Iowa State University
College of Engineering

ADSICS

Anomaly Detection System for Industrial Control Systems

SDMay22-38 : Alex Nicolellis, Muhamed Stilic, Pallavi Santhosh, Jung Ho Suh

Client : Dr. Manimaran Govindarasu

Advisors : Dr. Manimaran Govindarasu, Moataz Abdelkhalek

Problem Statement

- **Cyber attacks on power distribution companies** are debilitating and are now more common due to the **increased use of IoT devices** and the **inadequate security** on power grid systems

- To feel as secure in how our power is protected, our security measures must continue to advance with the constantly developing technology of cybercriminals.

Cybersecurity Threats

- **Domestic terrorists** have developed "credible, specific plans" to **attack the U.S. power grid** and view it as a "particularly attractive target given its interdependency with other infrastructure sectors," according to a security briefing **issued in January by Department of Homeland Security**

- **Past Attacks Globally:**
  - The European Network of Transmission System Operators for Electricity (2020)
  - Russian power grid (2019)
  - Saudi Aramco petrochemical plants (2017)

Solution Statement

ADSICS is a surveillance program that detects cyber attacks u**sing machine learning analysis.**

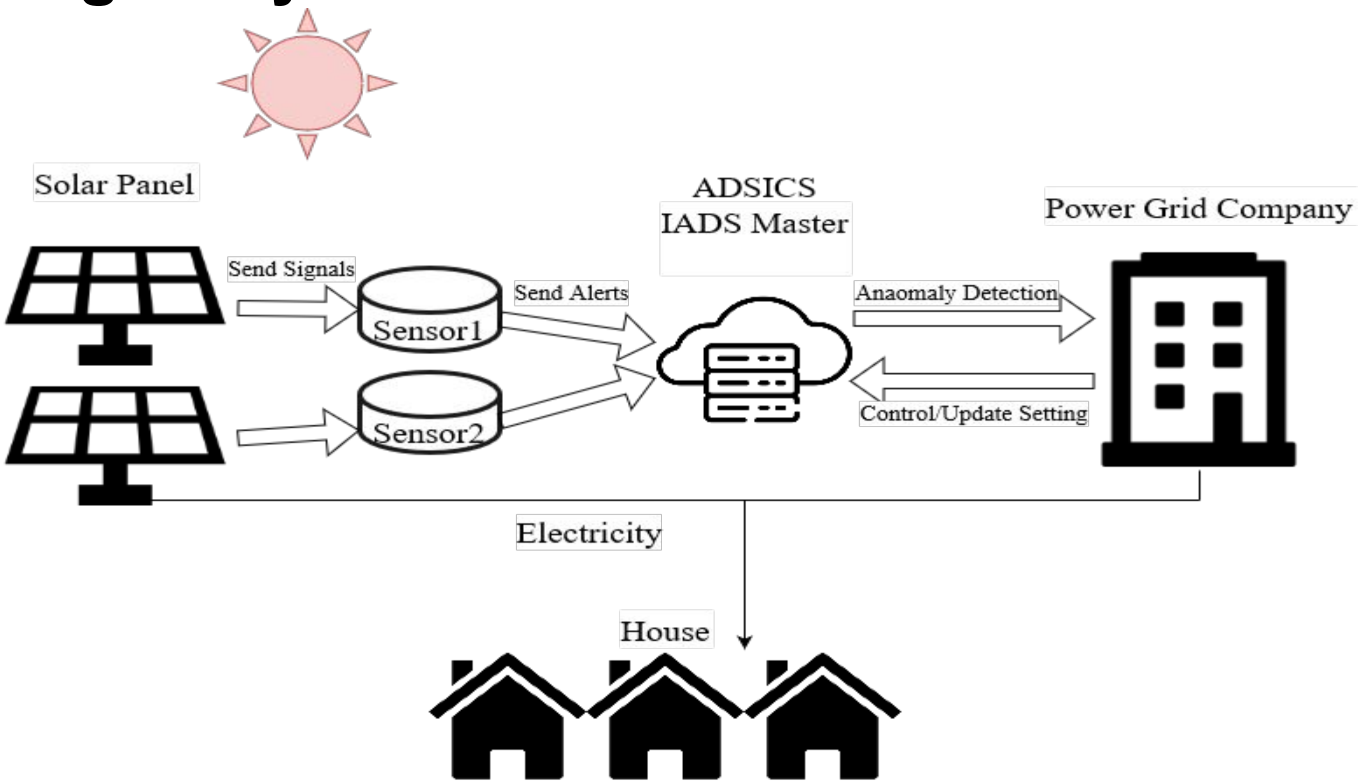

Solution Statement

- Traditional security measures easily disabled by user error; spear phishing
- Machine Learning Anomaly Detection System is needed
- Building our own ML ADS will cost a lot, use existing platform

elasticsearch

Elastics Stack

- Elastic Stack is an open source tool
- Modification to features is like a black box
- Three Projects that make the Elastic Stack
  - Elasticsearch - Search and Analytics Engine
  - Logstash(Beats) - Data ingestion Pipeline
  - Kibana - Data Dashboard and Visualizer

Project Description

- The end goal is described as **a detailed visual output** of temporal anomaly detection, tracking alert information, and performing **machine learning analysis** quickly and efficiently **using the desired platforms (Elastic Stack)**.

