



AmLight *EXP*

Americas Africa Lightpaths *Express & Protect*

Innovating the Network for Data-Intensive Science -INDIS, November 18, 2024



Leveraging In-band Network Telemetry for Automated DDoS Detection in Production Programmable Networks: The AmLight Use Case

*Hadi Sahin, PhD
Research Scientist*

Hadi Sahin, Jeronimo Bezerra, Italo Brito, Renata Frez, Vasilka Chergarova, Luis Fernandez Lopez, Julio Ibarra

Motivation

- In-band Network Telemetry (INT) has been available since 2015, providing rich network state information.
- We have been using INT in Amlight since 2018
- Research and practical deployments of INT for Distributed Denial of Service (DDoS) threat detection remain limited:
 - Existing studies primarily rely on data generated from simulation environments (e.g., Mininet)
 - In the industry, sFlow and NetFlow are generally used for traffic monitoring.
 - There is a lack of analysis comparing the performance of INT-based approaches to traditional monitoring tools, such as sFlow.

Key Contributions of This Paper

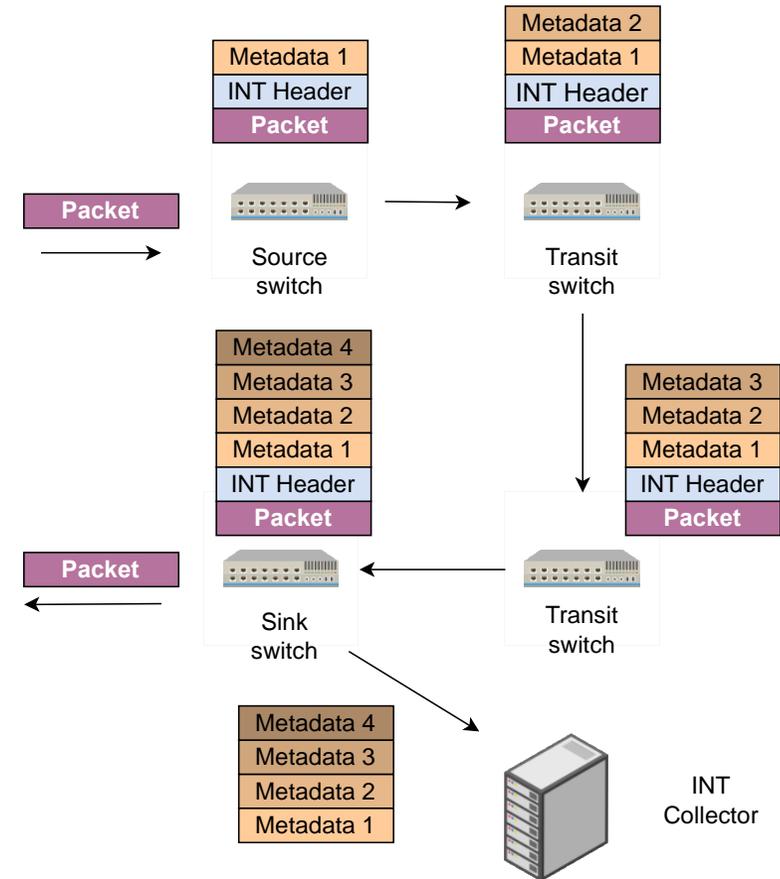
In this work, we leverage In-band Network Telemetry (INT) technology implemented in the AmLight network to detect Distributed Denial of Service (DDoS) attacks:

- Utilize real-world production INT data to detect DDoS attacks.
- Compare DDoS attack detection using INT-based analysis with traditional sFlow-based monitoring.
- Propose an automated, machine learning-driven approach for DDoS attack detection.

