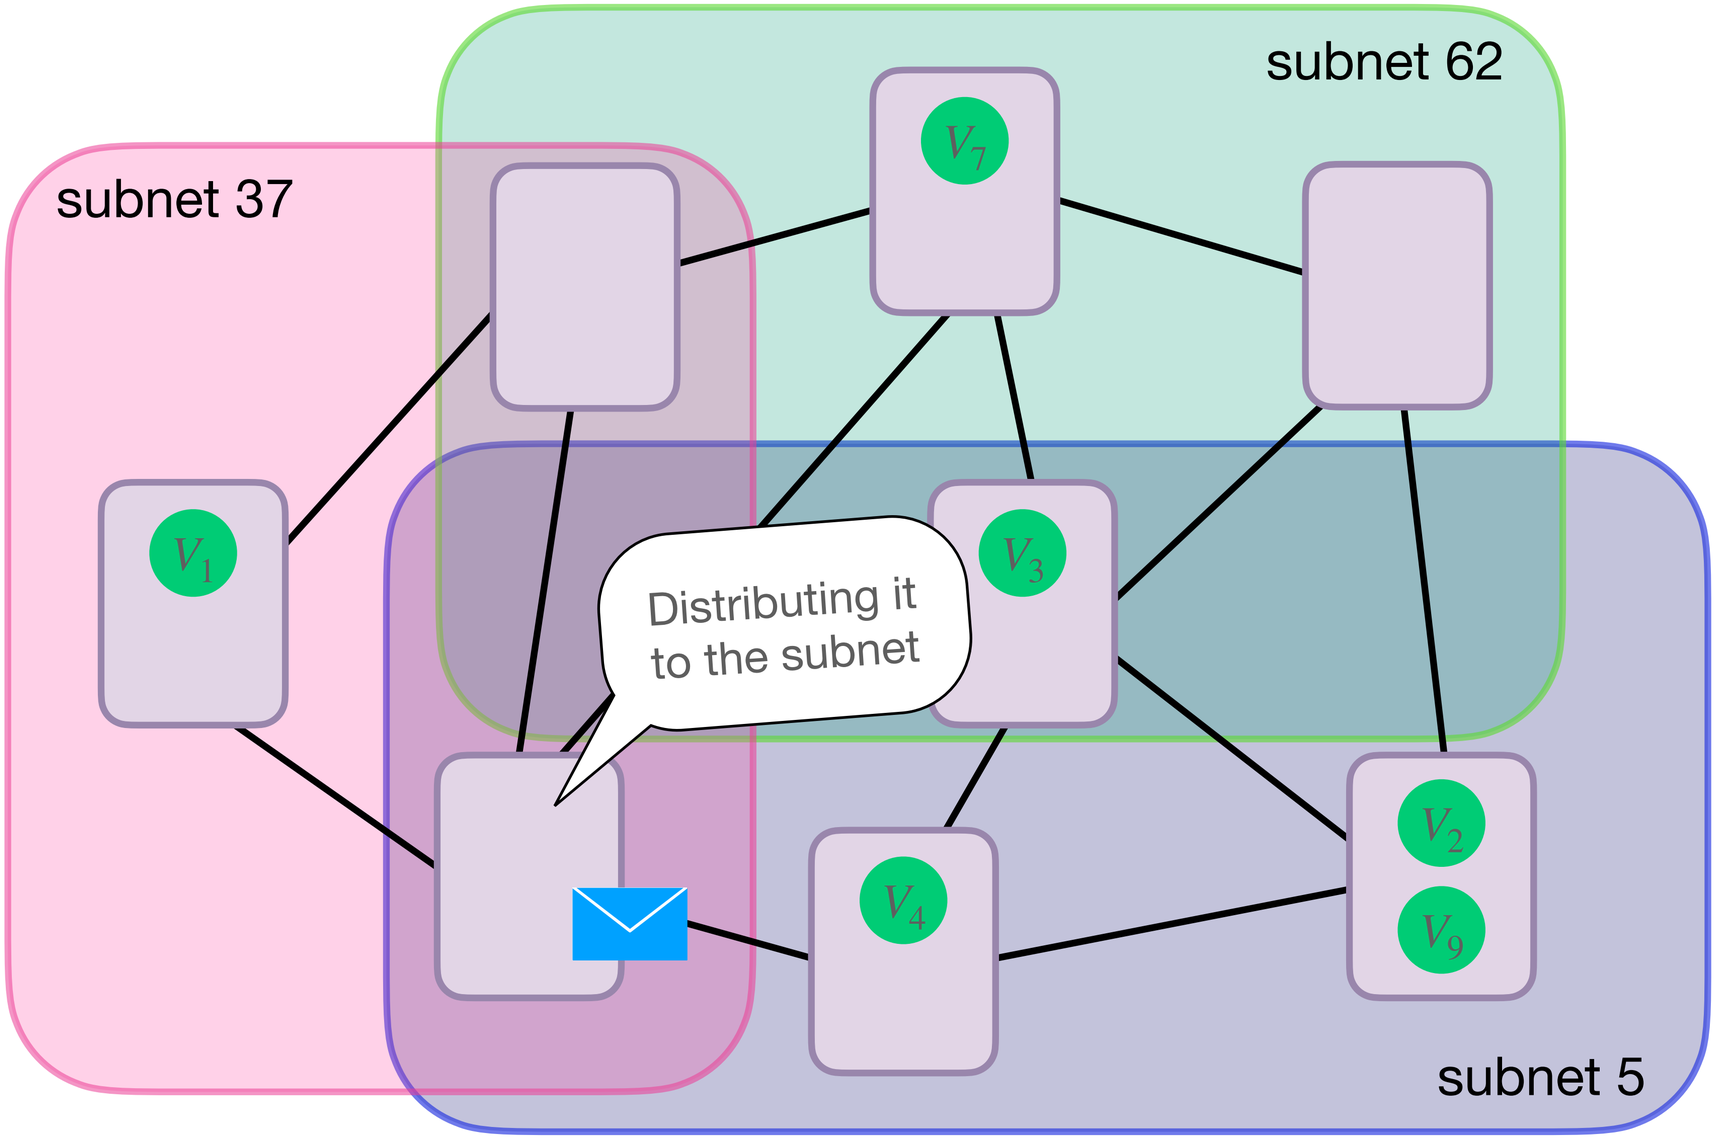
subnet 62
subnet 37
V_7
V_1
V_3
Distributing it
to the subnet
V_4
V_2
V_9
subnet 5

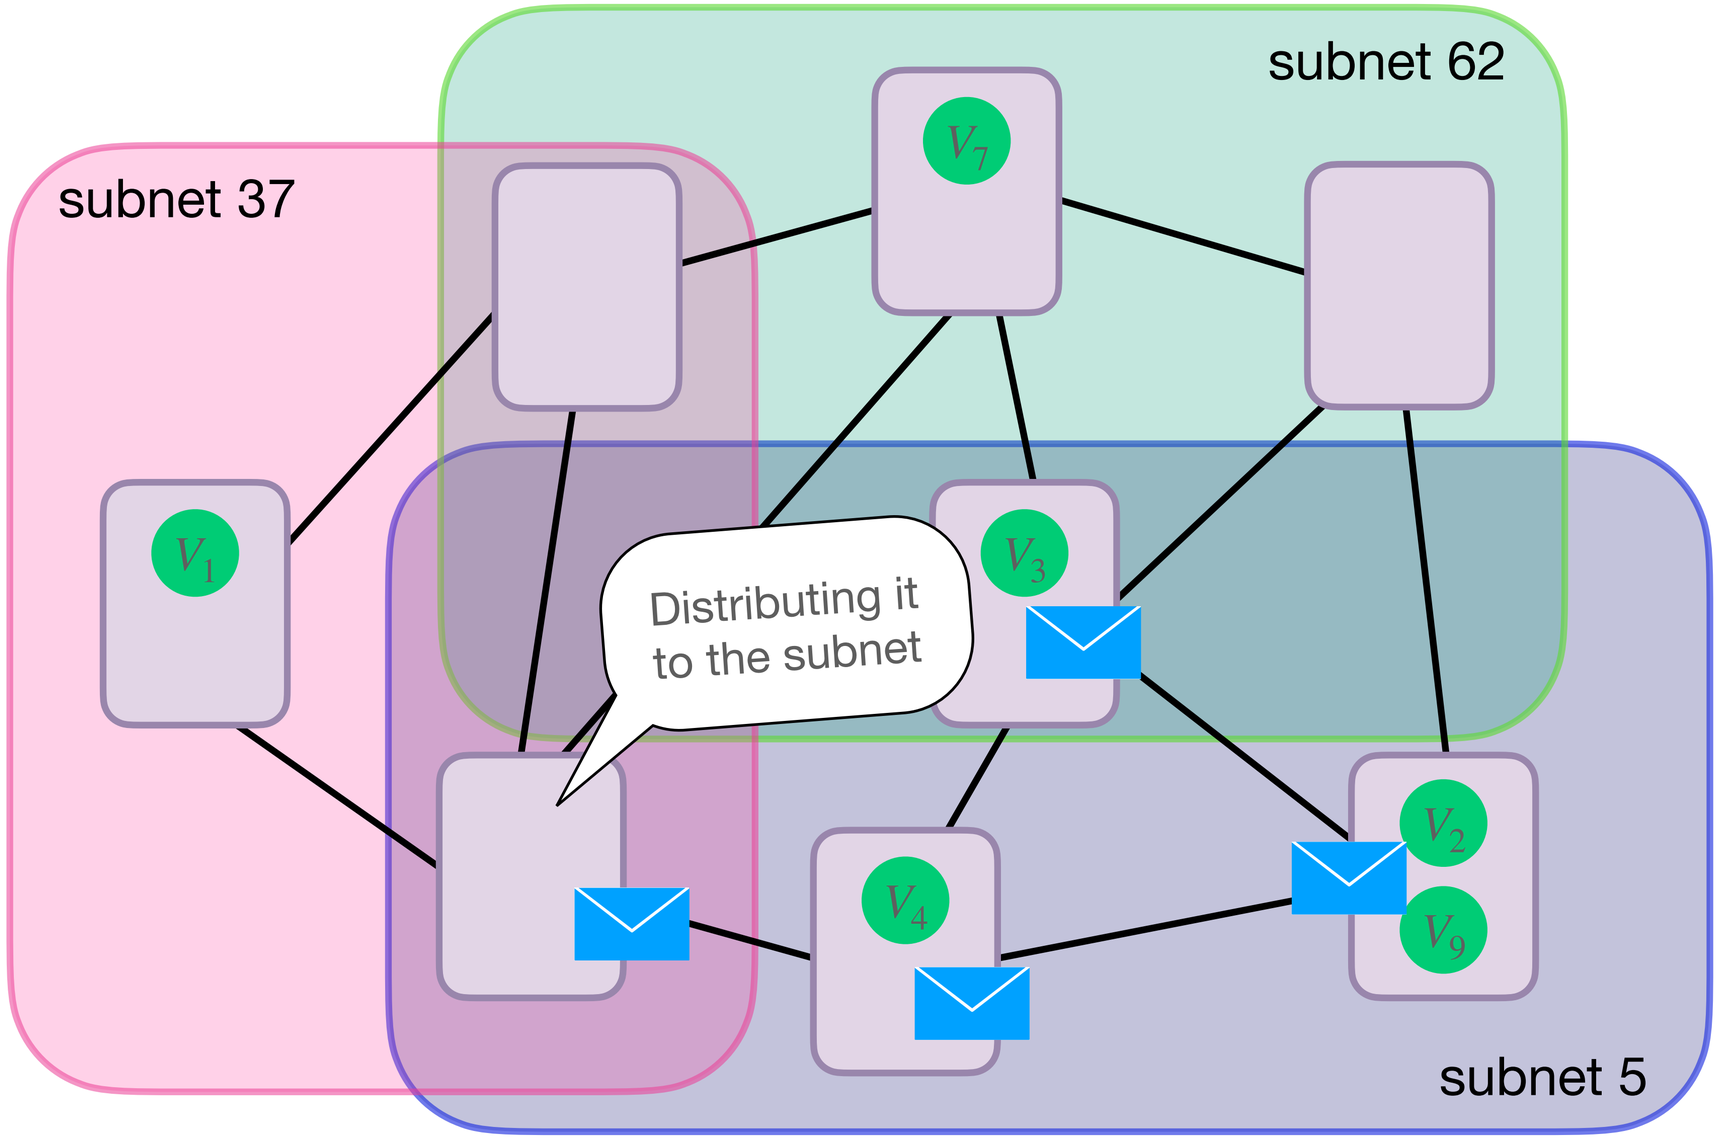
subnet 62
subnet 37
V_7
V_1
V_3
Distributing it
to the subnet
V_4
V_2
V_9
subnet 5

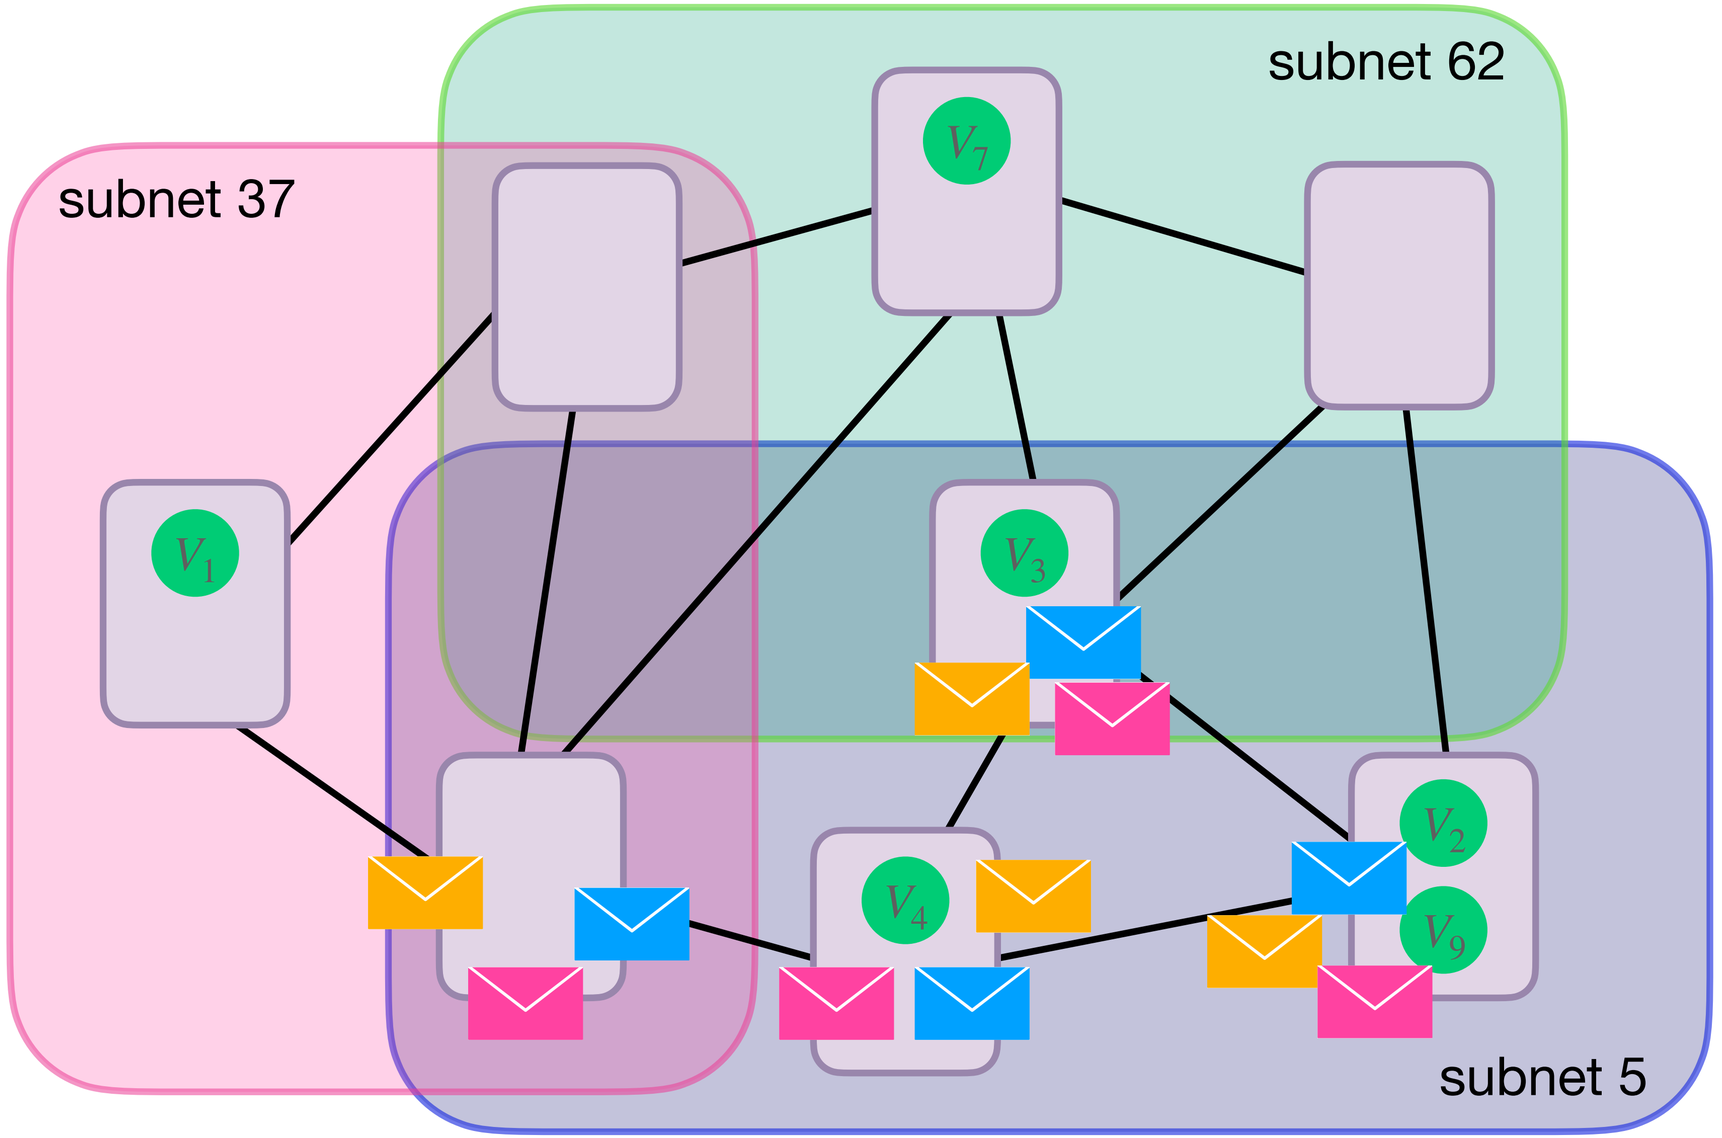
subnet 62
subnet 37
subnet 5
V_7
V_1
V_3
V_4
V_2
V_9

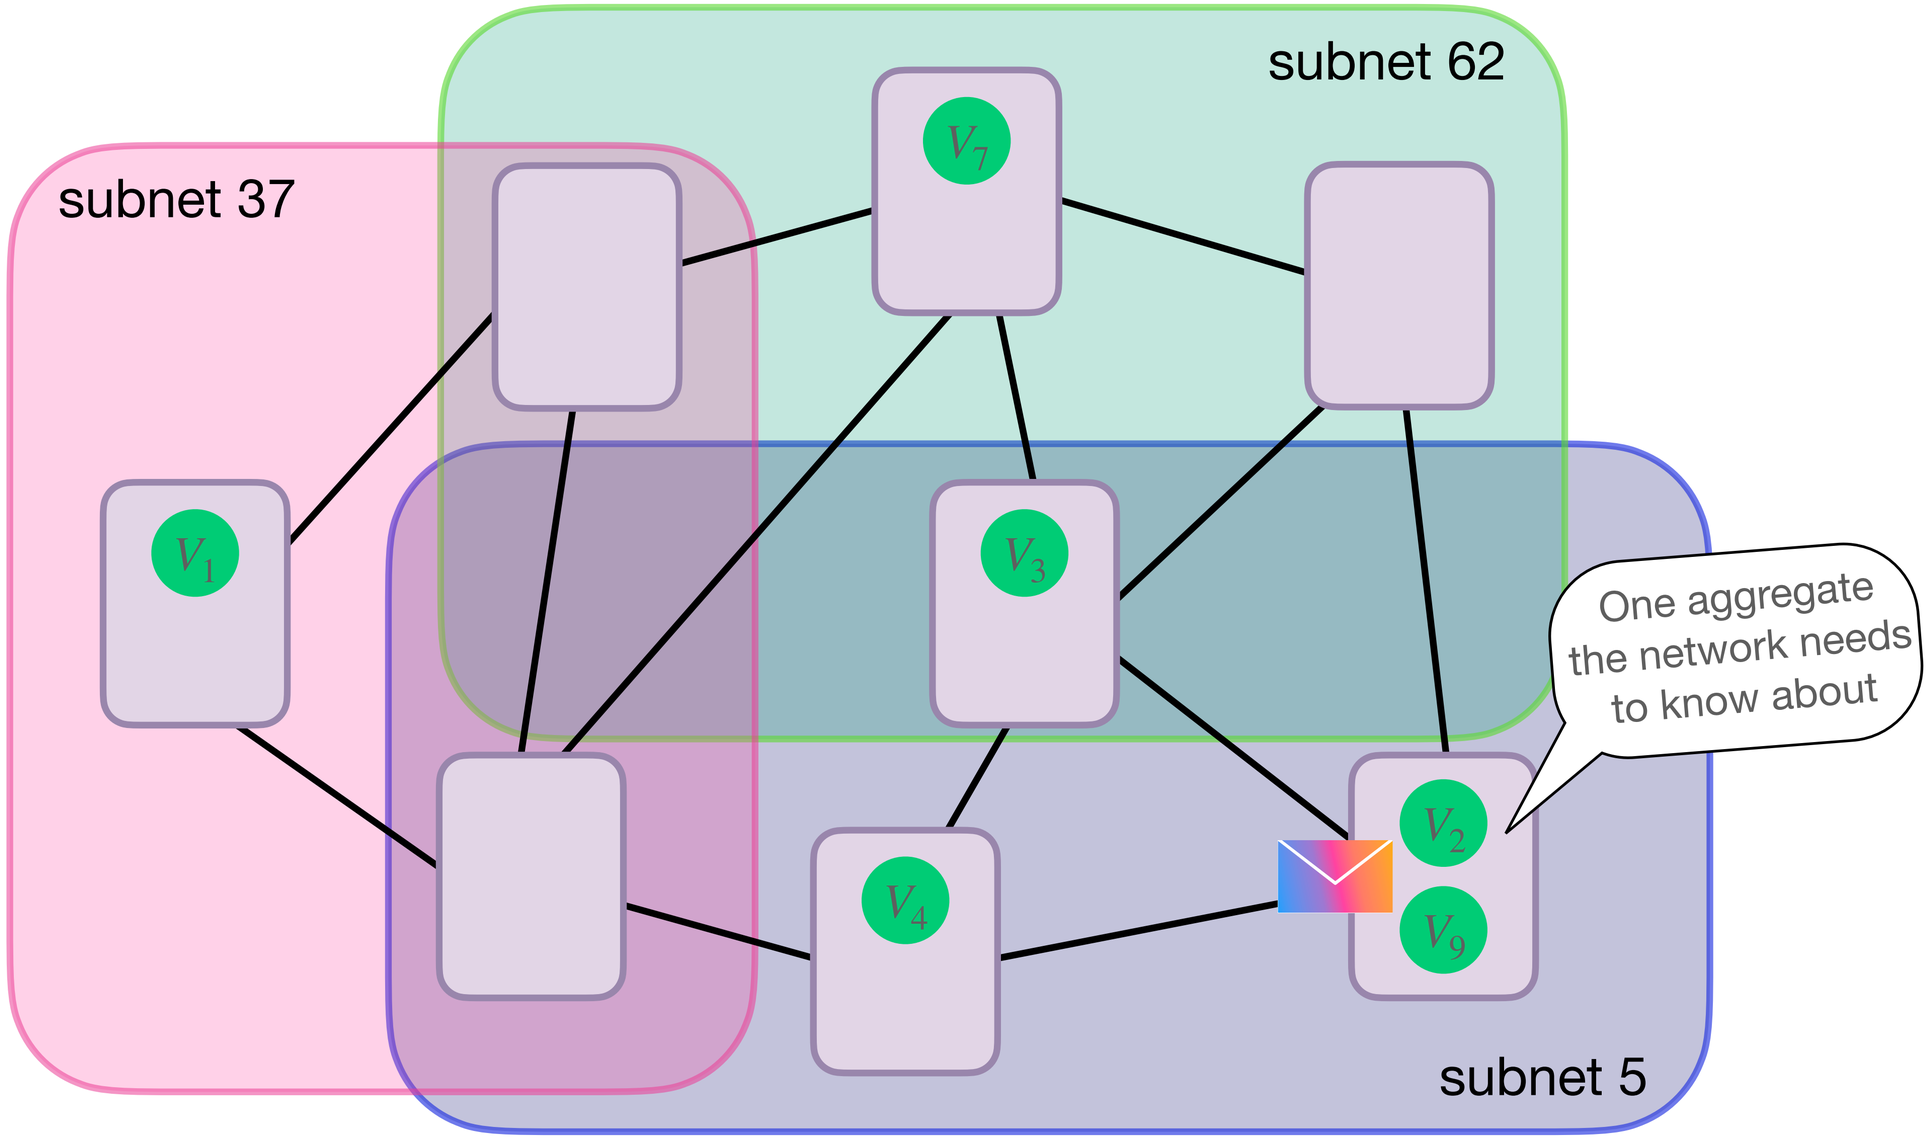

Do we give up privacy as a tradeoff for lower network load?