- Additionally, we built our own, **local machine learning model** to compare with the outputs from Elastic to determine whether or not this relatively “new”platform could be **relied upon solely** as a valid solution to **detect and analyze** various cyberattacks on the grid.

Requirements & Constraints

Functional Requirements:

- **Use machine learning to detect network anomalies**
- Verify incoming alerts and detect false positives
- Display alerts for easy human understanding
- Present temporal and spatial details for each alert
- Users should be able to add or remove data visualizations on the dashboard
- Alerts should be distinguishable from each other and labled

Non-Functional Requirements:

- **Alerts should be presented intuitively**
- **Alerts should be color coded by severity**
- The system should be able to handle a large volume of alerts
- The system should be reliable and consistent with analysis
- System should have an **accuracy of >95%** for detecting each alert

Constraints:

- One branch of the project must use **Elastic Stack tools** (specifically Elasticsearch, Kibana, and Logstash)

Datasets

- The main datasets we used for different types of attacks such as **DoS, Brute Force, and XSS** from https://www.unb.ca/cic/datasets/ids-2017.html (University of New Brunswick 2017).
  - Chosen for: variety of attacks and the length of an entire workday with network traffic
- Modbus dataset which includes **Modbus query flood, ping flood, tcp SYN flood attacks** from https://github.com/tjcruz-dei/ICS_PCAPS/releases\ (Frazão, I.)
  - Chosen for: OT protocol (common for power grid)

Datasets

- We used the **PCAP files to replay the network flow** along with attacks inside it and used **CSV files to classify attacks** with our VM ML model.

- The CSV datasets are each compiled of a data flow in the PCAP which was grouped together by CICFlowMeter.
  - These data flows consist of **78 features** ranging from Destination Port to Flow ID

- By analyzing the alert data of the PCAP files, a “Label” column was created and input with the various **“attack” types associated with each data flow**.

Datasets: Wednesday

Wednesday_Workingday dataset is the largest dataset with **13 gigabytes** and has four classified DoS attacks: **DoS slowloris, DoS Slowhttptest, DoS Hulk, DoS GoldenEye**.

Wednesday_Workingday			
DoS / DDoS			
Network Traffic from 9:00 A.M. to 5:00 P.M.			
	OS	**IP**	**Local IP**
Attacker	Kali	205.174.165.73	
Victim DDoS	WebServer Ubuntu	205.174.165.68	192.168.10.50
Victim Heartbleed	Ubuntu12	205.174.165.66	192.168.10.51
Type of Attack		**Start**	**End**
DoS Slowloris		9:47	10:10
DoS Slowhttptest		10:14	10:35
DoS Hulk		10:43	11:00
DoS GoldenEye		11:10	11:23

