Security .Days

<Prague.CZ>
<9-11 April 2024>

GÉANT

Exploring machine learning for DDoS mitigation

Jakub Man
CESNET

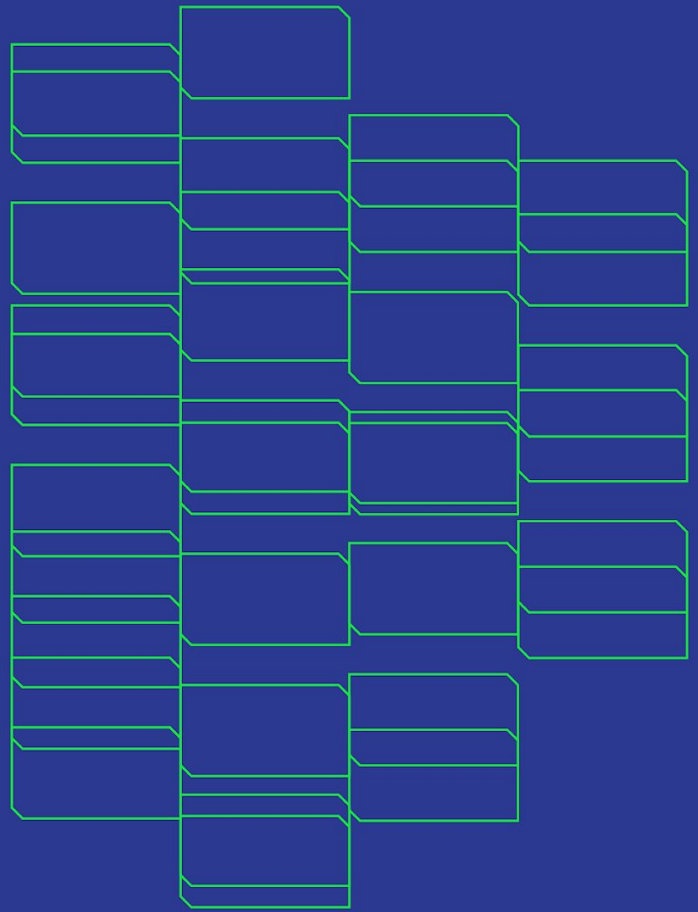

Co-funded by
the European Union

Requirements on ML

- Fast response to changing attacks
- Low false-positives
- Good performance

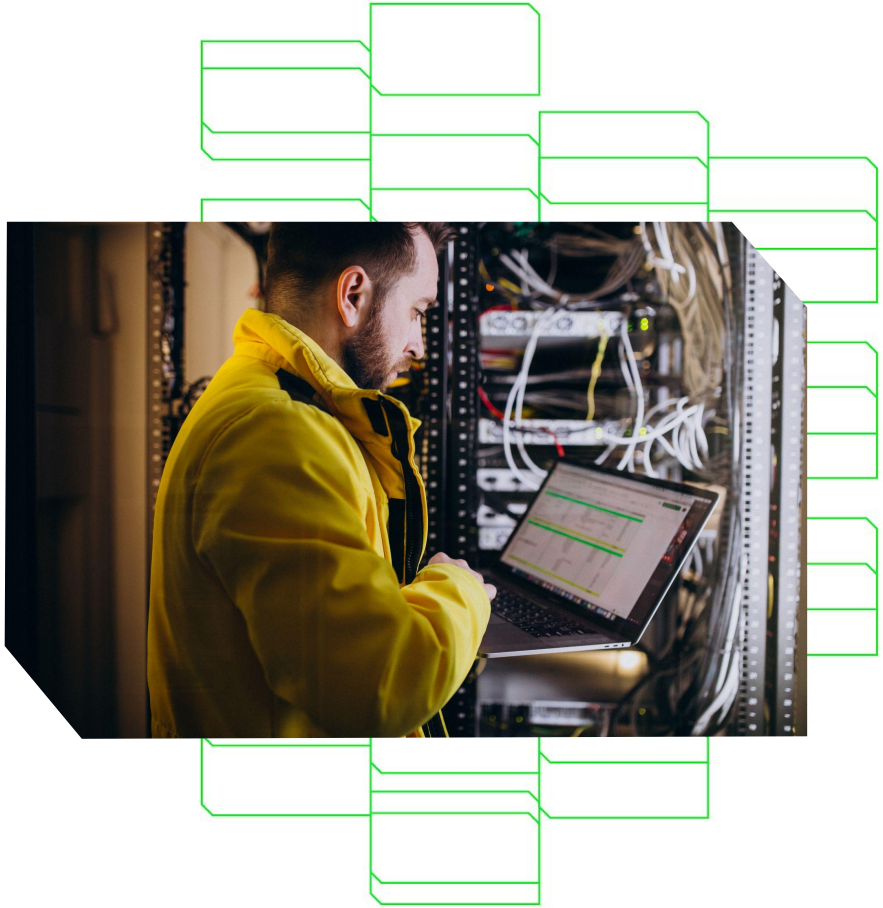

Image by senivpetro on Freepik

Autoencoder

- Artificial neural network
- Reconstructs it’s input on output (X = Y)
- Minimizes reconstruction error RMSE(X,Y)
- Trained in unsupervised manner (only legitimate traffic)