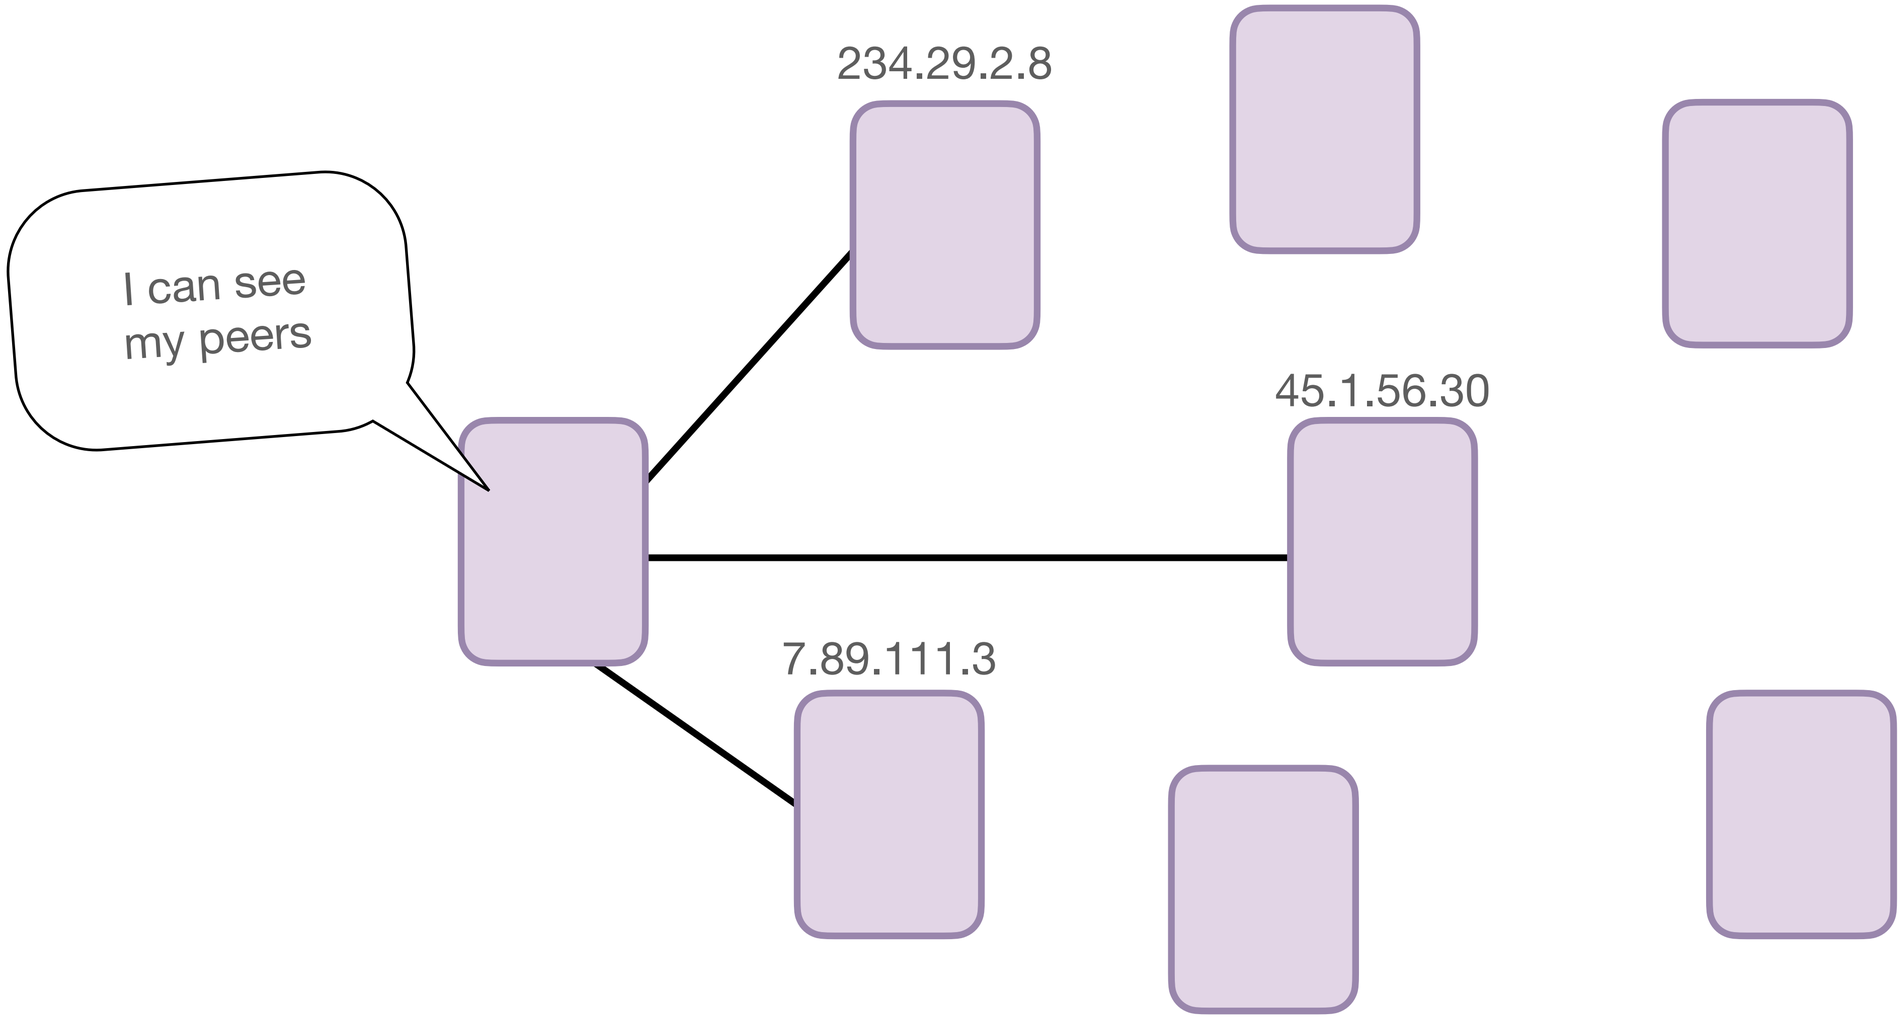
234.29.2.8
I can see
my peers
45.1.56.30
7.89.111.3

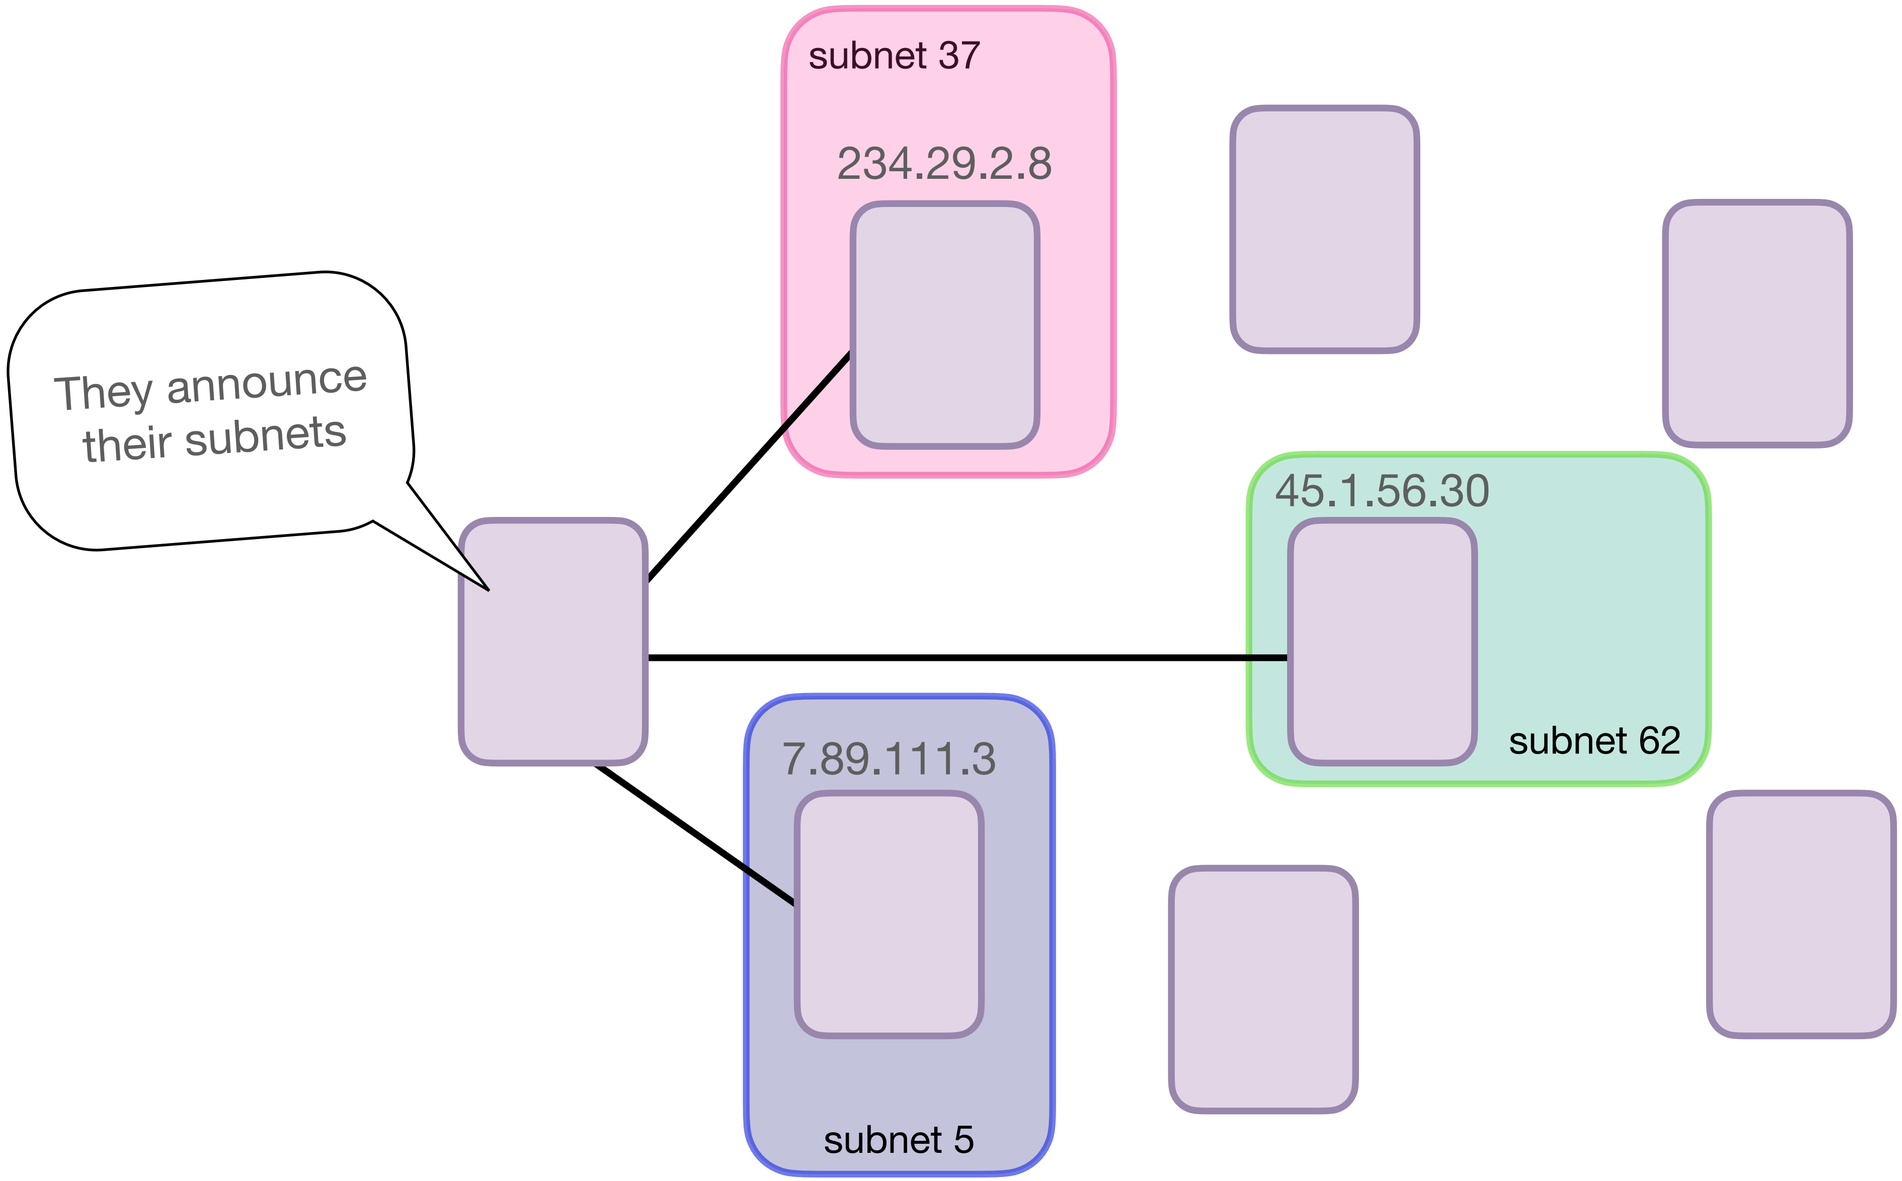
subnet 37
234.29.2.8
They announce
their subnets
45.1.56.30
subnet 62
7.89.111.3
subnet 5

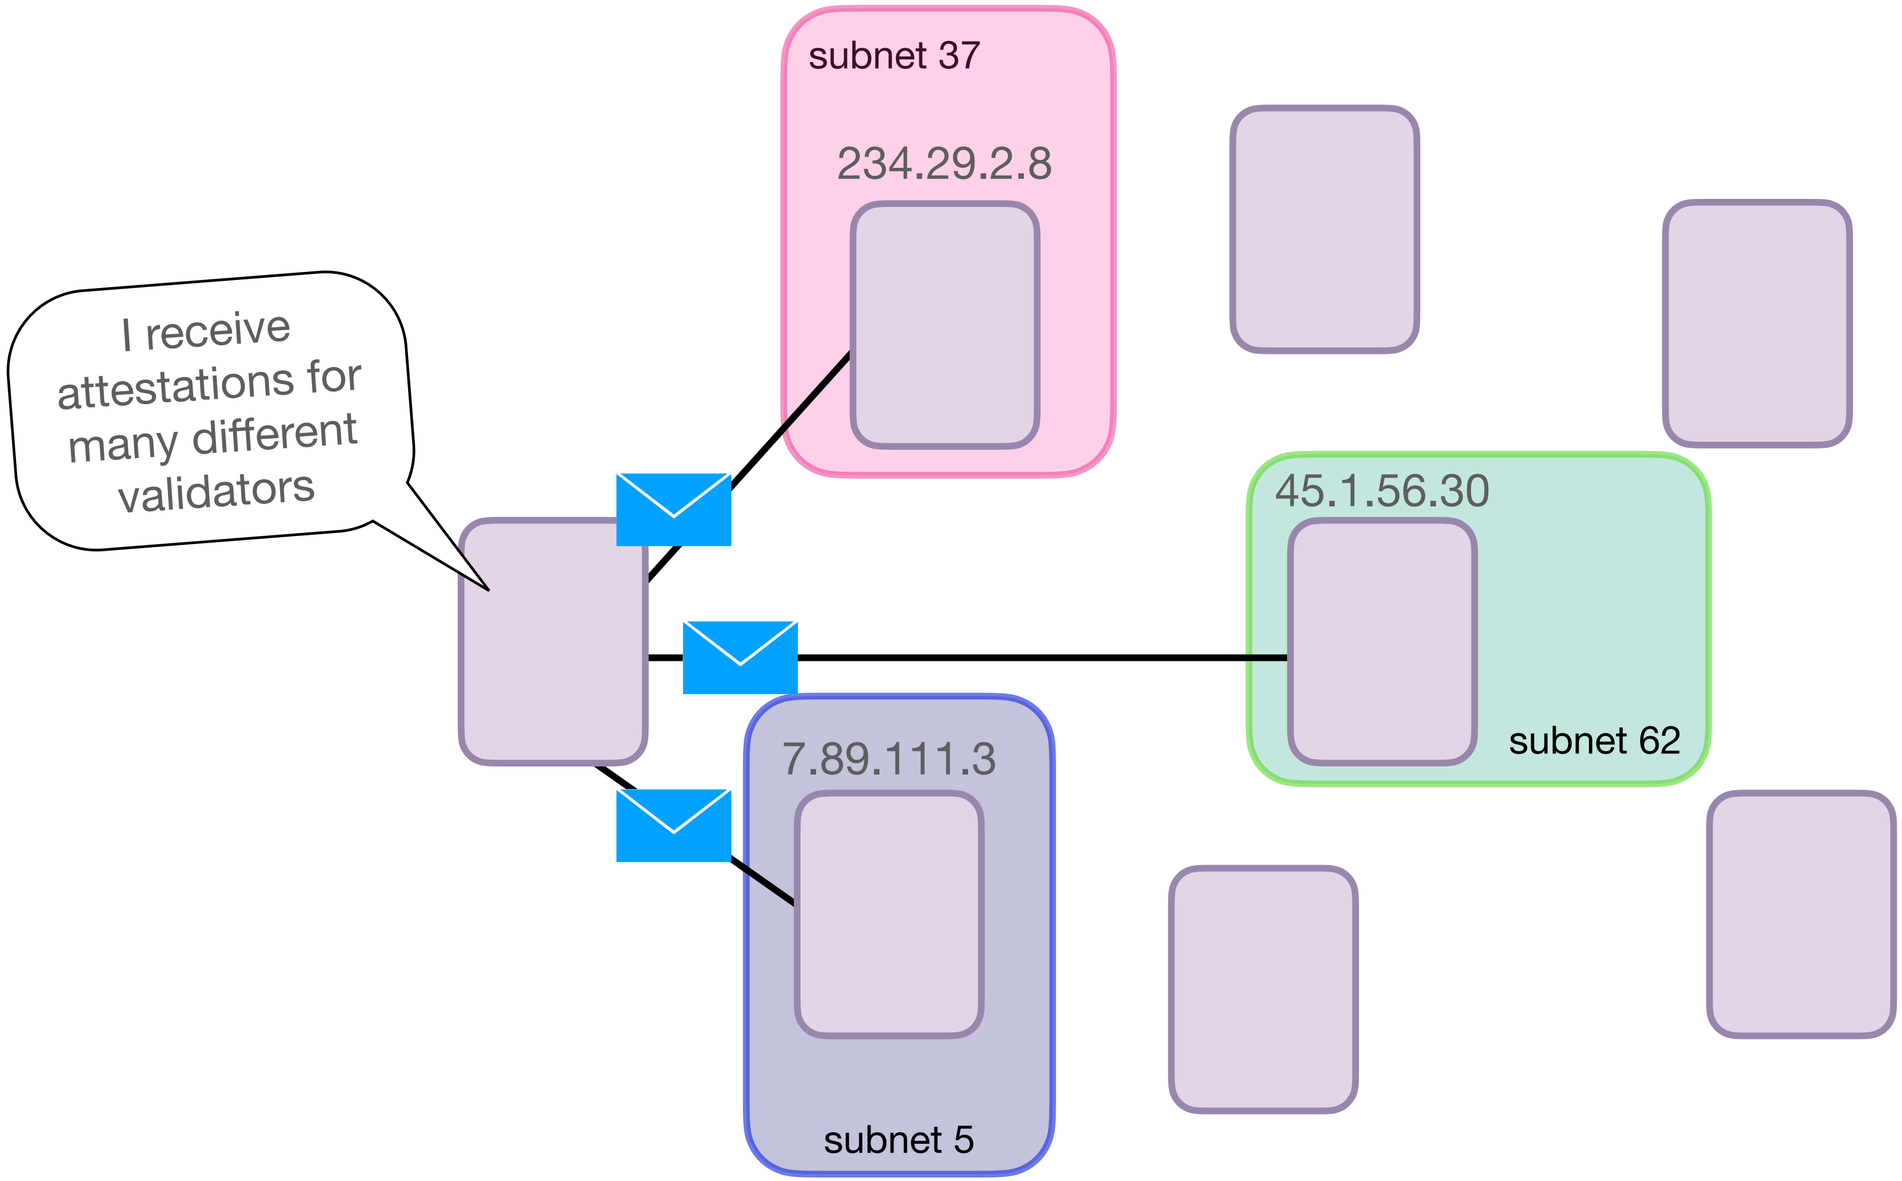
subnet 37
234.29.2.8
I receive attestations for many different validators
45.1.56.30
subnet 62
7.89.111.3
subnet 5

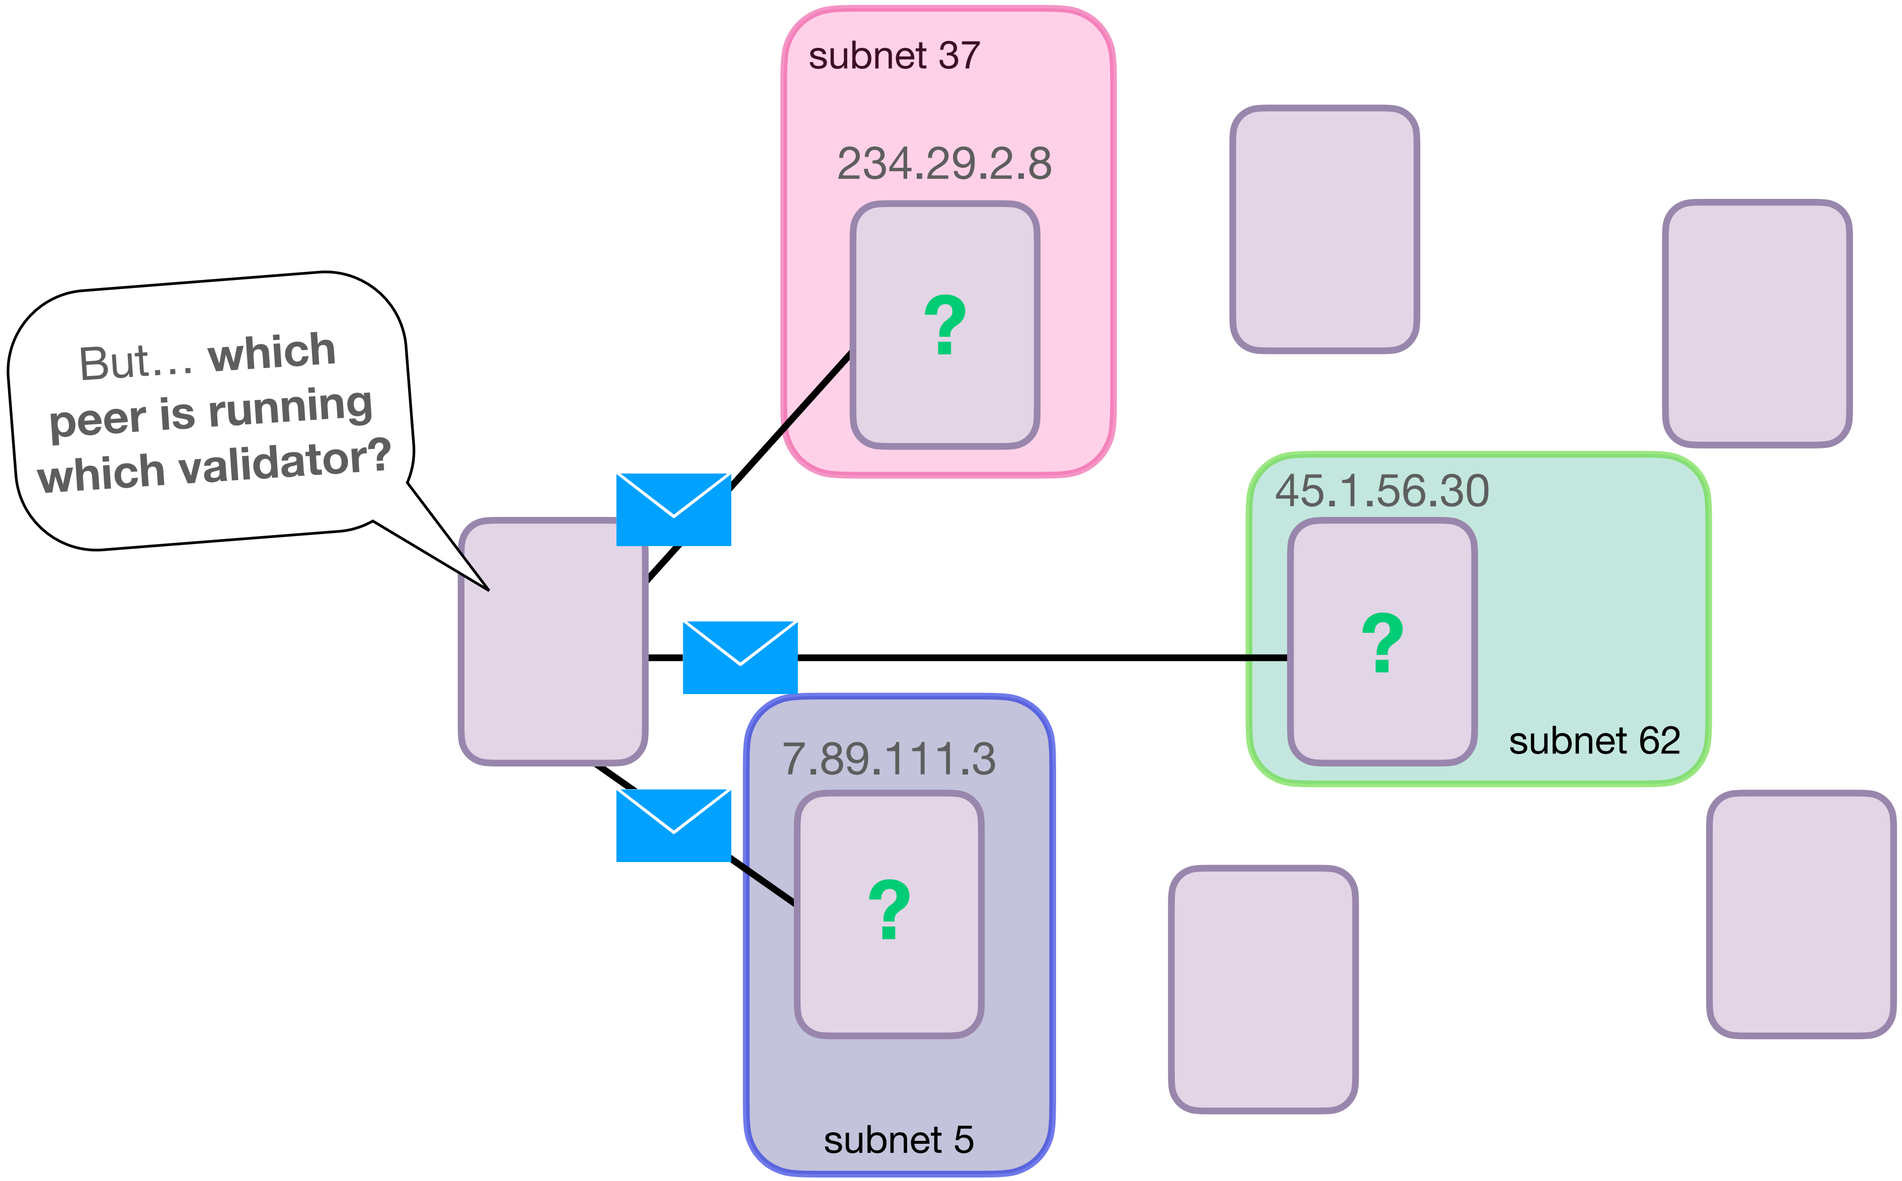
subnet 37
234.29.2.8
?
But… which peer is running which validator?
45.1.56.30
?
subnet 62
7.89.111.3
?
subnet 5

Why does anonymity matter?

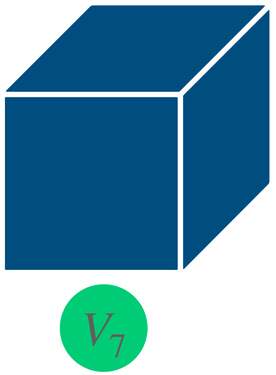

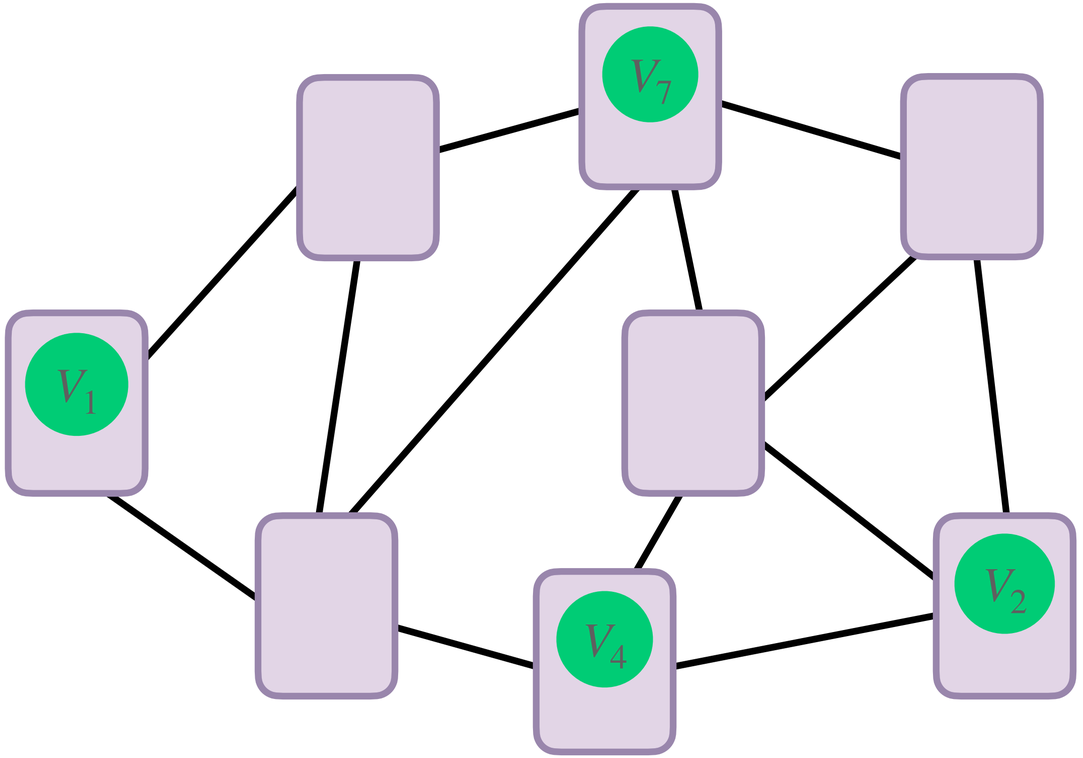

Why does anonymity matter?

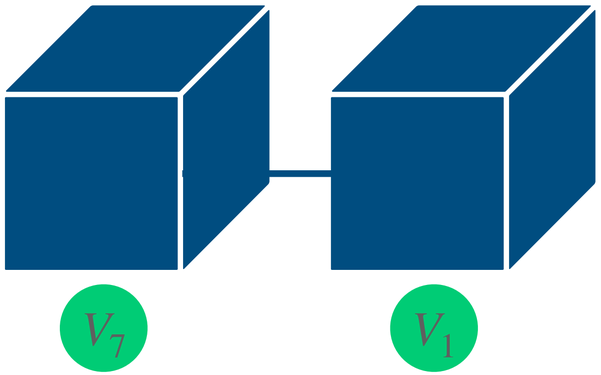

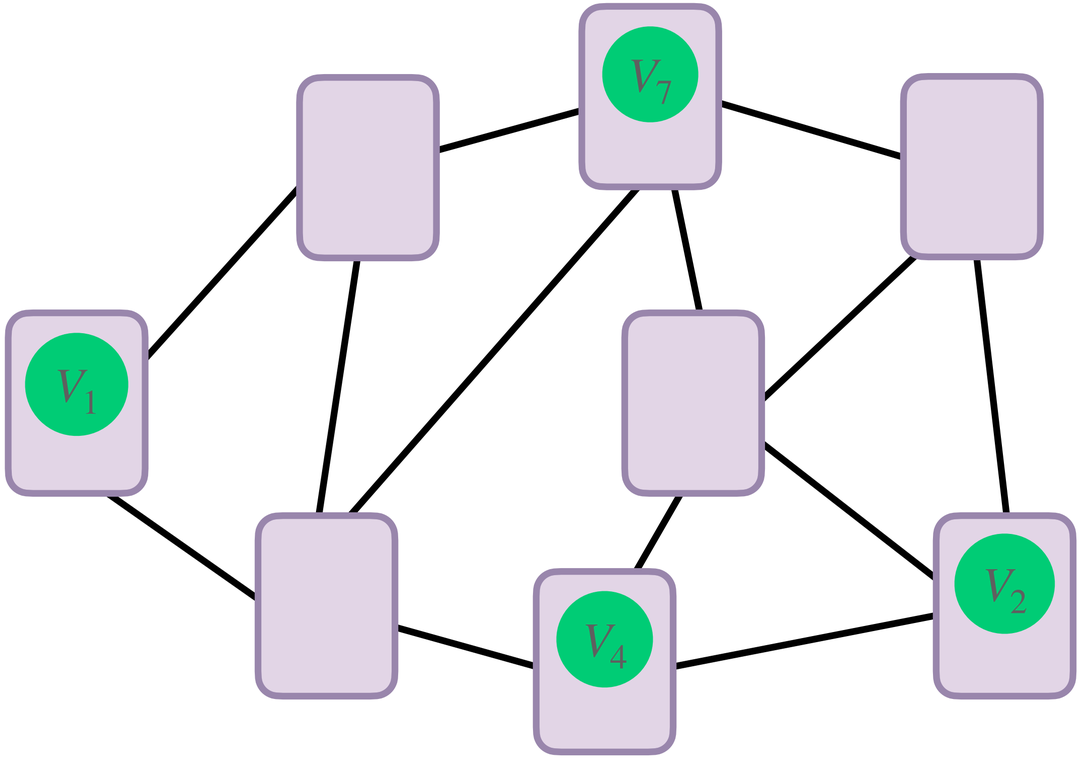

Why does anonymity matter?

