Heimdallr: Fingerprinting SD-WAN Control-Plane Architecture via Encrypted Control Traffic

Minjae Seo, Jaehan Kim, Eduard Marin, Myoungsung You, Taejune Park, Seungsoo Lee, Seungwon Shin, and **Jinwoo Kim**

Software-Defined WAN (SD-WAN)

- A new use case for efficiently operating a private WAN
  - To manage geographically distributed sites with a unified platform, i.e., controller
  - Can achieve network-wide optimization ➔ Used by many WAN operators, e.g., Google[1], Microsoft[2]

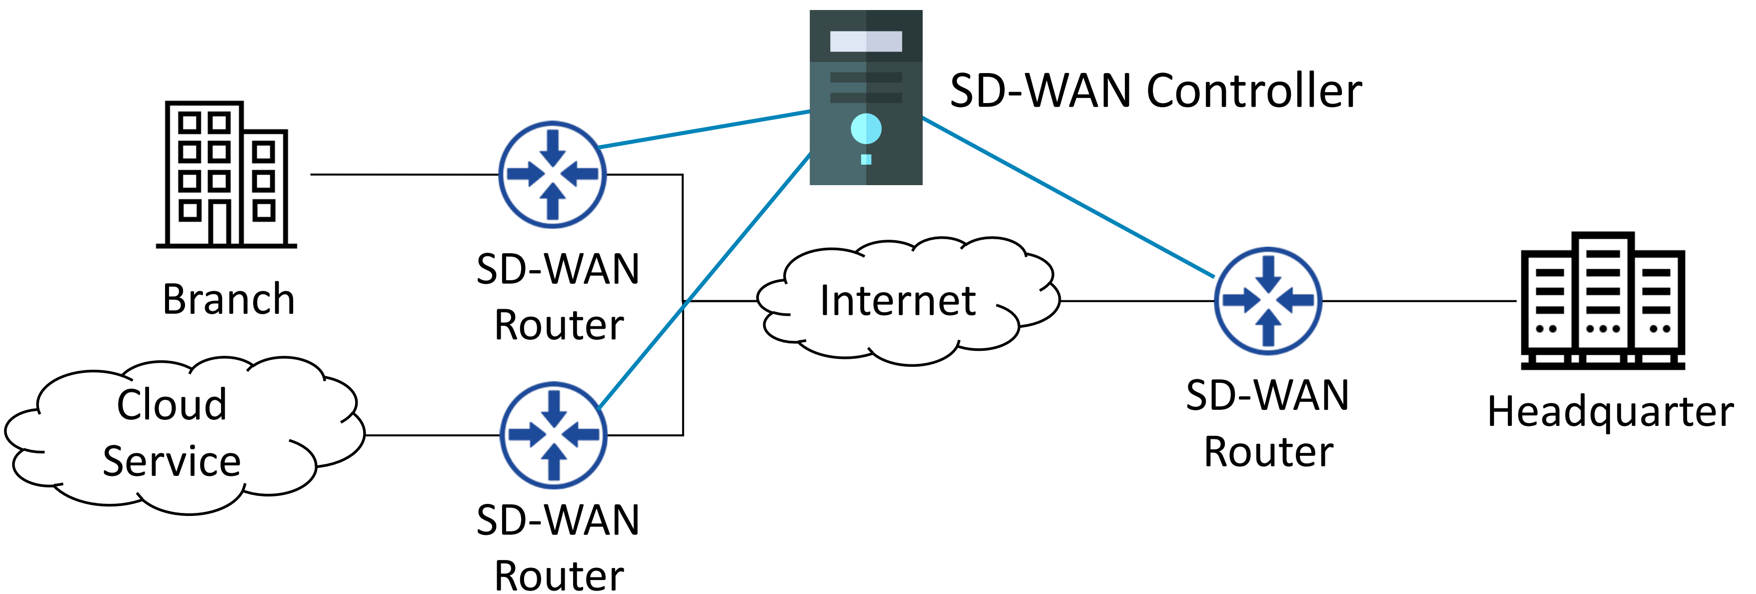

1 Orion: Google's Software-Defined Networking Control Plane, NSDI '21
2 Achieving high utilization with software-driven WAN, SIGCOMM '13

Control Plane: SD-WAN's Brain

- ***Single*** controller
  - Weak to a single point of failure
- ***Multiple*** controllers → cluster
  - Physically distributed for fault-tolerance and high-performance

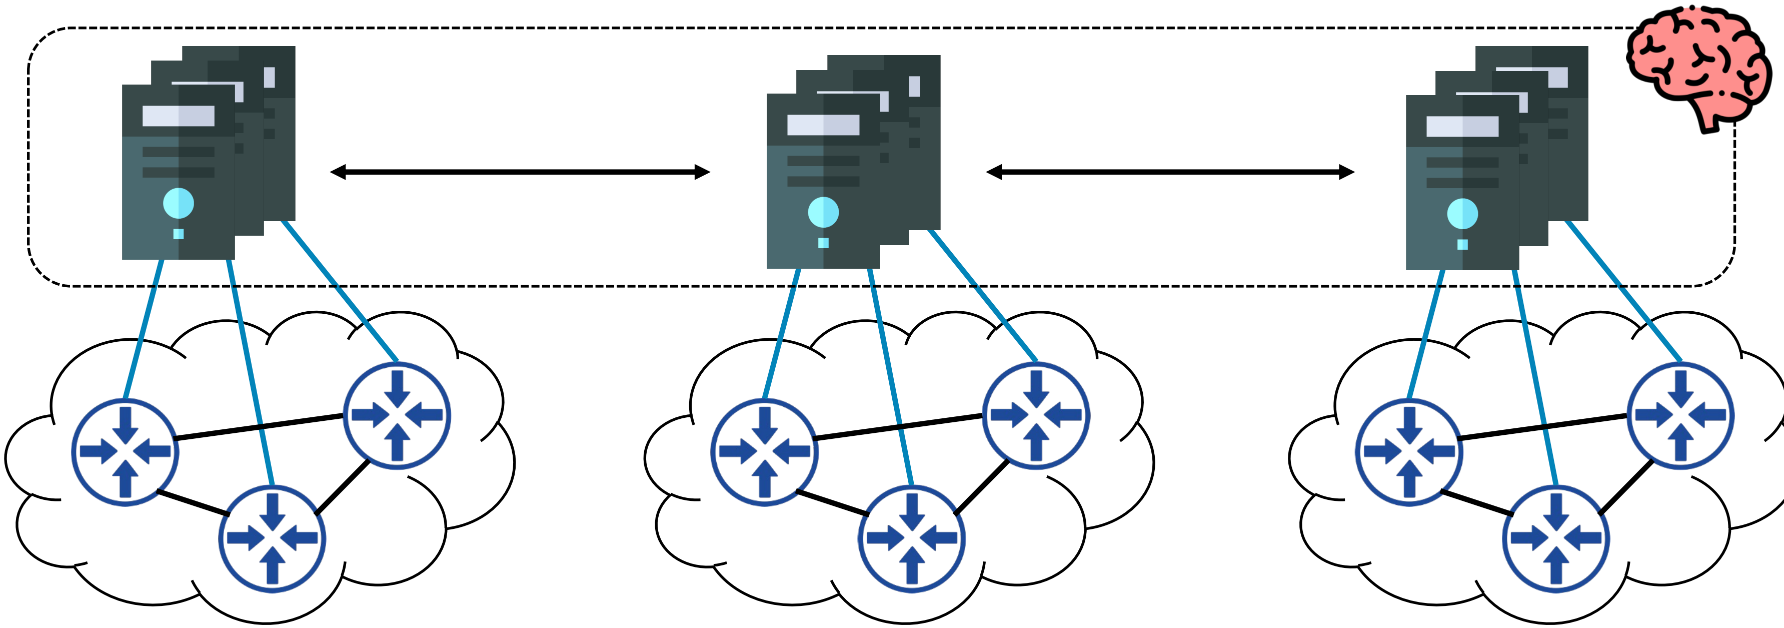

SD-WAN Control Traffic

- Exchanged between controllers/switches
  - To make a cluster keep consistent states
- Includes diverse cluster management protocols
  - E.g., consensus, membership, southbound
- Normally transmitted by a secure channel
  - E.g., SSL/TLS

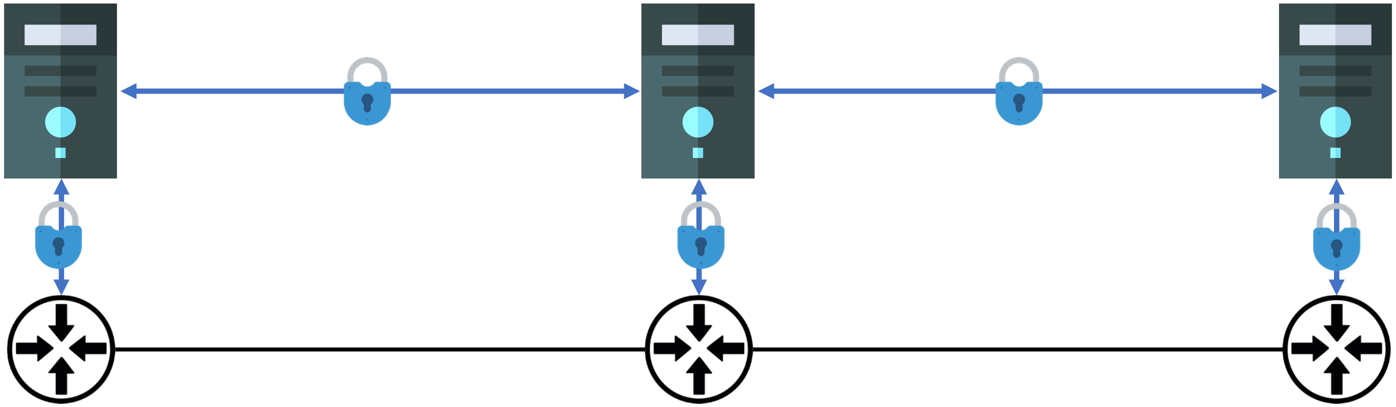

Threat Model: Eavesdropper

- Can illegally sniff WAN traffic in the middle
  - Ditto [NDSS '22][1]
- Local eavesdropper: router/link wiretapping[2]
- Network eavesdropper: BGP hijacking[3]