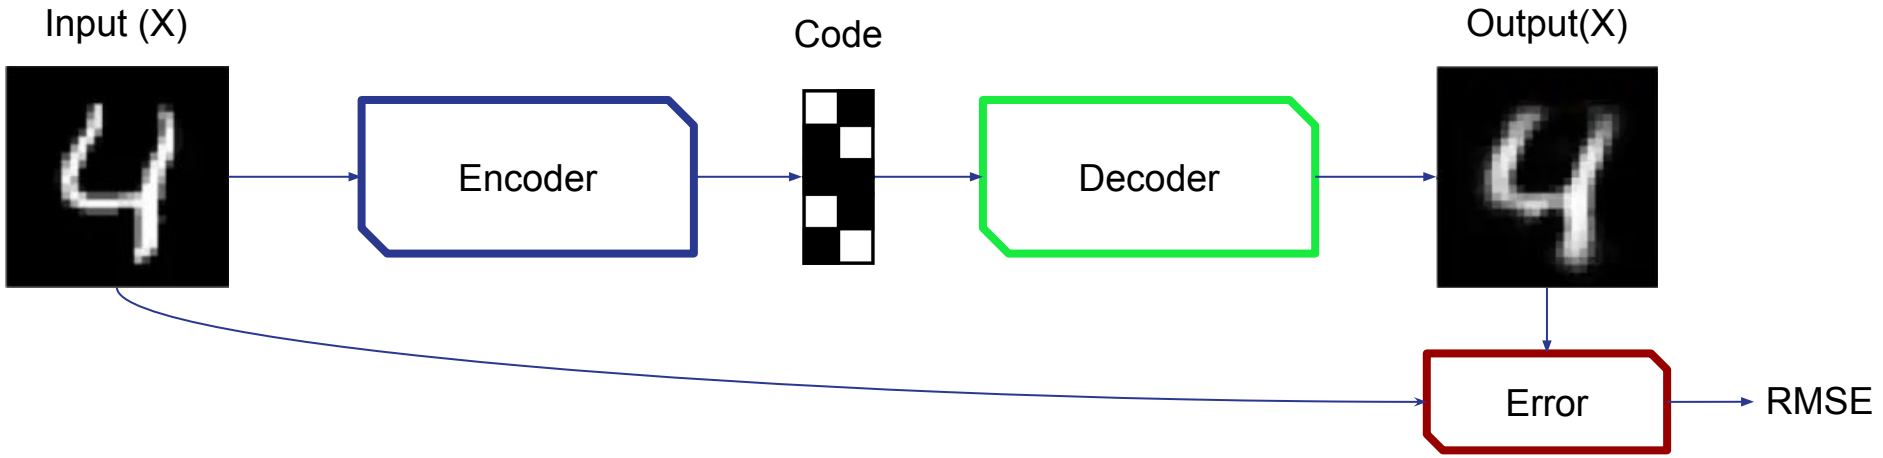

Autoencoder - anomaly detection in packets

Benign packets

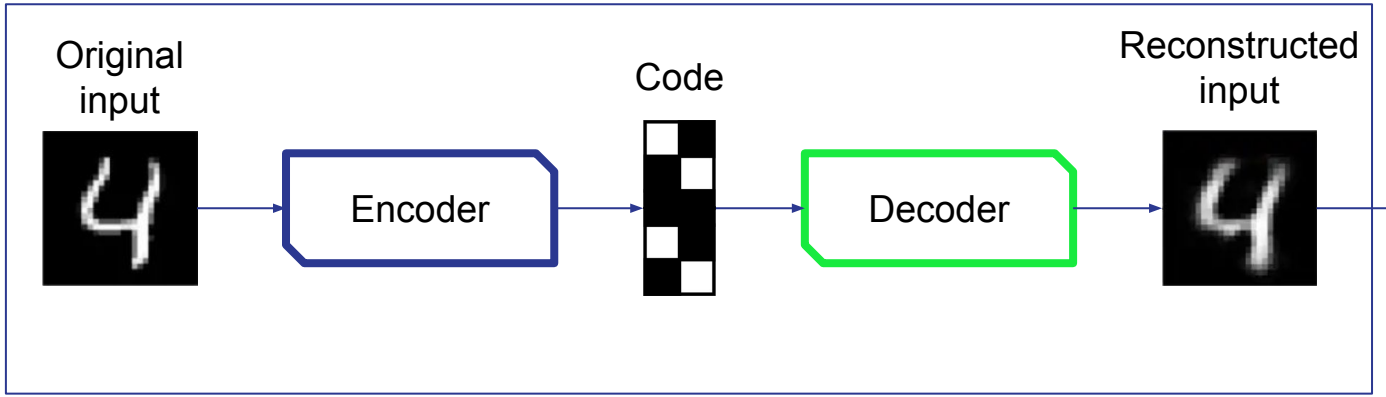

Small reconstruction error

Malicious packets

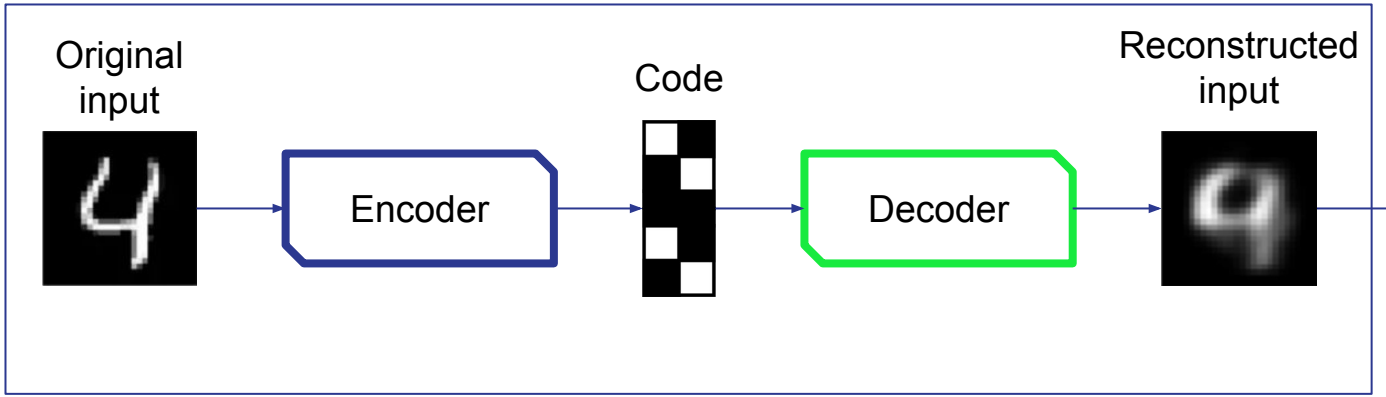

Large reconstruction error

Experimenting with autoencoders

- Uses sFlow (sampled packets)
- Aggregates “windows” of traffic by source IPs
- Features typical for attacks are extracted from the windows
  - Timestamp, L4 ports, payload length, …
- Result is a list of malicious IPs