Elastic: Theoretical Design

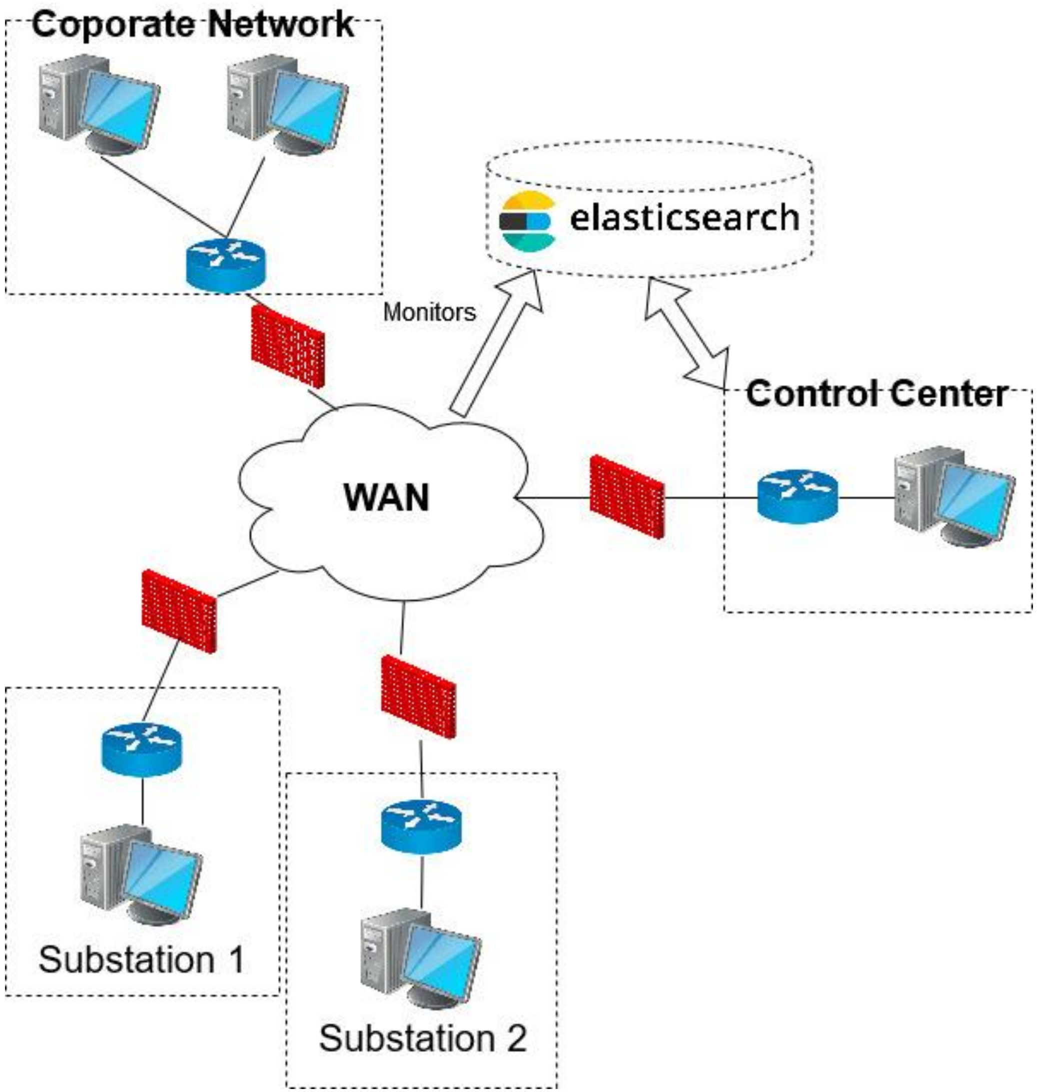

Elastic: Practical Design

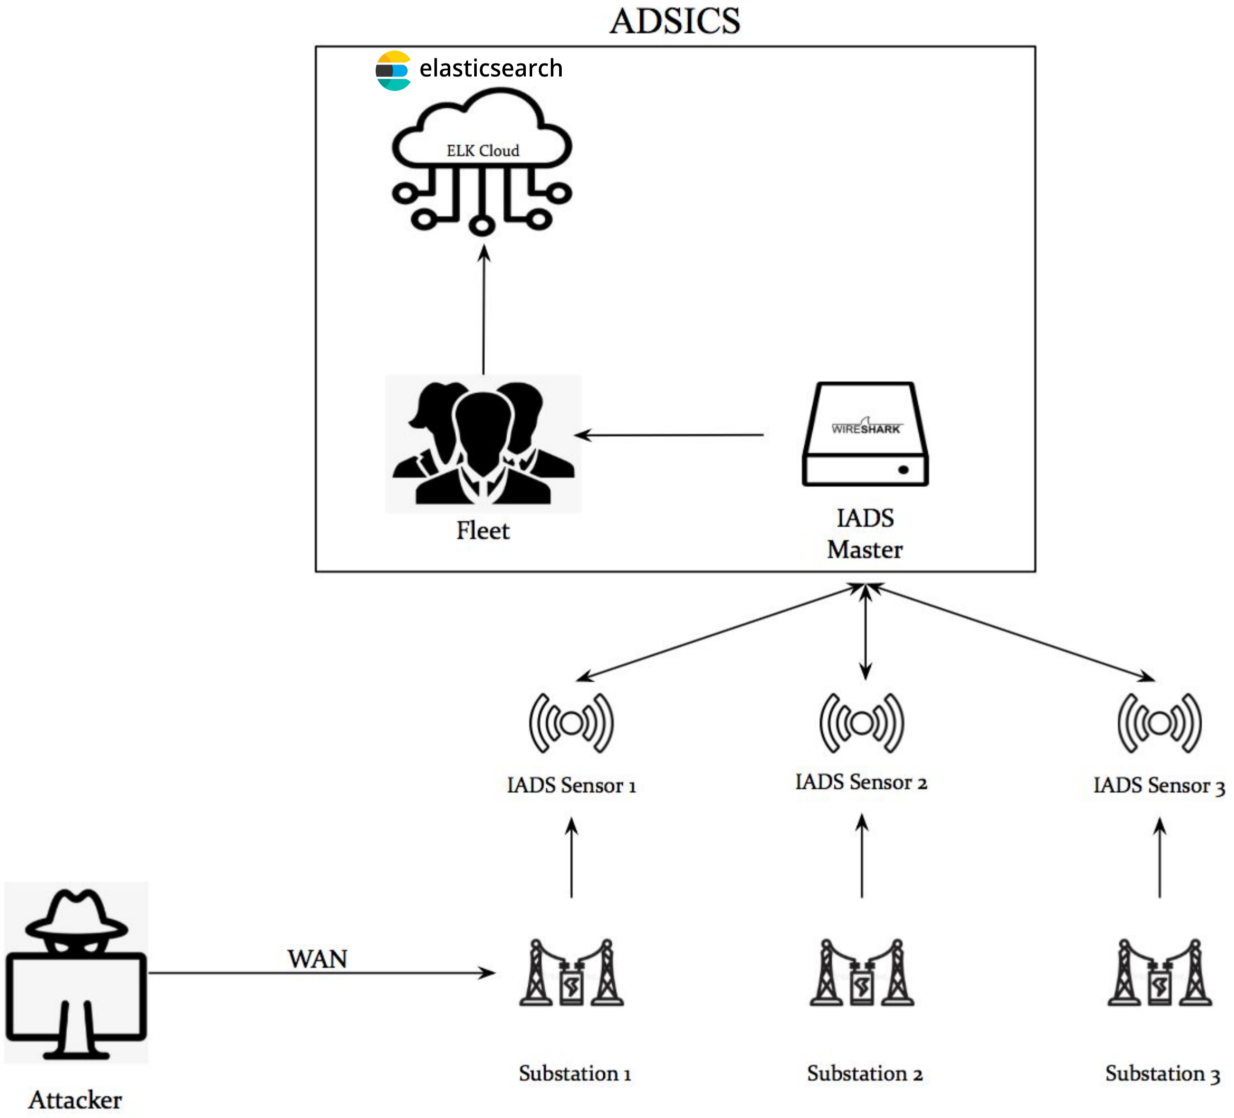

Elastic: System Architecture

Functional Decomposition - Operating Systems:

- Attacker - Kali Linux
- ELK_Agent, IADS_Sensor, IADS_MASTER - Linux
- Victim - Windows XP

Software Architecture:

- vSphere - VM Platform
- SIEM
  - Elastic Stack
    - Elastic Search
    - Logstash(Beats)
    - Kibana
- Tools
  - TCPReplay
  - Wireshark
  - Editcap

Elastic: Integrations