Background Information

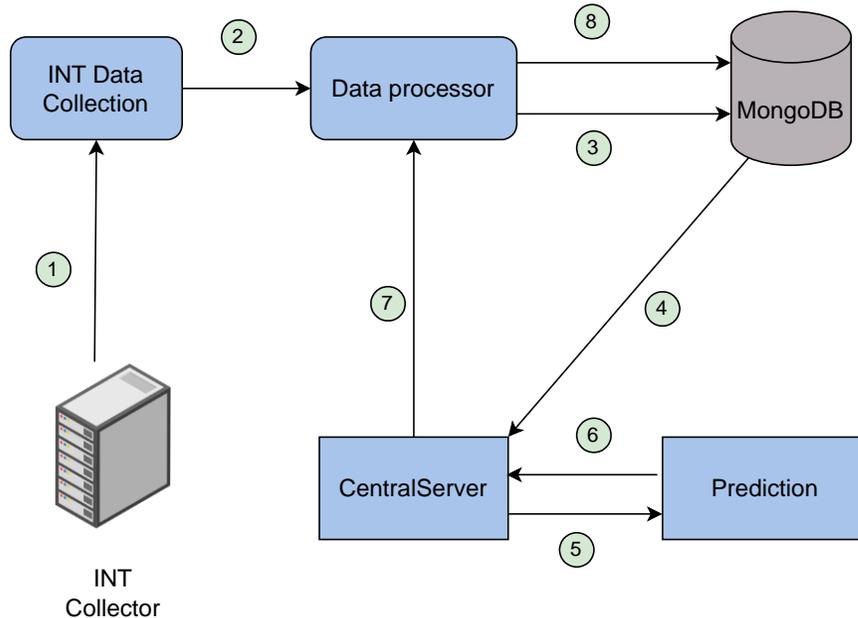
INT and sFlow tools

- INT technology combines data packet forwarding with network measurement.
- It embeds telemetry information into packets as they traverse the network.
- sFlow captures and samples packets across network devices (1/4096).



Automated DDoS Detection

Proposed Mechanism



1. Gather INT data.
2. Send INT data to the Data processor:
 - Flow ID: src/dst IP, src/dst ports, protocol.
 - Create flow-level features (e.g., Packets per second, Flows per second).
3. Save processed data to the database.
4. Retrieve processed data from database
5. Send data to the prediction model.
6. Receive predictions.
7. Send predictions to the Data processor for ensemble voting.
8. Save final prediction score to the database.

Evaluation Setup

Machine Learning Models

We use machine learning (ML) models to classify benign versus normal flows (binary classification). The following models are used:

- Random Forest (RF)
- K-Nearest Neighbors (KNN)
- Gaussian Naive Bayes (GNB)
- Neural Network (NN) with three hidden layers of 32, 16, and 8 neurons.

Evaluation Setup

Used Metrics

- Accuracy: Proportion of correctly classified instances among all instances. $(TP+TN)/(TP+TN+FP+FN)$
- Recall (R): Proportion of actual positives correctly identified. $TP/TP+FN$
- Precision (p): Accuracy of positive predictions. $TP/TP+FP$
- F1: Harmonic mean of Precision and Recall. $2P*R/(P+R)$

Note: True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN)

Evaluation Setup

Data Source

- Data were collected from a subnet of the AmLight network from June 6 to June 11, 2024
 - Benign flows were collected for all days.
 - Simulated attack flows were generated on June 10th and 11th (~60M packets per day, ~700 per second)
 - Attack types include SYN Flood, SYN scan, UDP scan, and SlowLoris

Attack Type	Date	Attack Episode
SYN Scan	06.10.2024	13:24:02 - 13:57:03
SYN Scan	06.10.2024	16:30:51 - 16:35:20
UDP Scan	06.10.2024	16:36:20 - 16:53:00
UDP Scan	06.10.2024	16:56:45 - 16:59:99
SYN Flood	06.10.2024	20:48:01 - 20:49:01
SYN Flood	06.10.2024	20:52:11 - 20:54:12
SYN Flood	06.11.2024	20:13:31 - 20:15:31
SYN Flood	06.11.2024	20:16:41 - 20:17:01
SYN Flood	06.11.2024	20:17:17 - 20:17:37
SlowLoris	06.11.2024	20:27:37 - 20:28:37
SlowLoris	06.11.2024	20:29:12 - 20:31:12