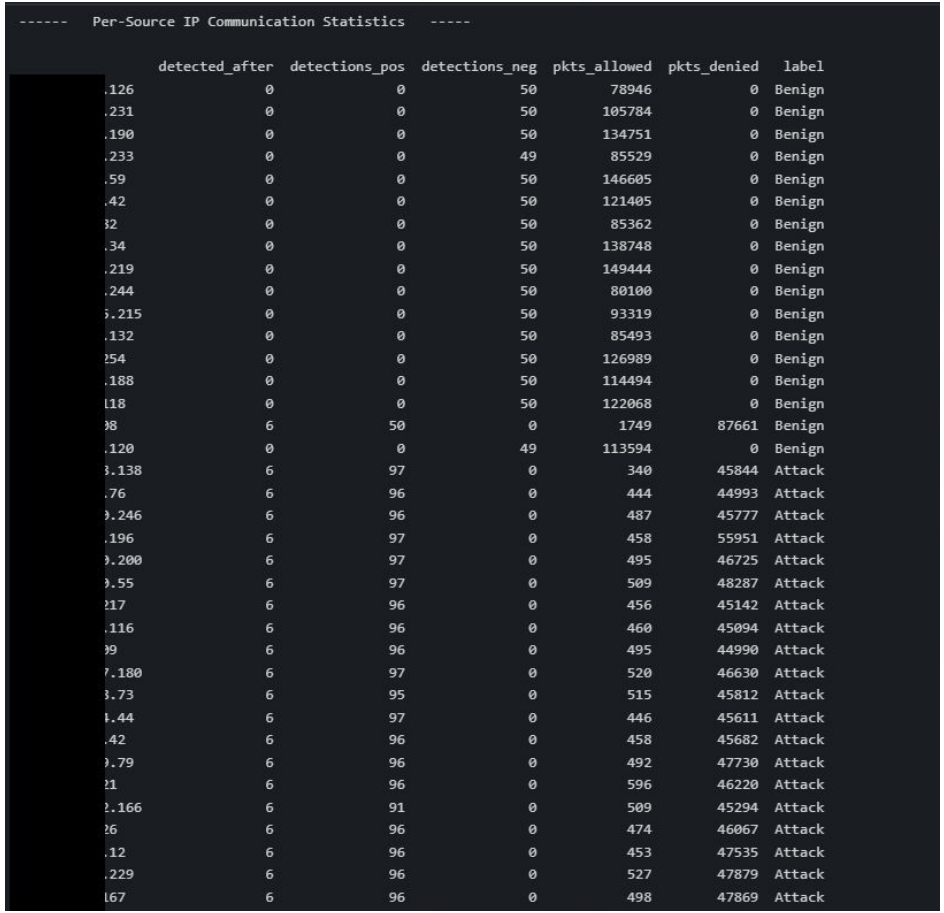

------ Per-Source IP Communication Statistics -----

	detected_after	detections_pos	detections_neg	pkts_allowed	pkts_denied	label
.126	0	0	50	78946	0	Benign
.231	0	0	50	105784	0	Benign
.190	0	0	50	134751	0	Benign
.233	0	0	49	85529	0	Benign
.59	0	0	50	146605	0	Benign
.42	0	0	50	121405	0	Benign
32	0	0	50	85362	0	Benign
.34	0	0	50	138748	0	Benign
.219	0	0	50	149444	0	Benign
.244	0	0	50	80100	0	Benign
.215	0	0	50	93319	0	Benign
.132	0	0	50	85493	0	Benign
54	0	0	50	126989	0	Benign
.188	0	0	50	114494	0	Benign
118	0	0	50	122068	0	Benign
98	6	50	0	1749	87661	Benign
.120	0	0	49	113594	0	Benign
.138	6	97	0	340	45844	Attack
.76	6	96	0	444	44993	Attack
.246	6	96	0	487	45777	Attack
.196	6	97	0	458	55951	Attack
.200	6	97	0	495	46725	Attack
.55	6	97	0	509	48287	Attack
217	6	96	0	456	45142	Attack
.116	6	96	0	460	45094	Attack
99	6	96	0	495	44990	Attack
.180	6	97	0	520	46630	Attack
.73	6	95	0	515	45812	Attack
.44	6	97	0	446	45611	Attack
.42	6	96	0	458	45682	Attack
.79	6	96	0	492	47730	Attack
21	6	96	0	596	46220	Attack
.166	6	91	0	509	45294	Attack
26	6	96	0	474	46067	Attack
.12	6	96	0	453	47535	Attack
.229	6	96	0	527	47879	Attack
167	6	96	0	498	47869	Attack

Autoencoder test results

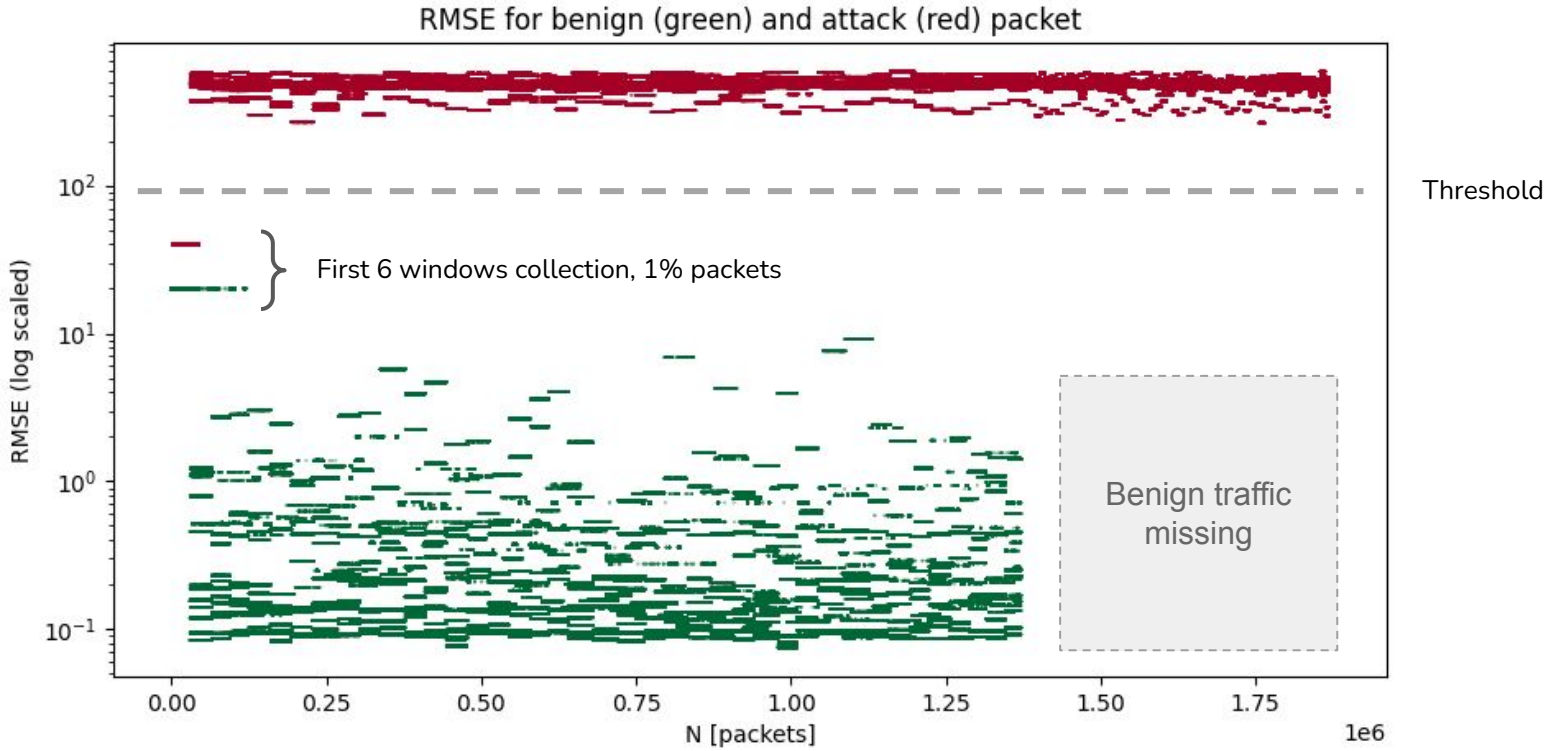

Threshold: 100; Recall: 0.99; Precision: 1.00; Accuracy: 0.99

How do we help networking professionals to configure this correctly?