- Integrations
  - Fleet
    - Agents
    - Packet beat(Data ingestion)
  - Network Packet Capture
  - Security Integrations(Jobs)

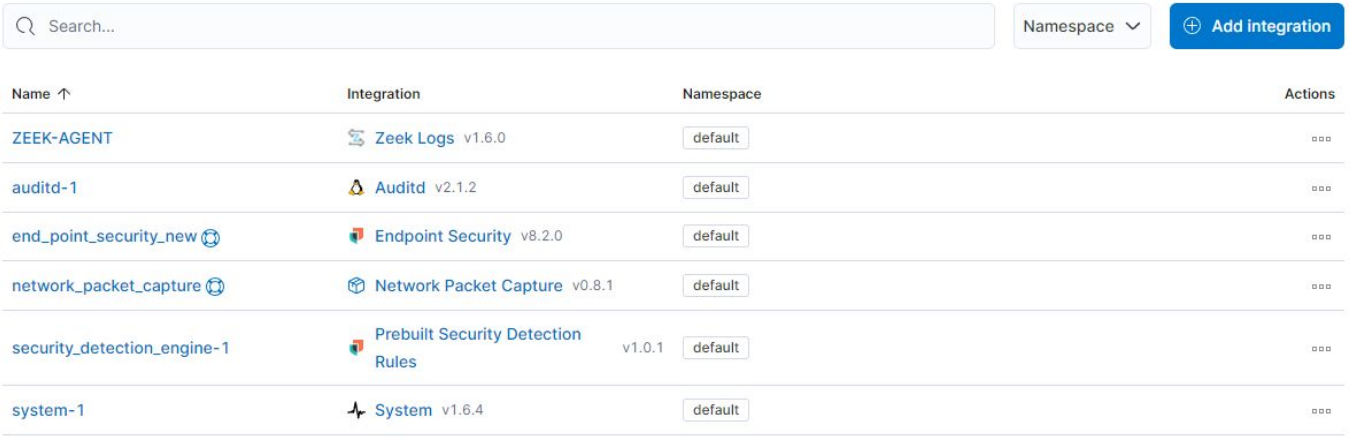

Name ↑	Integration	Namespace	Actions
ZEEK-AGENT	Zeek Logs v1.6.0	default	
auditd-1	Auditd v2.1.2	default	
end_point_security_new	Endpoint Security v8.2.0	default	
network_packet_capture	Network Packet Capture v0.8.1	default	
security_detection_engine-1	Prebuilt Security Detection Rules v1.0.1	default	
system-1	System v1.6.4	default	