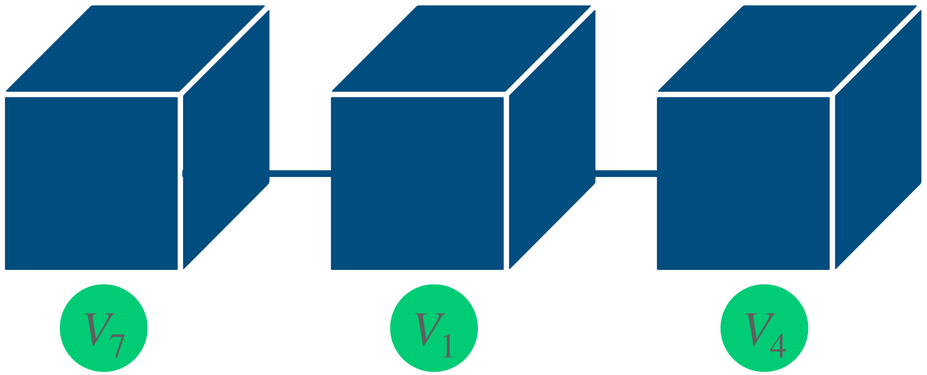

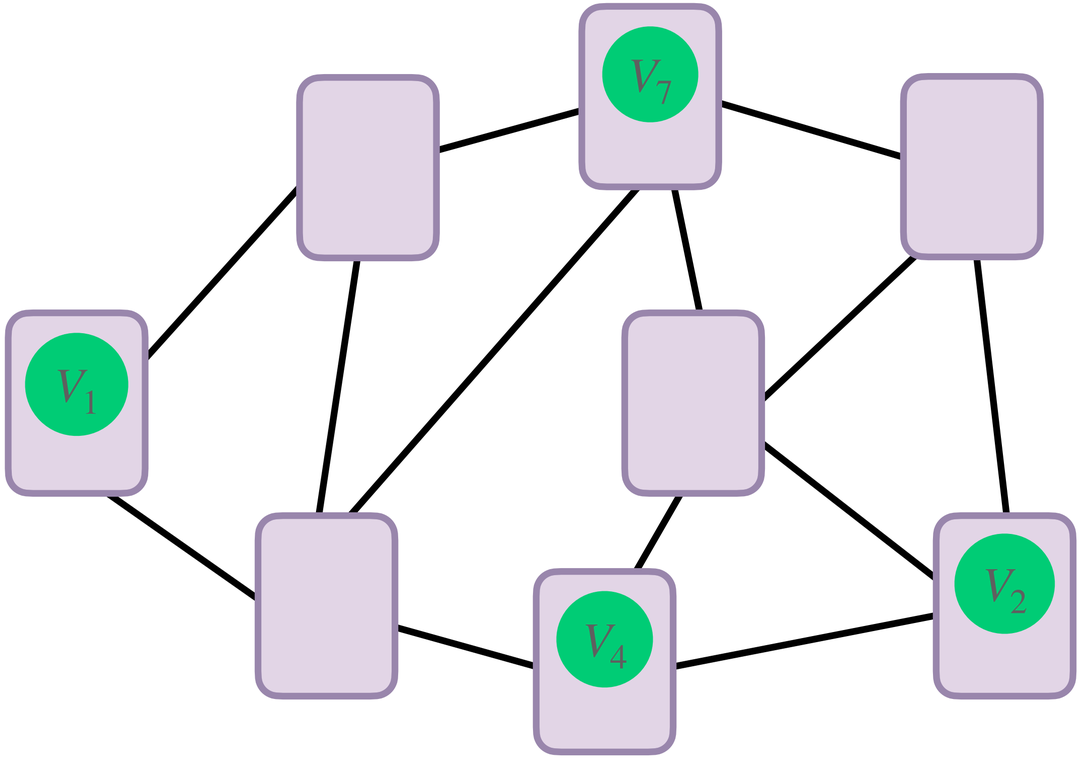

Why does anonymity matter?

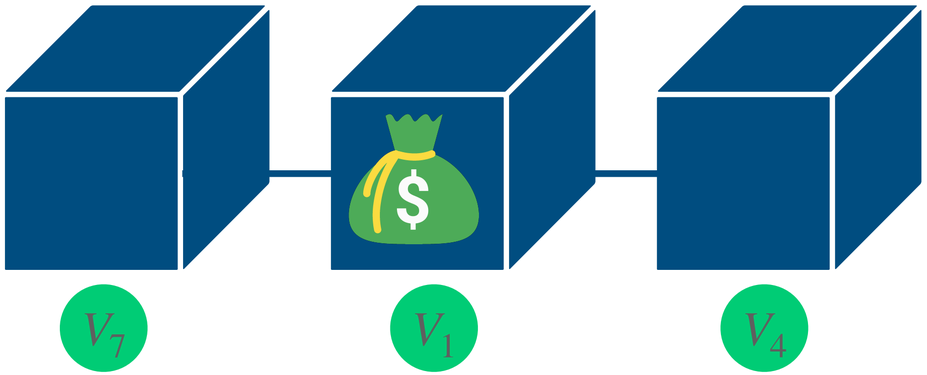

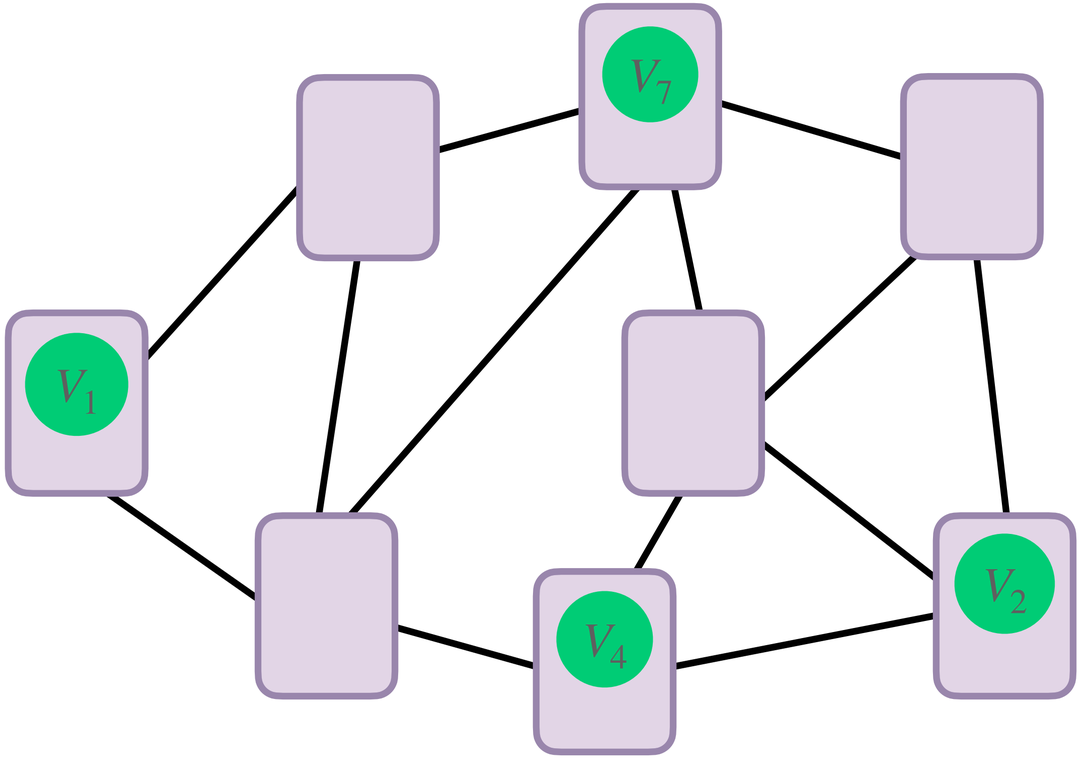

Why does anonymity matter?

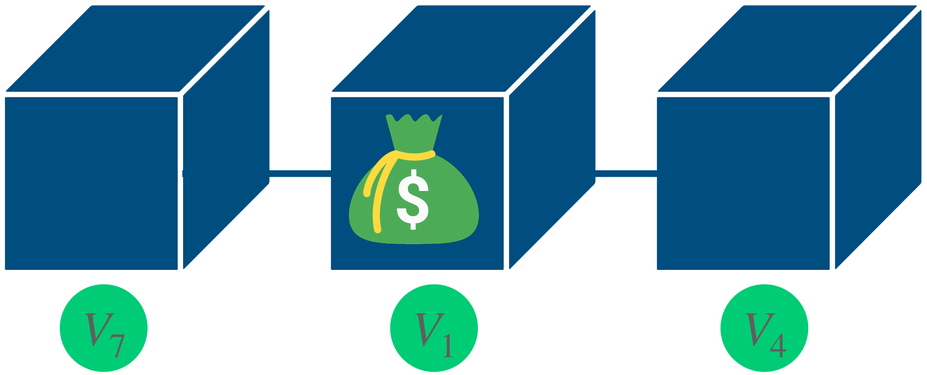

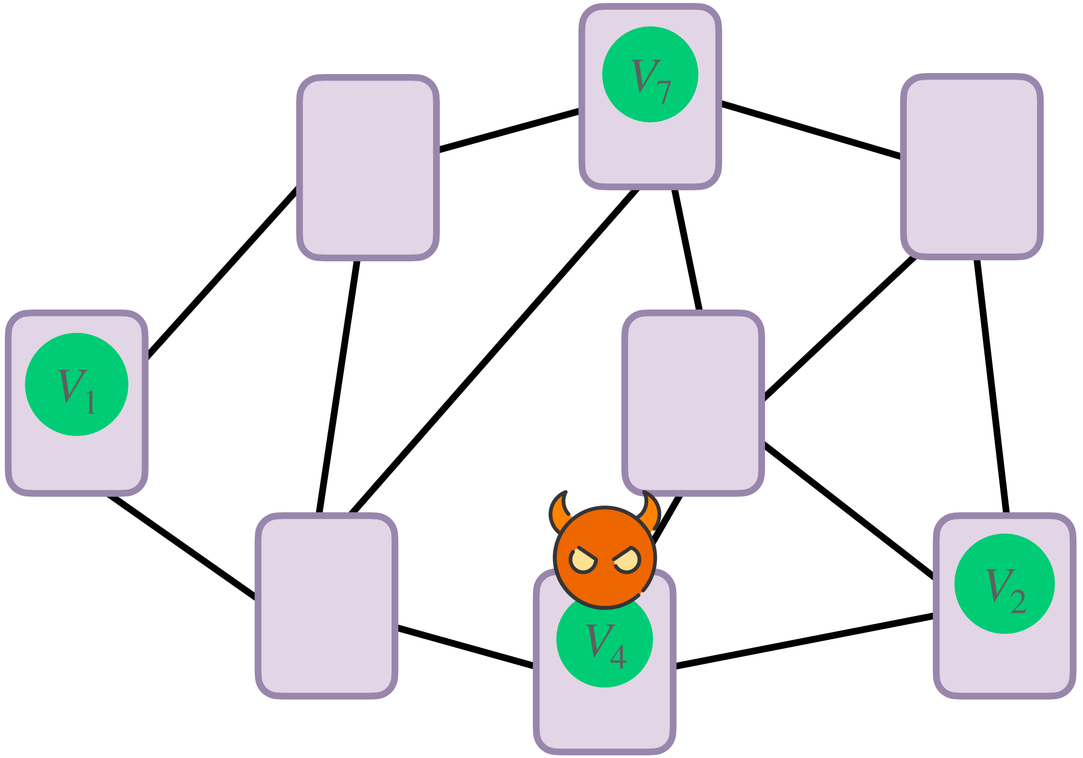

Why does anonymity matter?

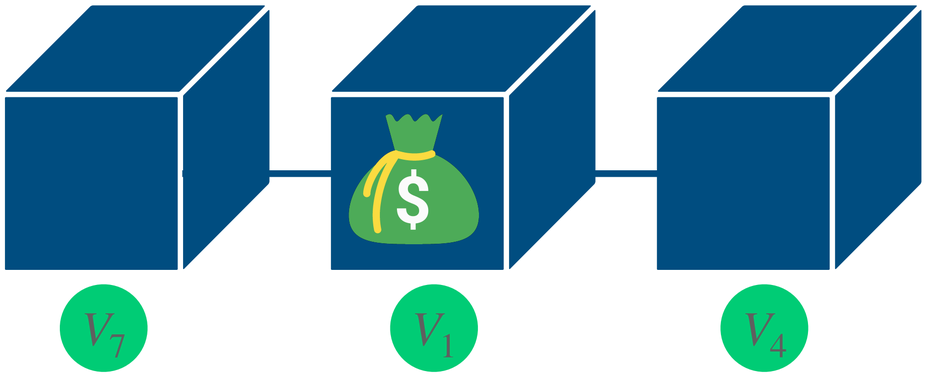

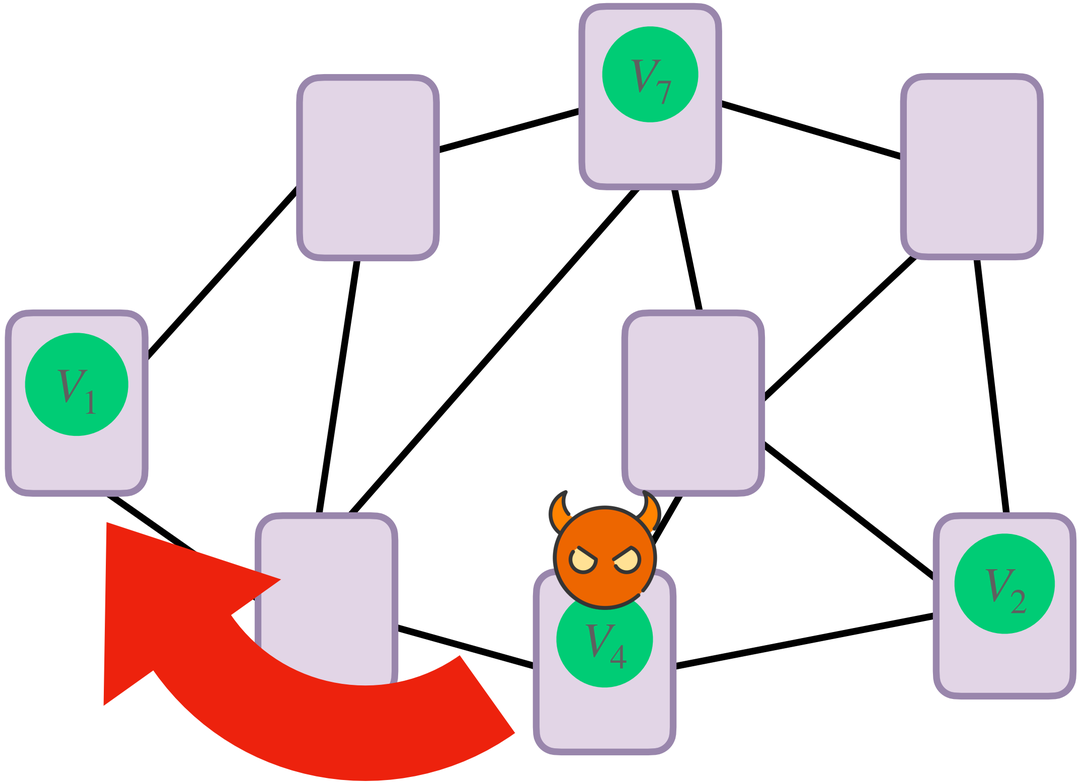

Why does anonymity matter?

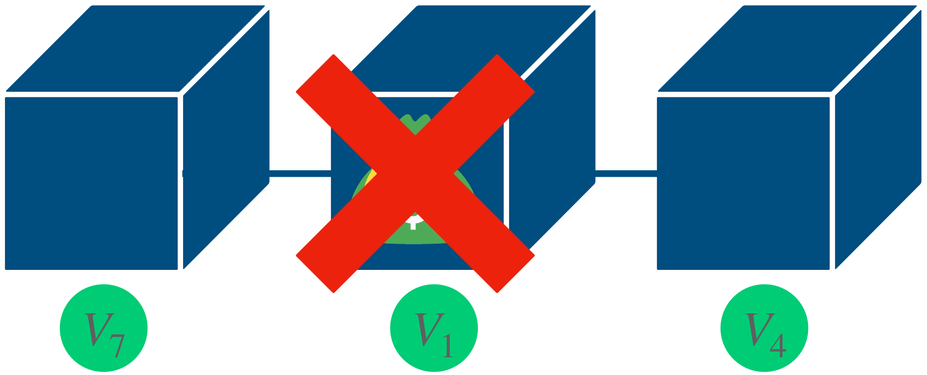

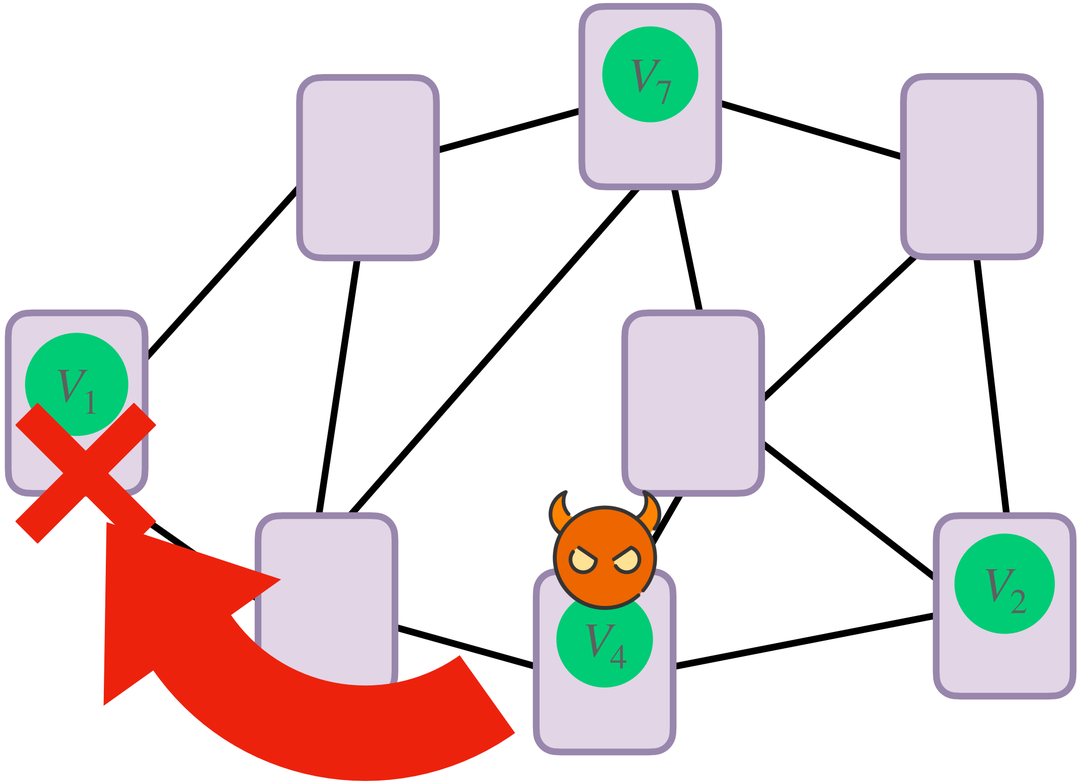

Why does anonymity matter?

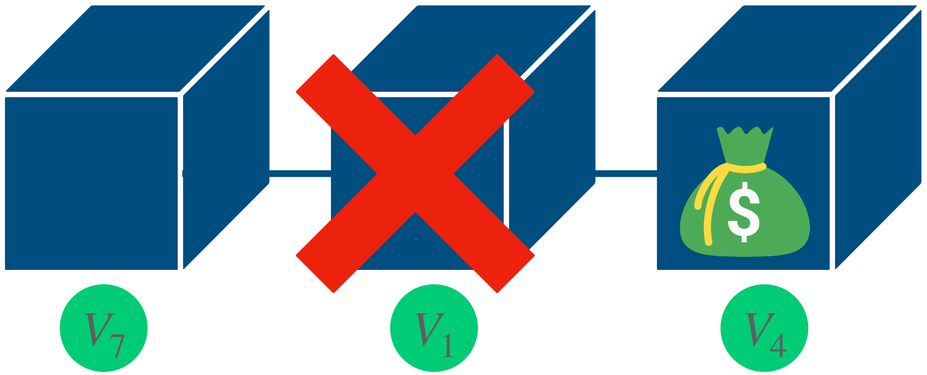

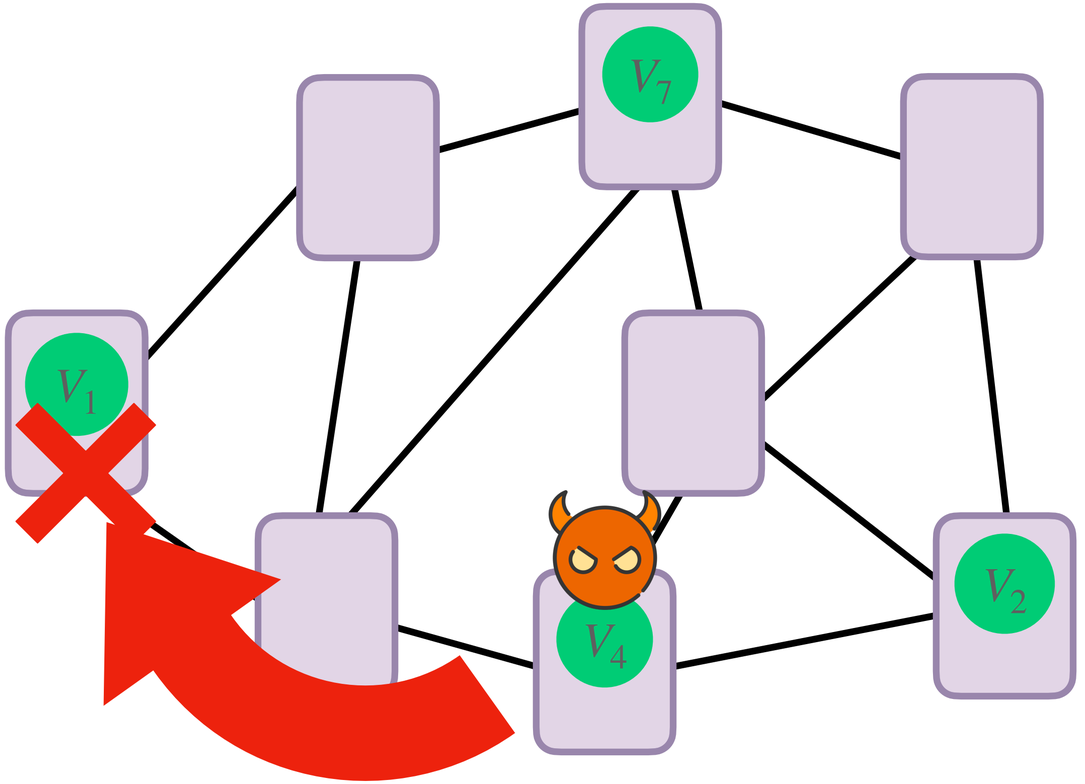

Implications of de-anonymizing validators:

De-anonymization Methodology


I am in
all subnets!
?

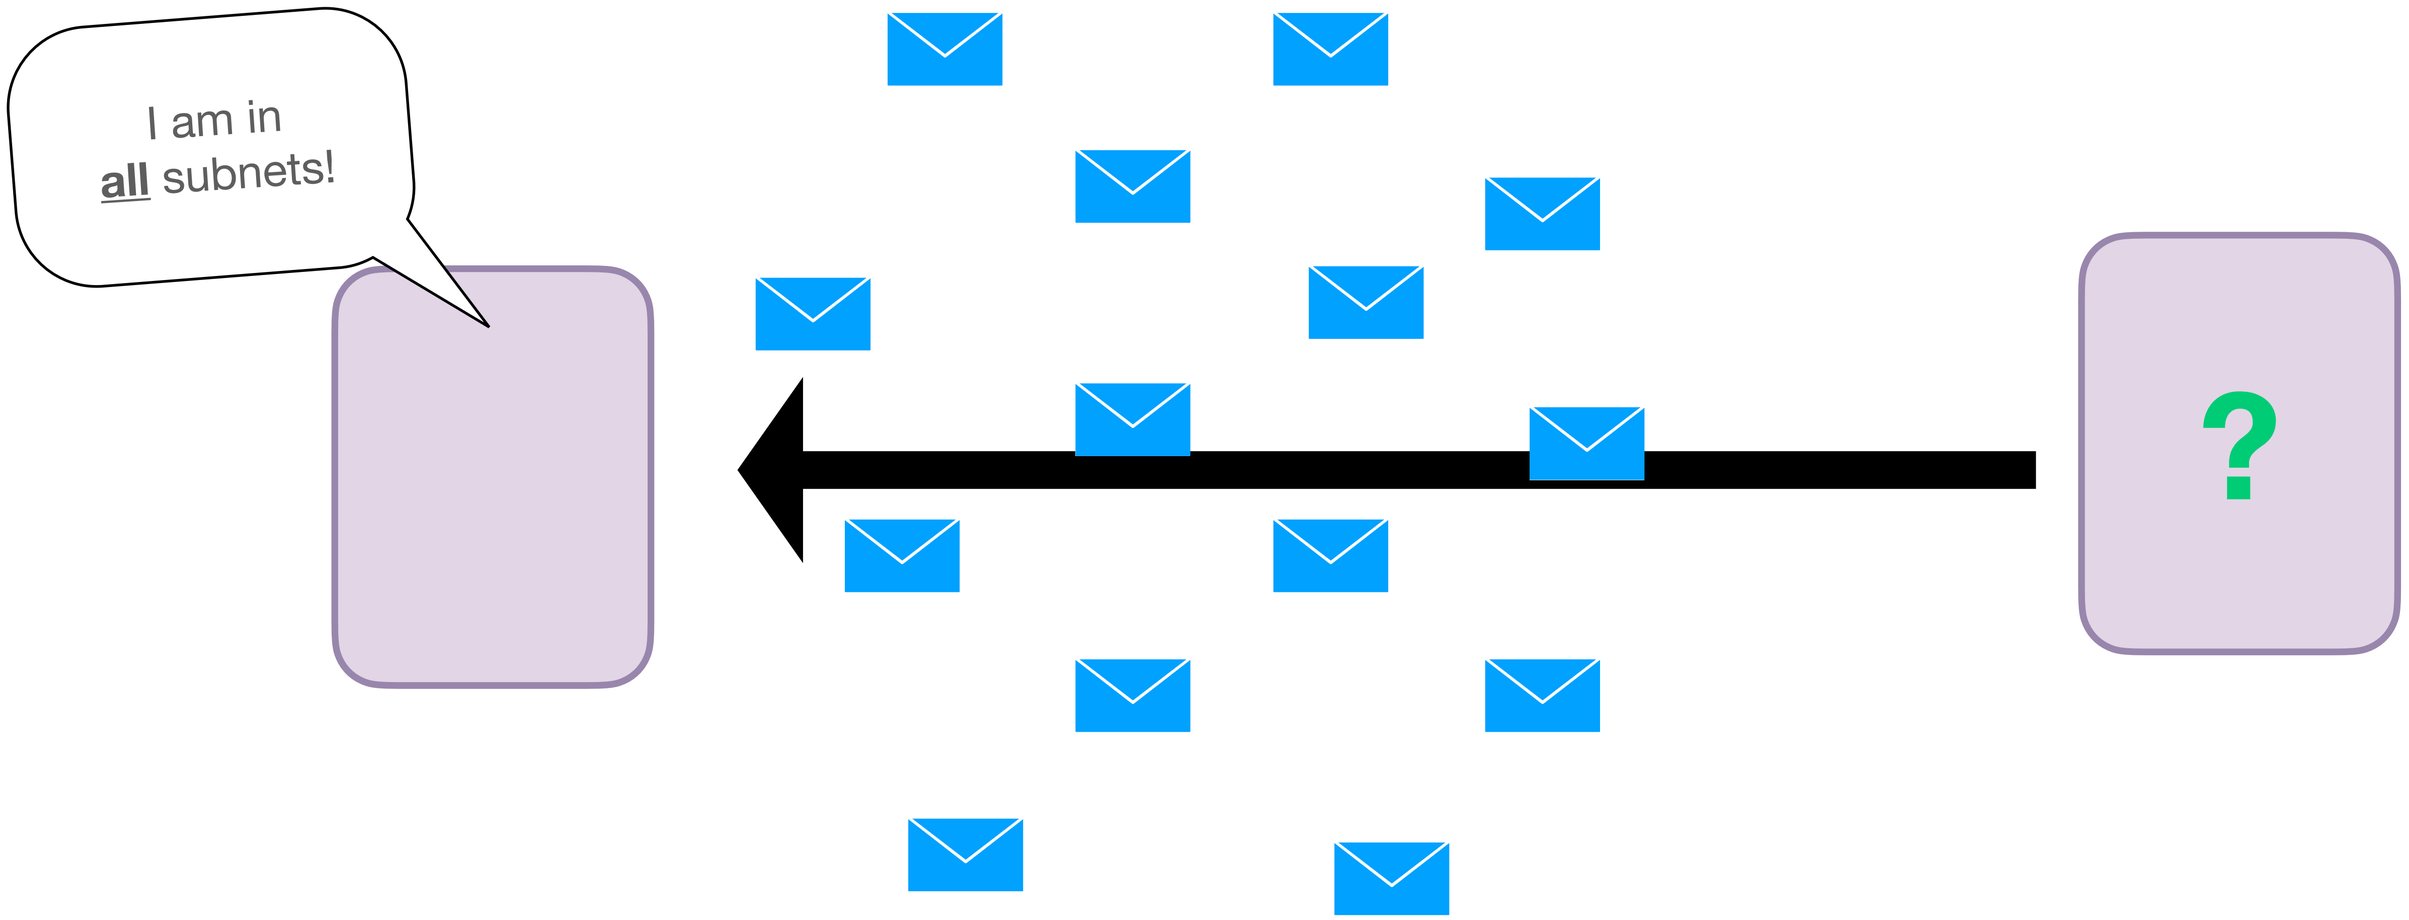
I am in
all subnets!
?