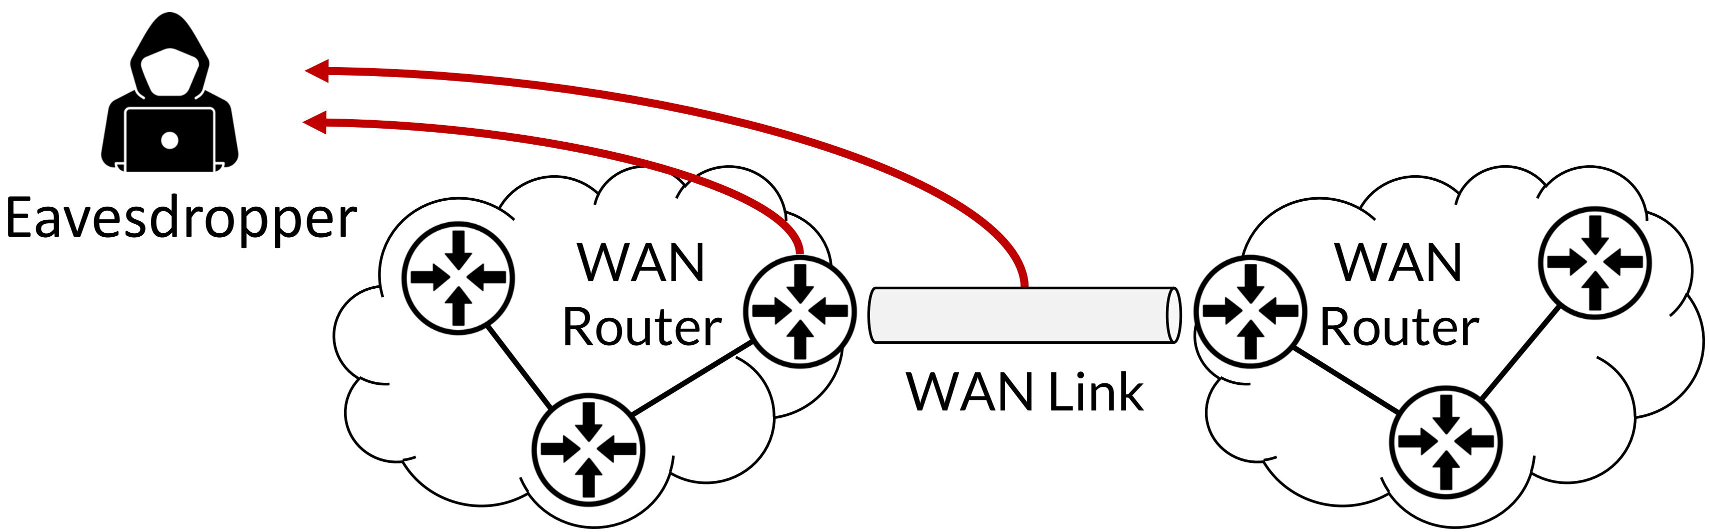

1 ditto: WAN Traffic Obfuscation at Line Rate, NDSS '22
2 "The Creepy, Long-Standing Practice of Undersea Cable Tapping", The Atlantic '17
3 RAPTOR: Routing attacks on privacy in tor, USENIX Security '15

In-band Control Channel

- Shares the same link between the control and data traffic[1]
  - Can be wiretapped by an eavesdropper

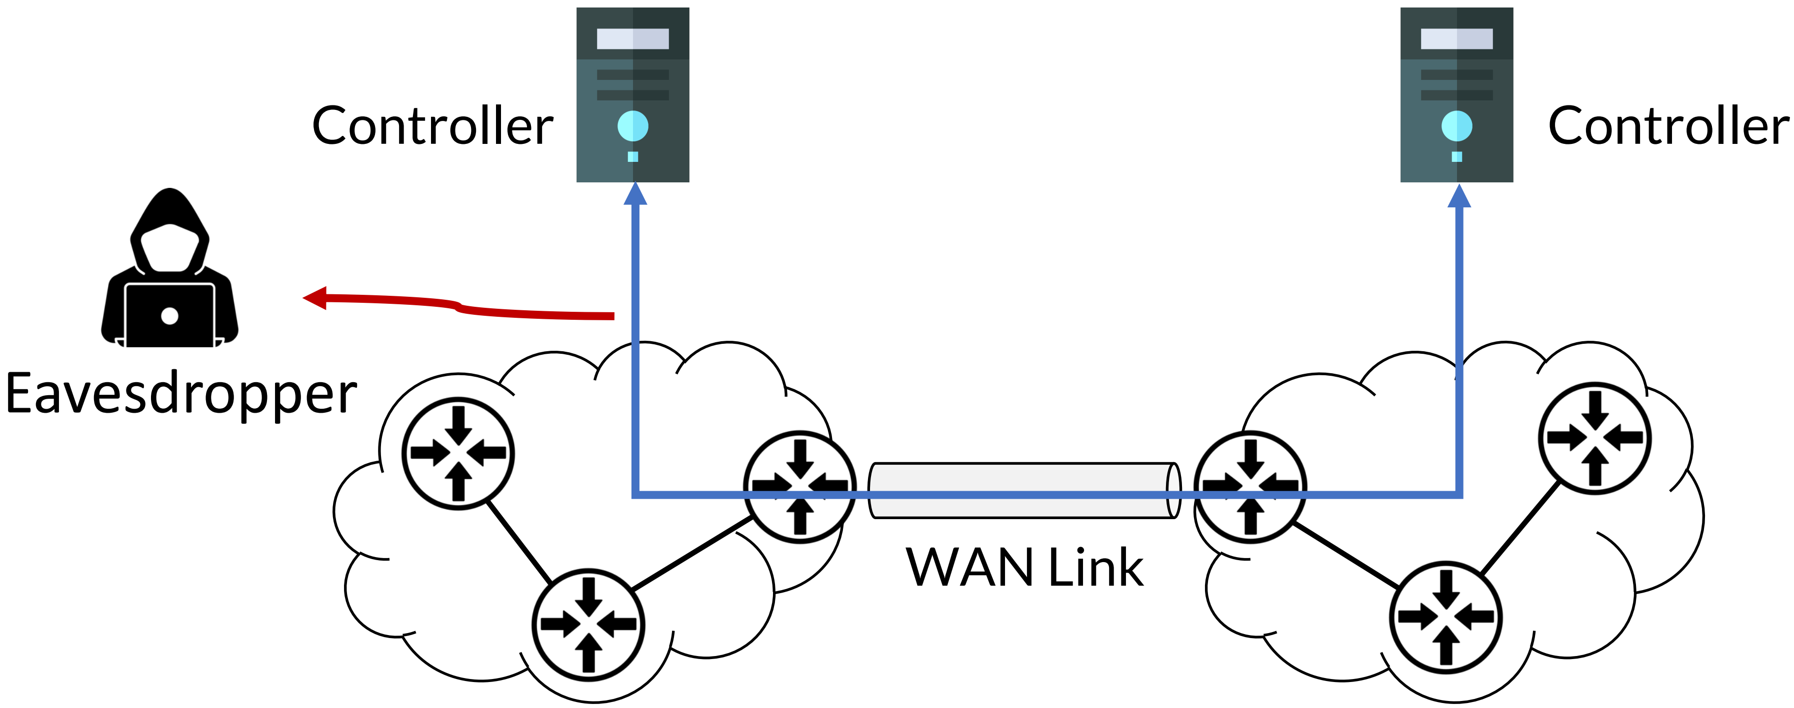

1 The CrossPath Attack: Disrupting the SDN Control Channel via Shared Links

Research Question

- “Can an eavesdropper fingerprint the confidential SD-WAN information by analyzing encrypted control traffic?”

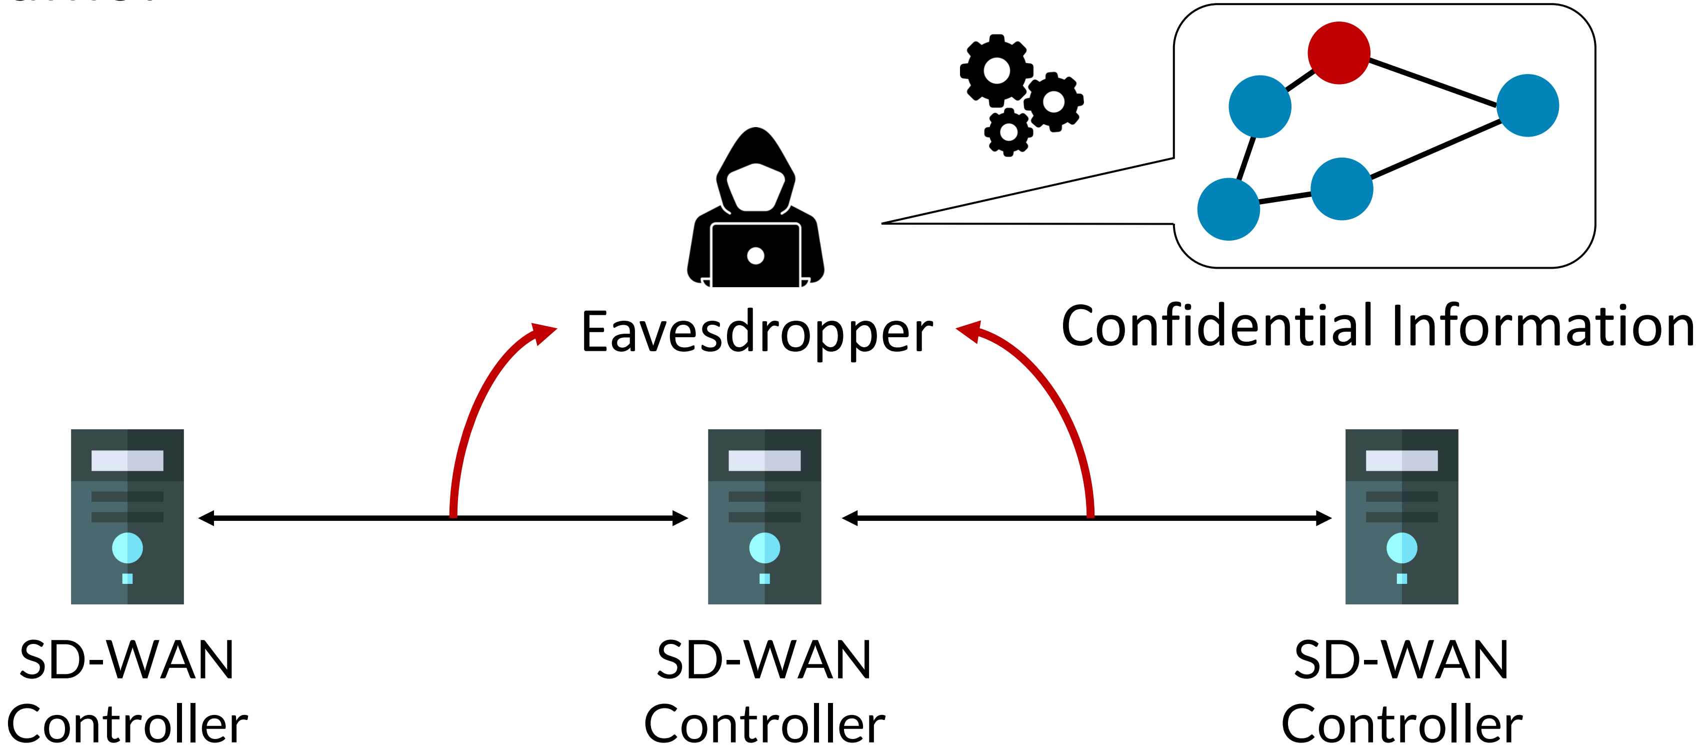

Related Work

- Aiming to leak confidential information from SDN
  - Shin and Gu [HotSDN ‘13] → Fingerprinting SDN architecture
  - Sonchack et al. [ACSAC ‘16] → Fingerprinting SDN policies
  - Achleitner et al. [SOSR ‘17] → Fingerprinting SDN policies
  - Cao et al. [RAID ‘19] → Fingerprinting SDN applications
- …using control traffic analysis