Elastic: Kibana

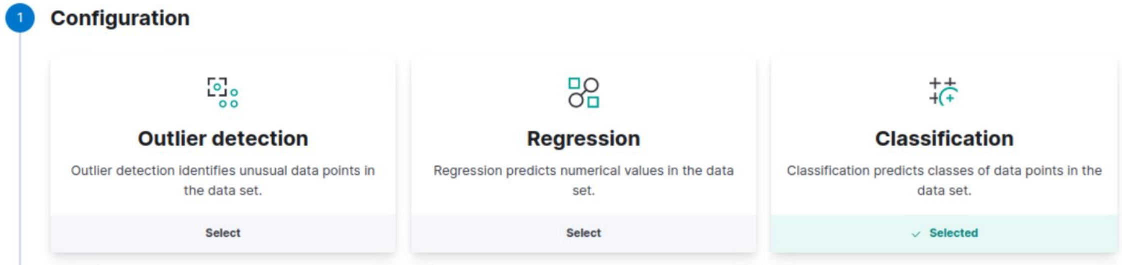

- Machine Learning - Anomaly Detection
  - High Count Jobs
  - Rare Processes
- Machine Learning - Data Frame Analytics
  - CSV Jobs

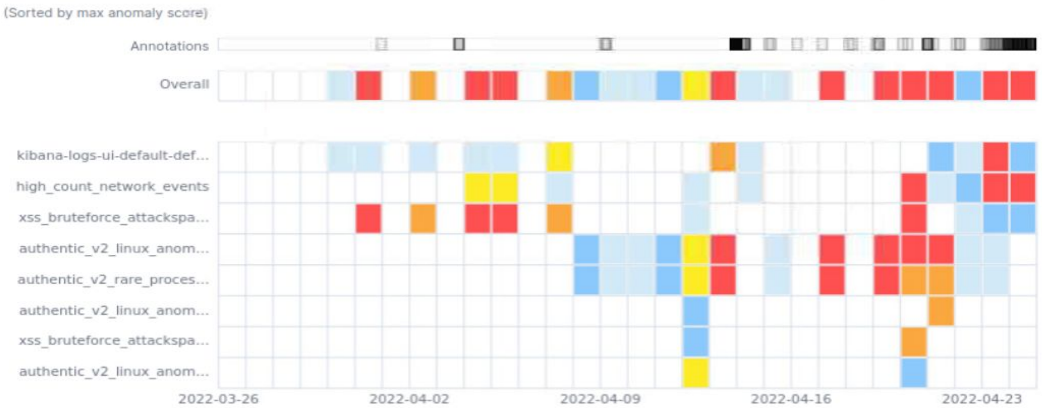

ID ↑	Description	Memory status	Source index
benign_then_attac	test	ok	test_benign_th...
heart_disease		ok	heart_2020
mitm_15h_6h_test	test	ok	mitm_15m_6h
modbusqueryflooc	test	ok	modbusqueryfl...
no_header_thursd.	test	ok	no_headers_th...
pingflooddos_15	test	ok	pingflooddos...
tcpsynflooddos_:	test	ok	tcpsynfloodd...
test_attack_then_t	test	ok	test_attack_be...
test_dos_first_try	test	ok	test_dos_first
test_dos_last_try	test	ok	test_dos_last

Elastic/Network Issues

- No modbus port integrations
  - Cannot edit Network Packet Capture
- Near real time analytics
  - Data ingested in real time but analyzed in near-real time
- TCPReplay
  - MTU Size