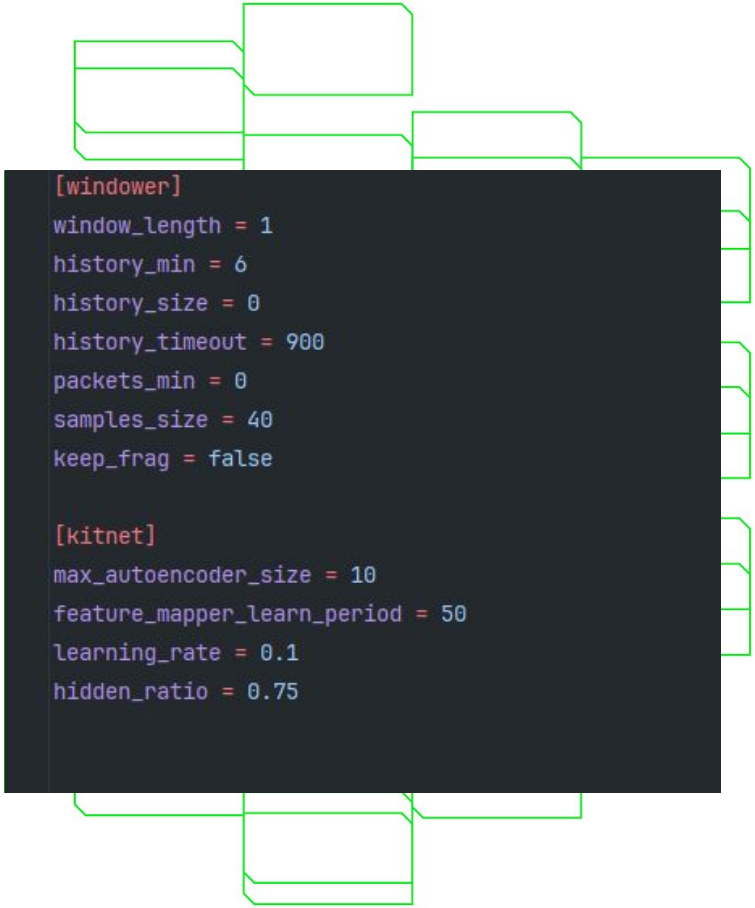

Configuation

Administrator

Configuation

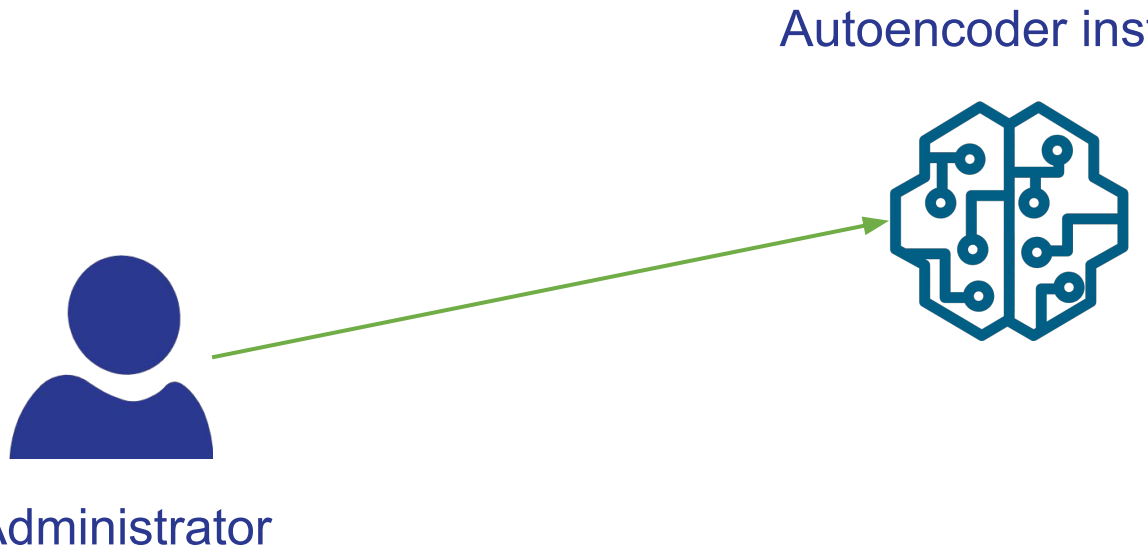


```
autoencoder_instance:
  detection_threshold: 50
  name: "Primary instance"
  ip_set_id: 1
```

Configuation

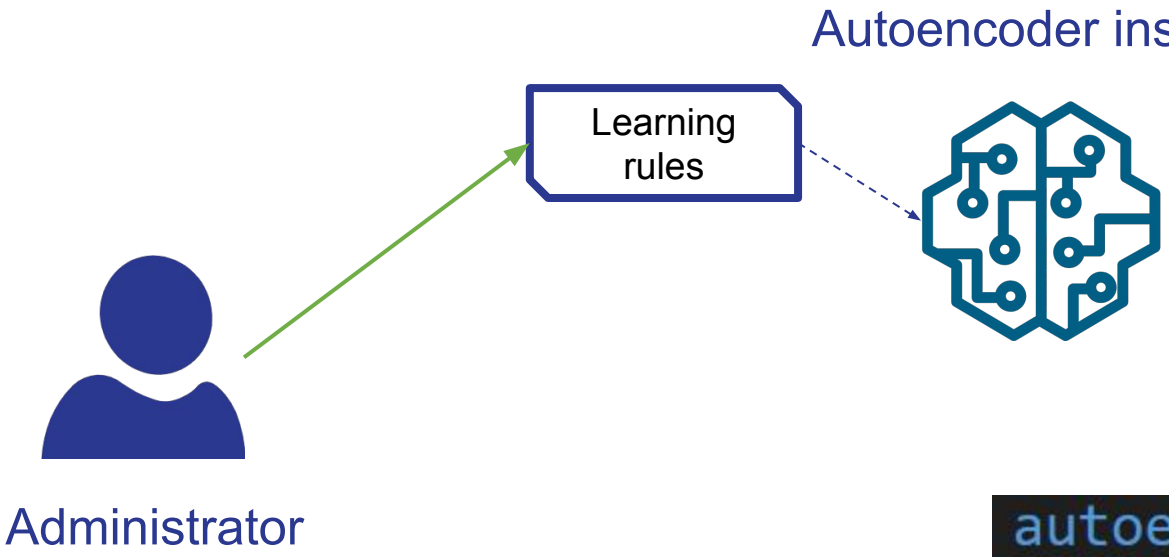


```
autoencoder:
  mode: learn
  ip_src: ['192.168.0.0/16']
  vlan: 100
  instance: "Primary instance"
```

Configuation

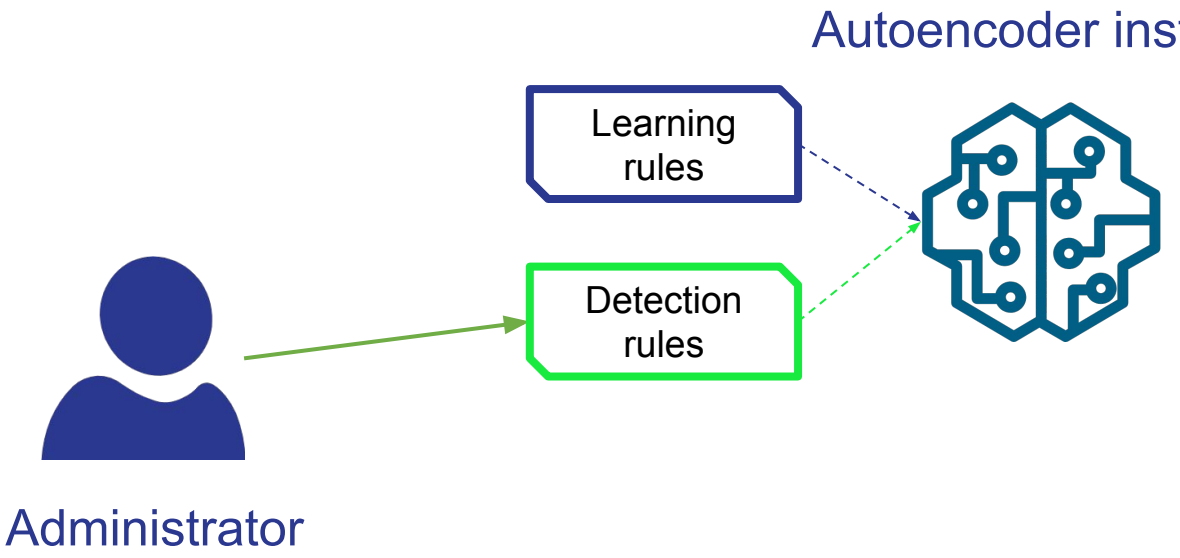


```
autoencoder:
  mode: detect
  ip_src: ['10.50.0.0/16']
  instance: "Primary instance"
```