None of them focuses on fingerprinting ***SD-WAN***

Heimdallr

- A system that fingerprints SD-WAN control plane information
  - Collects traffic and extracts features automatically
  - Learns traffic patterns using a deep learning model
  - Infers confidential information on SD-WAN control-plane

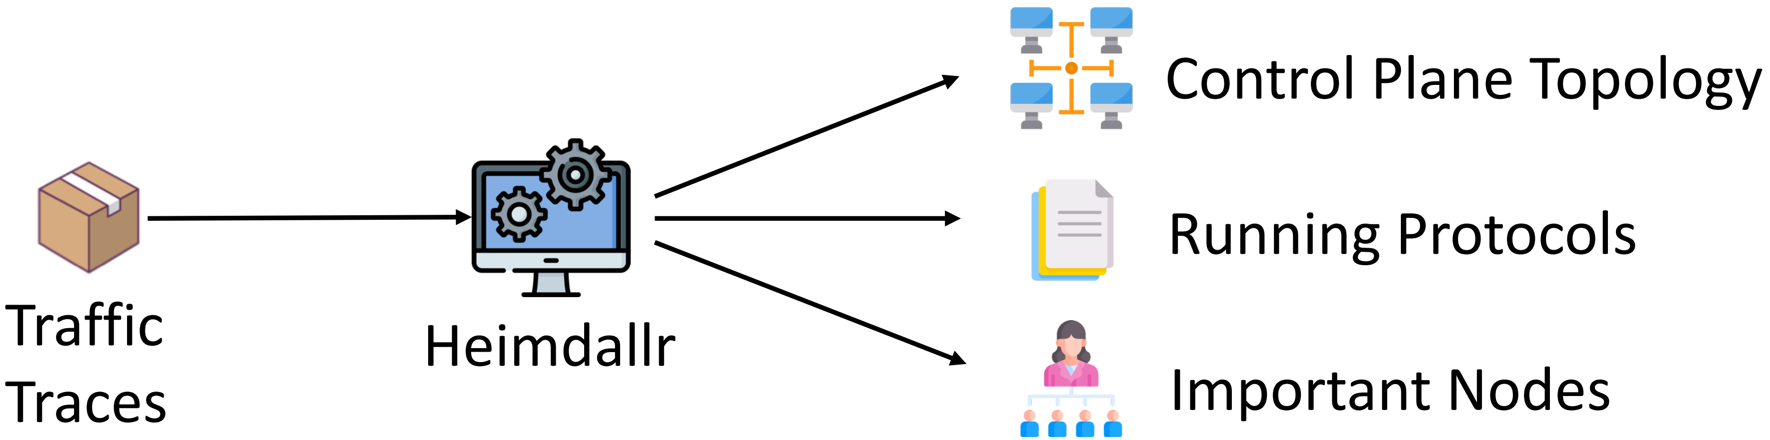

Confidential Information?

- What information might an eavesdropper have an interest in?
  - No clear definition so far
  - We define three representative types

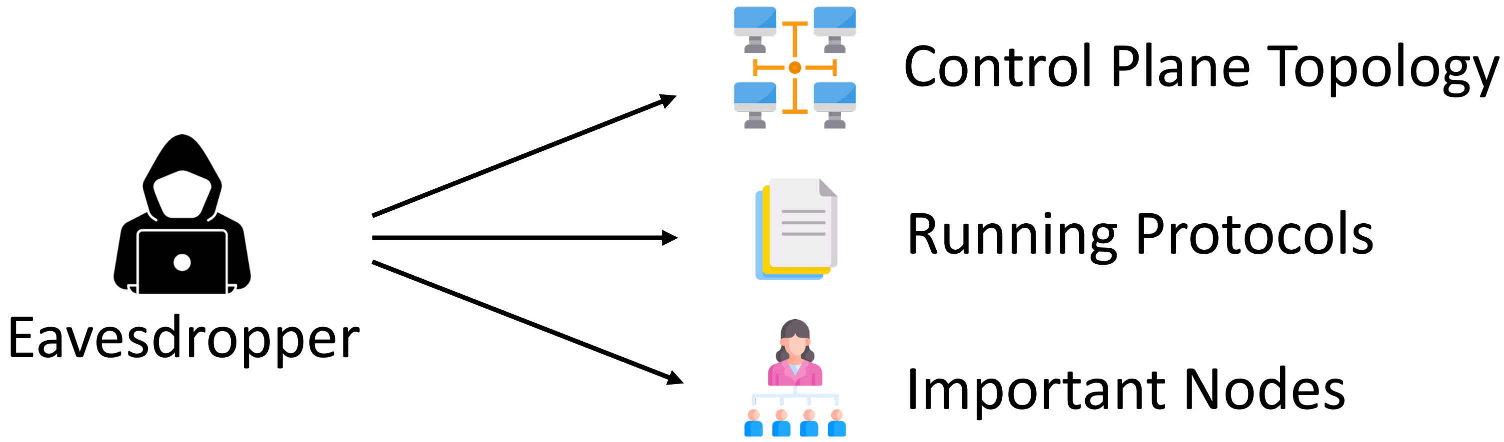

Control Plane Topology

- How a cluster is (logically) structured?
  - Controller-to-controller link?
  - Controller-to-switch link?
- What if attacker targets a specific connection?
  - E.g., The CrossPath Attack[1]

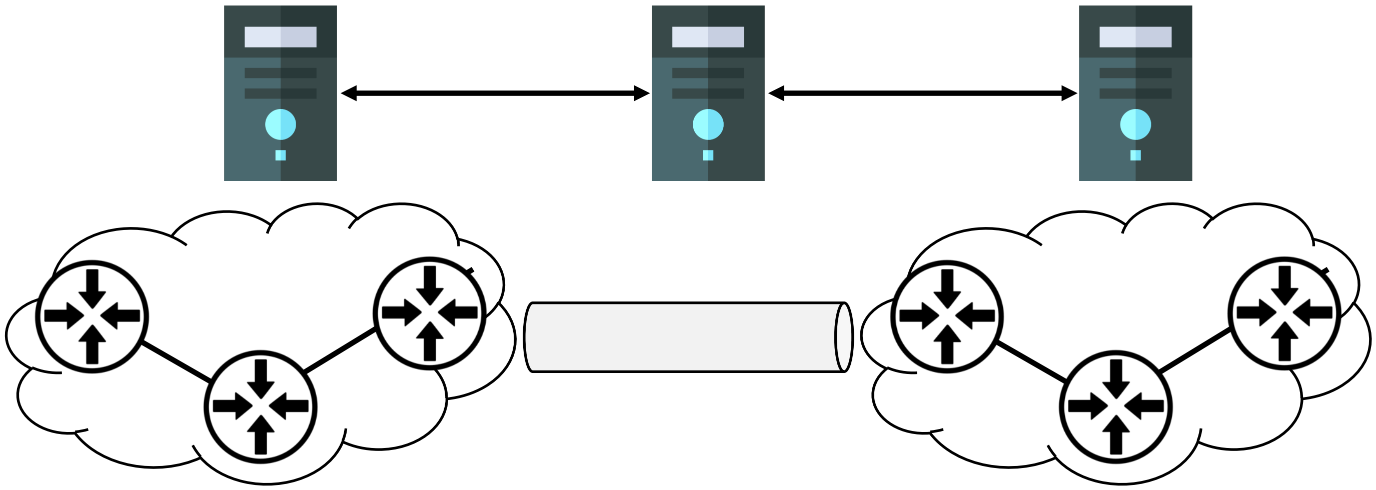

1 The CrossPath Attack: Disrupting the SDN Control Channel via Shared Links

Cluster Management Protocols

- What protocols are being used?
  - Consensus: synchronizes states between controllers
  - Membership: checks whether a controller is alive
  - Southbound: communicates with switches
- What if attacker abuses a protocol vulnerability?

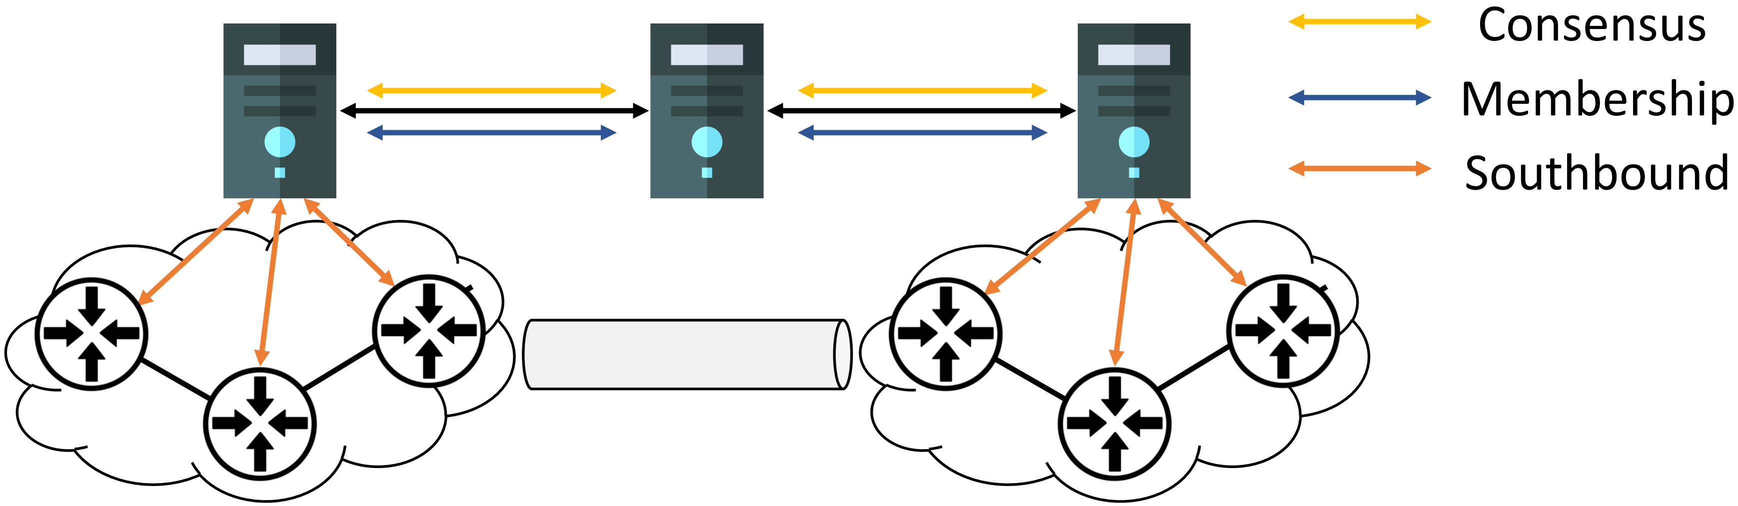

Node Roles

- Which controller is a primary role?
  - Which controller is a leader for consensus?
  - Which controller is a master for southbound?
- What if attacker targets the primary?