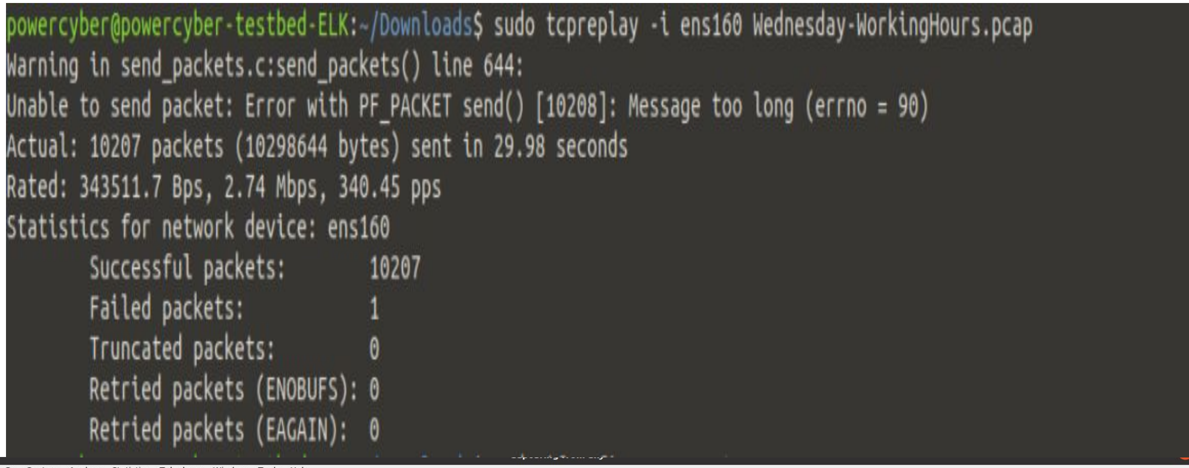

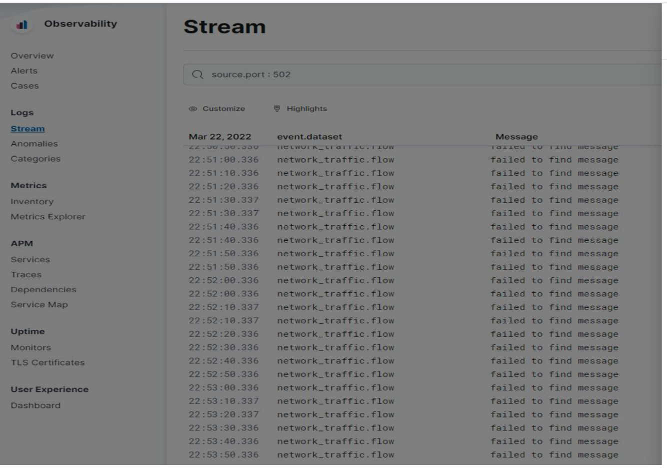

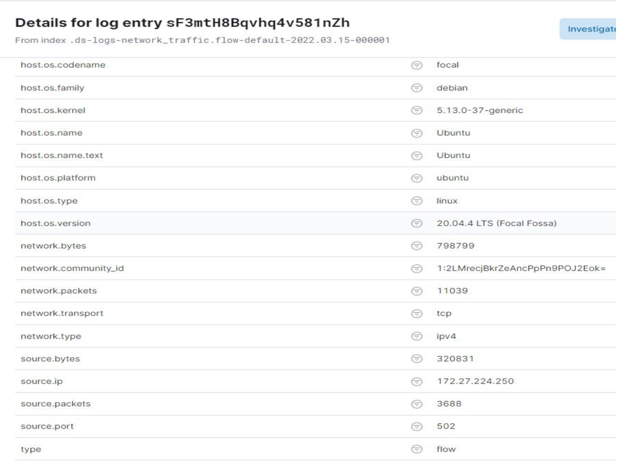

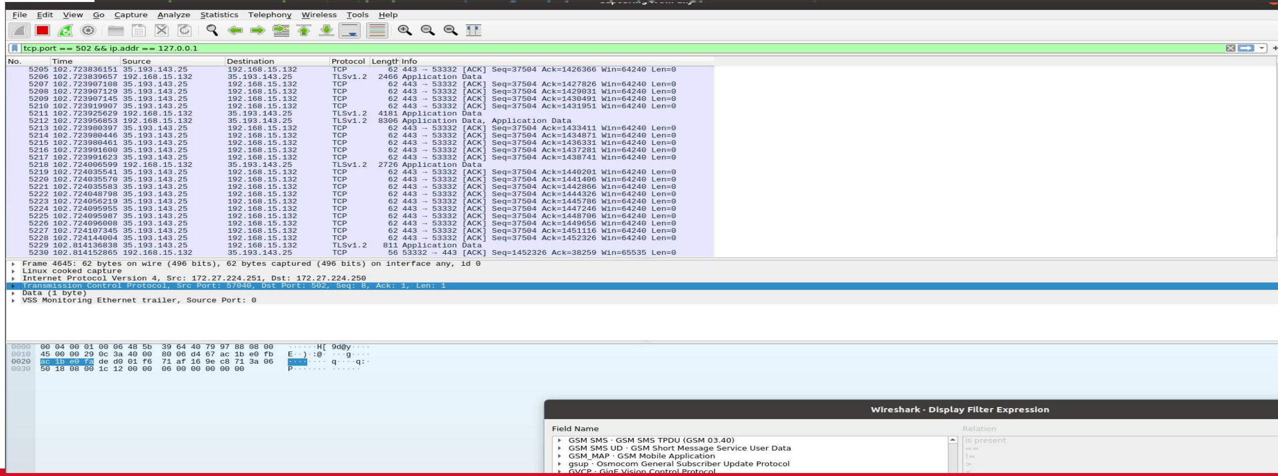

Elastic: Dashboard UI

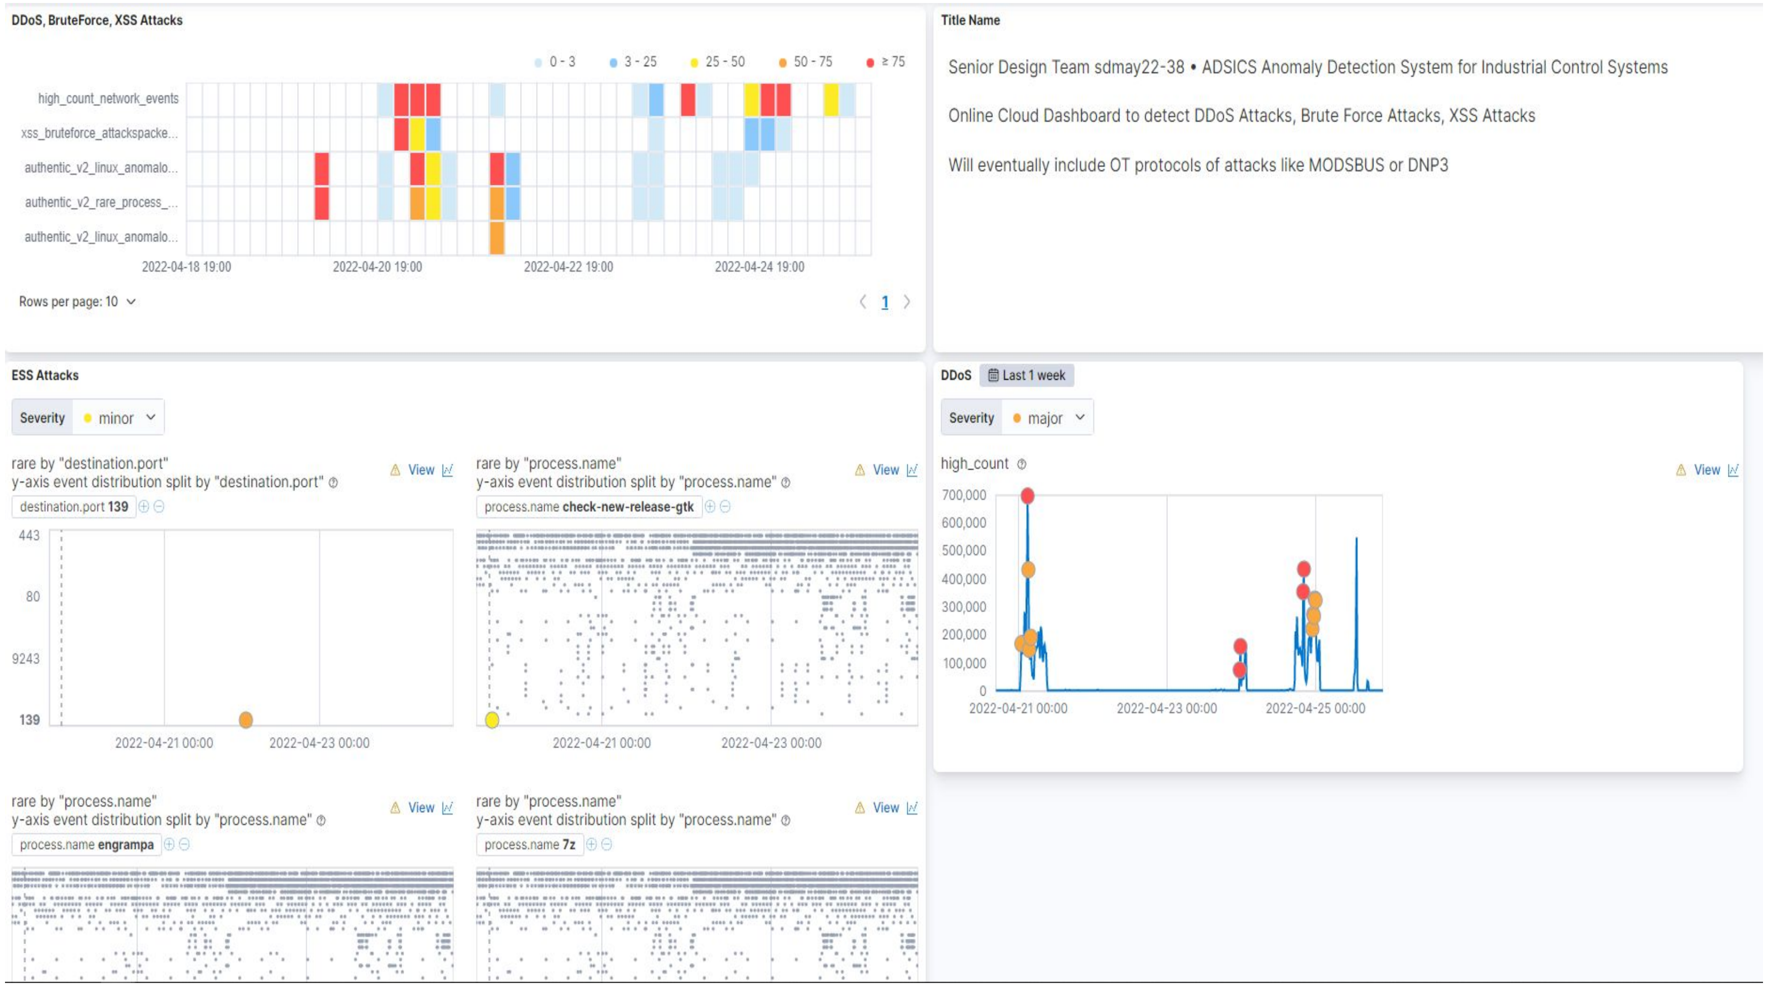

Elastic: Machine Learning - Anomaly Detection

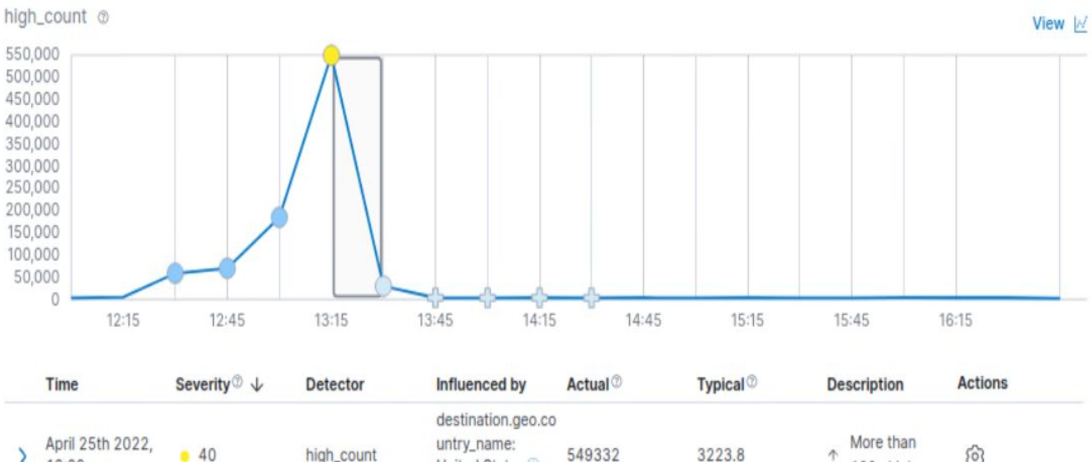

Elastic: Machine Learning Data Frame Analytics

Actual class

	BENIGN	DoS Hulk	DoS Slowhttptest	DoS slowloris