Configuation

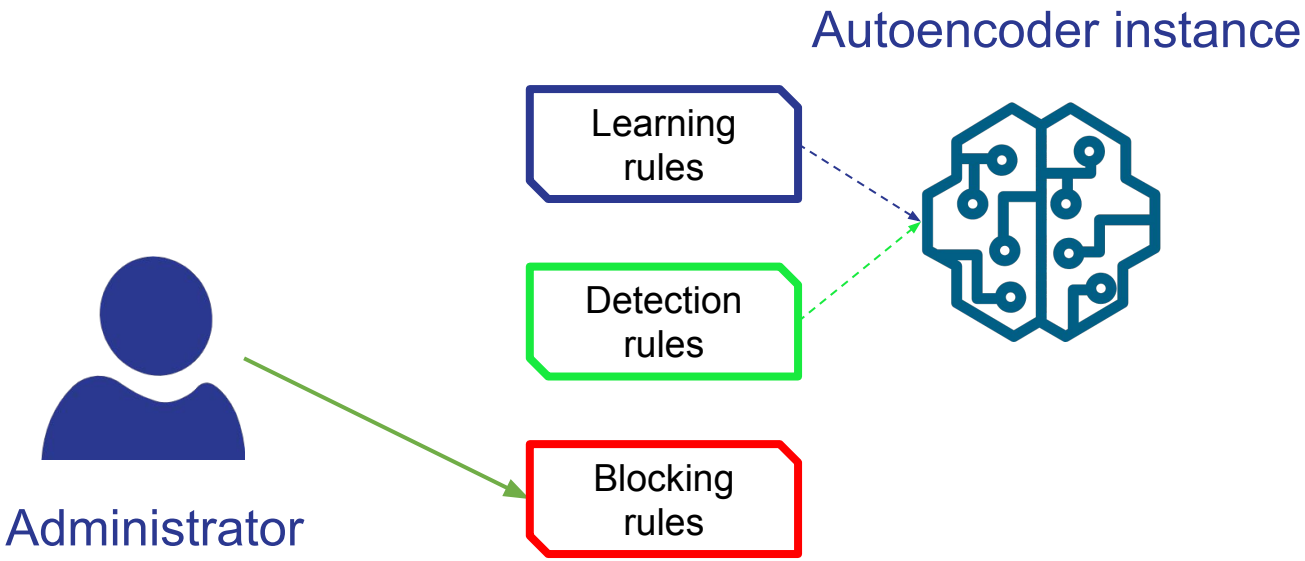


```
exact_match:
  ip_set_id: 1
  ip_dst: ['10.60.0.0/16']
```