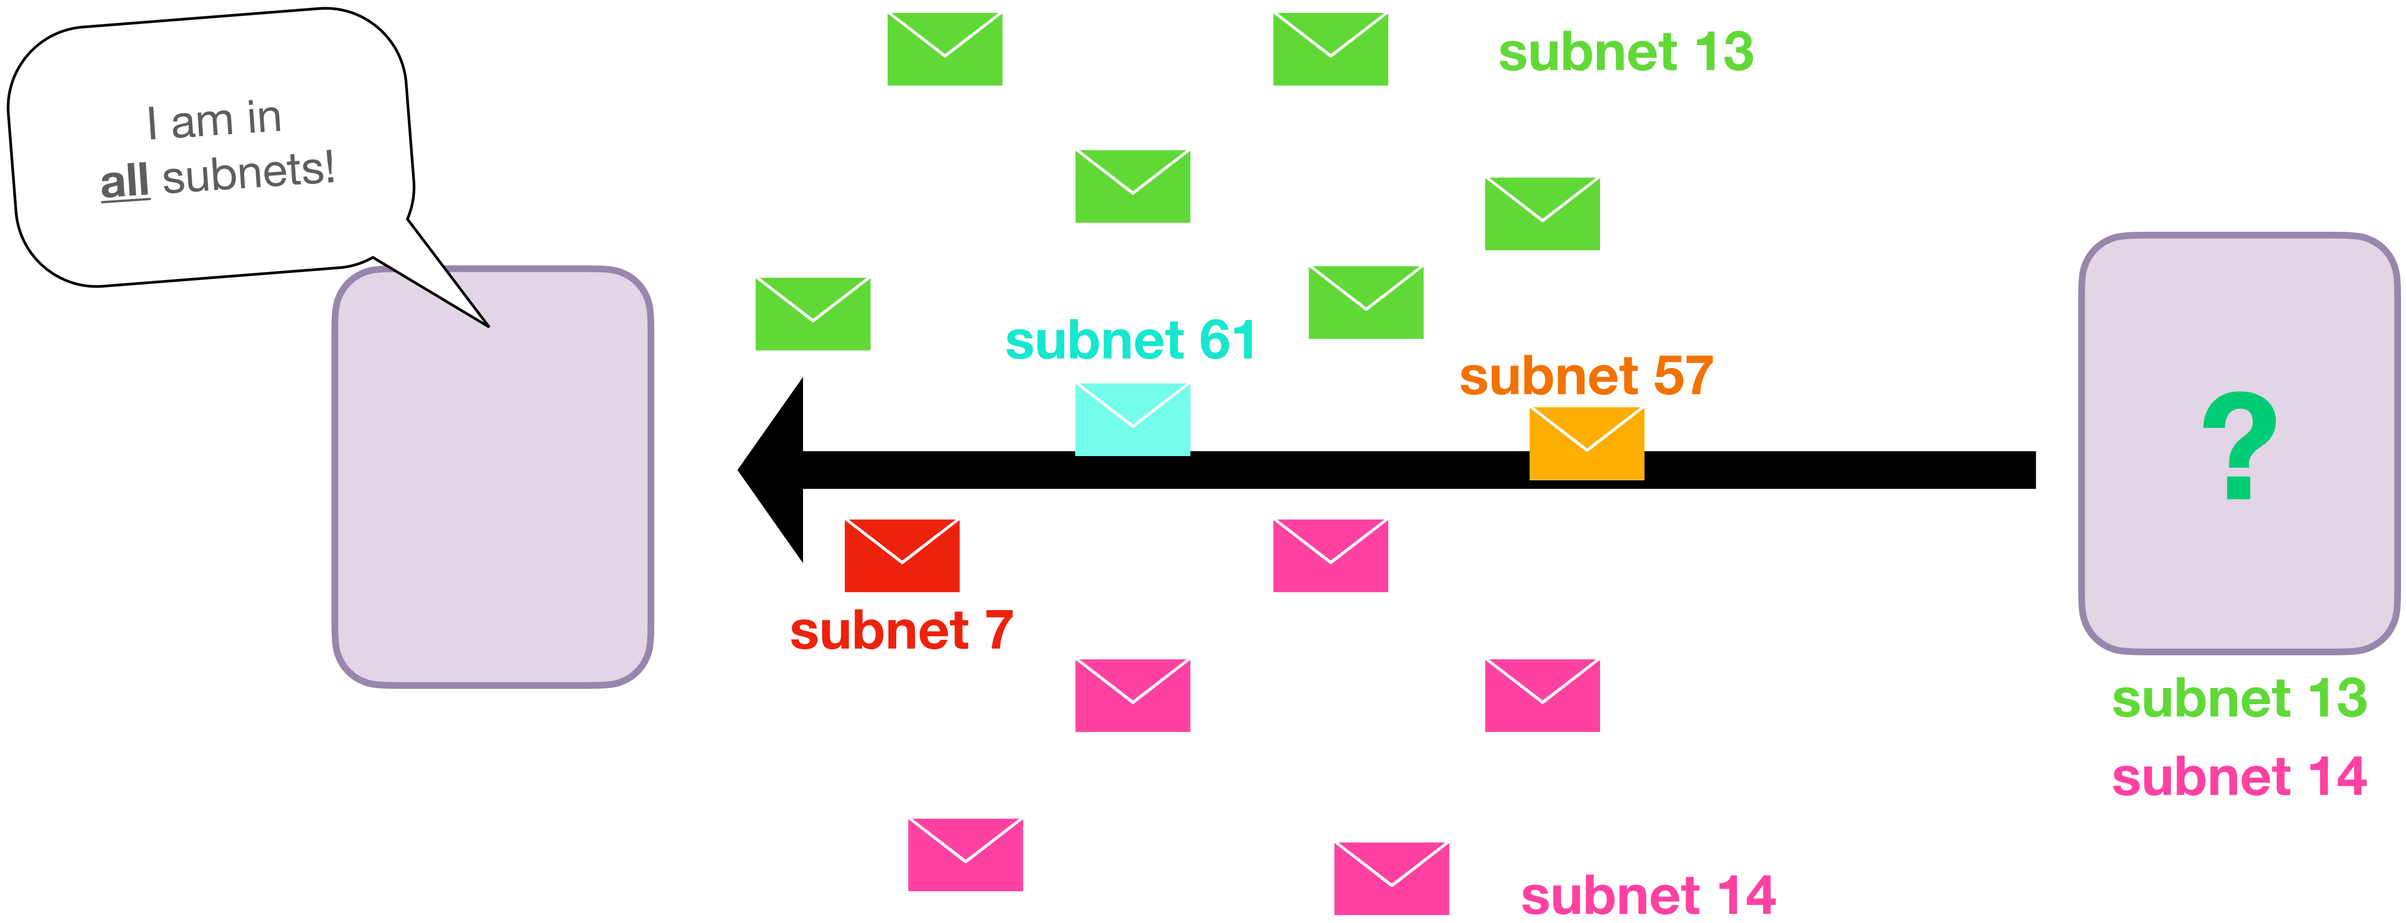
I am in
all subnets!
subnet 13
subnet 61
subnet 57
subnet 7
subnet 14
?
subnet 13
subnet 14

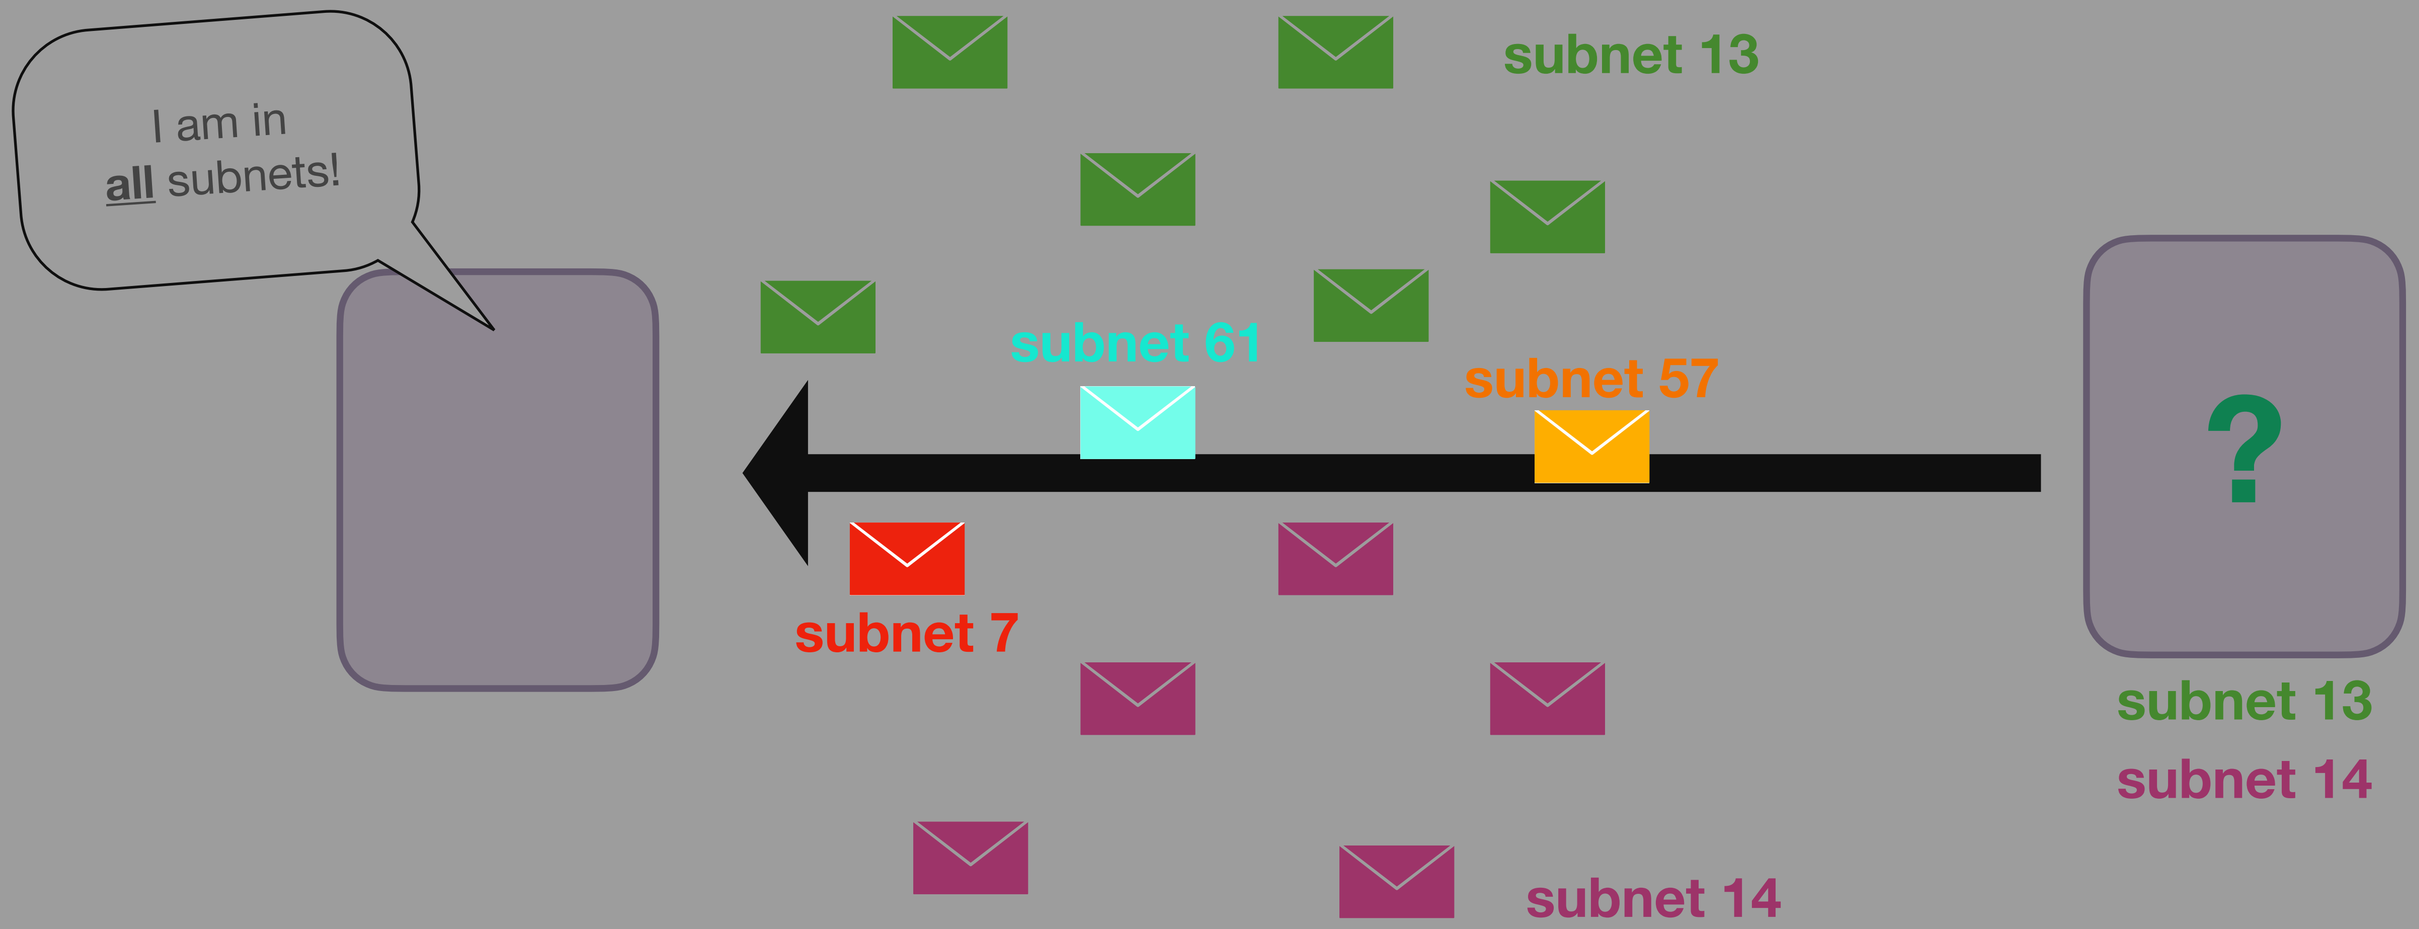
I am in
all subnets!
subnet 13
subnet 61
subnet 57
subnet 7
subnet 14
?
subnet 13
subnet 14

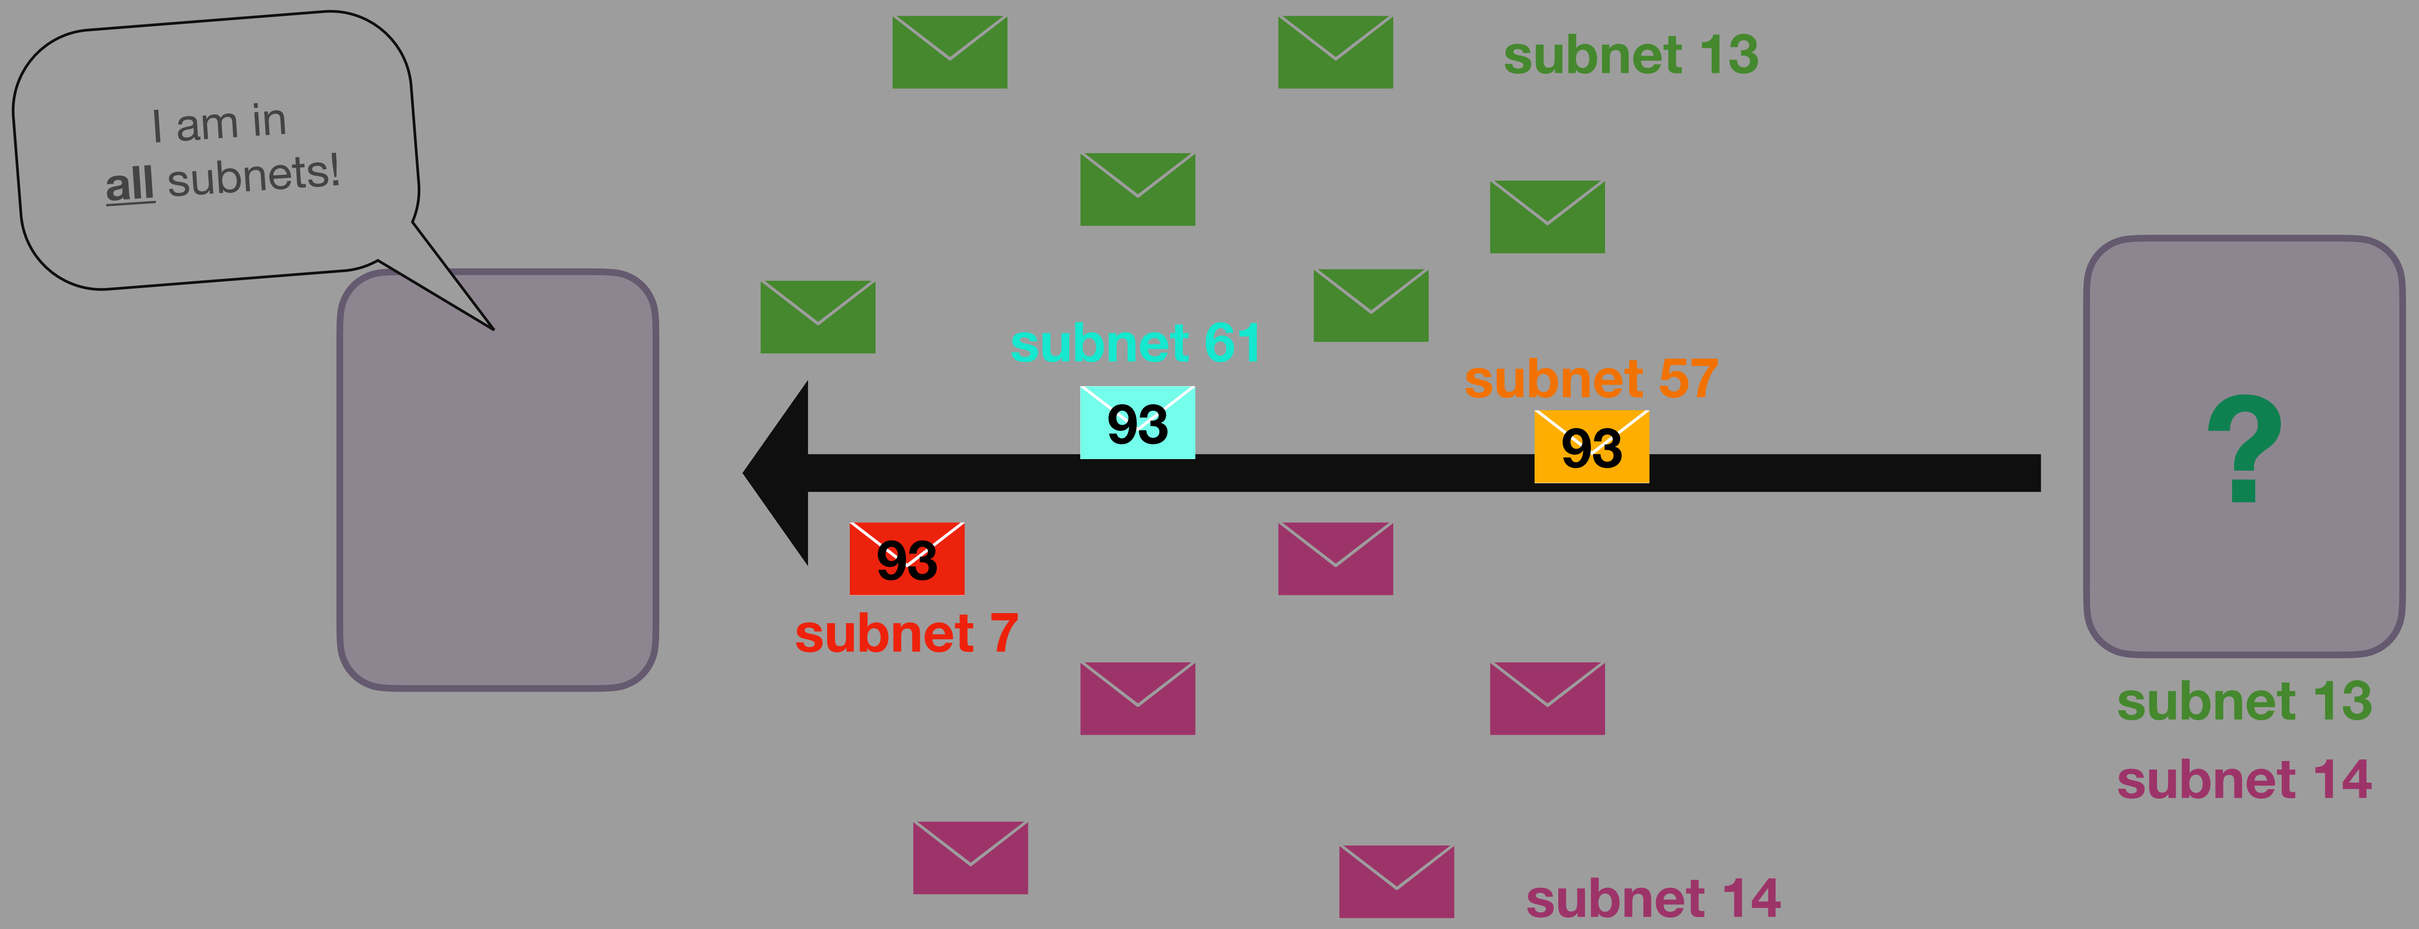
I am in
all subnets!
subnet 13
subnet 61
93
subnet 57
93
93
subnet 7
?
subnet 13
subnet 14
subnet 14

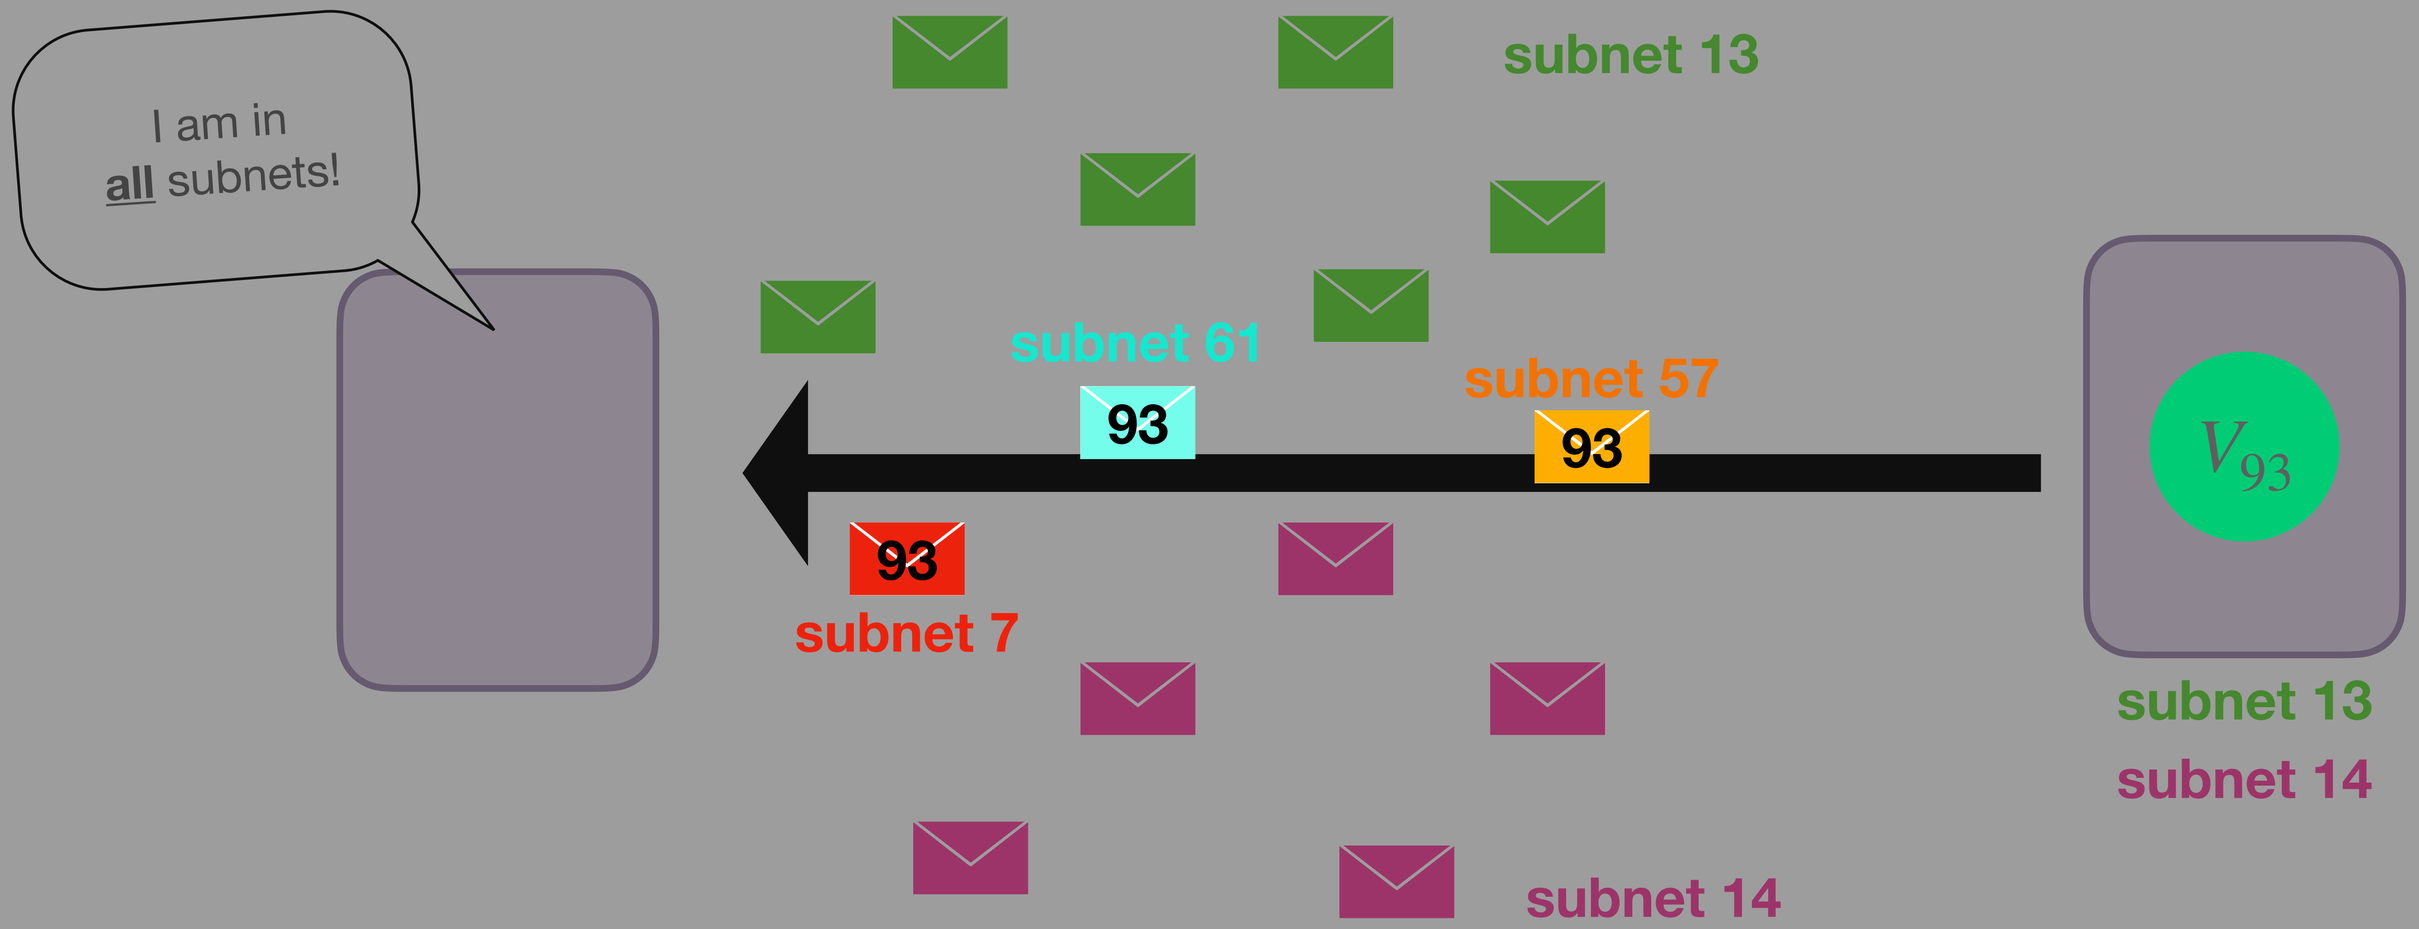
I am in
all subnets!
subnet 13
subnet 61
93
subnet 57
93
93
subnet 7
V_{93}
subnet 13
subnet 14
subnet 14