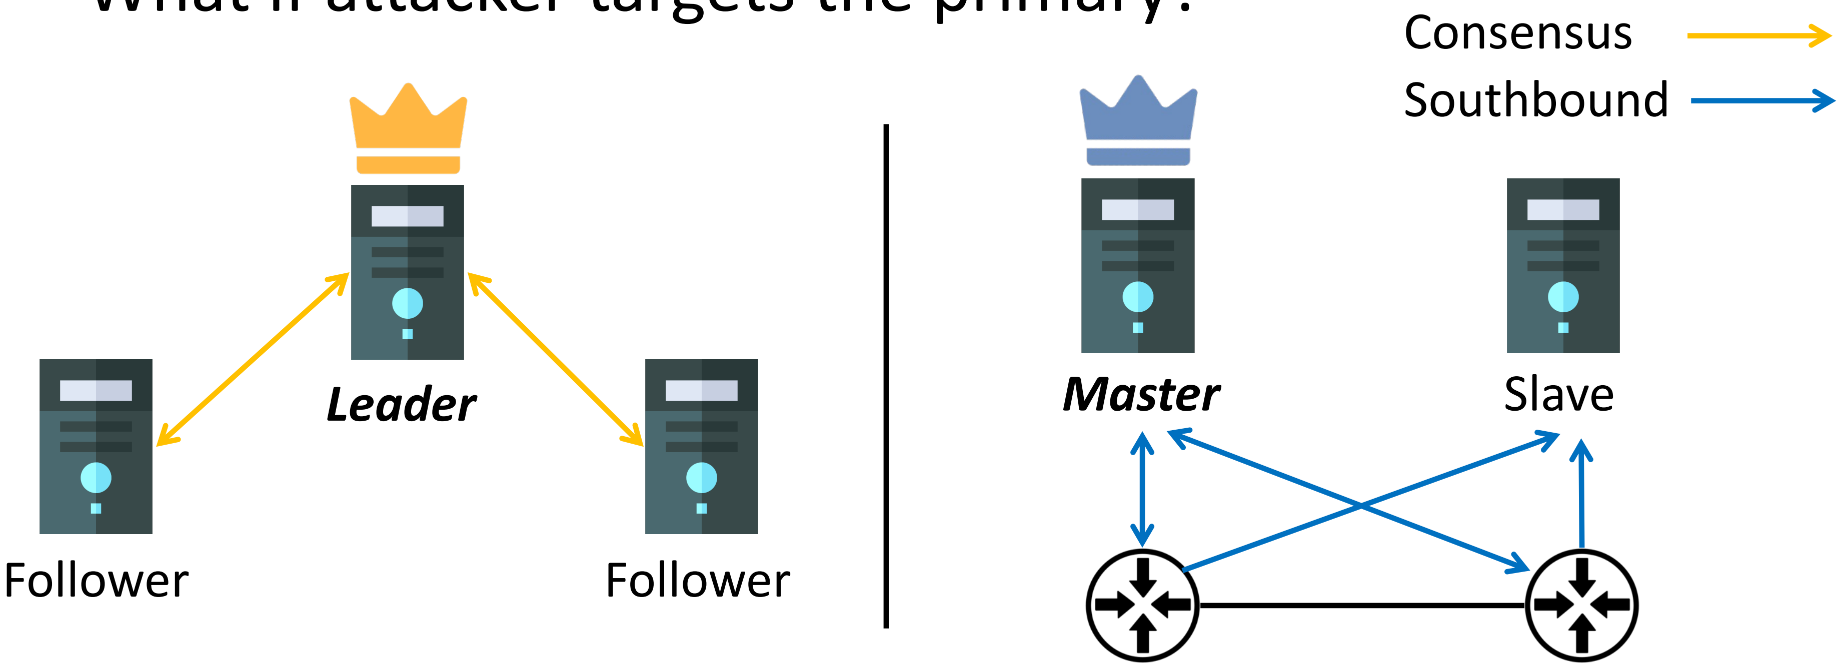

Challenges

- How to distinguish control traffic from data traffic?
  - Many traffic types in the wild
- How to distinguish cluster protocols?
  - All packets mixed in the similar connection
- How to distinguish a role for each node?
  - No information available from encrypted packets