Evaluation Setup

Feature Selection

Features	INT	sFlow
Protocol	✓	✓
Packet Size*	✓	✓
Number of packets	✓	✓
Queue Occupancy*	✓	×
Hop Latency*	✓	×
Inter Arrival Time*	✓	✓
Flow rate (Gbit/s)	✓	✓
Packet rate (Packet/s)	✓	✓

- Features from INT data differ from those of sFlow in *Queue occupancy* and *Hop latency*
- For (*) we include *cumulative, average, and standard deviation of the variables.*
 - *Cumulative IAT => flow duration*
 - *Flow rate = Total packet size/flow duration*
 - *Packet rate = Total # of packets/ flow duration*

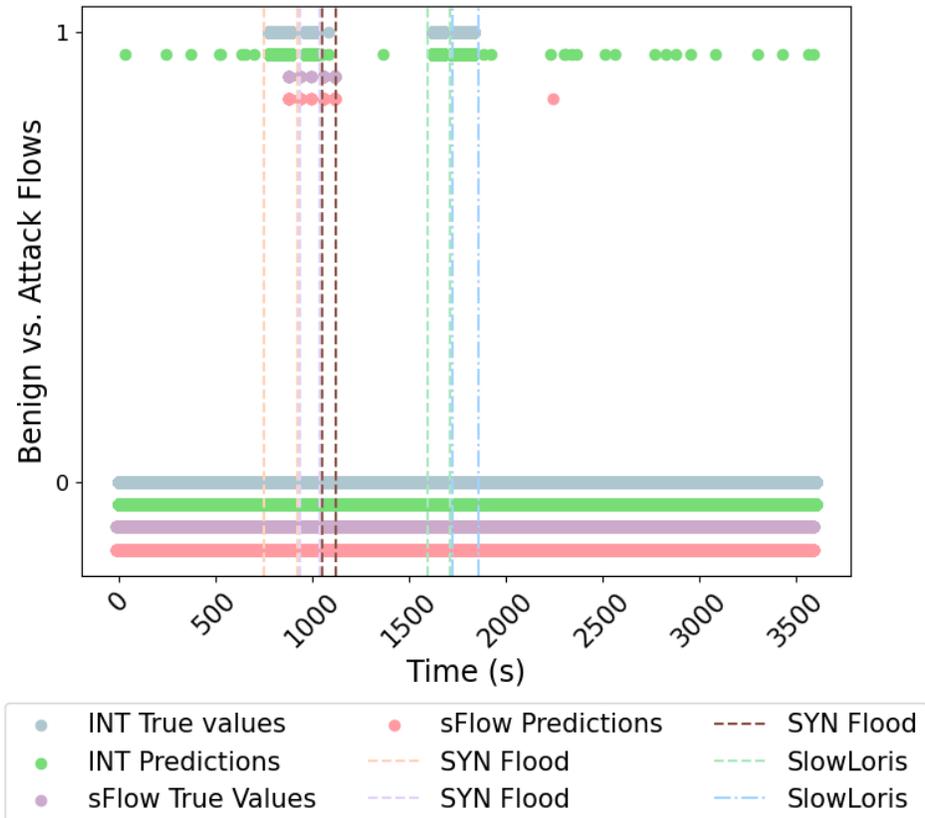
Experimental Results I

- We use data flows from June 11, 2024, as the test set to evaluate the models.
 - This includes both benign and attack flows.
 - It also contains SYN Flood (seen) and SlowLoris (unseen attacks) attacks.
- For the RF and KNN models INT and sFlow data show comparable performance.