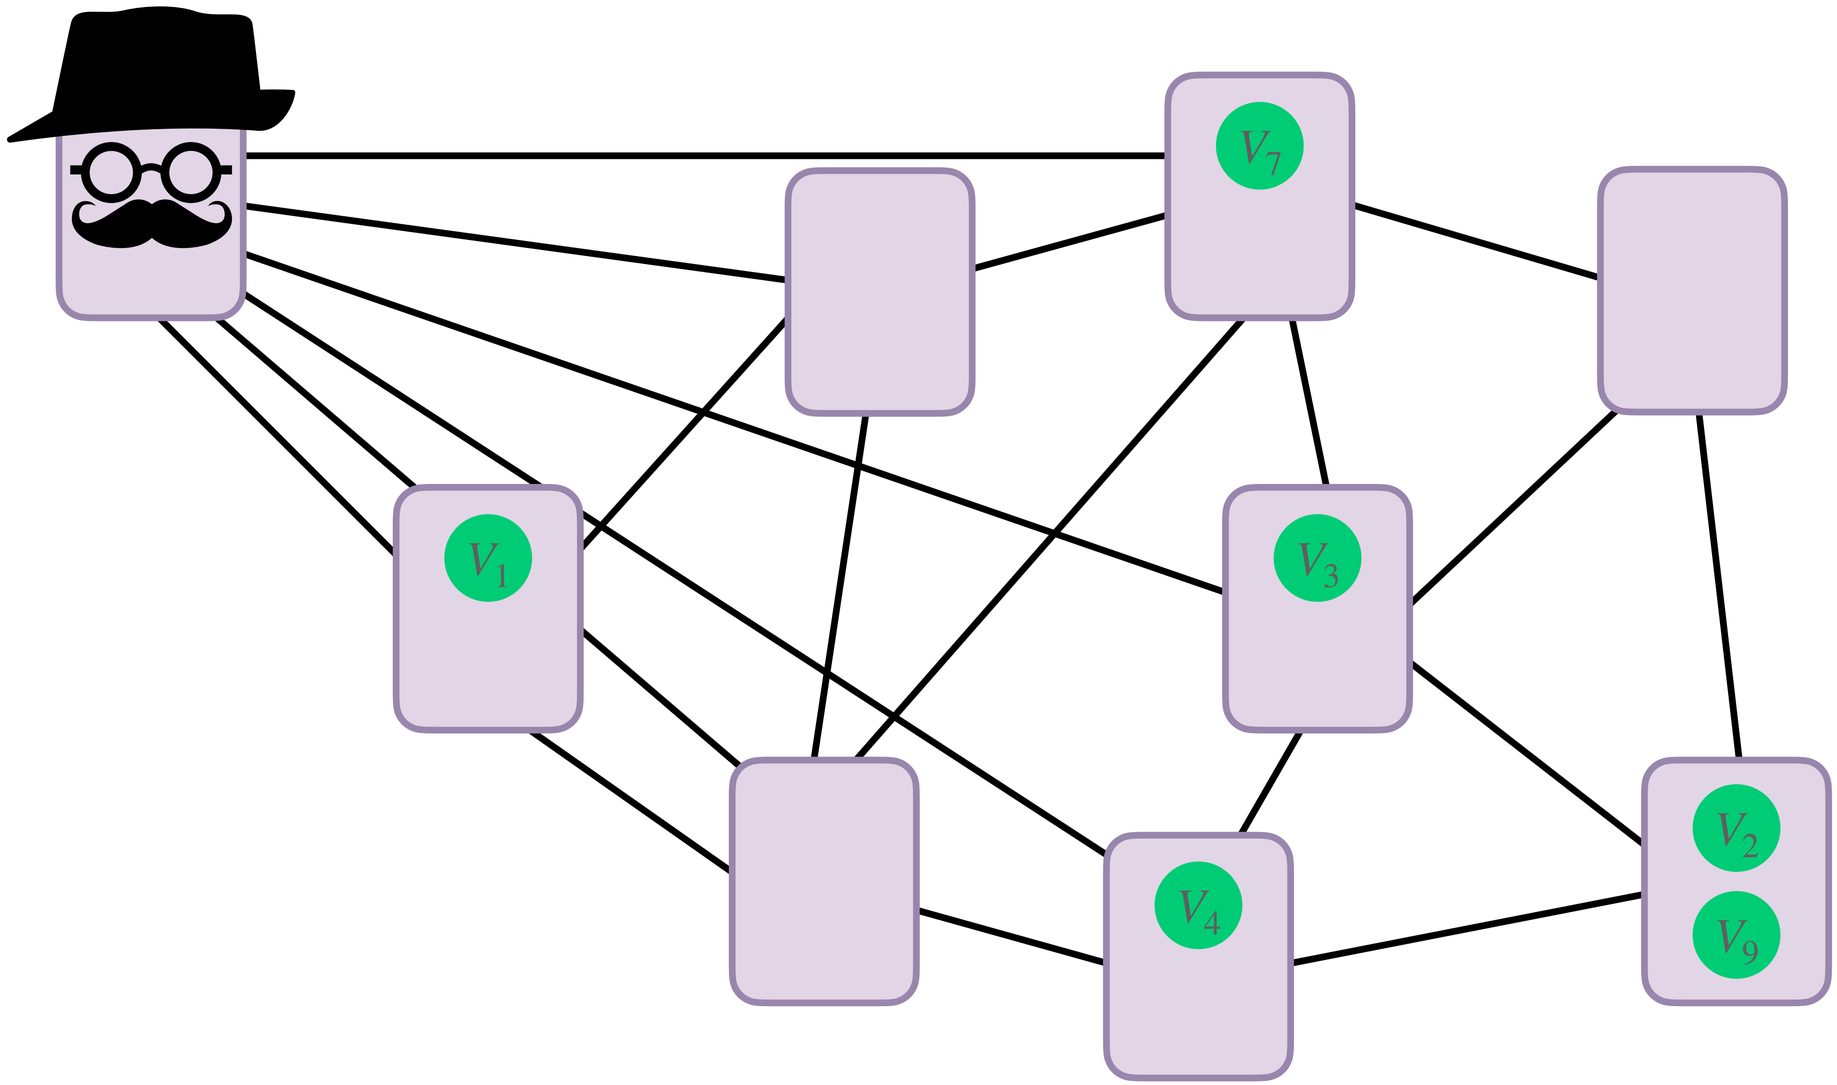
V_7
V_1
V_3
V_4
V_2
V_9

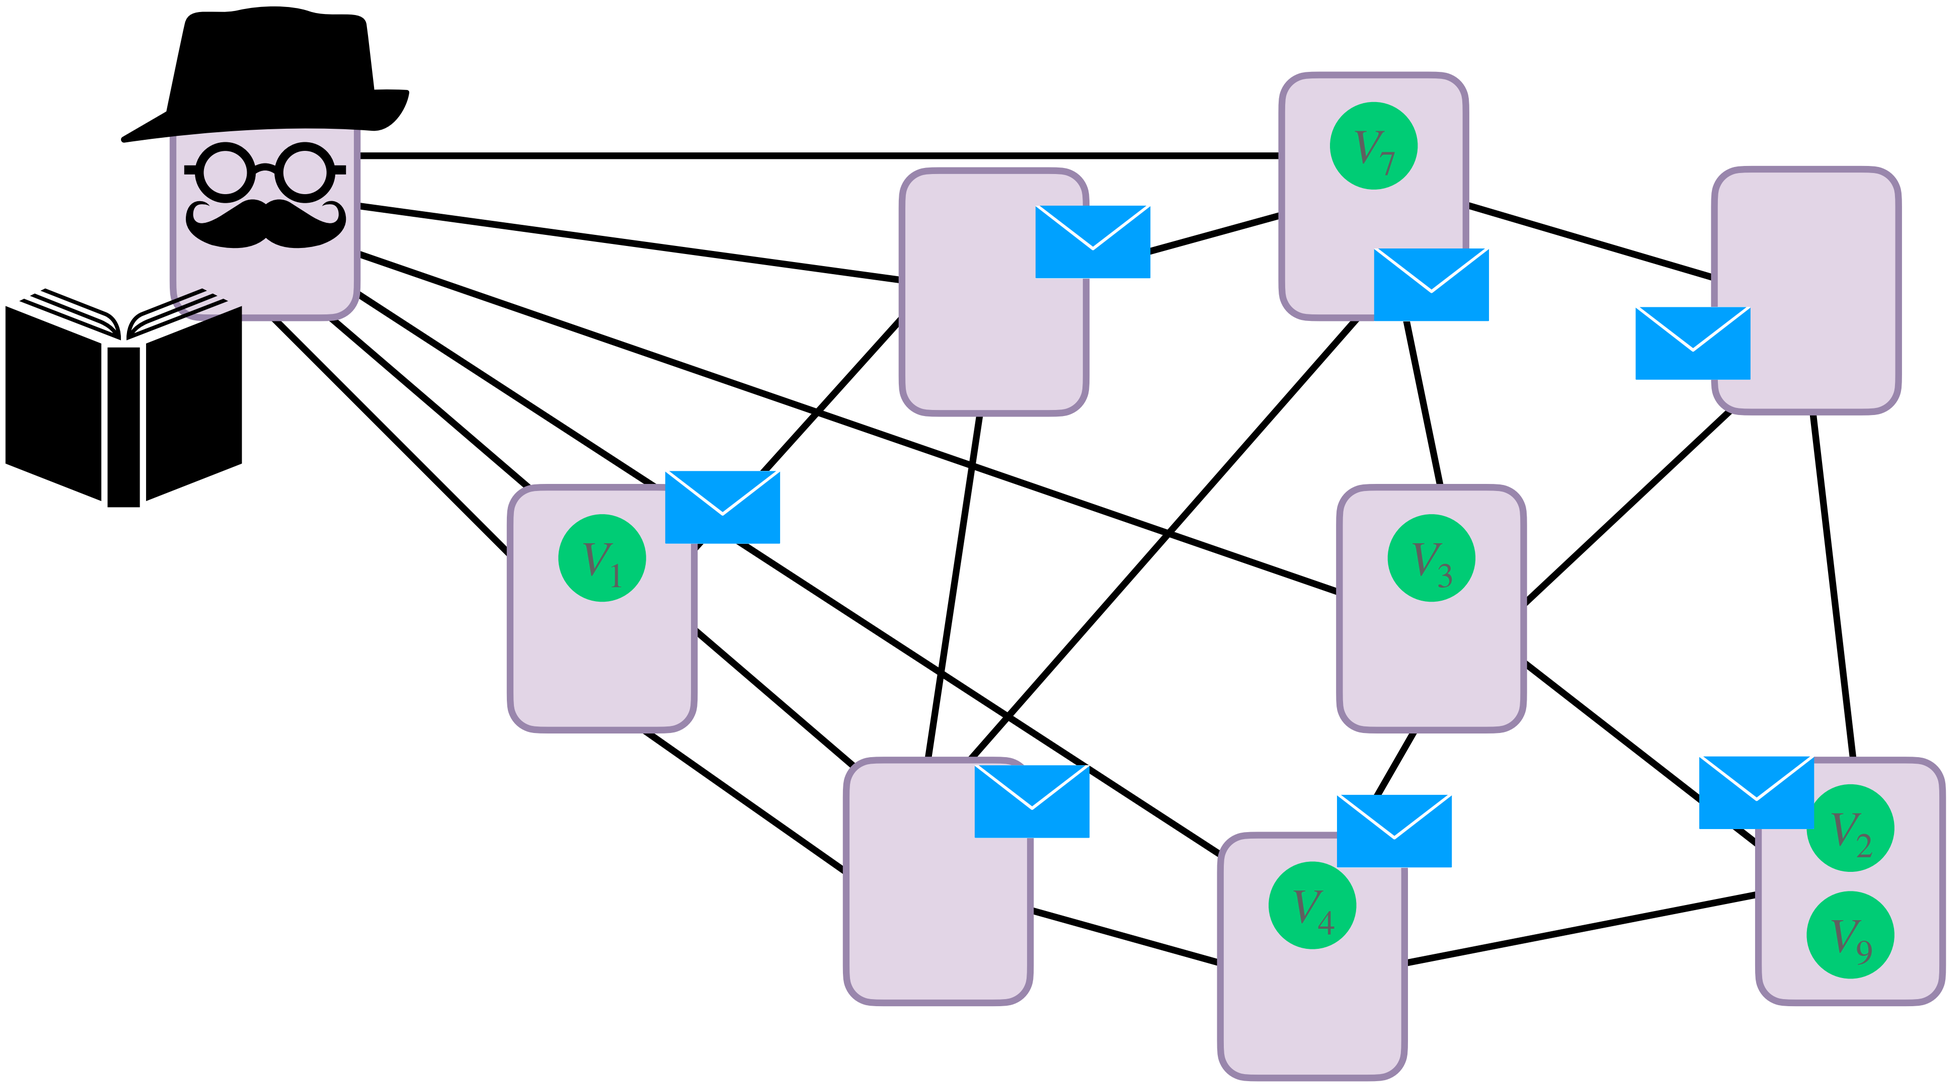
V_7
V_1
V_3
V_4
V_2
V_9

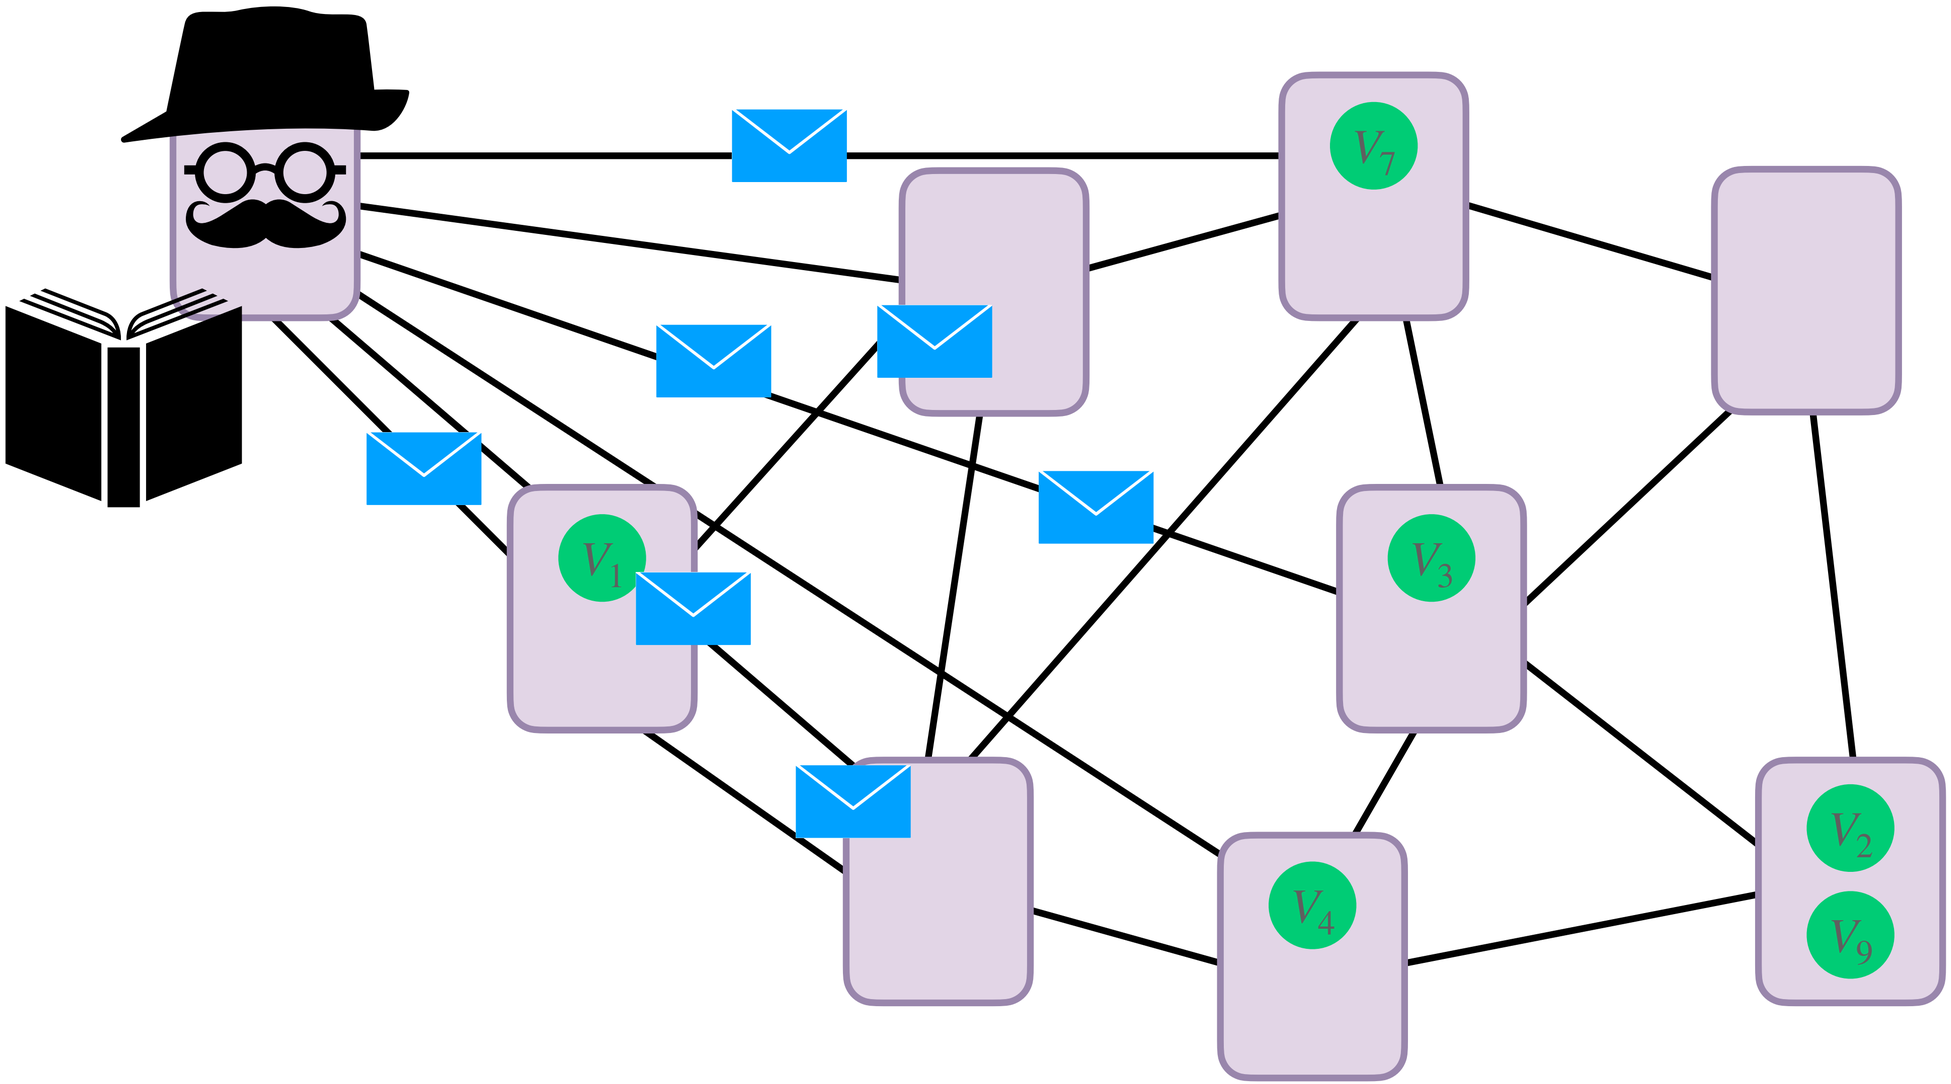
V_7
V_1
V_3
V_4
V_2
V_9

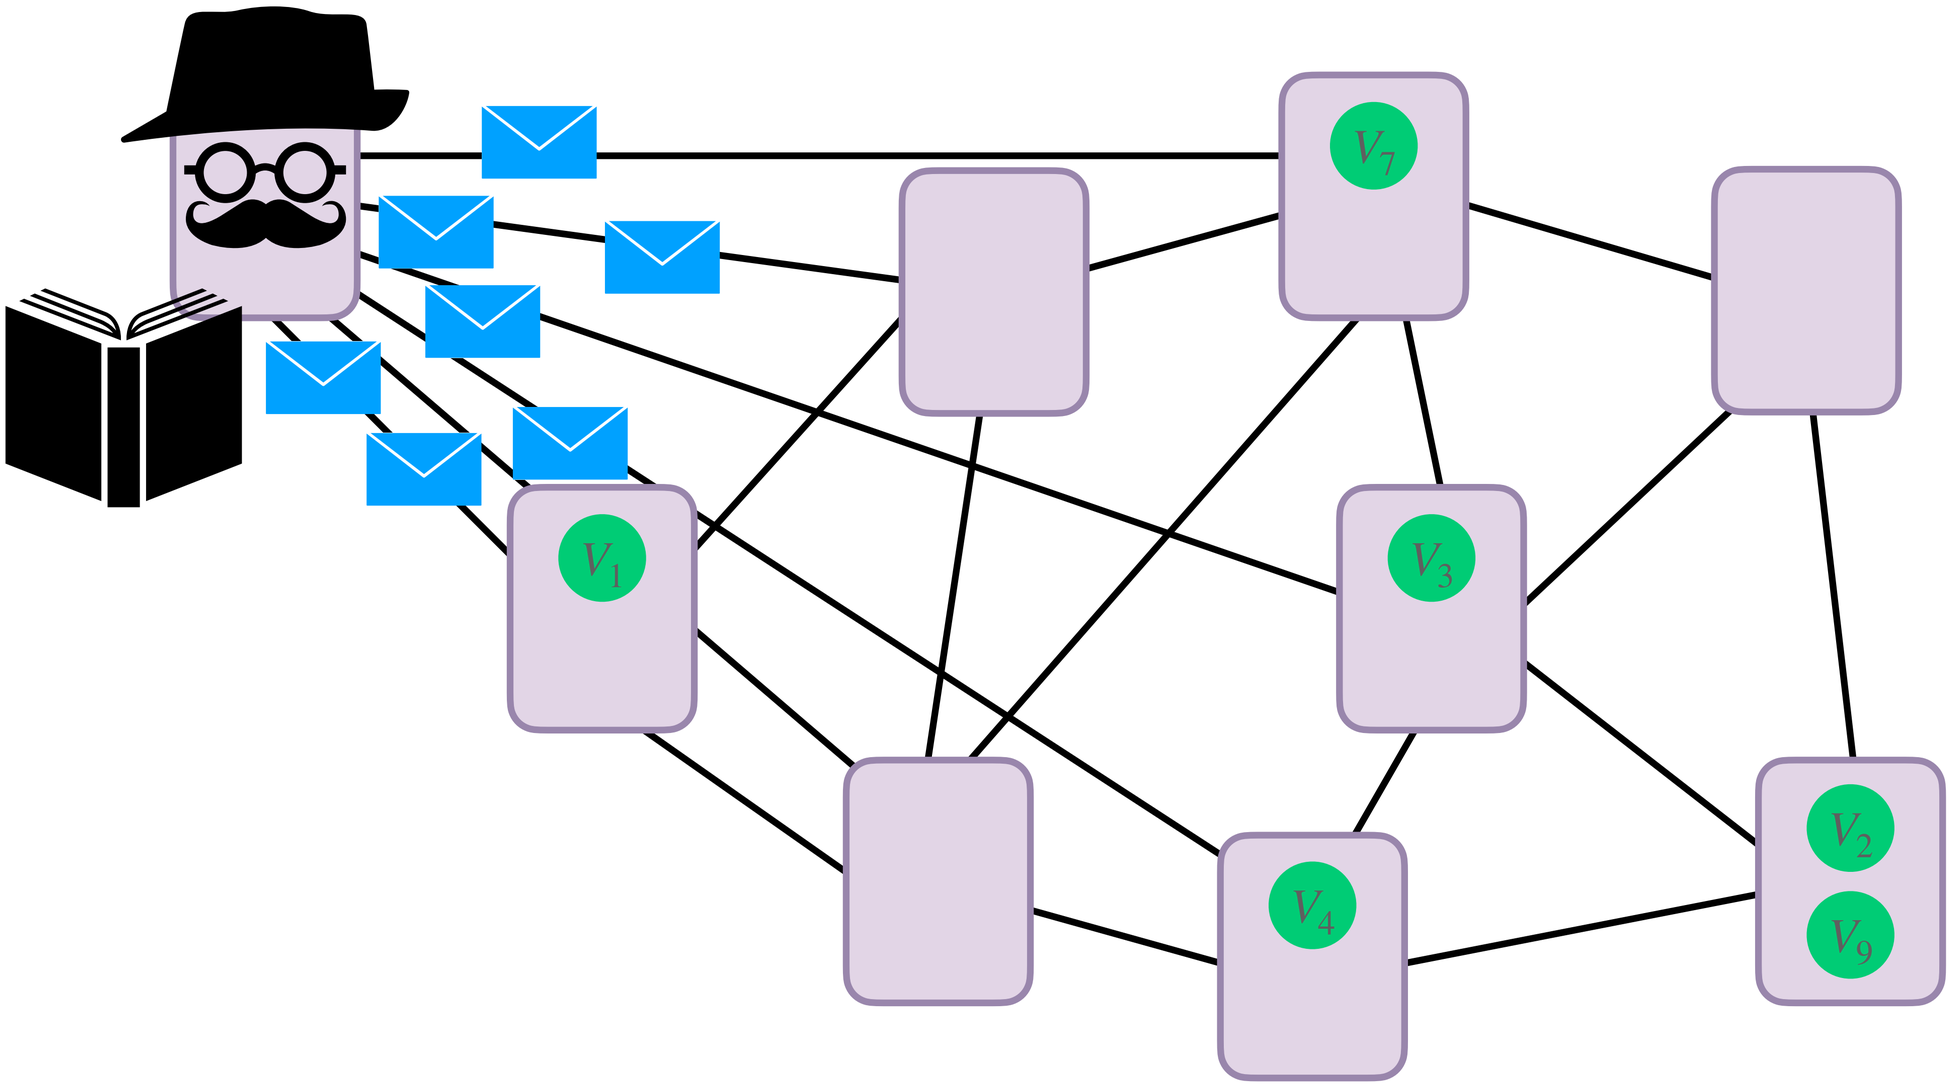
V_7
V_1
V_3
V_4
V_2
V_9

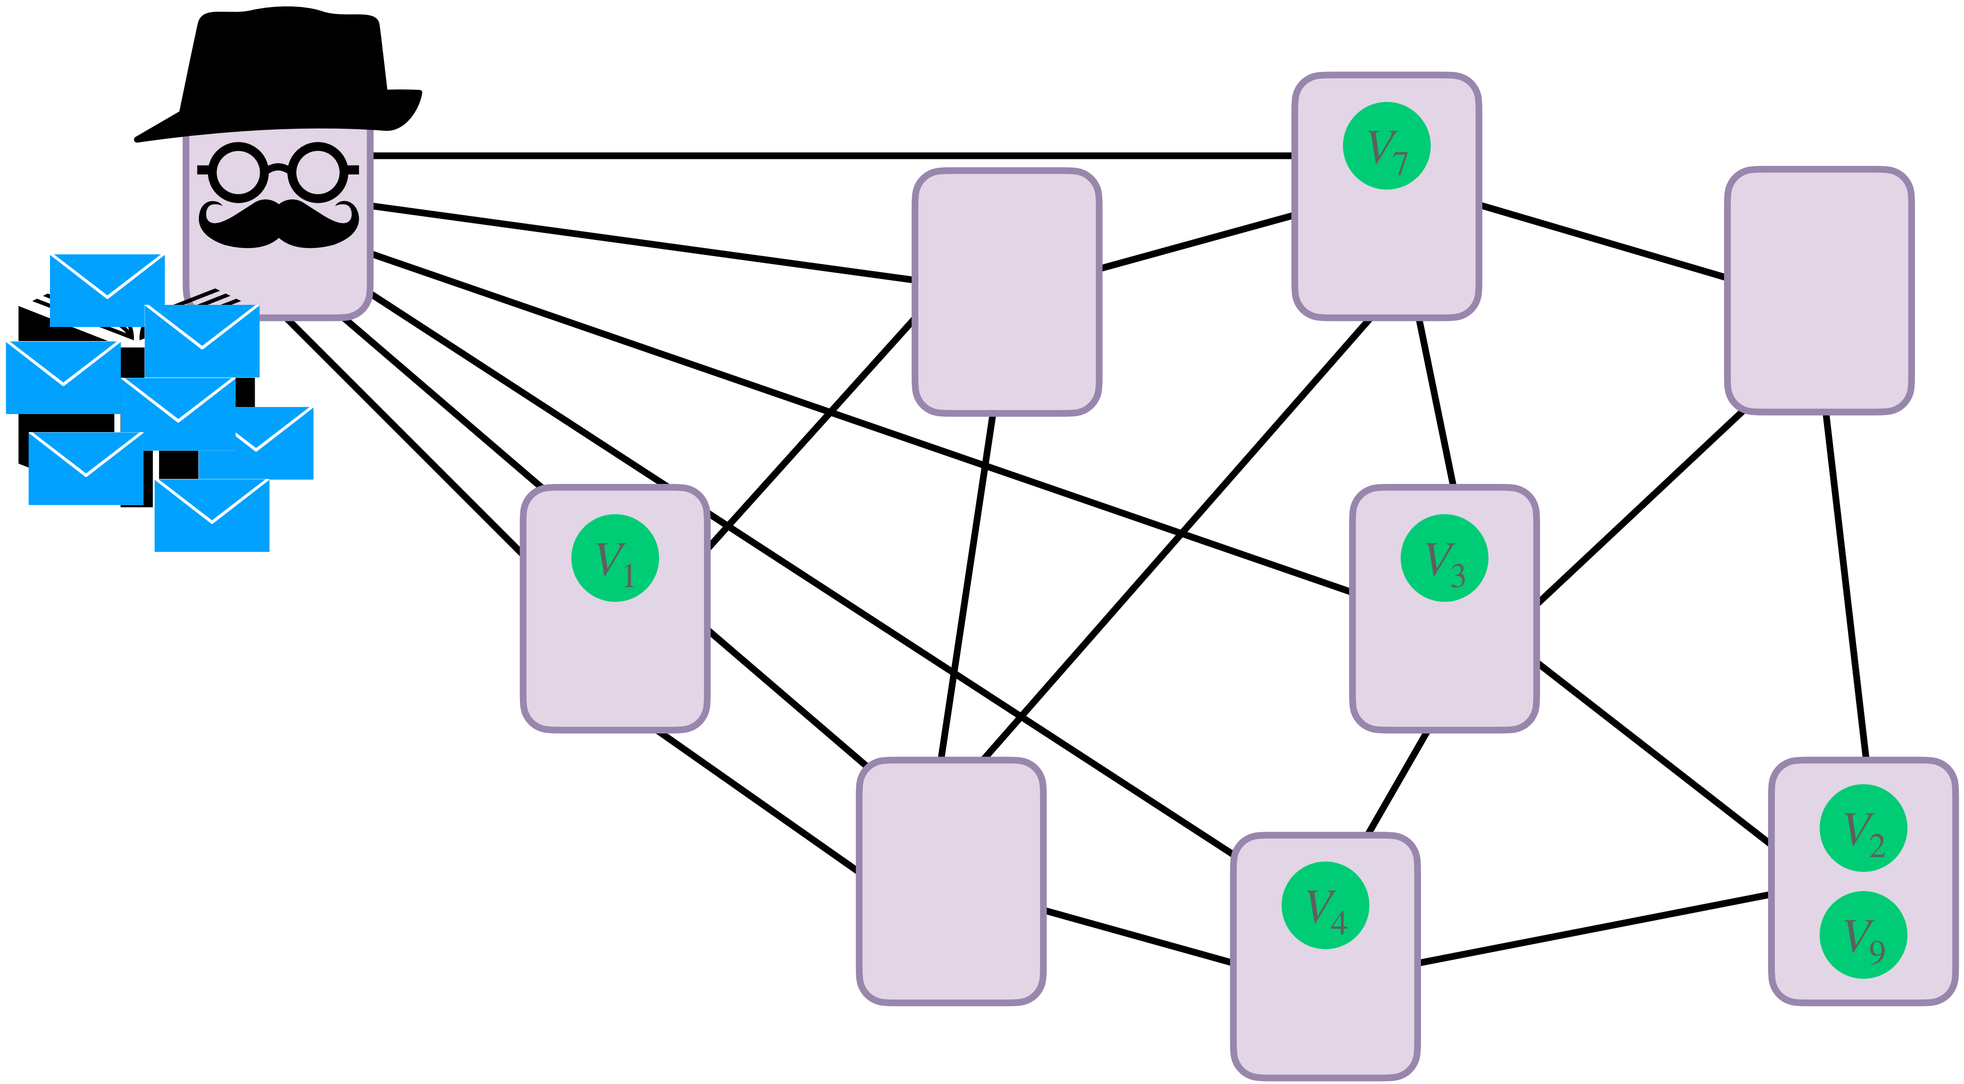
V_7
V_1
V_3
V_4