Insight 1: Periodical Pattern

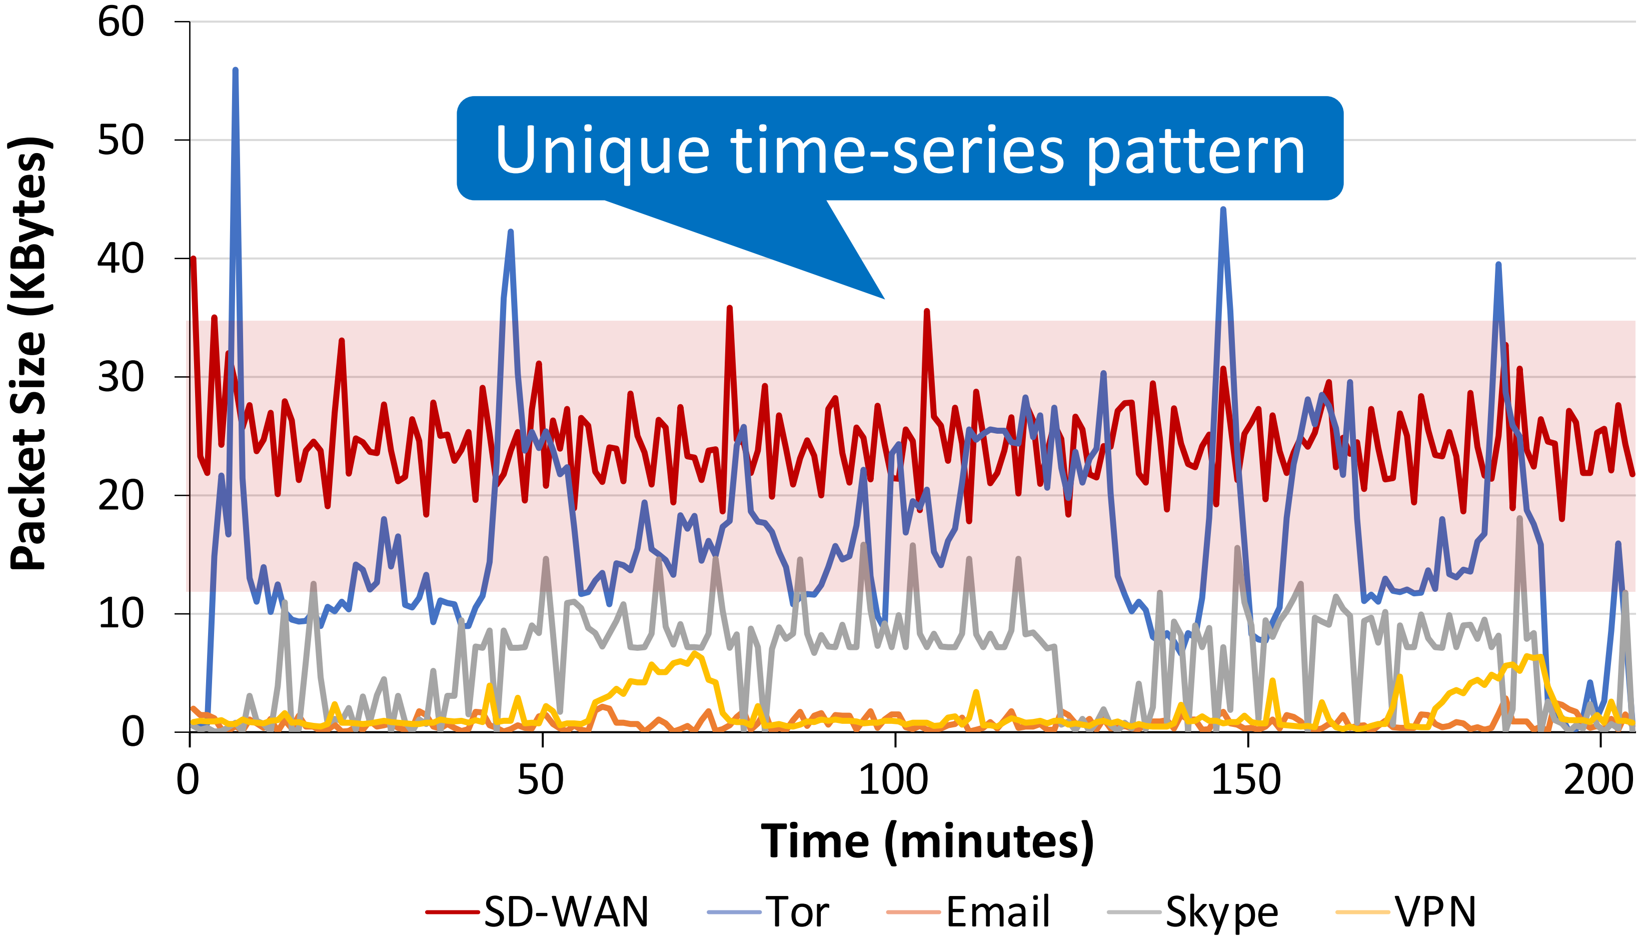

Insight 2: Directional Pattern

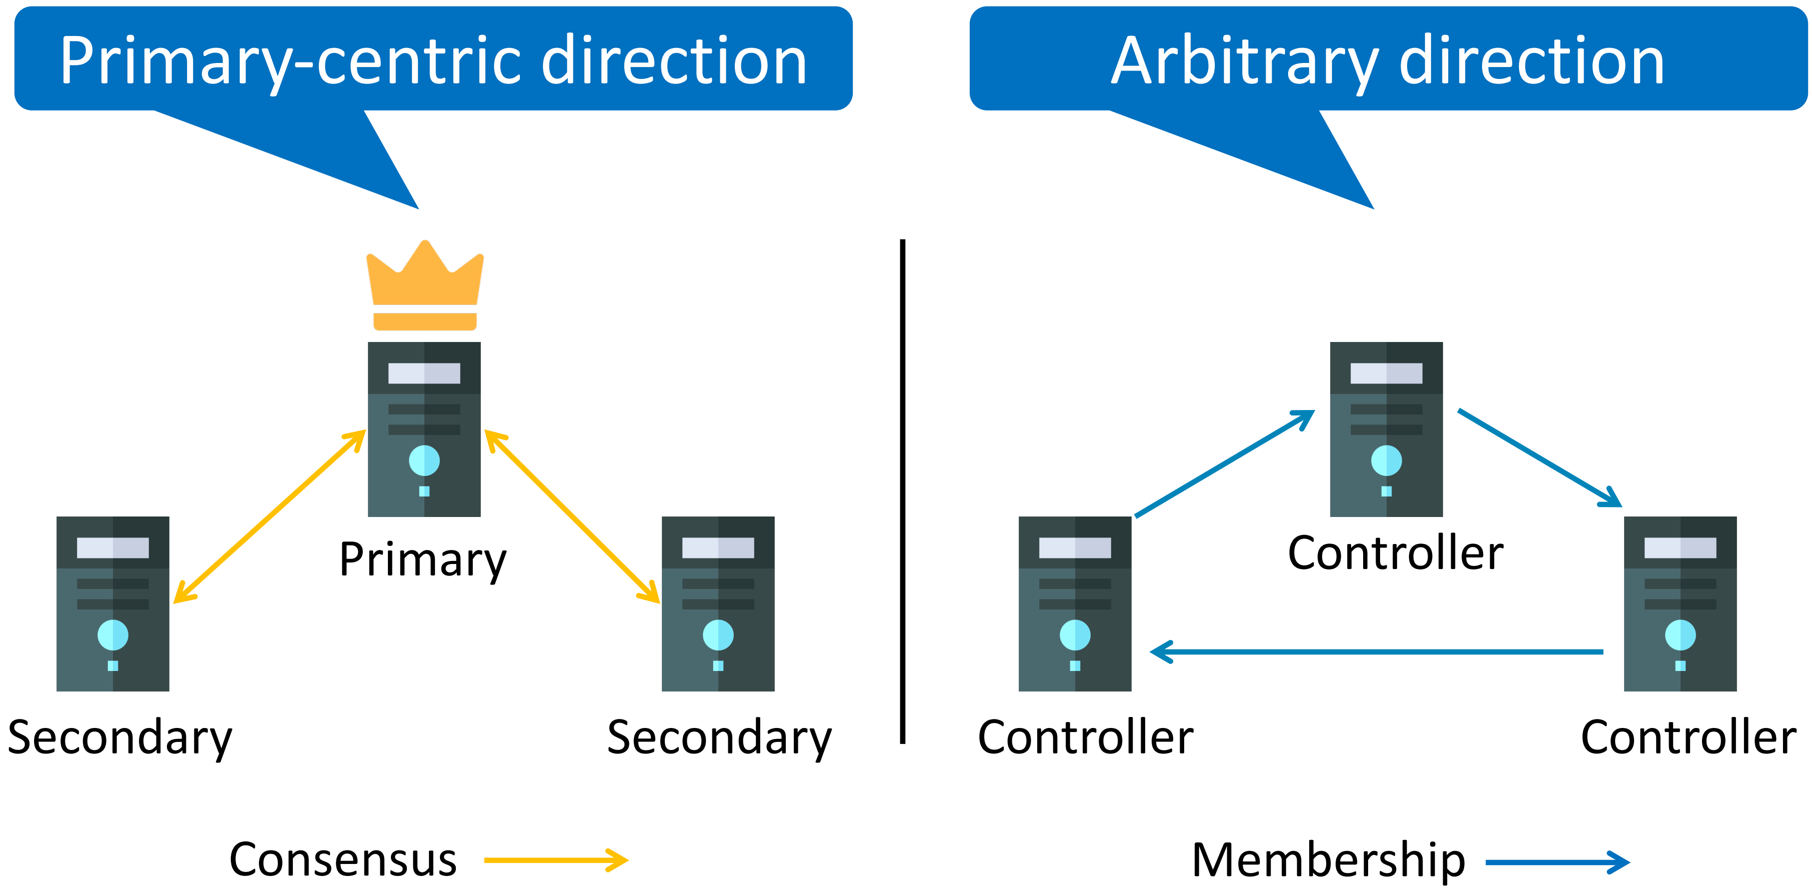

Insight 3: Traffic Distribution

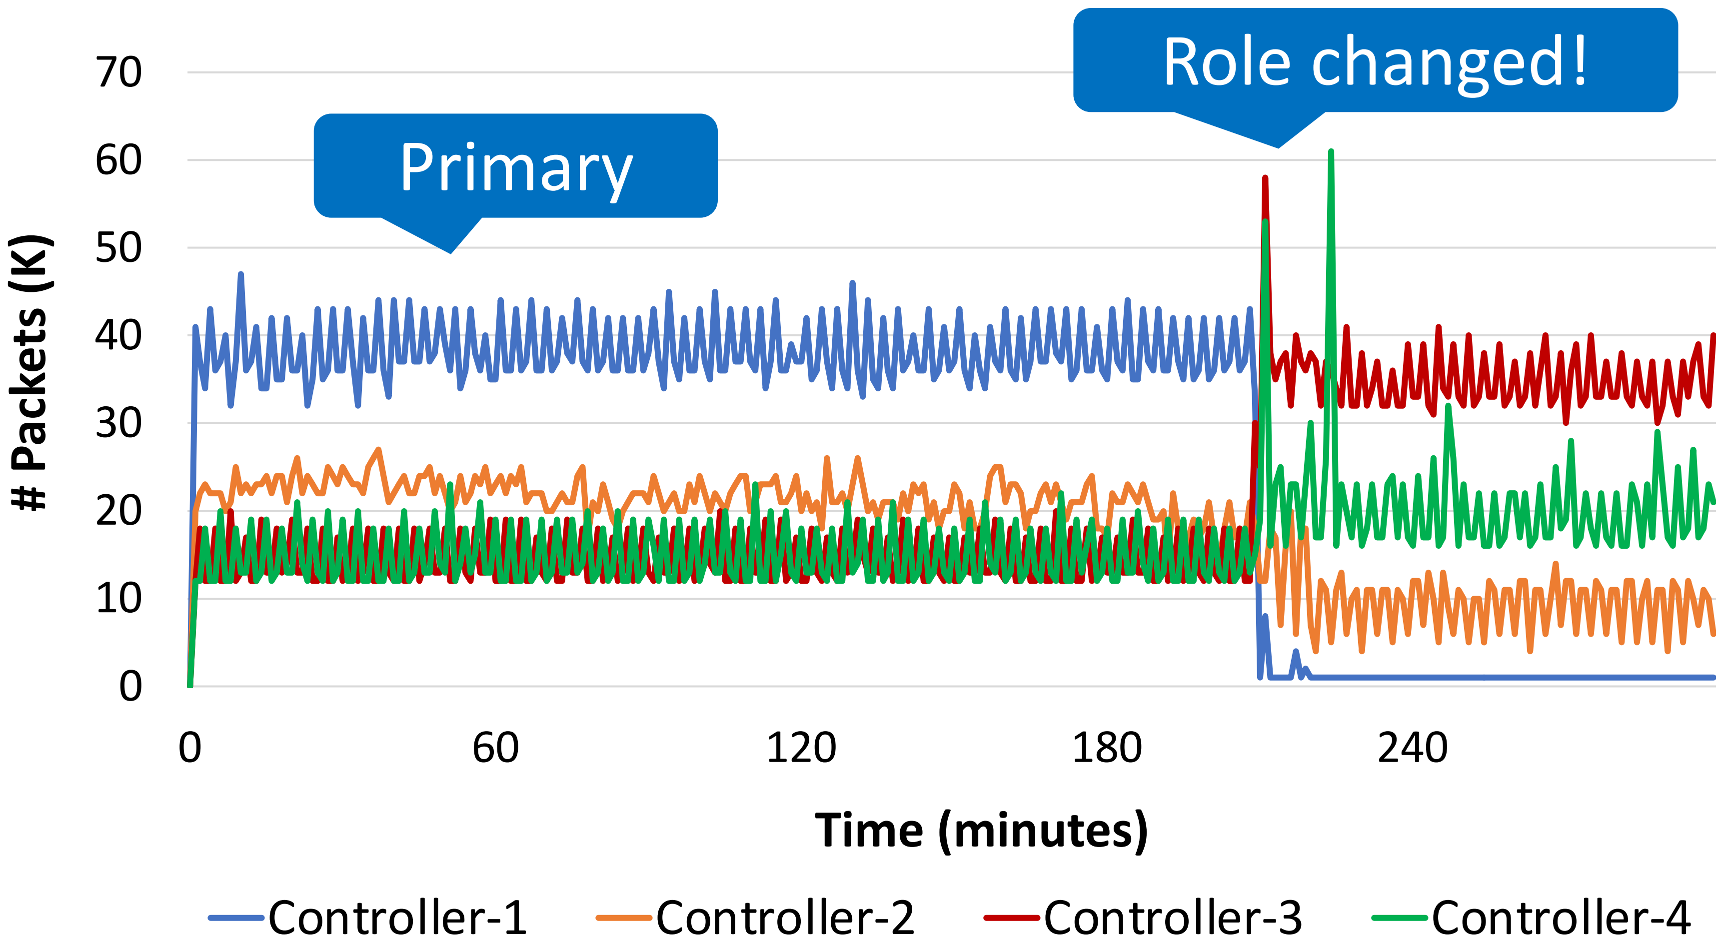

KWANGWOON UNIVERSITY

1st Phase: Identifying SD-WAN Control Traffic

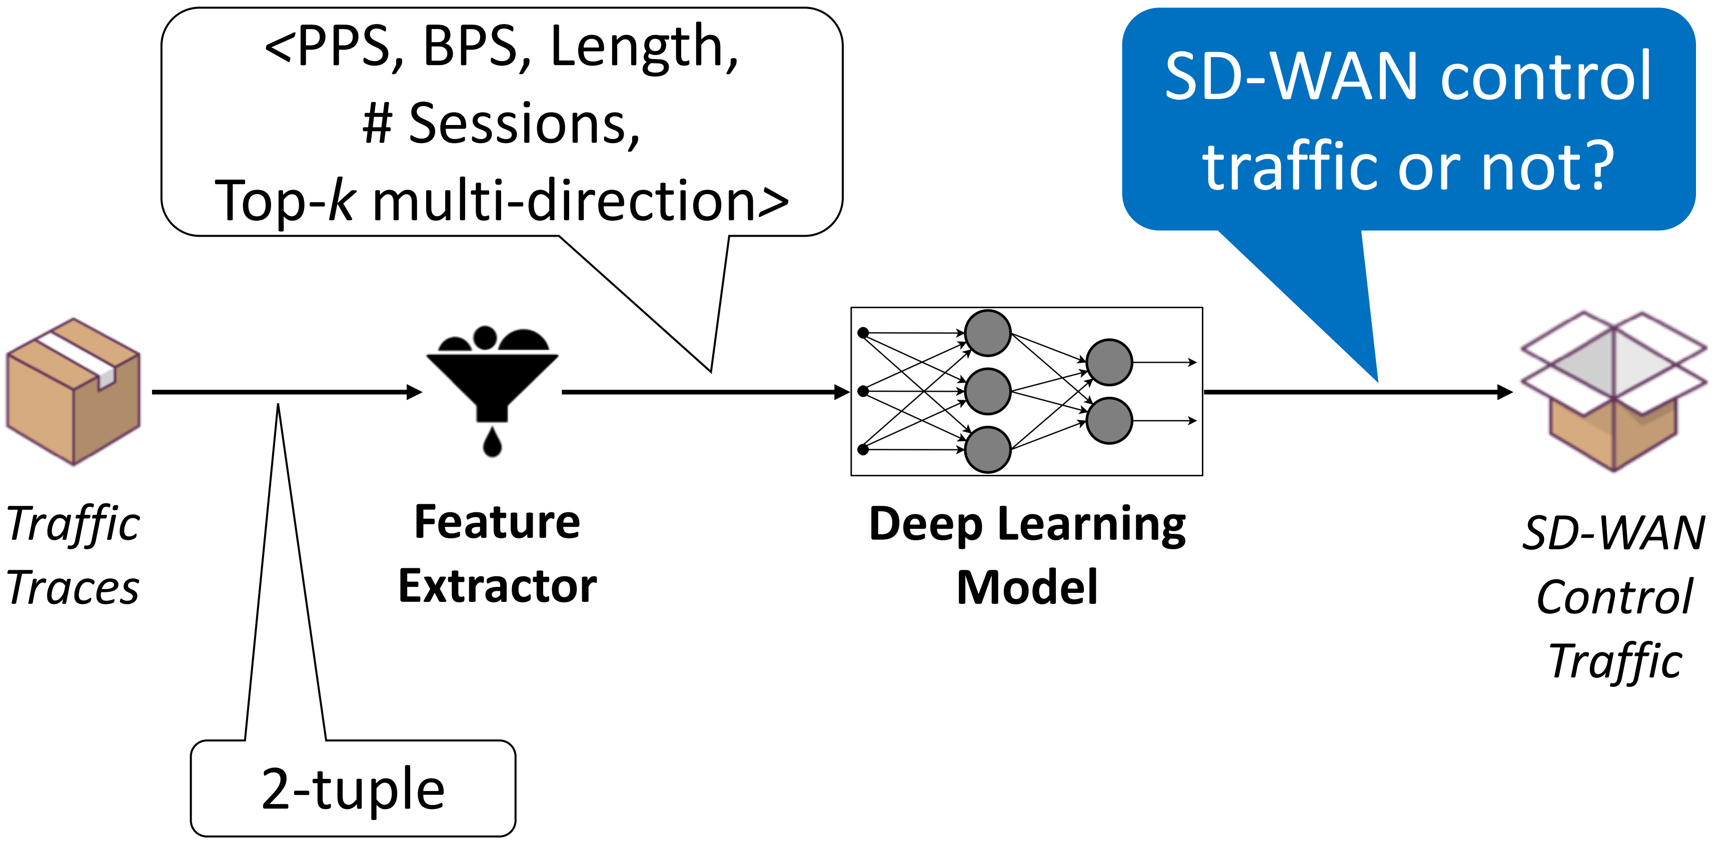

2nd Phase: Identifying Cluster Management Protocols

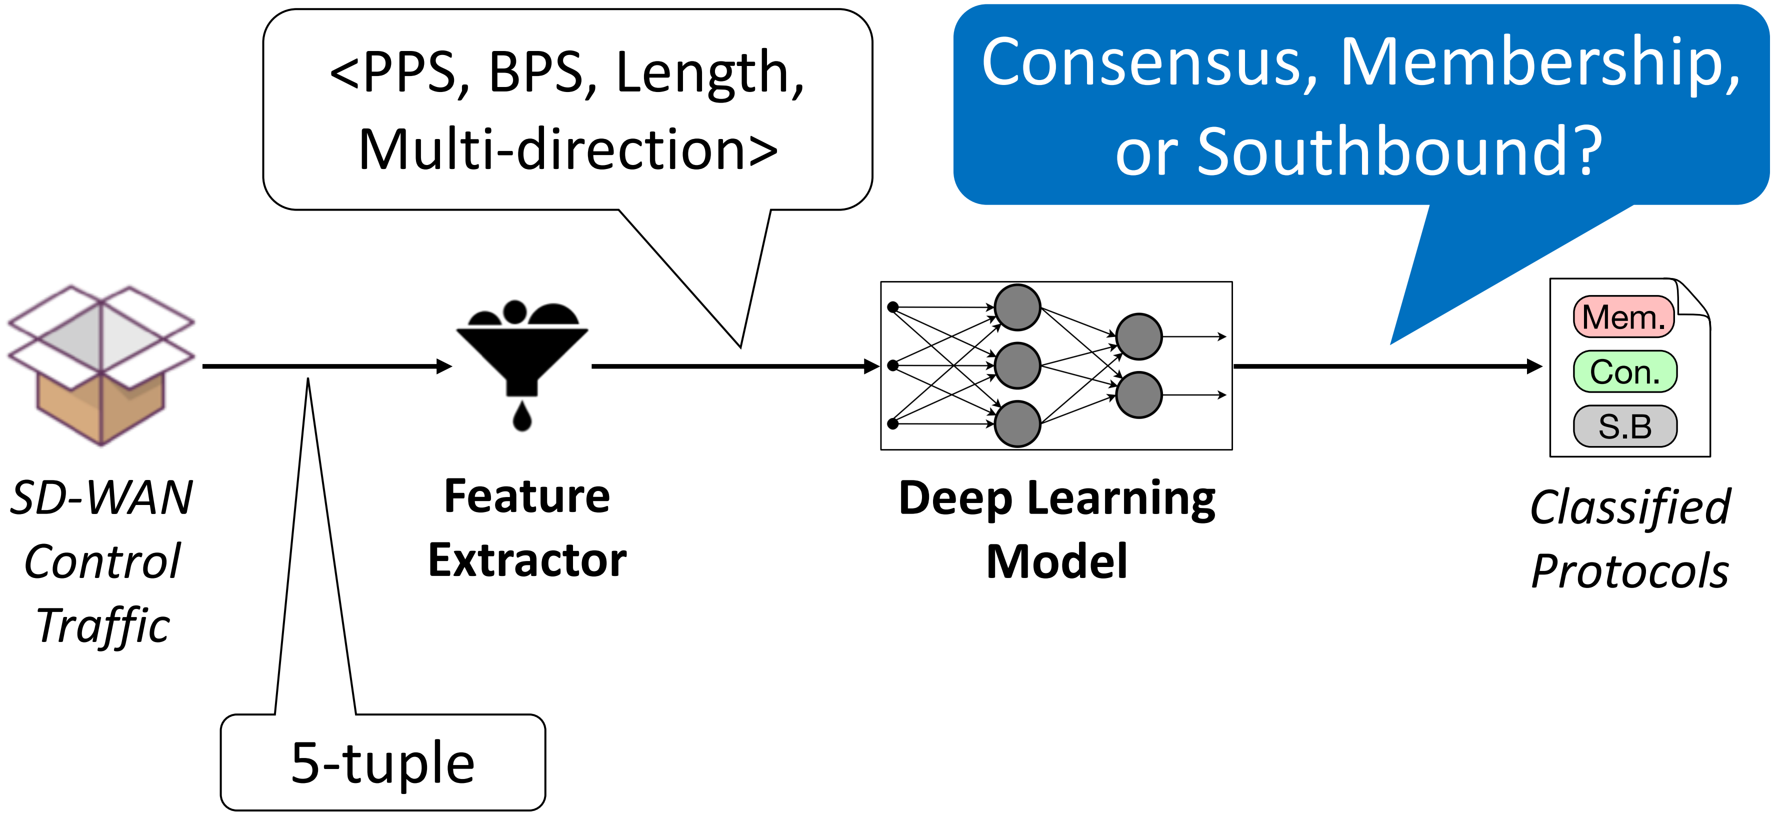

Classification Task

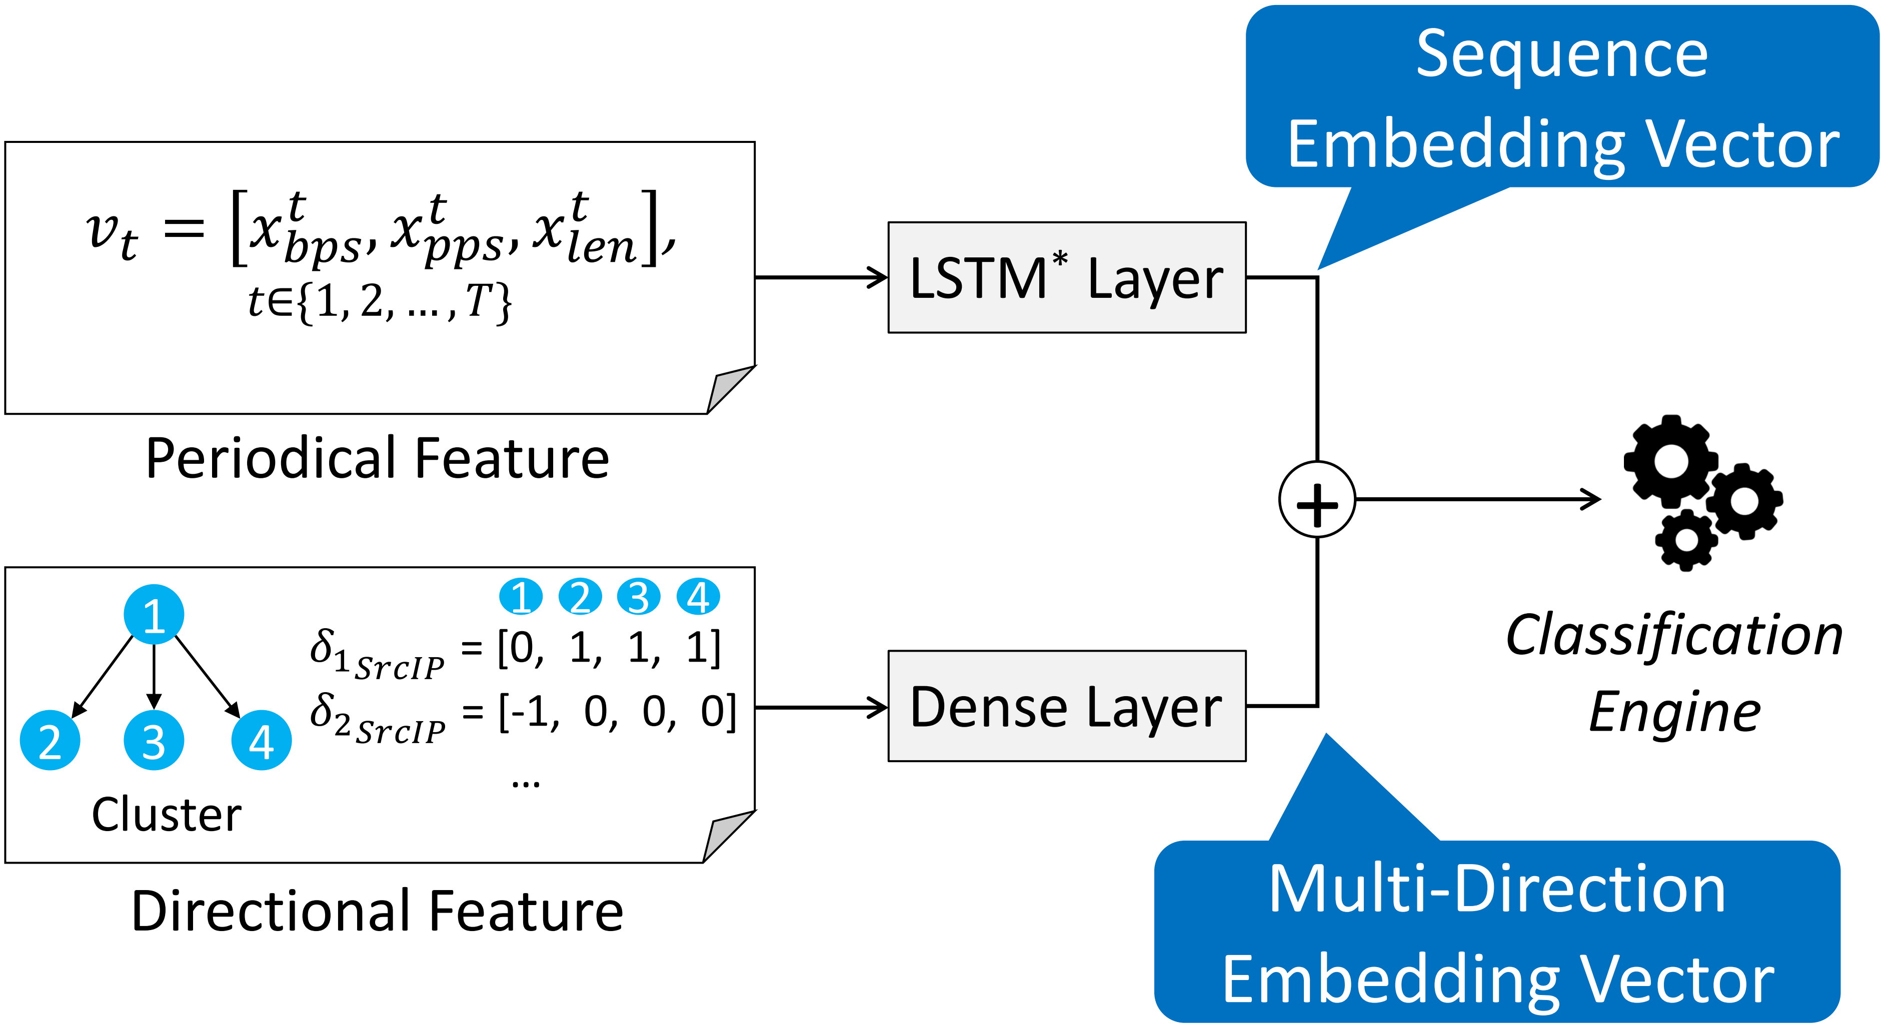

*Long Short-Term Memory

3rd Phase: Identifying Roles and Control Plane Architecture

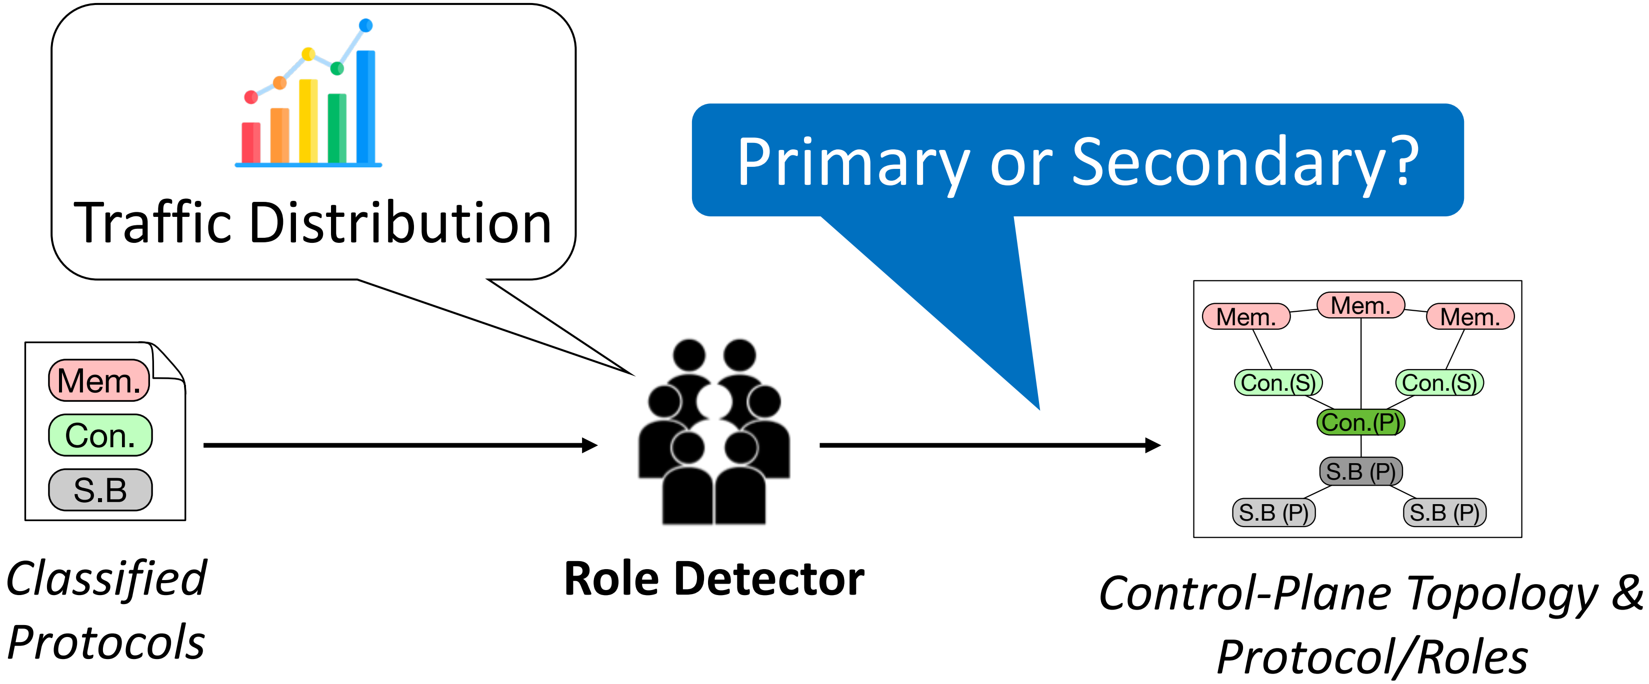

Inferring Roles with Z-Score

- Utilizes z-score of traffic amount to identify an outlier