Data	Model	Accuracy	Recall	Precision	F1-score
INT	RF	1.0000	1.0000	0.9999	1.0000
sFlow	RF	0.9999	1.0000	0.9907	0.9953
INT	GNB	0.9919	1.0000	0.9959	0.9959
sFlow	GNB	0.9959	1.0000	0.6057	0.7544
INT	KNN	0.9988	0.9993	0.9984	0.9988
sFlow	KNN	0.9997	1.0000	0.9550	0.9770
INT	NN	0.9996	1.0000	0.9992	0.9996
sFlow	NN	0.9937	0.0000	0.0000	0.5000

Experimental Results I

A closer Look at the Predicted Flows



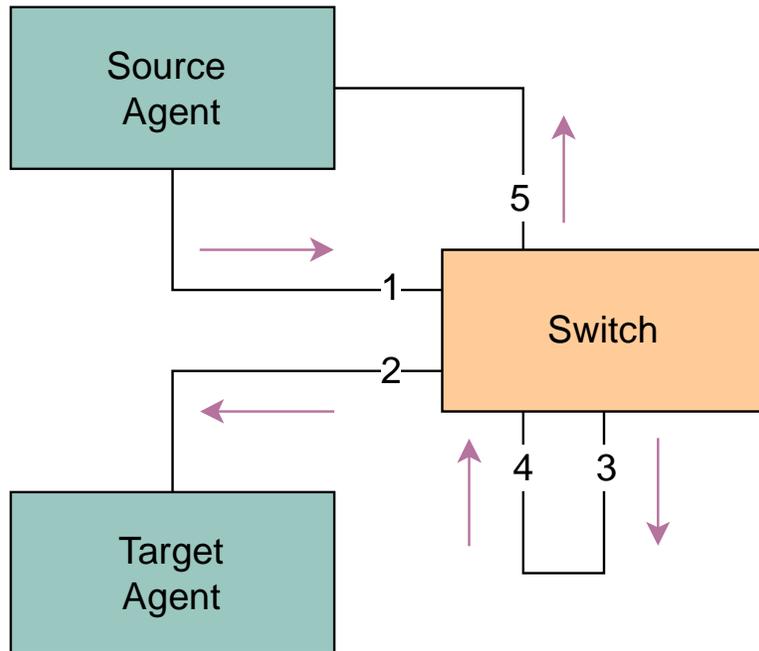
- Vertical lines: attack episodes
 - SYN Flood, SYN Flood, SYN Flood, SlowLoris, SlowLoris, respectively
- Colors indicate:
 - Gray: INT True Values
 - Green: INT Predictions
 - Purple: sFlow True Values
 - Pink: sFlow Predictions
- Y-axis: Predictions
 - 0: Benign flows
 - 1 : Attack Flows

sFlow may not capture all attack flows due to sampling.

- As a result, predictions based on sFlow data might miss certain threat episodes.

Experimental Results II

The INT Testbed



- `tcpreplay -i <interface> -p <number of packets> <pcap file path>`
 - Benign flows 3 packet/s
 - Attack flows up to 100 packet/s
- Flows go from the source to target
 - Switches 3 and 4 act as source and sink.
 - INT metadata is removed from the packet load at switch 4.
 - INT data is gathered from 5.

Experimental Results II

Attack Type	Accuracy	Misclassified/ Number of Predicted Packets	Average Prediction Time (s)
UDP Scan	0.9947	14/2628	0.12
SYN Scan	0.9961	10/2542	0.44
SYN Flood	0.9984	27/2814	0.09
SlowLoris	0.9795	16/779	0.05
Benign	0.9417	136/2331	103.14

- We achieved over 97% accuracy in predicting all attack types, with an average response time of under 2 seconds
- The accuracy for benign flow is slightly lower
- The average prediction time is also longer
 - This is due to bottlenecks in registering new flows, updating, and I/O operations.

Conclusion

- INT data proved effective in detecting DDoS attacks for both known and unknown attack patterns, with an F1 score consistently above 99% across all models.
- sFlow performed similarly to INT for RF and KNN models but did not perform well for GNB and NN models. sFlow may miss data due to sampling, which results in missing some attack flows.
- Automated detection, addressing bottlenecks, can be achieved in under 2 seconds.
- Efficiently storing, processing, and analyzing INT data requires substantial computational resources and optimized techniques.