V_2
V_9

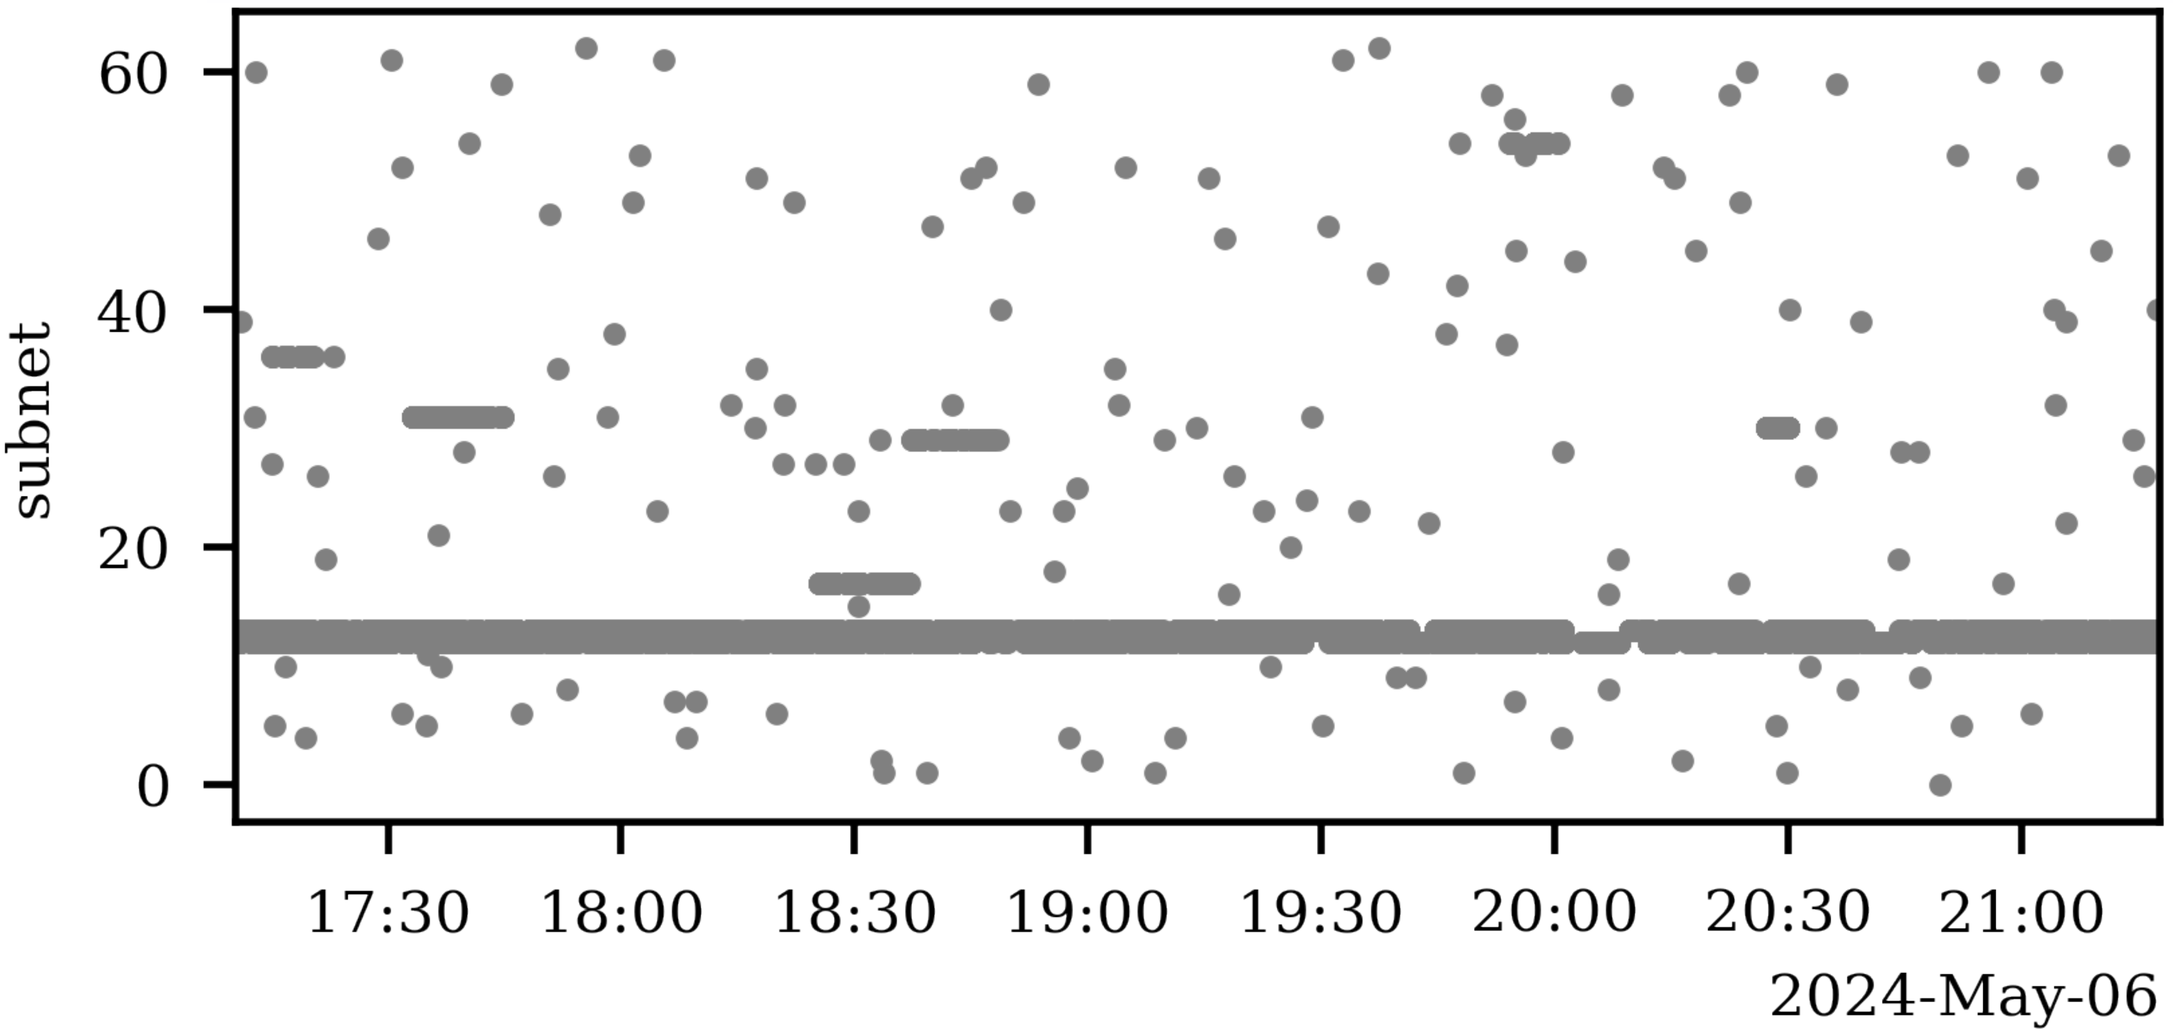
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

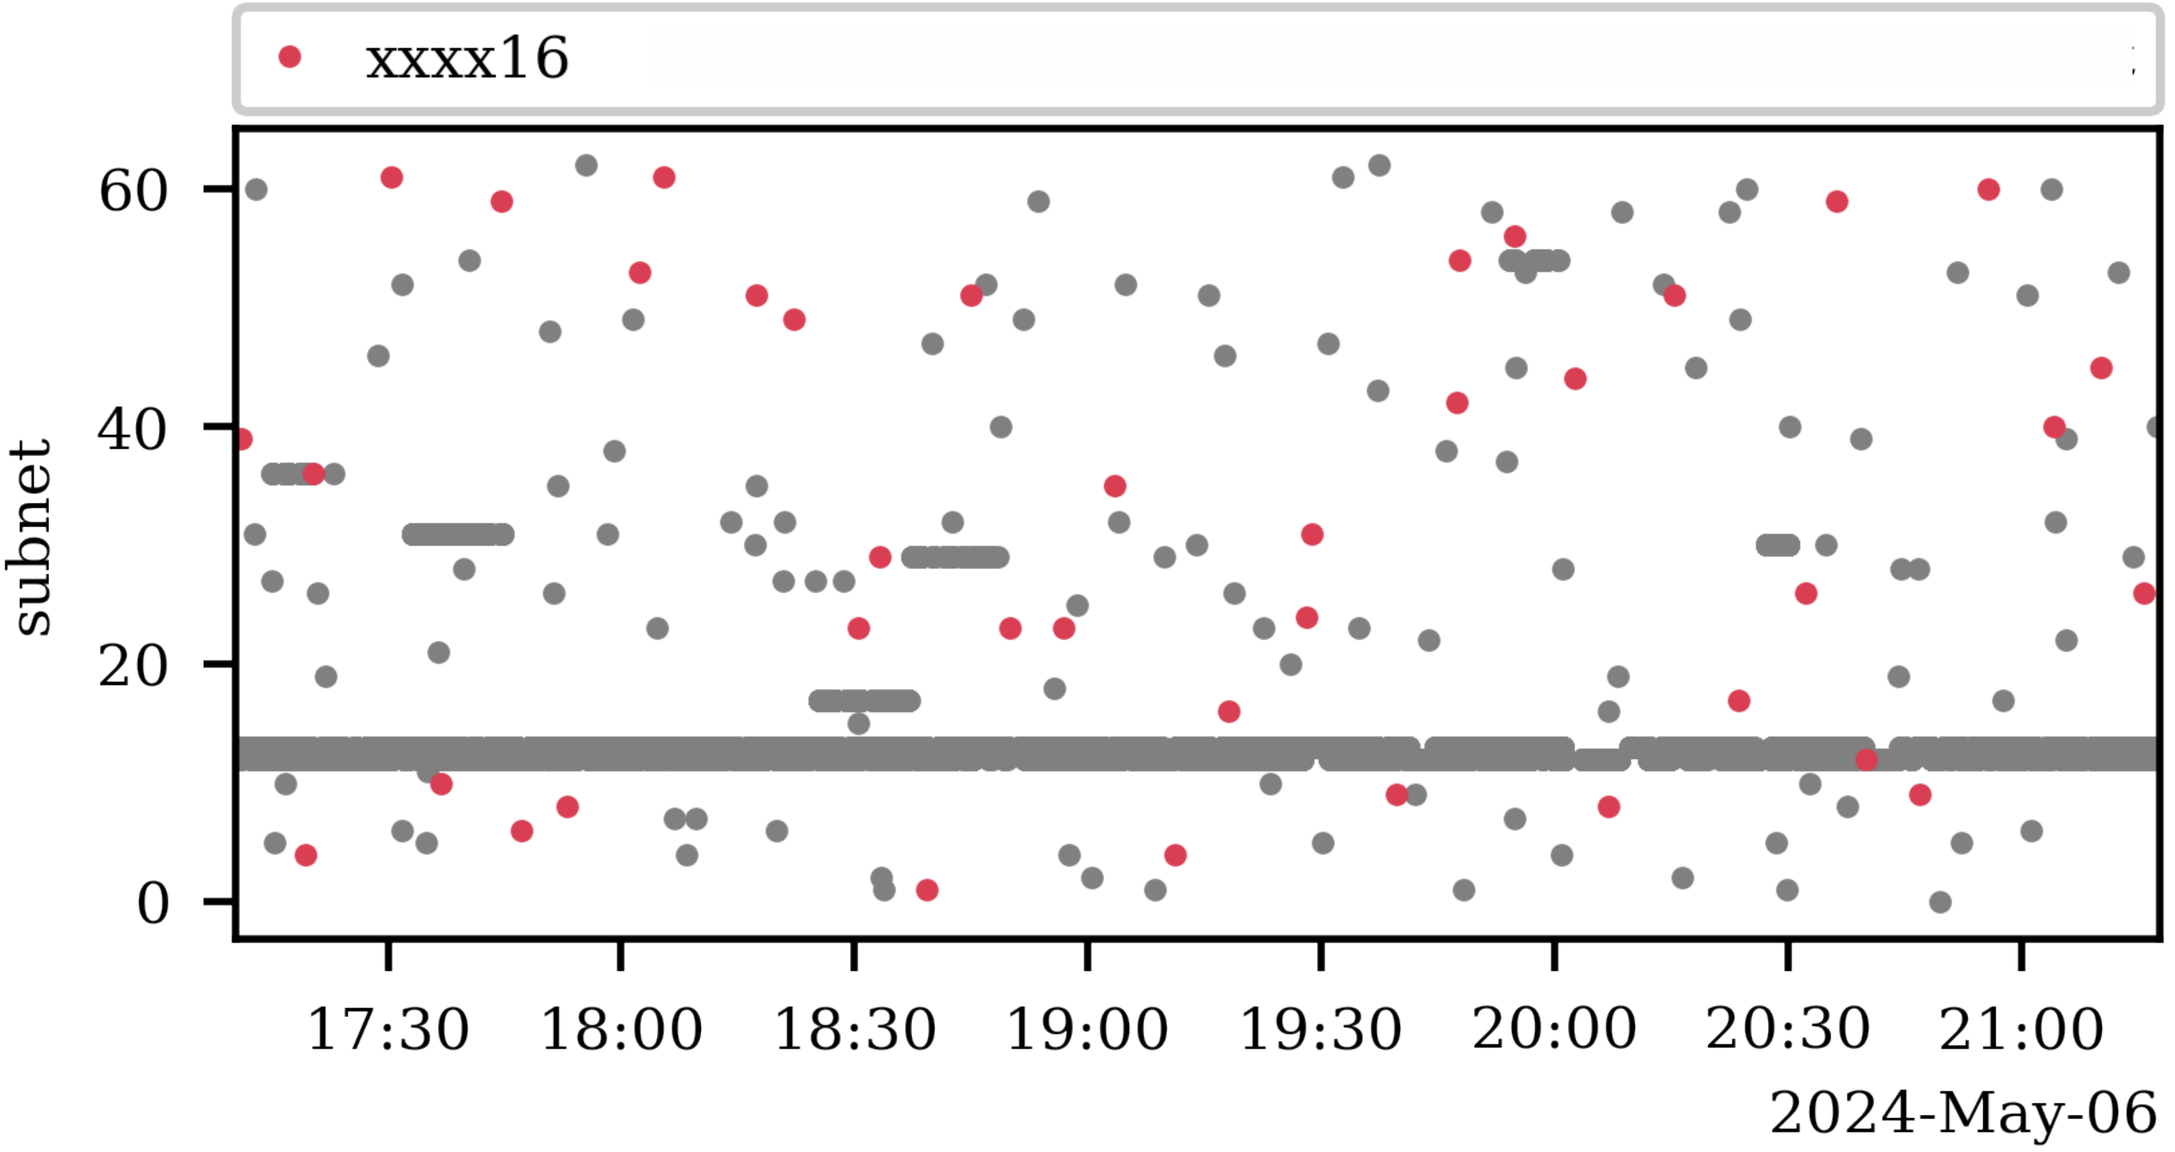
xxxx16
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

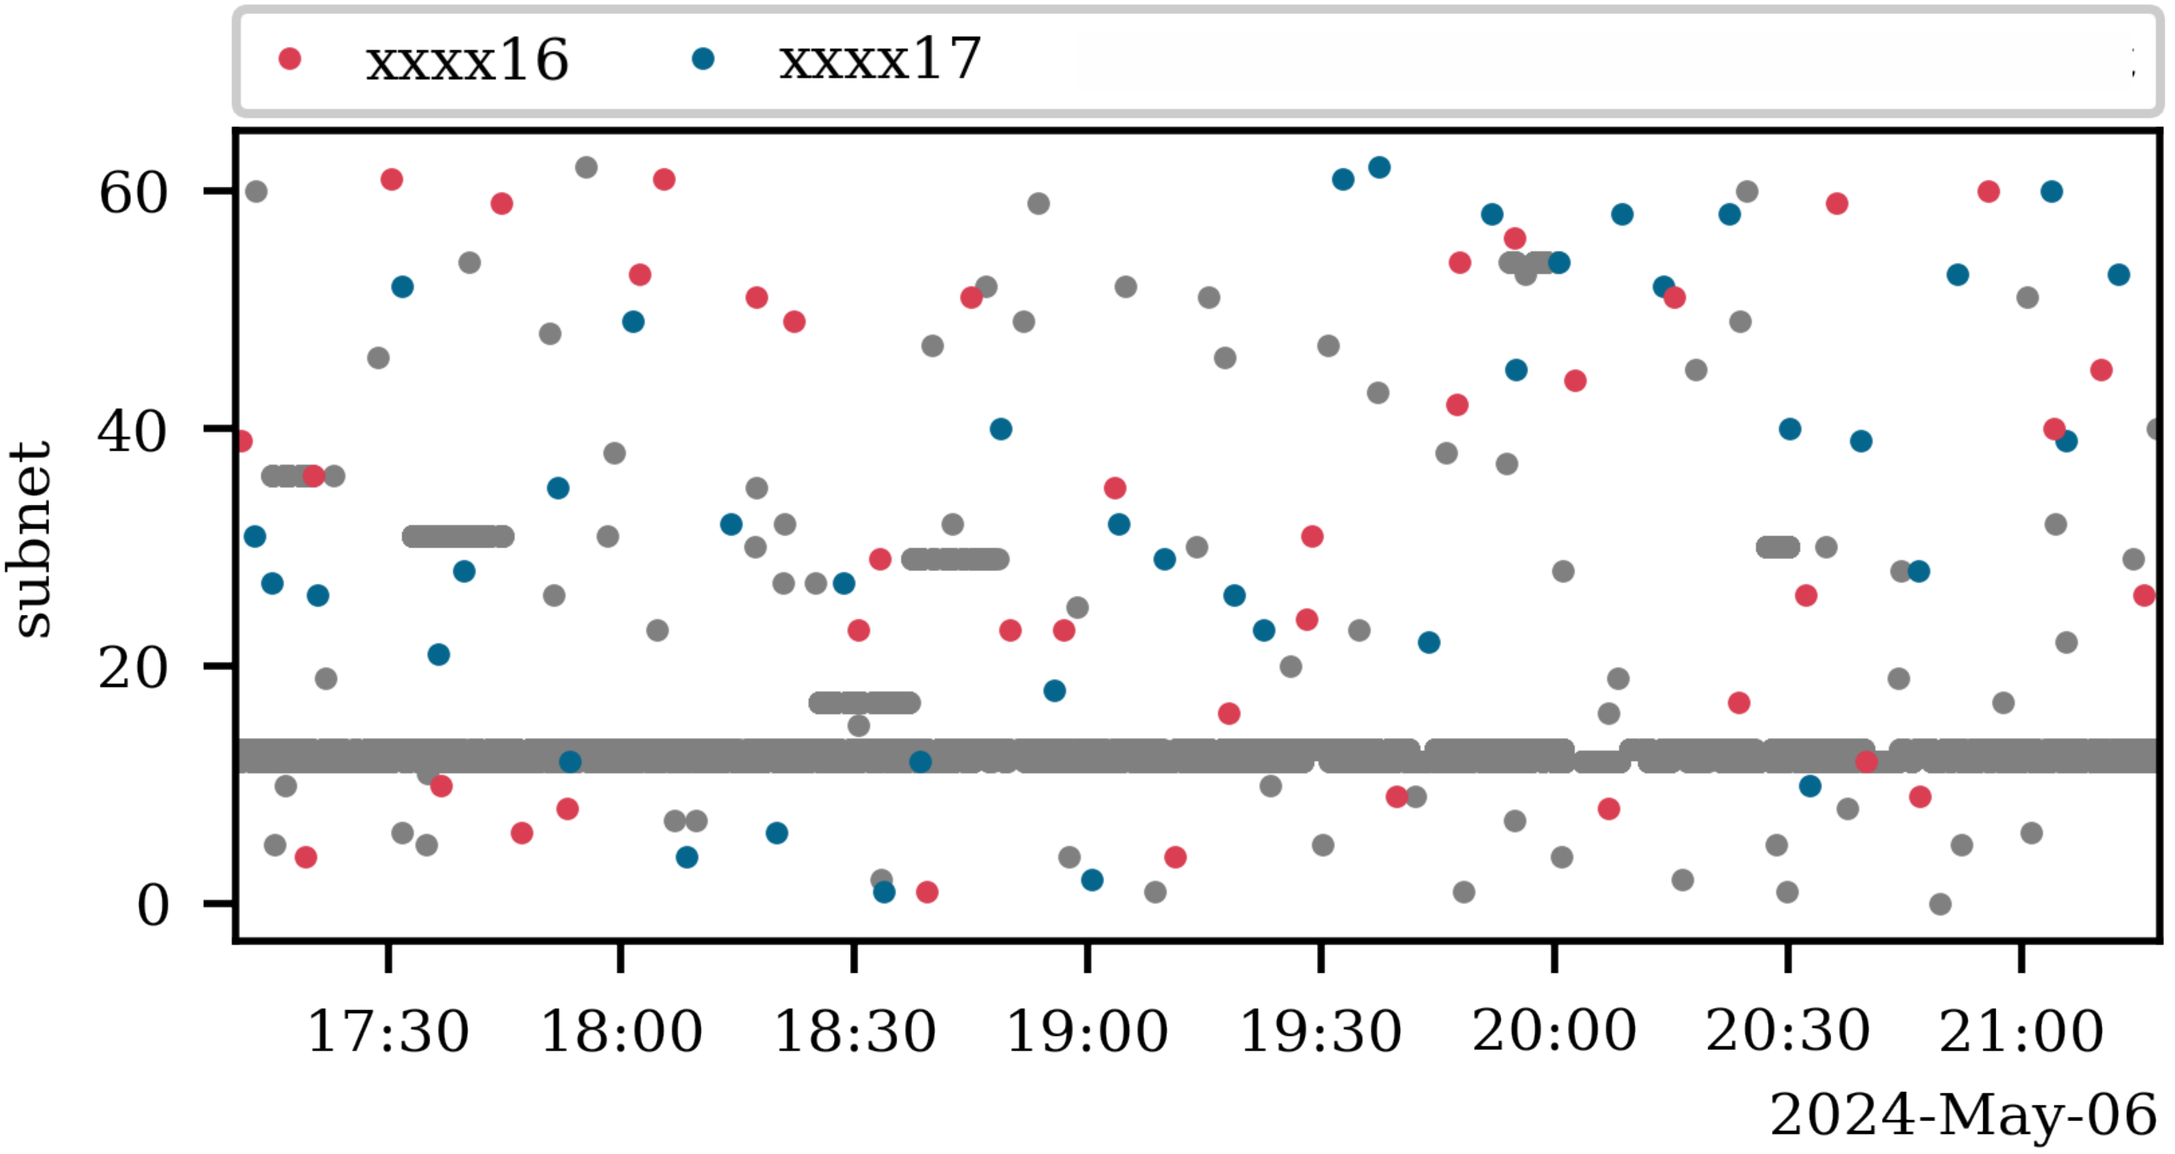
xxxx16
xxxx17
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