  - Outlier whose $BPS_z \geq \theta_z$ ➔ likely to be a primary role
- How to determine a threshold θ_z?
  - Based on the analysis of traffic distribution

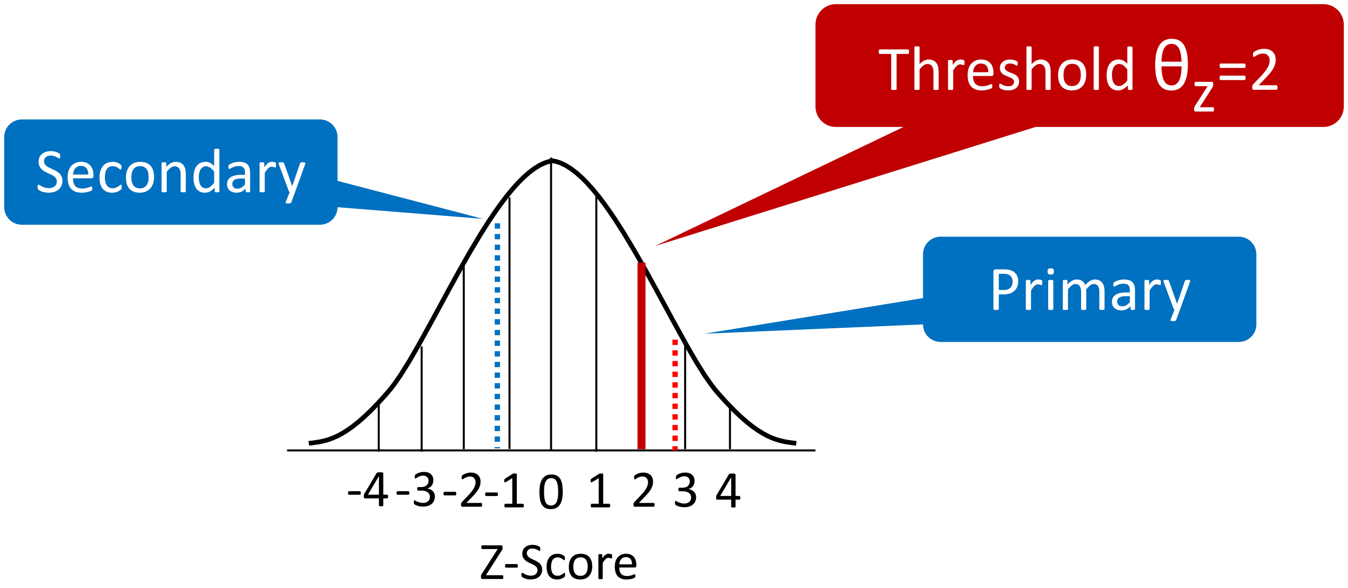

Evaluation

1. Can Heimdallr perform each fingerprinting task accurately?

2. Can Heimdallr infer SD-WAN control plane topology?

3. What is best-suited deep learning algorithm to perform fingerprinting?

4. Is Heimdallr robust to defense systems?

Evaluation

1. Can Heimdallr perform each fingerprinting task accurately?

2. Can Heimdallr infer SD-WAN control plane topology?

3. What is best-suited deep learning algorithm to perform fingerprinting?

4. Is Heimdallr robust to defense systems?

Please read our paper