Configuation

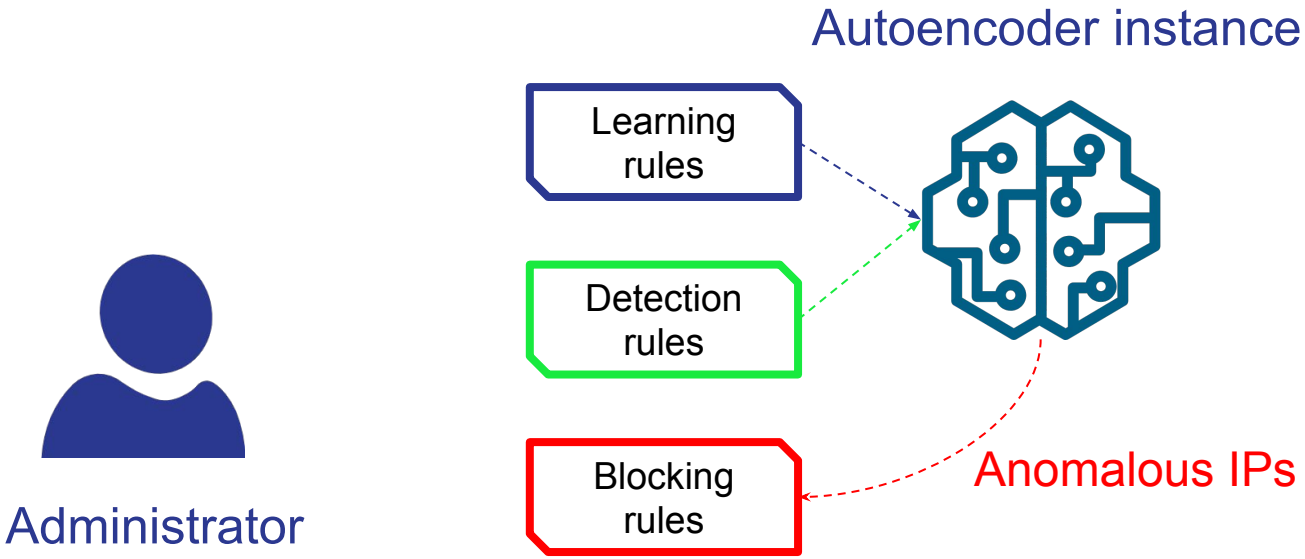

Integration into our system

Our DDoS Protector

All traffic

Autoencoder instance

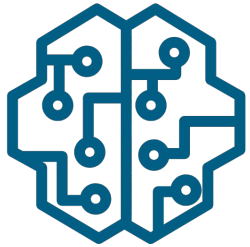

Integration into our system

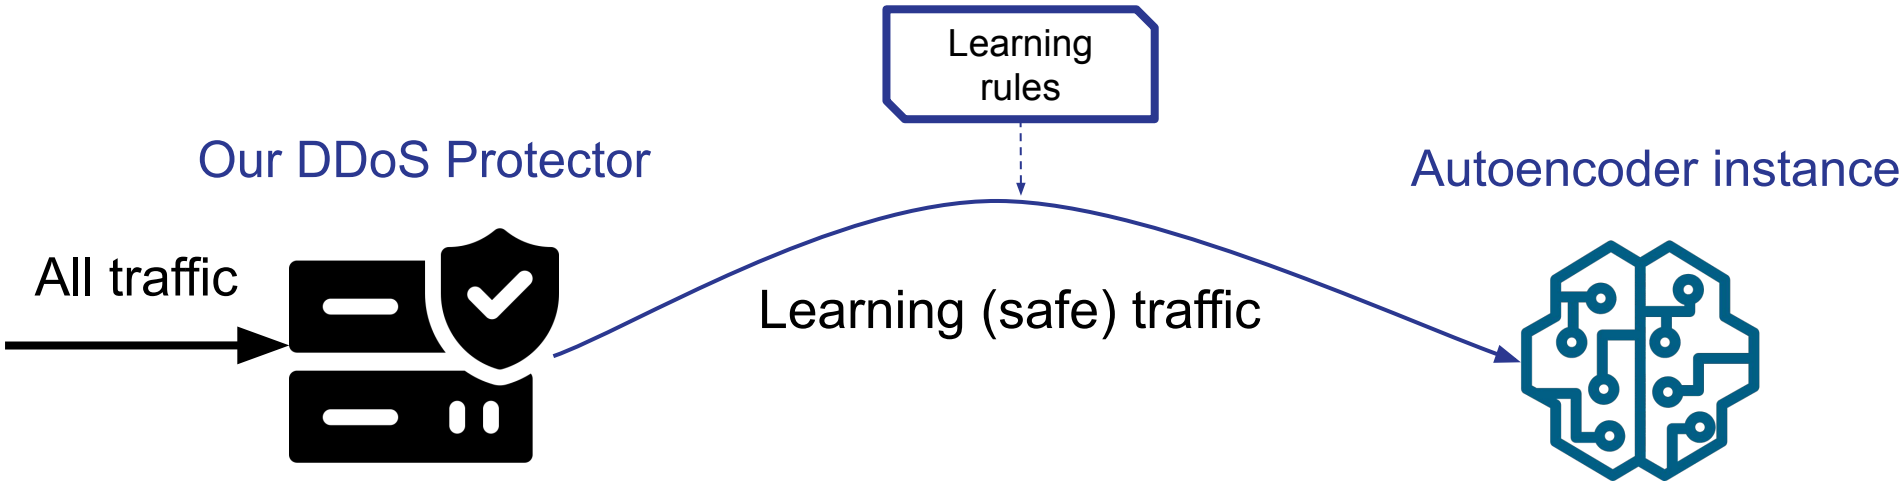

Integration into our system