Future Work

- We want to implement automated DDoS detection on part of our production network, with some form of mitigation capabilities.
 - To do this, we need to handle production volume and speed (at least 2.5 million packets per second across 50,000 to 70,000 unique flow IDs).
 - Currently, we use C/C++ code for parsing and processing, and Kafka for storing and streaming data.
 - Our goal is to scale this solution and deploy it across the entire production network.
- At AmLight, we want to continue utilizing INT data to address network issues that benefit from fine-grained data, such as jitter, packet loss, microbursts, congestion, and detailed traffic analysis.

THANK YOU



Additional Slides

Simulating Attack Flows

We used *hping3* to simulate SYN and UDP scans, as well as SYN Flood attacks, and a Python script to run SlowLoris attacks.

- TCP SYN scan against the host <host name>, all ports aggressive scan
- UDP Scan against the host <host name>, common ports
- TCP SYN Flood against the host <host name>, port <port name> around 5000ps
- SlowLoris against the host <host name>, port <port name>

Experimental Results I

sFlow Predictions

```
Accuracy: 0.9937
Confusion Matrix:
[[16587   0]
 [  106   0]]
Recall: 0.0000
Precision: 0.0000
F1-score: 0.0000
auc: 0.5000
-----
Prediction Time: 0.2350 seconds
```

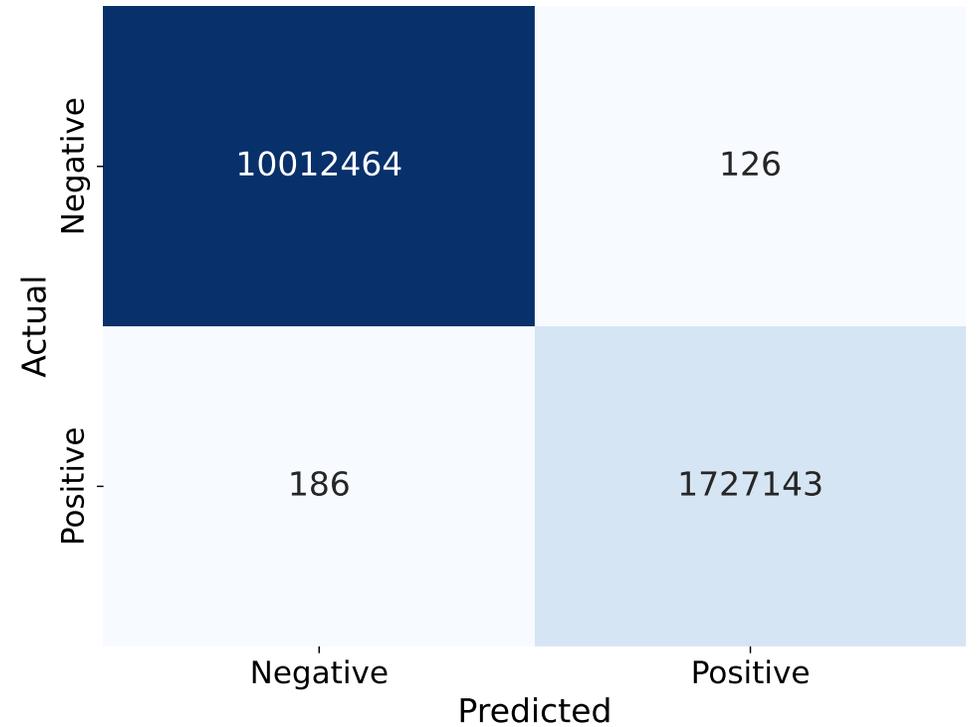
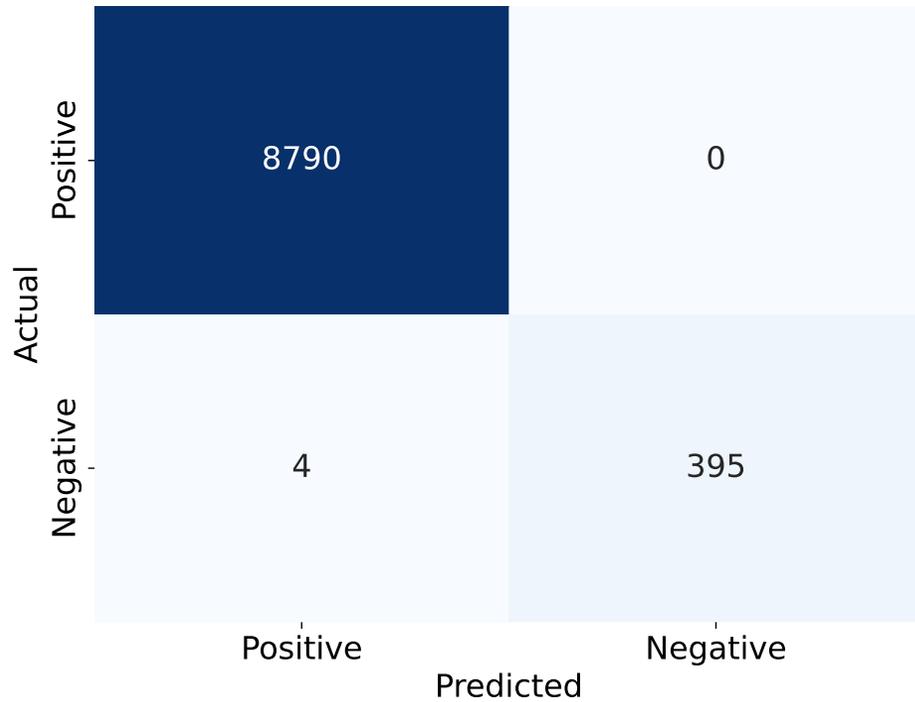
- Neural Network

```
Accuracy: 0.9979
Confusion Matrix:
[[16552   35]
 [   0  106]]
Recall: 1.0000
Precision: 0.7518
F1-score: 0.8583
auc: 0.9989
-----
Prediction Time: 0.0083 seconds
```