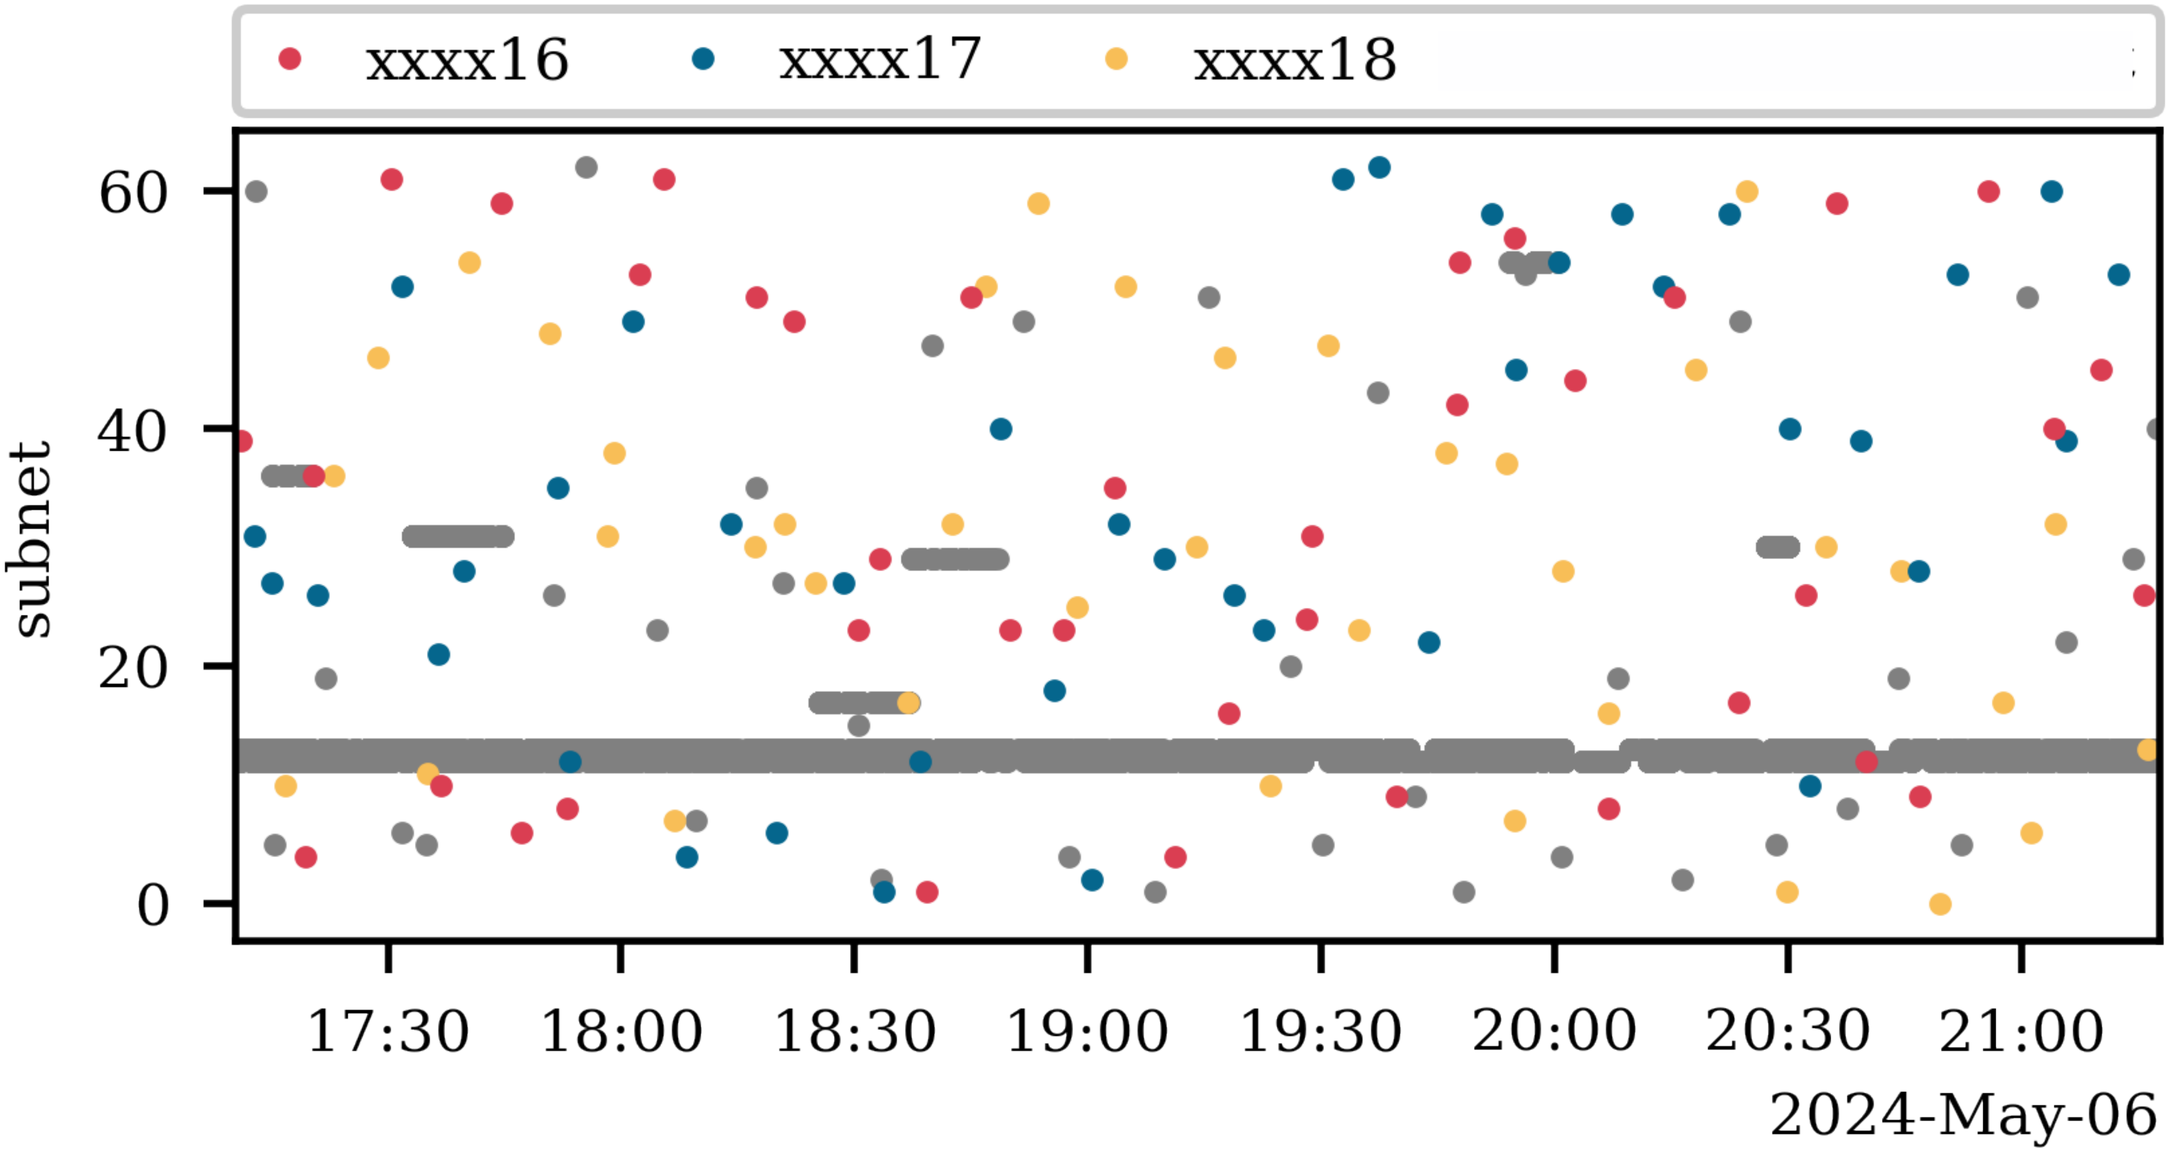
xxxx16
xxxx17
xxxx18
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

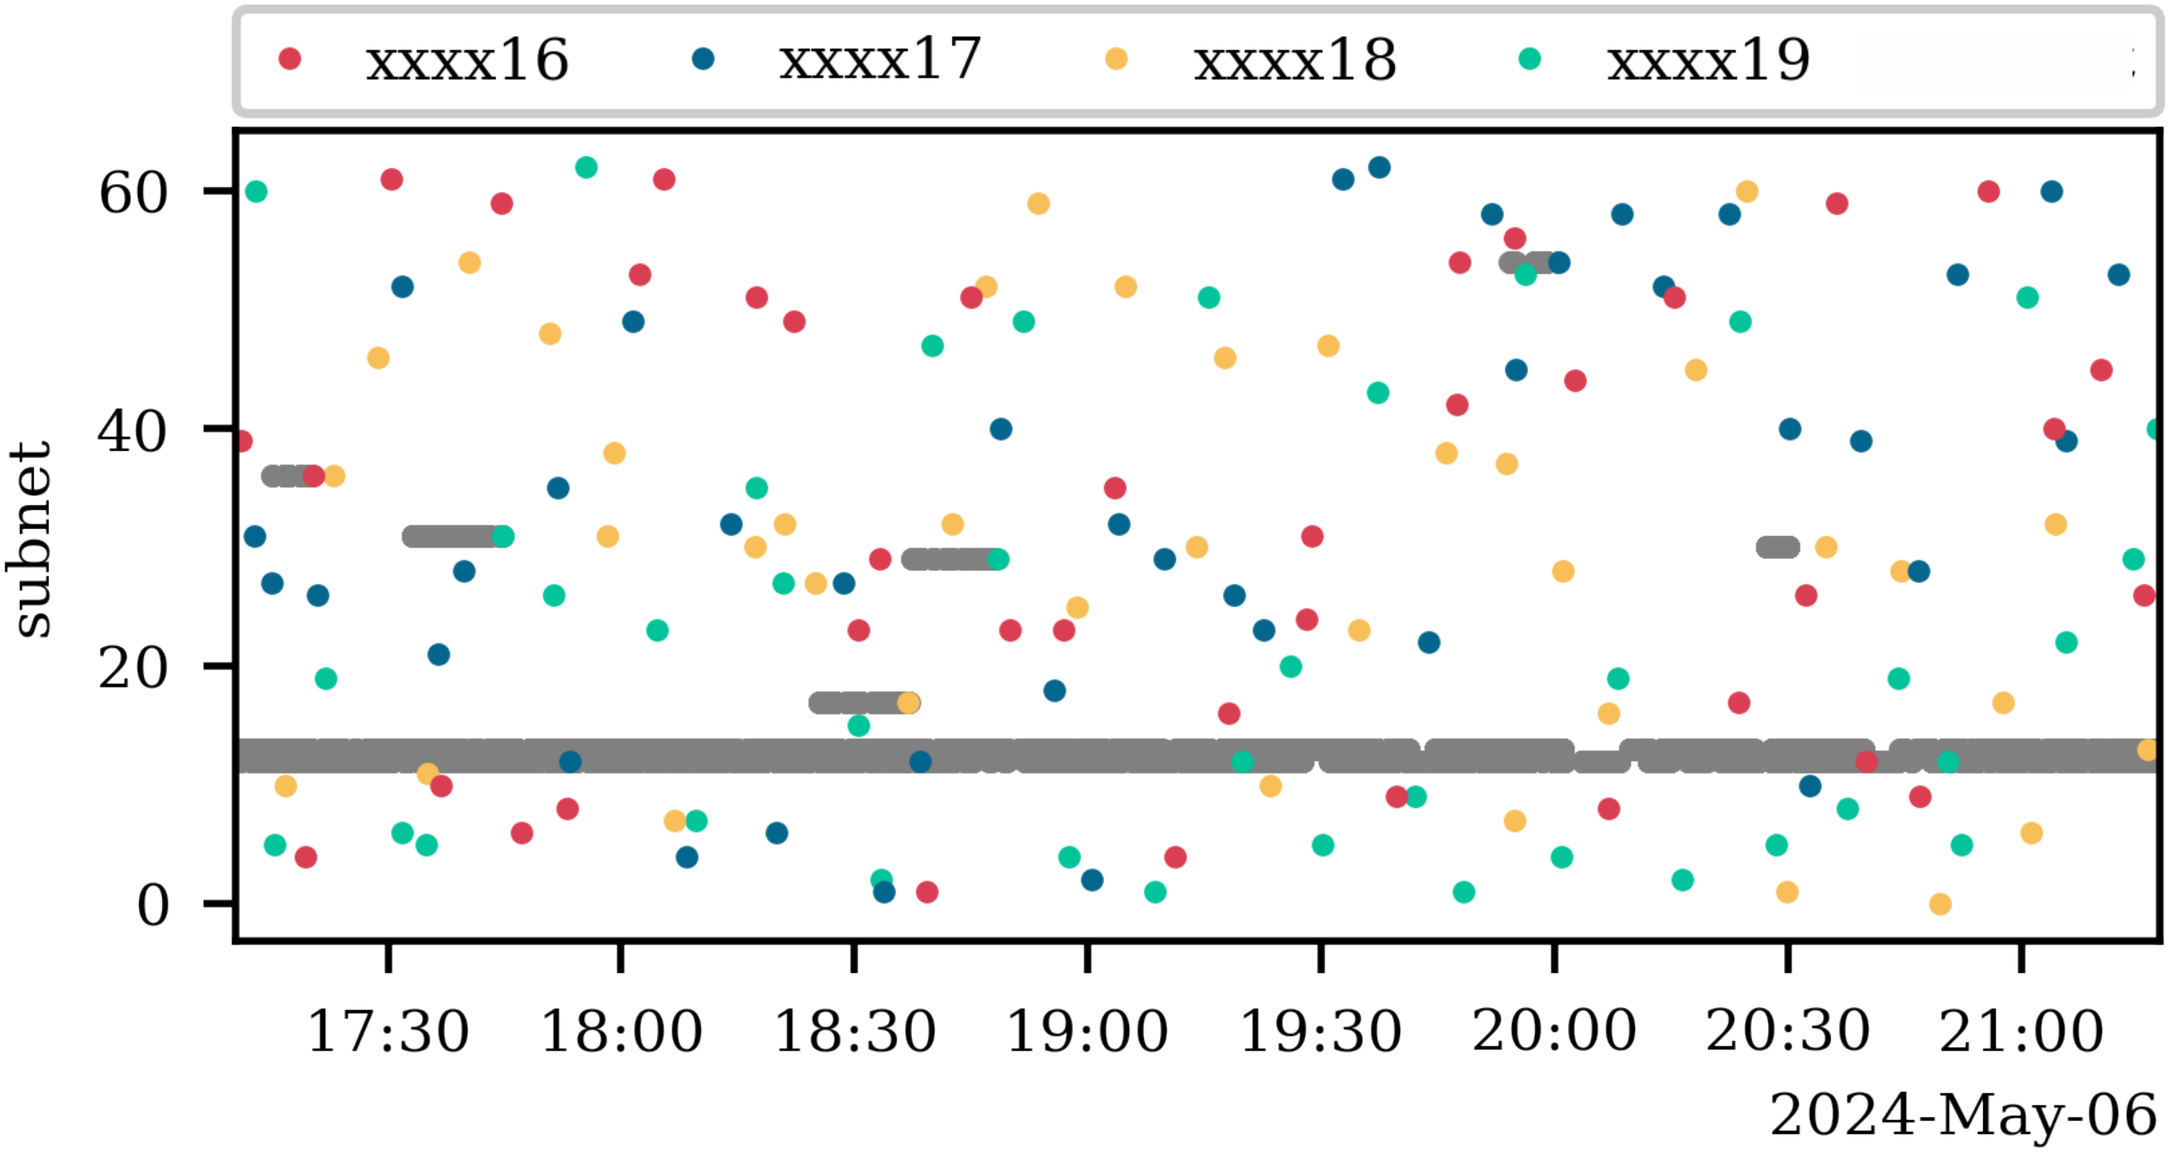
xxxx16
xxxx17
xxxx18
xxxx19
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

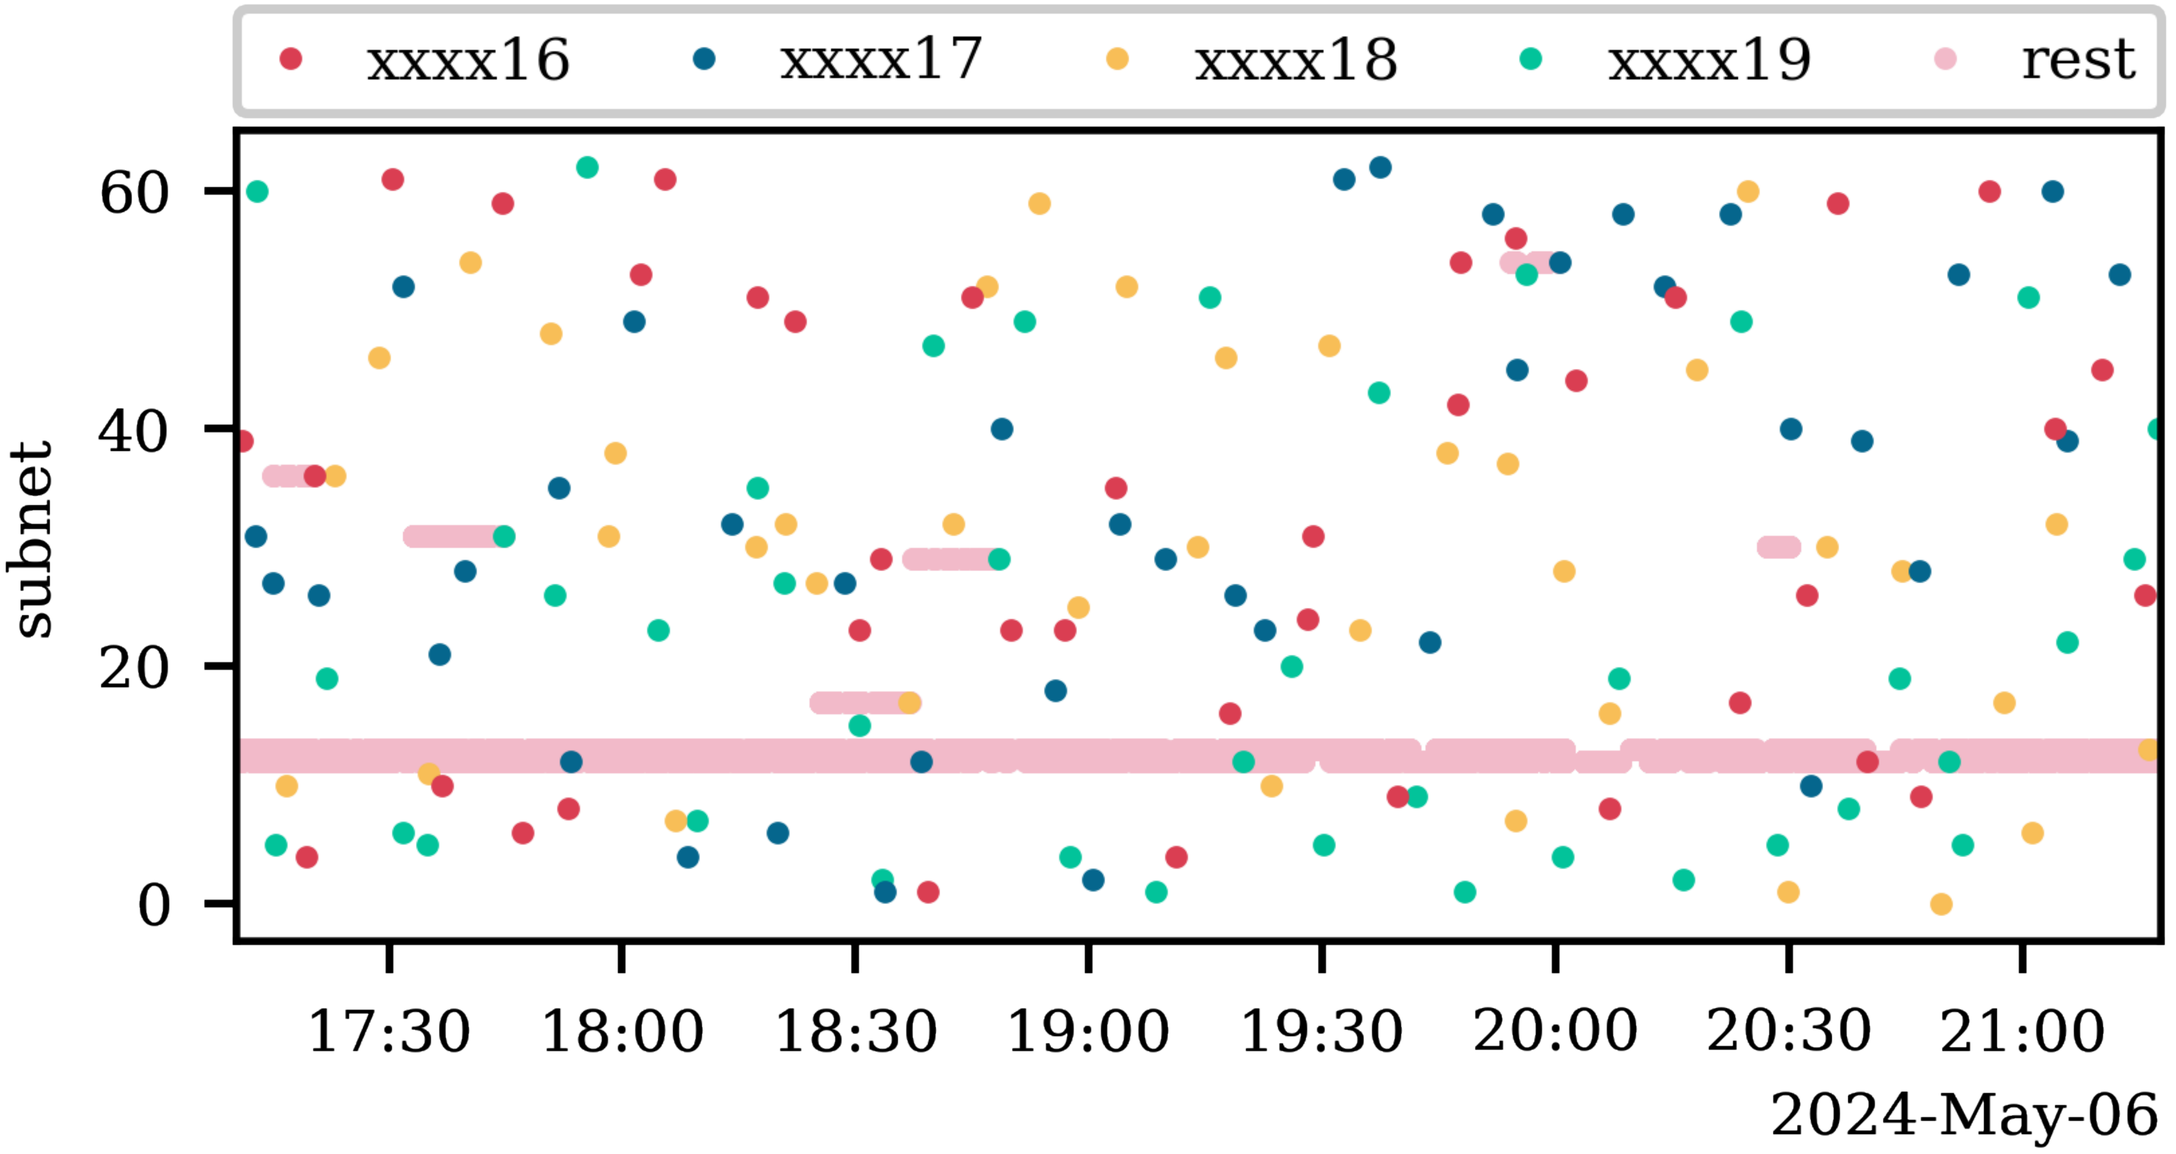
xxxx16
xxxx17
xxxx18
xxxx19
rest
subnet
60
40
20
0
17:30
18:00
18:30
19:00
19:30
20:00
20:30
21:00
2024-May-06

Measurement Methodology

Methodology

- Create a logging client: RAINBOW

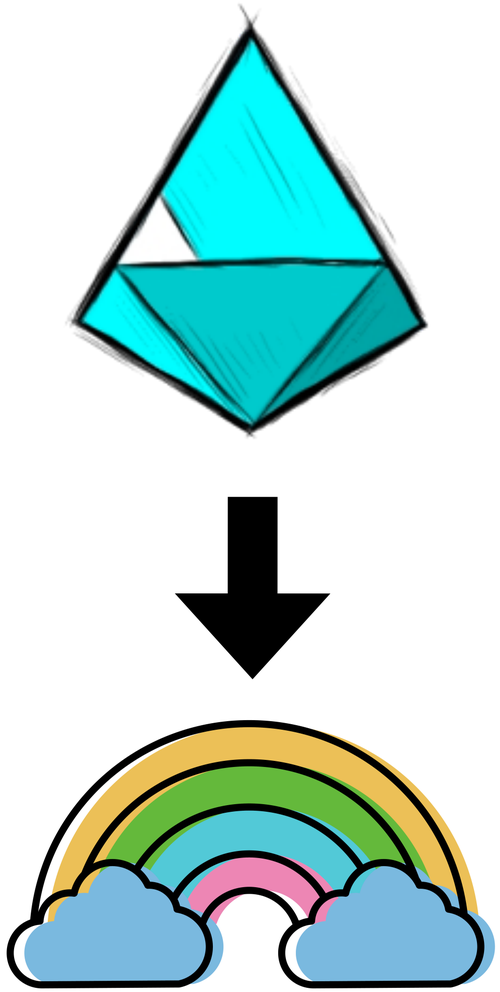

Methodology

- Create a logging client: RAINBOW
- Collect data on multiple locations for 3 days in May 2024

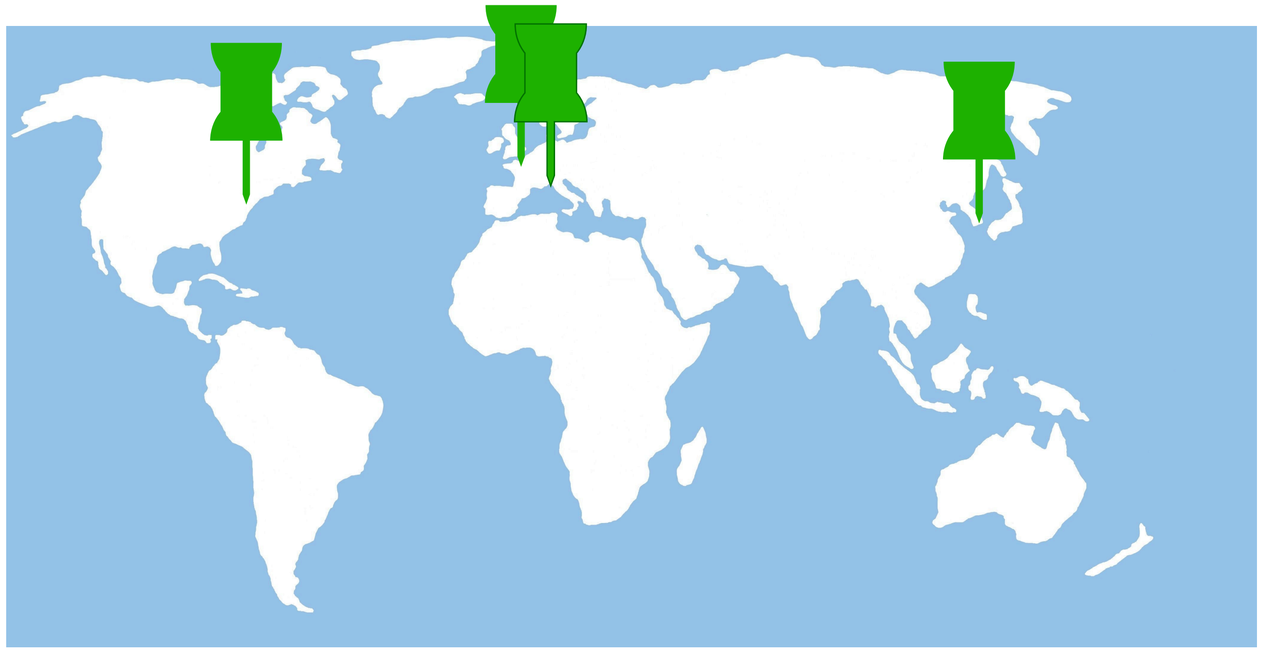

Methodology

- Create a logging client: RAINBOW
- Collect data on multiple locations for 3 days in May 2024
- Run a heuristic approach to deanonymization

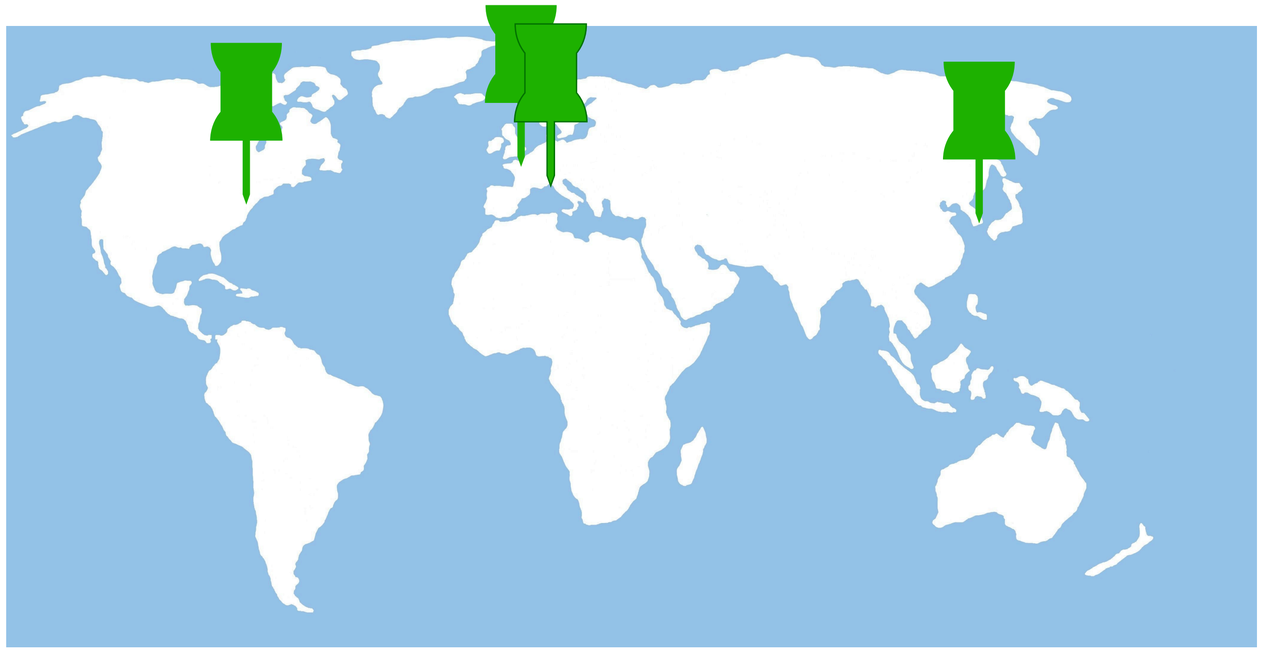

Methodology

- Create a logging client: RAINBOW
- Collect data on multiple locations for 3 days in May 2024
- Run a heuristic approach to deanonymization

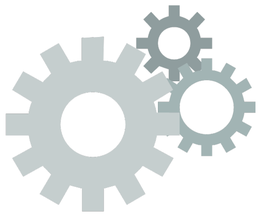

	peers		
	seen	with established connections	with long connections
FR	7,656	6,975	1,017
SO	7,816	7,122	1,142
VA	10,213	9,821	2,207
ZH	9,578	7,784	1,942
overall	11,219	10,785	4,372

Results

Results

Zurich, bare-bones, 1942 long connections

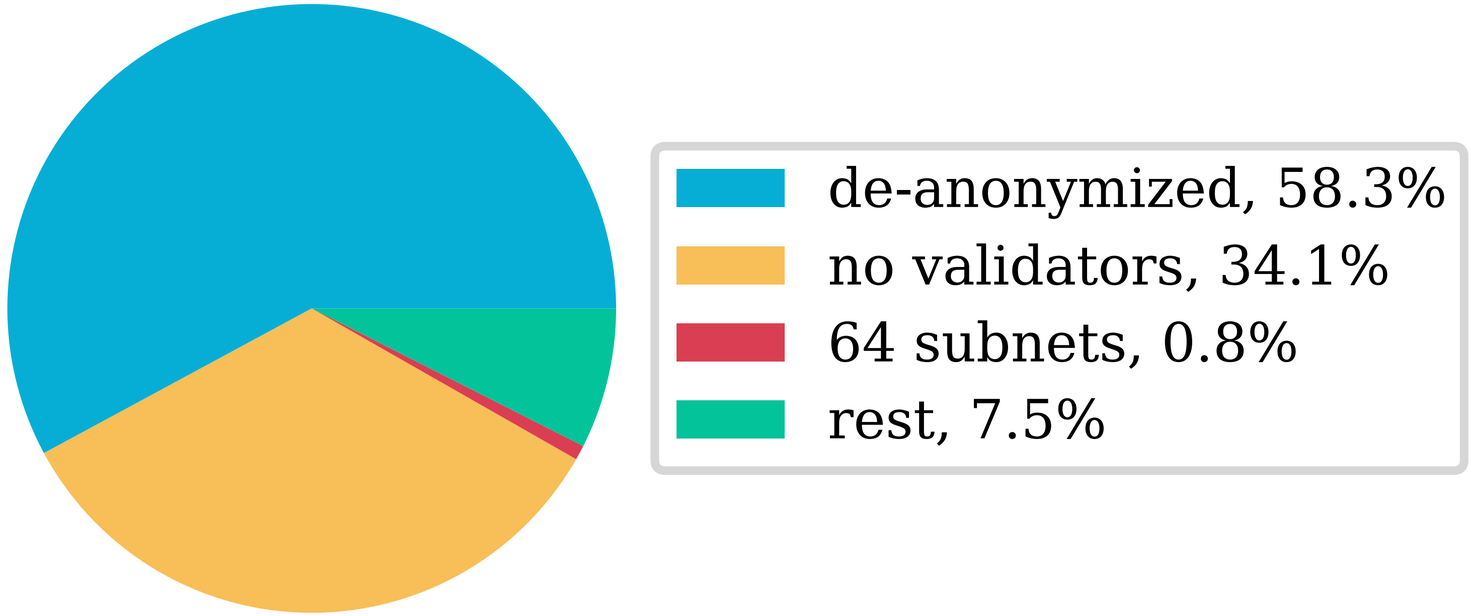

154’591

validators were deanonymized.