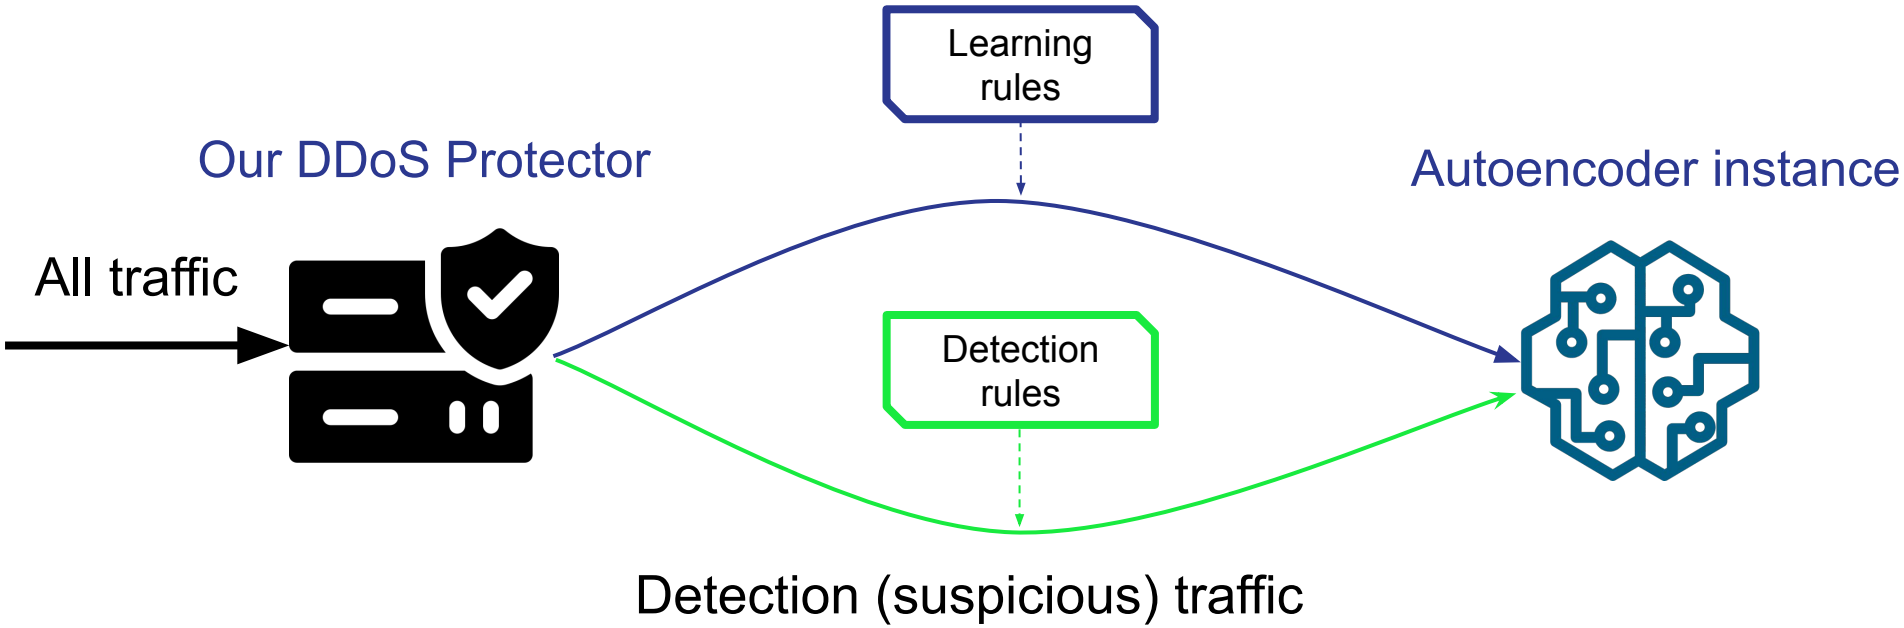

Integration into our system

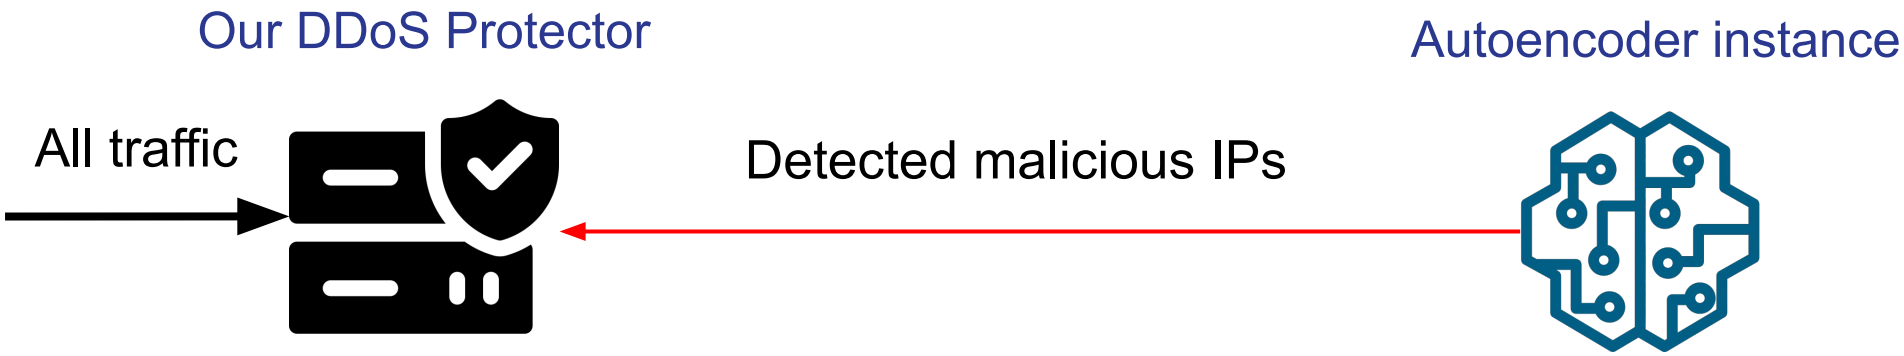

Integration into our system

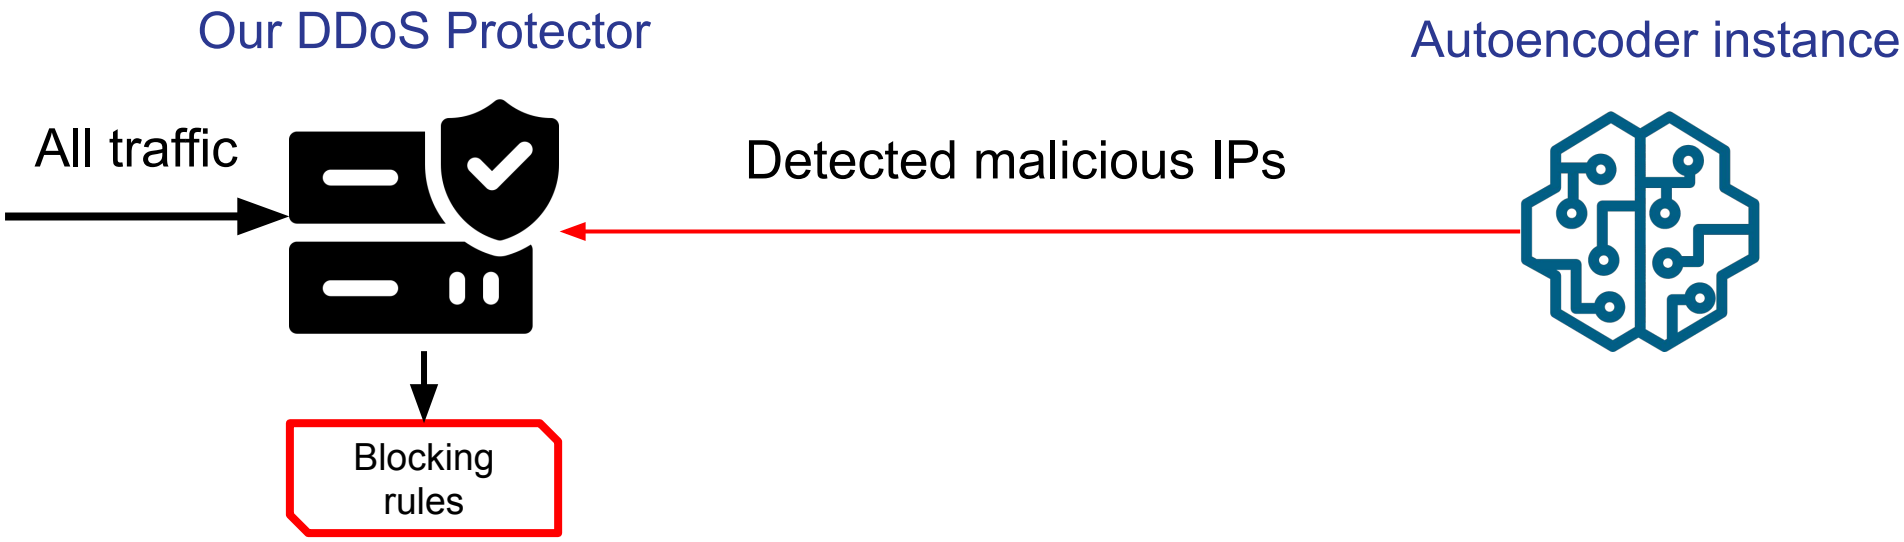

Integration into our system

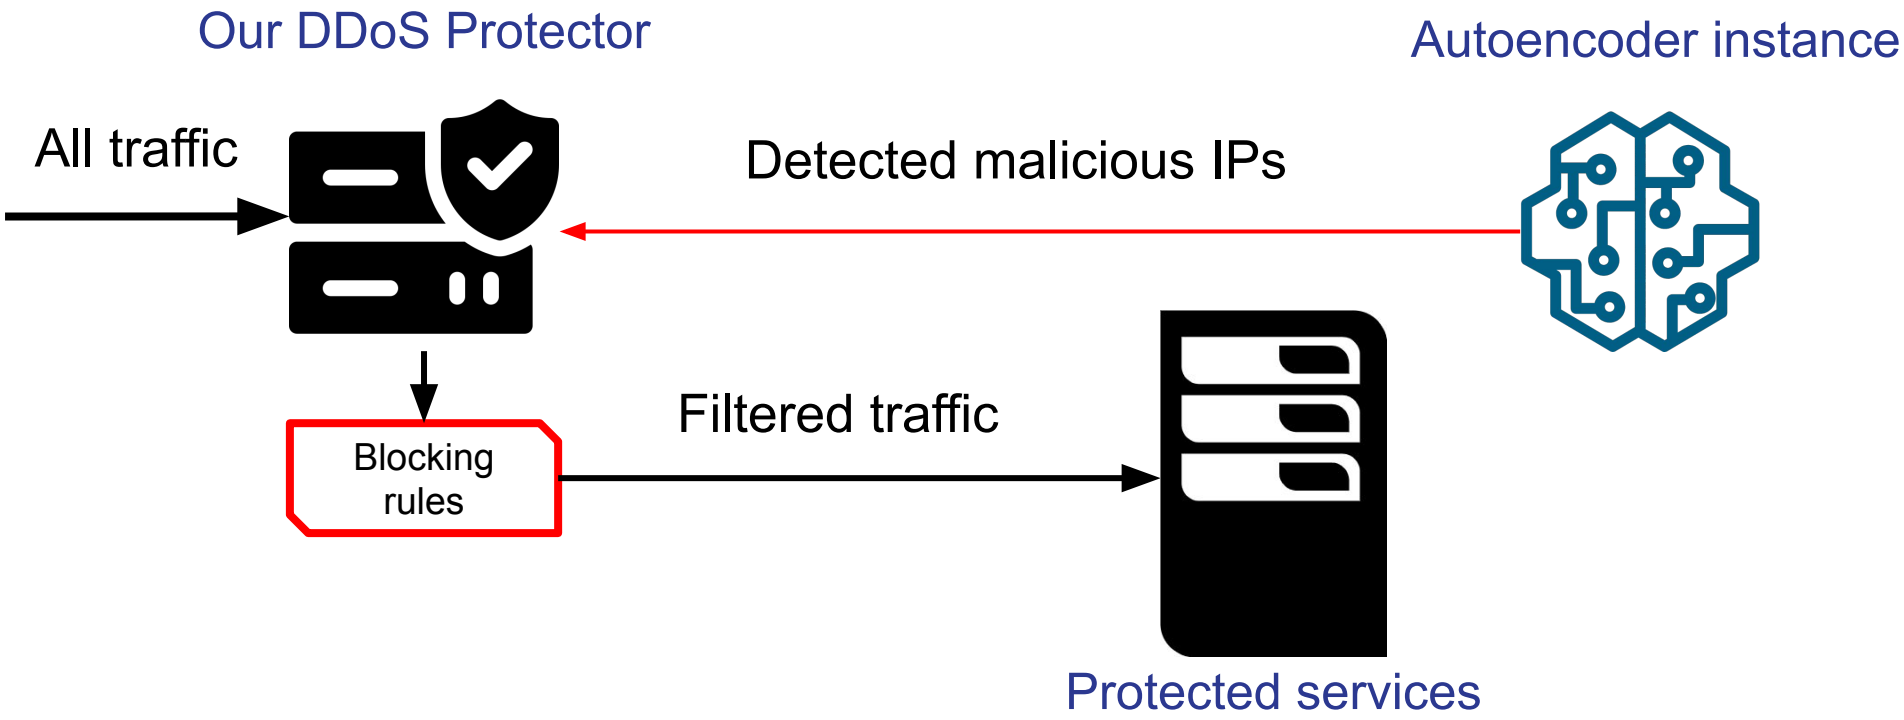

Experimental integration