- Gaussian Naïve Bayes

Experimental Results I

Confusion Matrix



Experimental Results I

Top Five Most Important Features

- The most important features for detecting DDoS attacks are *Inter-Arrival Time, Packet Size, Queue Occupancy, and Protocol*.
- Variants of these features, such as individual values, cumulative statistics, averages, and standard deviations, and their ranking differ in importance across ML models.

Features	RF	GNB	KNN	NN
Inter Arrival Time _{cum}	✓	✓	-	✓
Inter Arrival Time _{std}	✓	-	✓	✓
Packet Size	-	✓	✓	-
Packet Size _{avg}	✓	✓	✓	✓
Packet Size _{std}	✓	-	-	-
Queue Occupancy _{avg}	✓	-	✓	✓
Queue Occupancy _{std}	-	✓	-	-
Protocol	-	✓	✓	✓

Feature Permutation

- The table shows the ranked feature importance for the RF model.
- Permutation importance:
 - For each feature, the library shuffles its values across all samples while keeping other features unchanged.
 - It then observes how the outcome changes.

	Feature	Importance
7	mPktSize	1.090649e-01
5	clat	1.205070e-02
12	sQueueOccupancy	1.369430e-03
10	sPktSize	1.314149e-04
11	slat	1.126754e-04
9	mQueueOccupancy	1.025390e-04
2	pkt_size	9.252193e-05
8	mlat	4.906337e-05
6	cQueueOccupancy	2.982985e-05
4	cPktSize	2.361175e-05
0	protocol	1.829655e-05
3	iat	1.028968e-05
1	queue_occ	-4.599691e-07
13	flowRate	-5.621845e-07
14	pktRate	-1.056225e-06