16%

of all validators were deanonymized.

Verification

Verification

- **Consistency** of validators
  - Same staking pool
  - Same deposit address
  - Same fee recipient address
  - Consecutive IDs

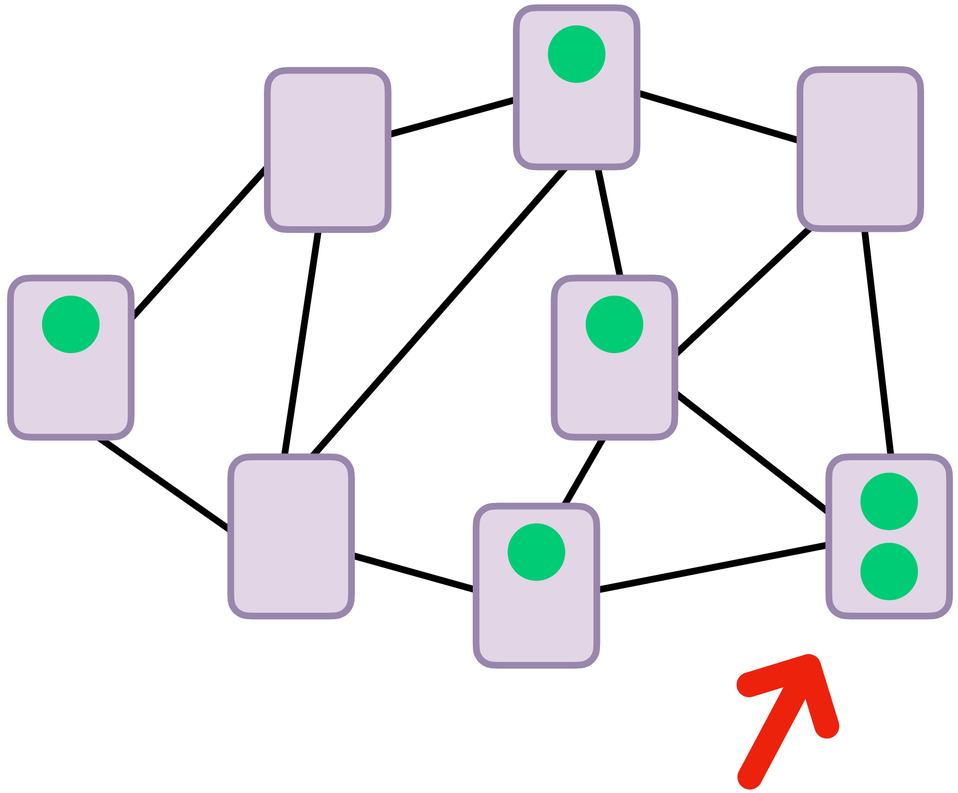

Verification

- **Consistency** of validators
  - Same staking pool
  - Same deposit address
  - Same fee recipient address
  - Consecutive IDs
- **Uniqueness** of Validator-IP Mapping

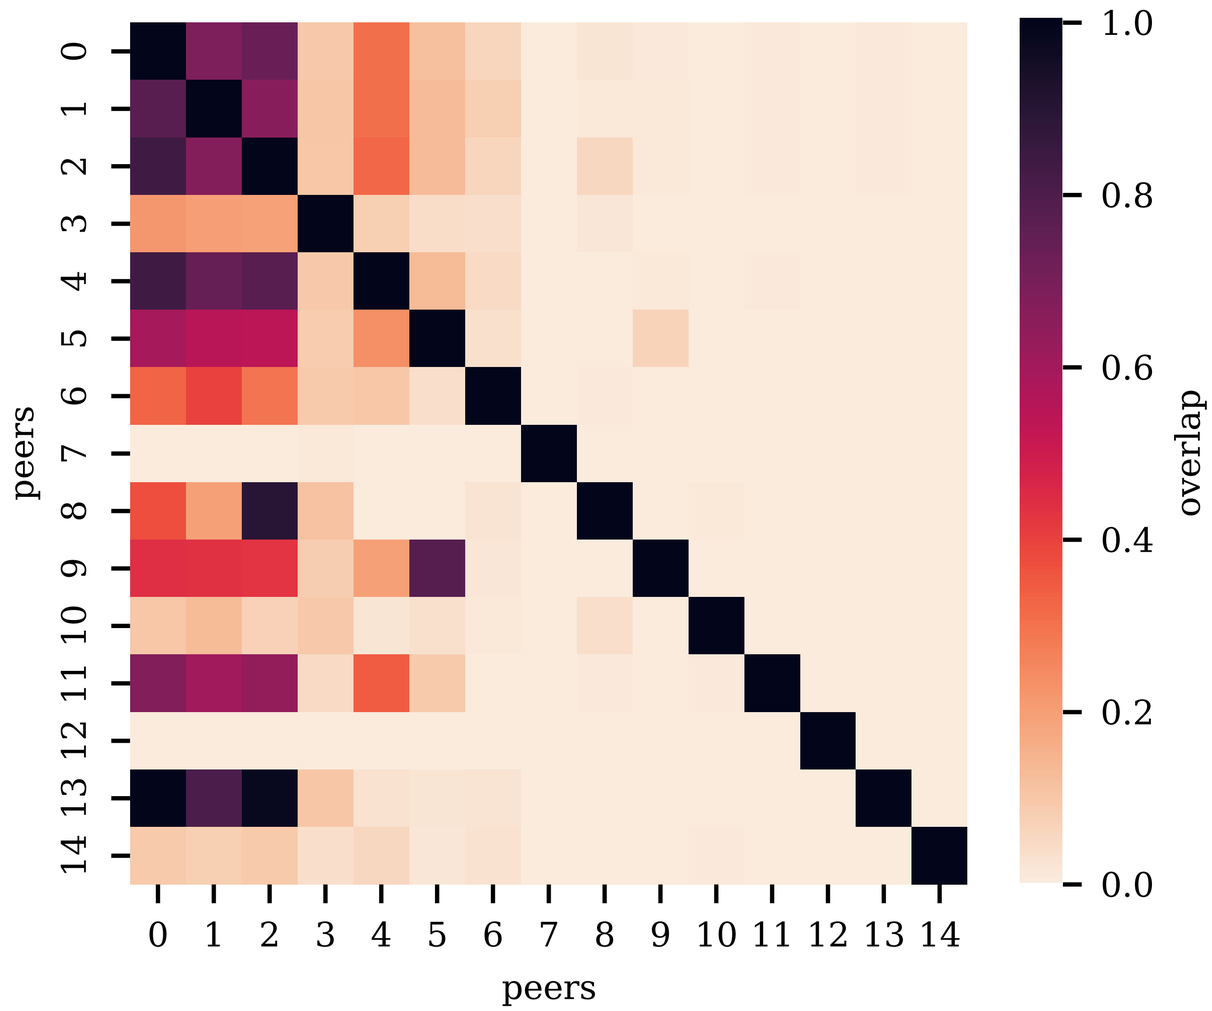

Verification

- **Consistency** of validators
  - Same staking pool
  - Same deposit address
  - Same fee recipient address
  - Consecutive IDs
- **Uniqueness** of Validator-IP Mapping
- **Similarity** of De-anonymizations

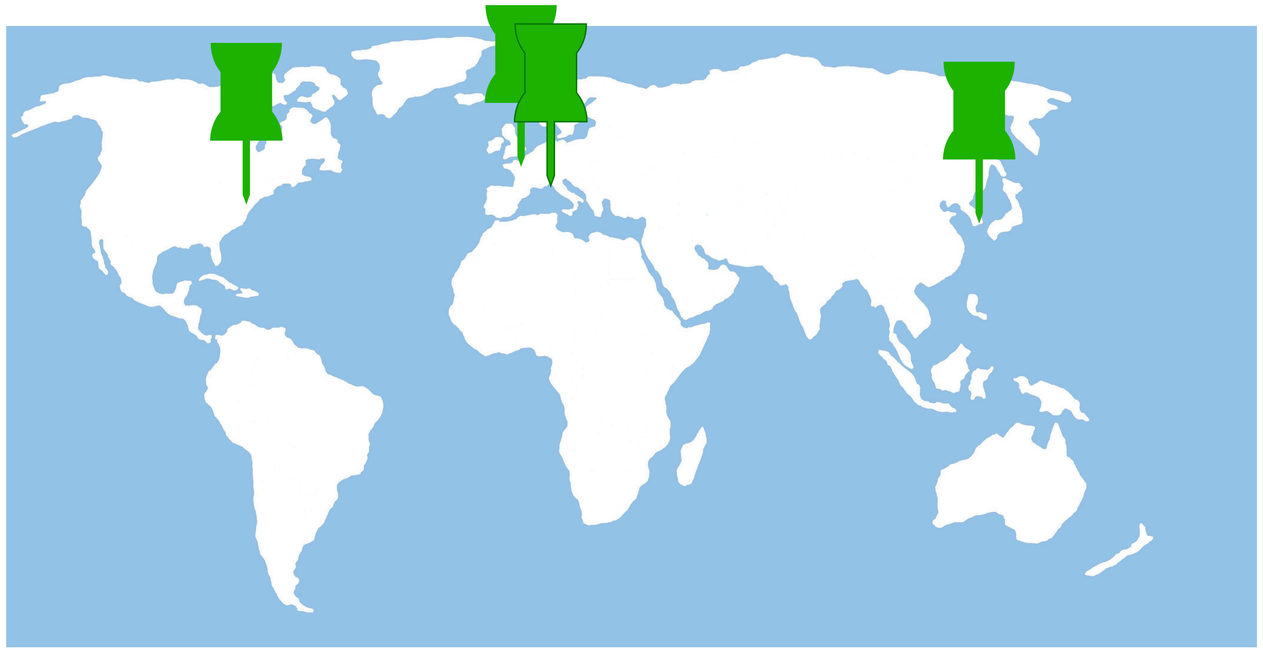

Take-aways

Validators per Peer

Overall

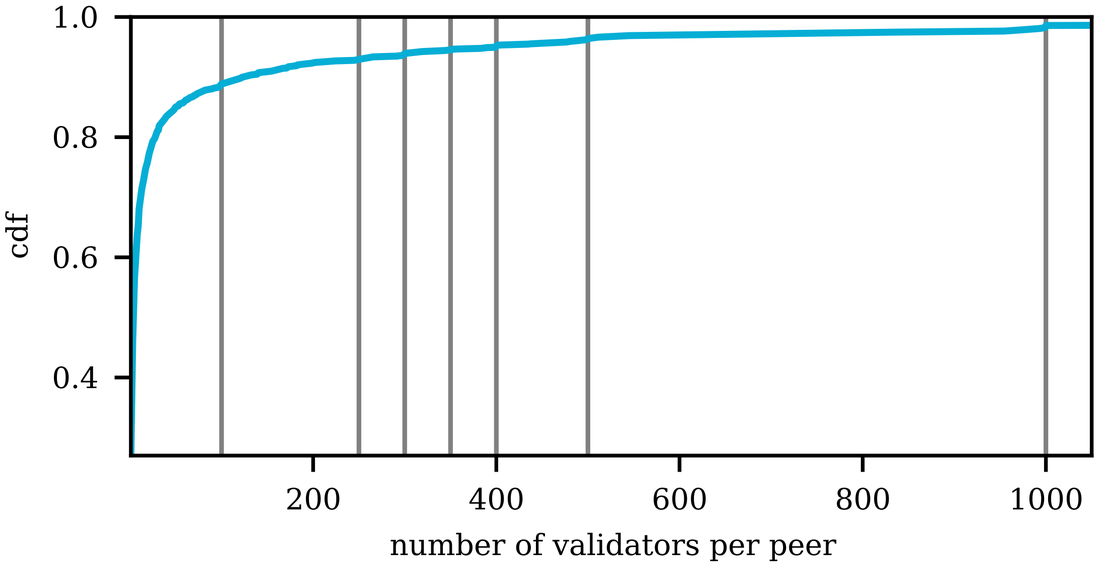

Five Largest Staking Pools

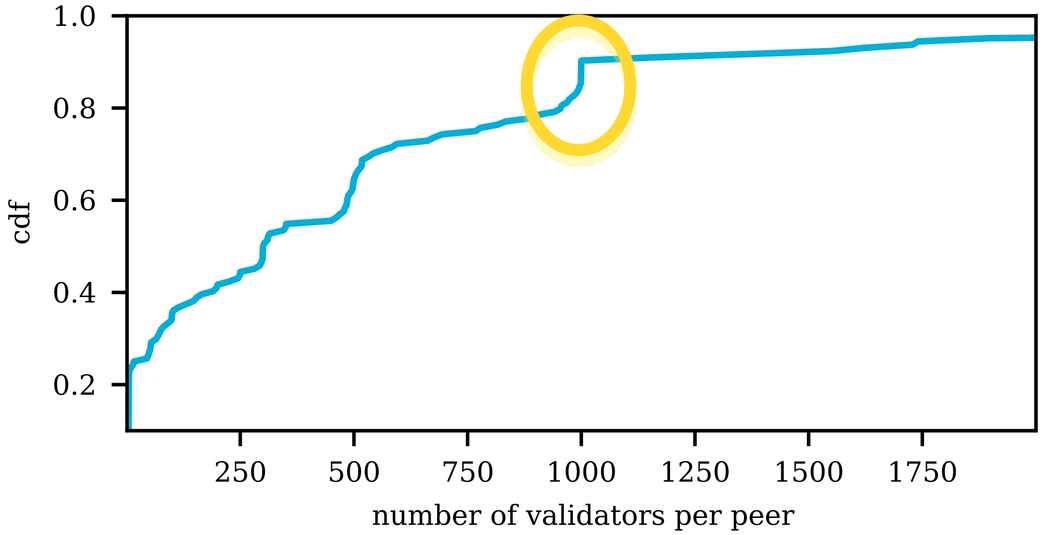

Location

Peers

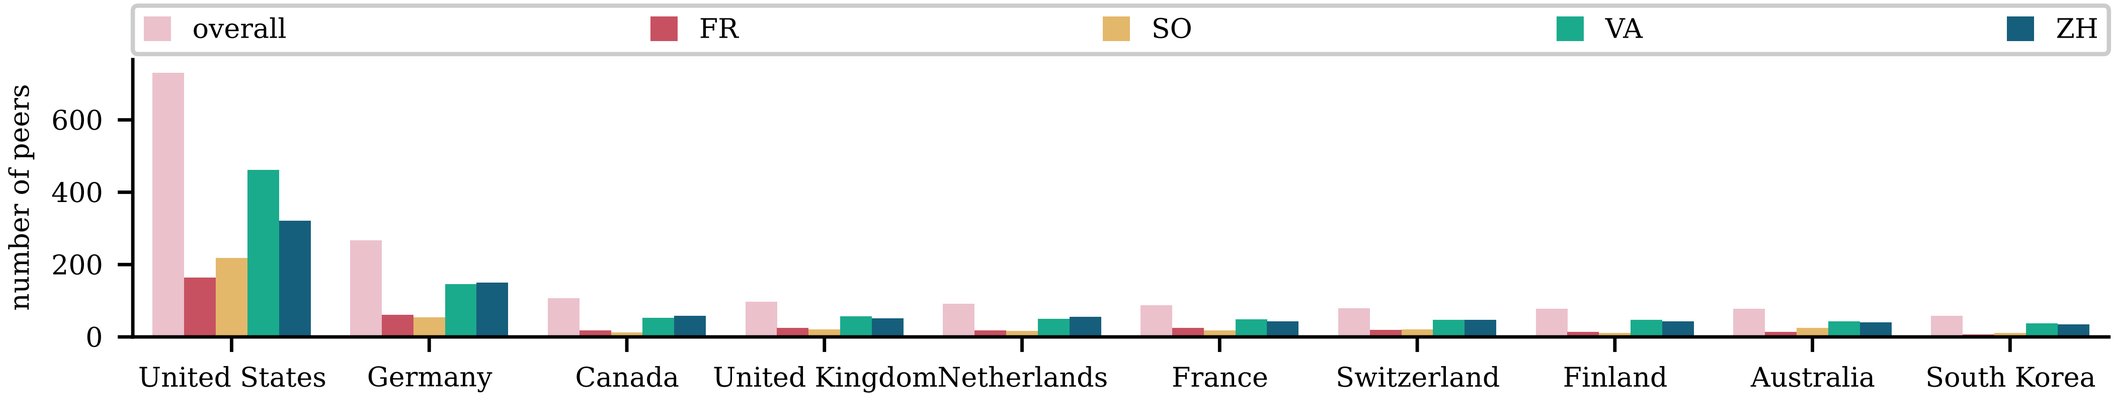

Validators

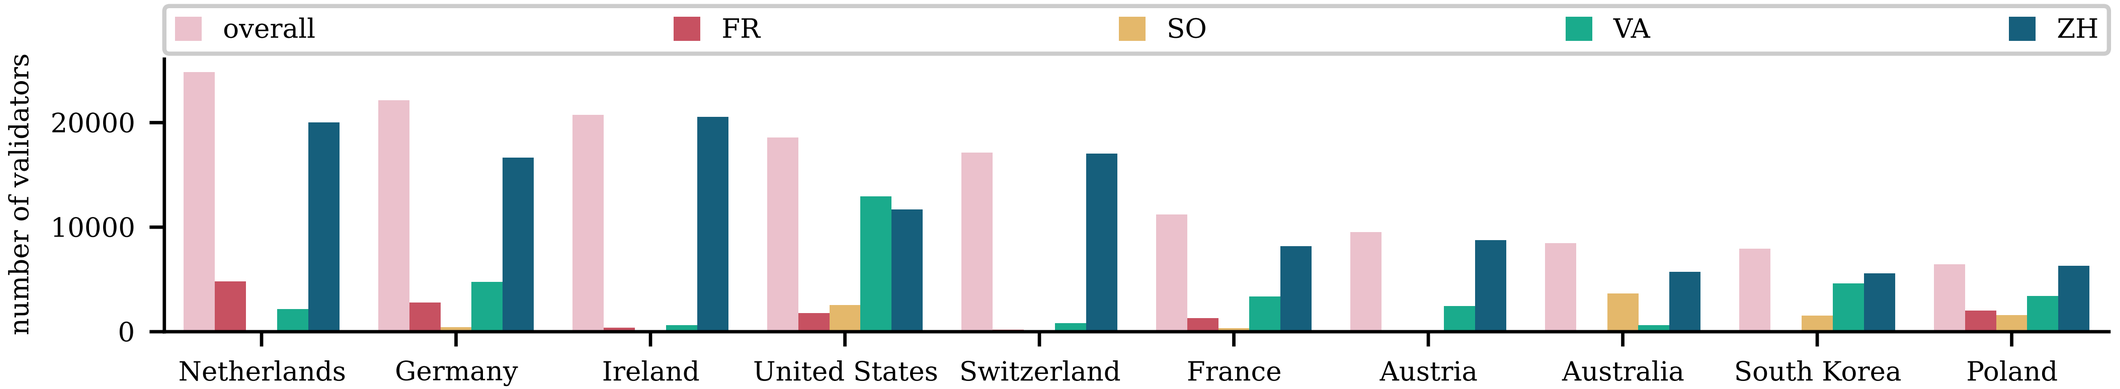

Organizations

Peers

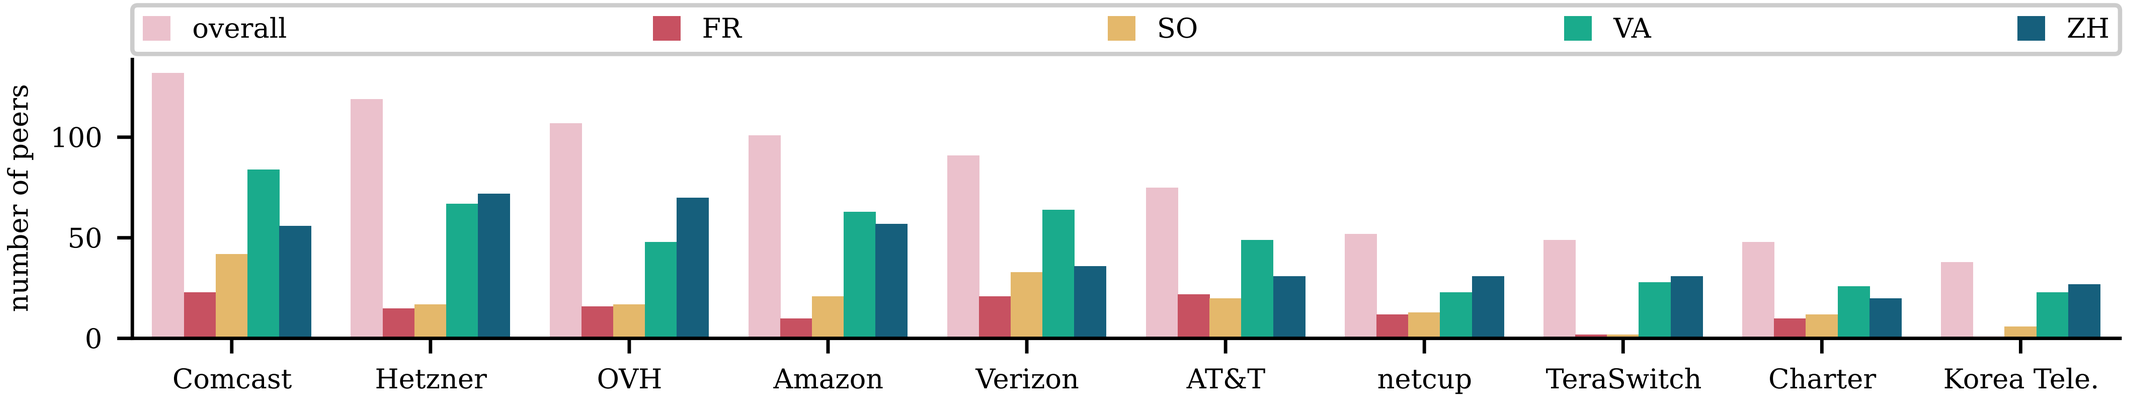

Validators

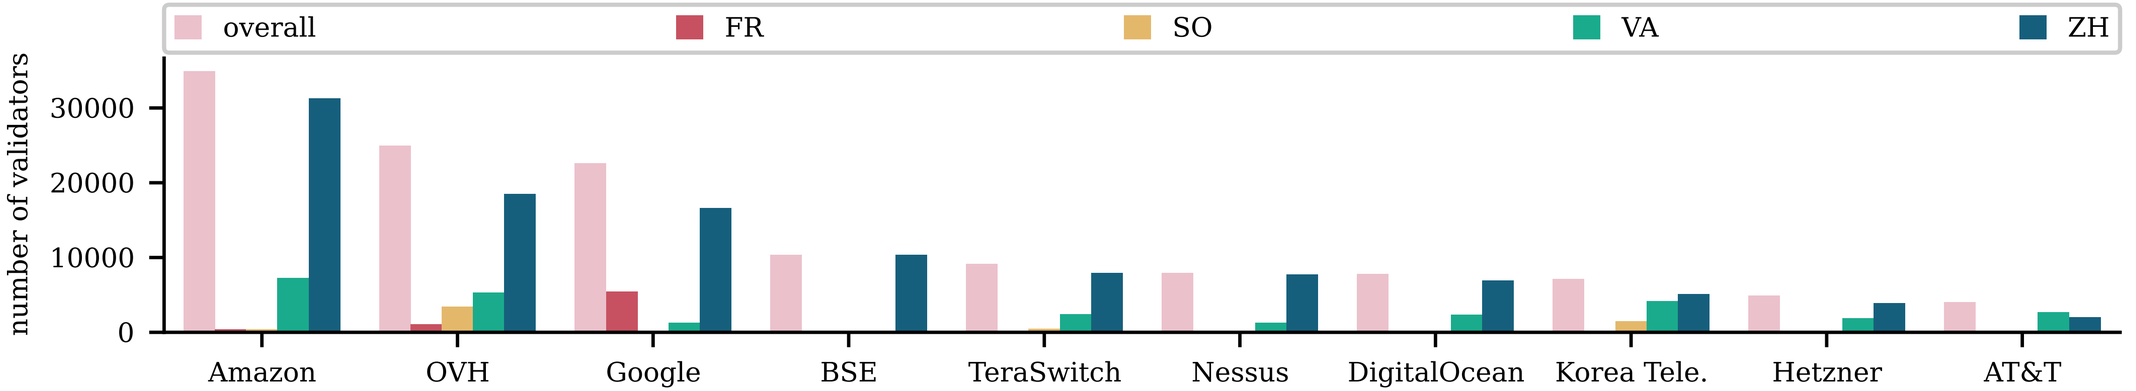

Summary

- RAINBOW implementation to show the feasibility of de-anonymizing Ethereum validators
  - low-cost, high-accuracy

Summary

- RAINBOW implementation to show the feasibility of de-anonymizing Ethereum validators
  - low-cost, high-accuracy
- Mitigations come at more cost, complexity, and/or latency
  - run more subnets, more nodes, more cryptography
  - increase anonymity set with friends
  - anonymous gossiping

Summary

- RAINBOW implementation to show the feasibility of de-anonymizing Ethereum validators
  - low-cost, high-accuracy
- Mitigations come at more cost, complexity, and/or latency
  - run more subnets, more nodes, more cryptography
  - increase anonymity set with friends
  - anonymous gossiping
- Reported attack to Ethereum Foundation + a grant for followup work on the gossip protocol

Thanks!

paper

institute imdea networks

ETH zürich