Experimental Environment

- A realistic SD-WAN testbed
  - Built over 2 campus and 1 enterprise networks
  - Consists of 4 sites where controllers and switches run
    - ONOS controller and EdgeCore/Pica switches

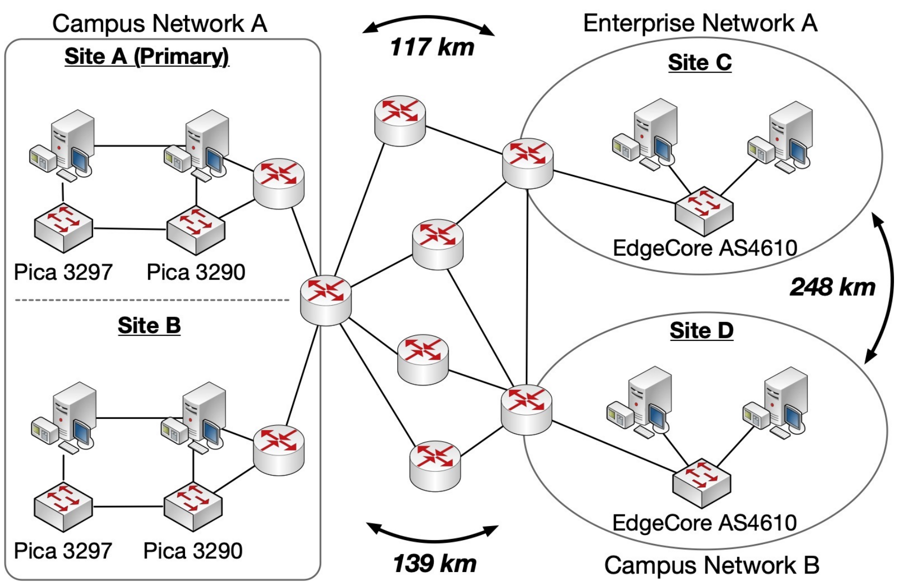

Dataset

- Collected about 53 million packets
  - Run SDN applications for control traffic and various services for data traffic
  - 70% for training and 30% for testing
- Divided into test cases for each threat model

Dataset Description
SD-WAN Control Traffic
CAIDA Backbone Traffic
Blockchain Management Traffic (Hyperledger)
Distributed Synchronization Service Traffic (ZooKeeper)
Commercial Traffic (Skype, Email, Video Streaming, etc.)