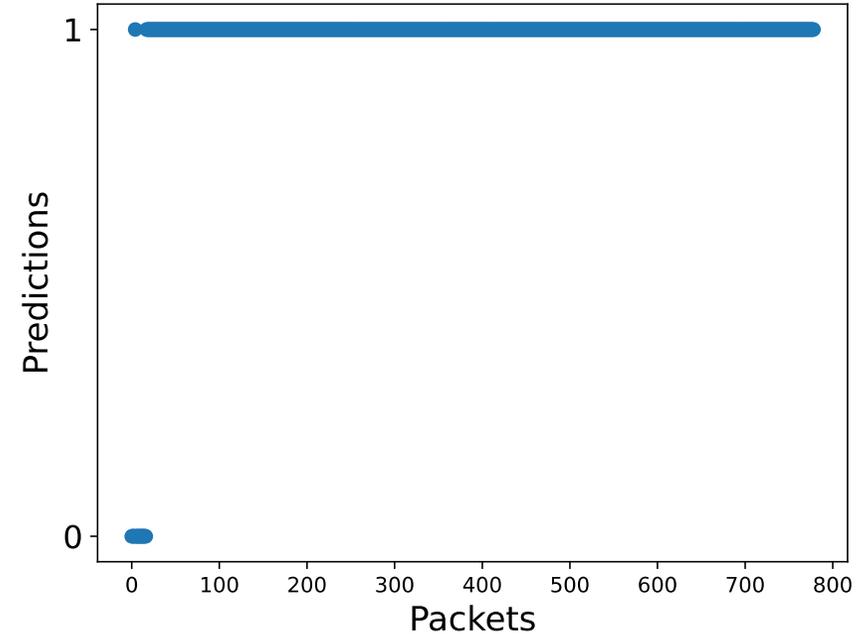
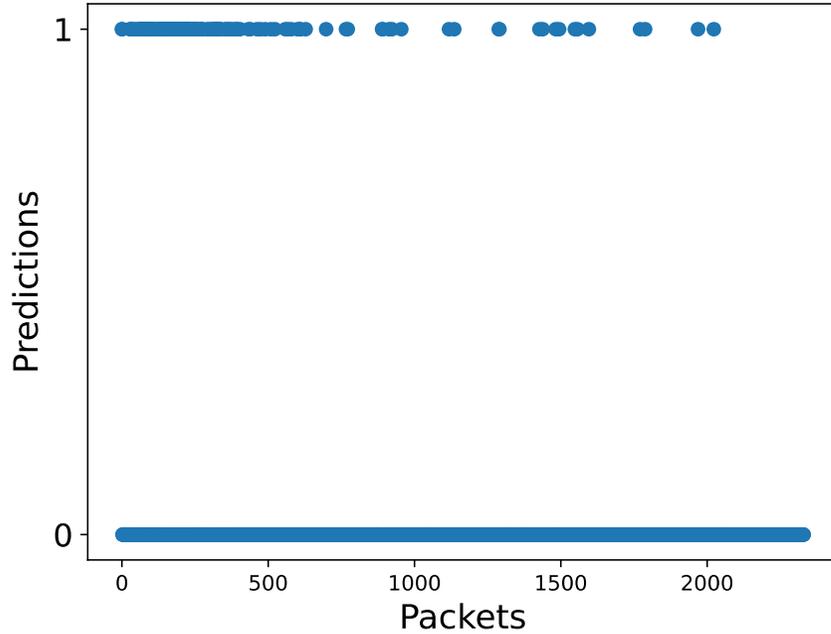
The INT Testbed

Hardware Configurations Overview

- The source and target servers powered by dual AMD EPYC 7451 24-core processors and 128GB of RAM. Each server utilizes a Mellanox ConnectX-5 network card capable of 100Gbps throughput.
- The switch is an Edgecore Wedge DCS800

Experimental Results II

A Closer Look at Predictions



- Misclassifications occur in the initial instances of a new flow.