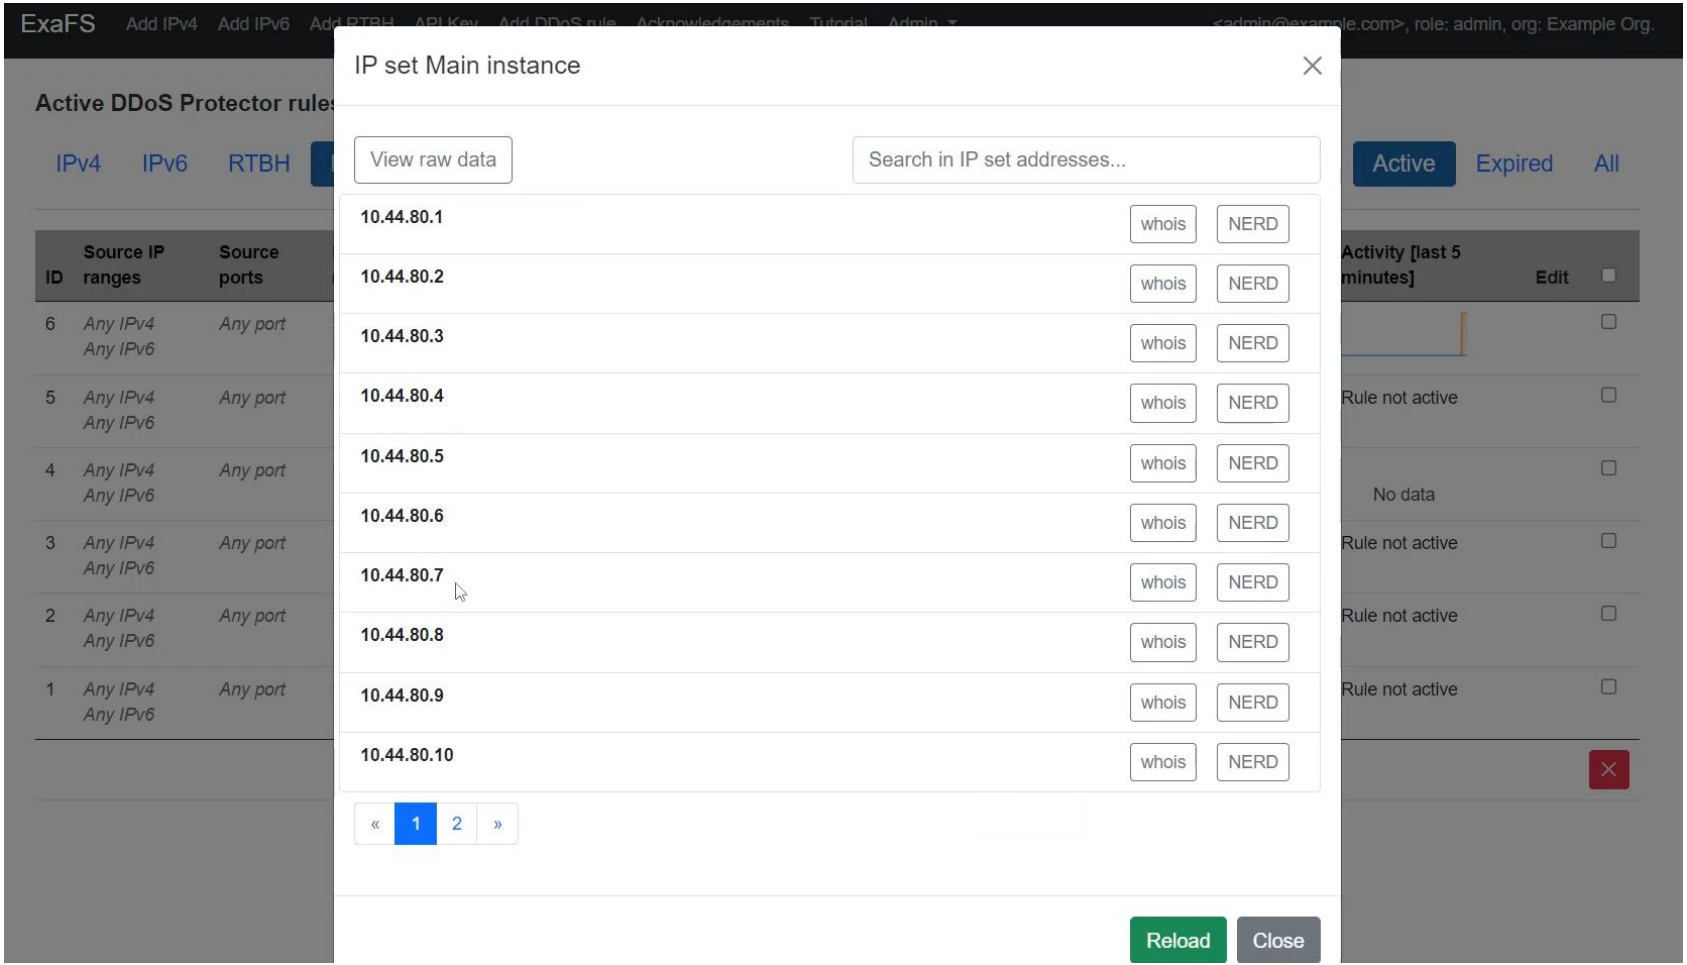

Conclusion

- Traditional methods are still needed as a pre-filter
- More complex attacks can be handled by machine learning
- Performance needs to be improved for production

Thank you

Any questions?

Jakub.Man@cesnet.cz
ddp@cesnet.cz

Co-funded by
the European Union