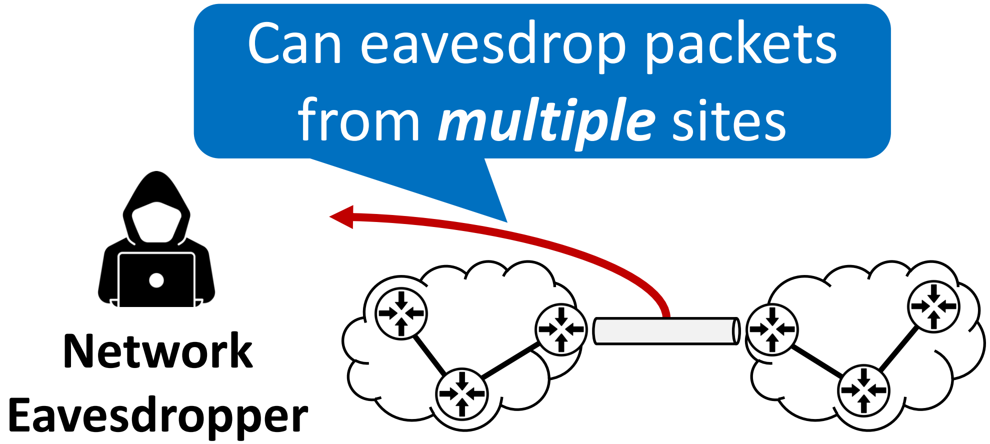

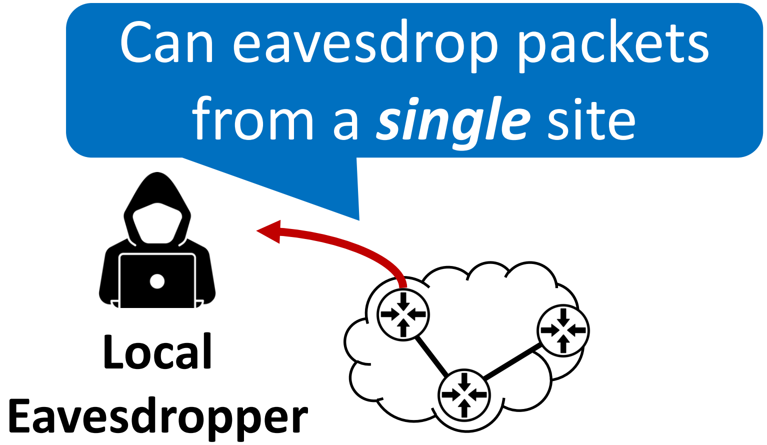

Performance of Control Traffic Classification (1st Phase)

- Uses an LSTM-based model for a classifier
  - To learn time-series features
- Can classify control traffic with ≥ 93% F1-score
  - Even by the local eavesdropper

	Traffic Type	Precision (%)	Recall (%)	F1-Score (%)
Network Eavesdropper	SD-WAN Control Traffic	96.73	95.57	96.08
	Data Traffic	99.70	99.78	99.32
Local Eavesdropper	SD-WAN Control Traffic	93.04	93.74	93.14
	Data Traffic	99.89	99.88	99.82

Performance of Cluster Protocol Classification (2nd Phase)

- To verify if Heimdallr can classify cluster protocols
  - I.e., Raft, Swim, OpenFlow
- Can classify protocols with at least ≥ 75% F1-score
  - Low F1-score due to small amount of collected packets

	Traffic Type	Precision (%)	Recall (%)	F1-Score (%)
Network Eavesdropper	Raft	81.67	78.39	80.73
	Swim	78.28	85.18	81.92
	OpenFlow	86.04	95.57	90.78
Local Eavesdropper	Raft	78.92	76.15	77.95
	Swim	76.01	72.24	74.68
	OpenFlow	84.21	95.19	89.13

Effectiveness of Role Detection (3rd Phase)

- To verify if Heimdallr can identify a role for each node
  - Leader-follower roles in Raft with a threshold θ_z=3
- Can distinguish them accurately
  - Except for the random eavesdropper (see our paper)

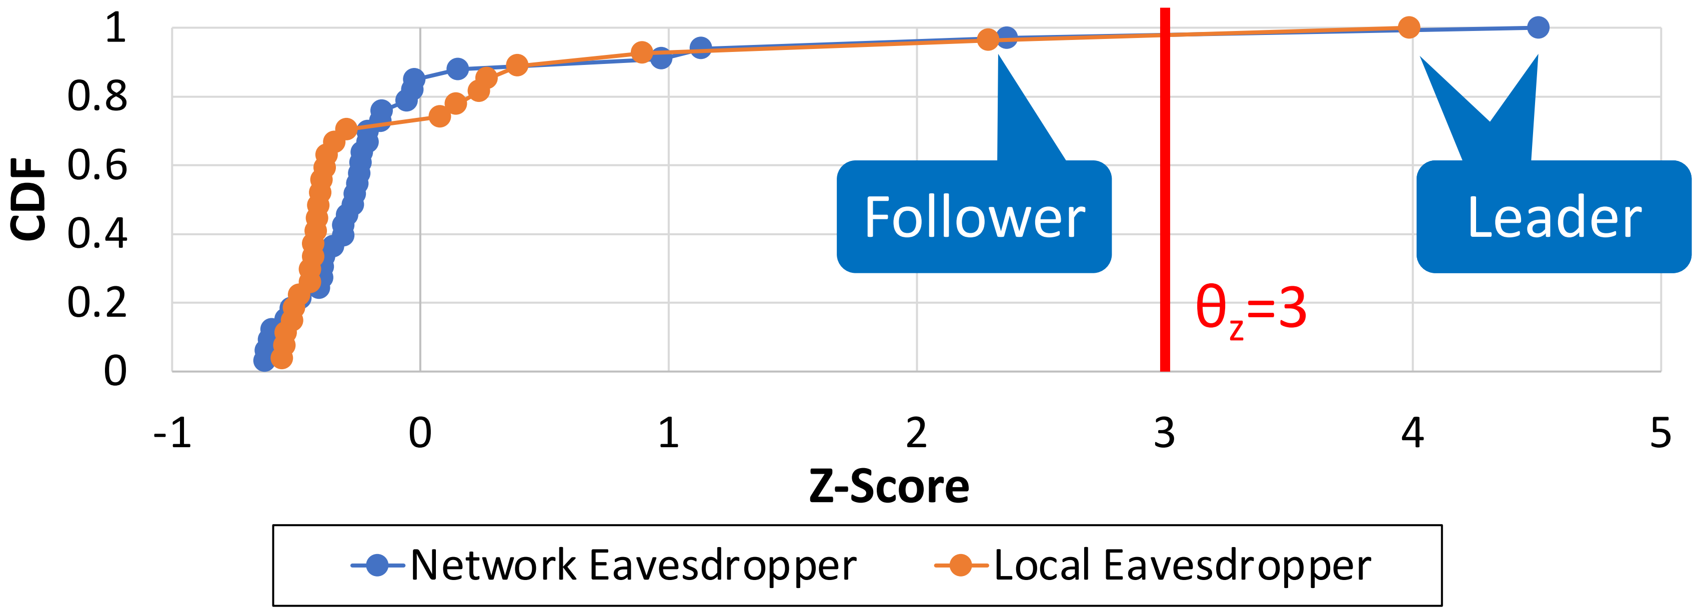

KWANGWOON UNIVERSITY

Similarity of Inferred Control Plane Topology

- Measured ***similarity*** between G_{inf} and G_{ori} using graph edit distance (GED)
  - G: a graph whose vertex V is protocol/role and edge E is their relationship[1]
  - $Similarity(G_{inf}, G_{ori}) = 1 - \frac{GED(G_{inf})}{|G_{inf}|+|G_{ori}|}$

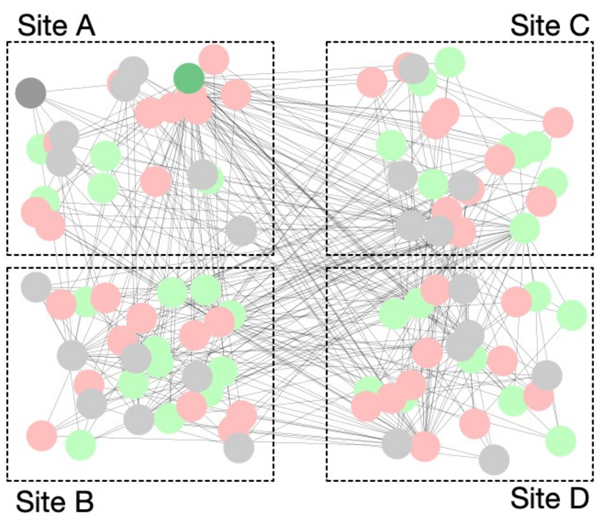

- 82% for network eavesdropper
- 70% for local eavesdropper

1 Fast Control Plane Analysis Using an Abstract Representation, SIGCOMM '16

Conclusion

- Software-Defined WAN (SD-WAN)
  - Widely deployed to operate private WANs efficiently
  - Employs multiple controllers for fault-tolerance and high-performance
  - Vulnerable to control traffic analysis attacks

- **Heimdallr**: a system for fingerprinting SD-WAN
  - Learns control traffic patterns systematically
  - Infers protocols, roles, and control-plane topology with a reasonable accuracy

Thank you for listening

(jinwookim@kw.ac.kr)

ACKNOWLEDGMENTS: This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2022-00166401). The research leading to these results have received funding from the European Union's Horizon 2020 research and innovation programme under grant agreements No 871793 (Accordion), No 101016509 (Charity) and No 101070473 (FLUIDOS).