BENIGN	99%	1%	0%	0%
DoS Hulk	0%	100%	0%	0%
DoS Slowhttptest	1%	0%	99%	1%
DoS slowloris	0%	0%	0%	100%

Evaluation quality metrics

0.998 Overall accuracy

0.994 Mean recall

Per class recall and accuracy

Class	Accuracy	Recall
BENIGN	0.998	0.993
DoS Hulk	0.998	1
DoS Slowhttptest	1	0.985
DoS slowloris	1	0.996

Total feature importance

Feature importance docs

Total feature importance values indicate how significantly a field affects the predictions across all the training data.

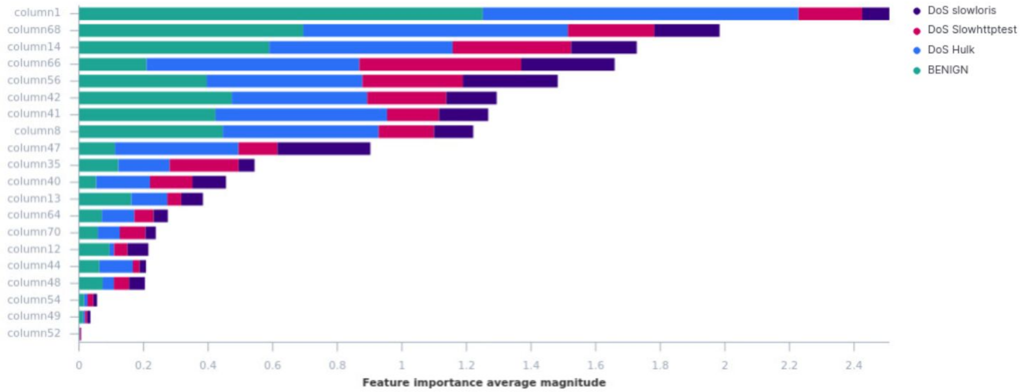

Model evaluation

Job status: stopped

docs evaluated: 170366

Normalized confusion matrix for entire dataset

Predicted class

Columns

Actual class

	BENIGN	Web Attack � Brute F...	Web Attack � Sql Inje...	Web Attack � XSS
BENIGN	100%	0%	0%	0%
Web Attack � Brute Force	1%	19%	0%	80%
Web Attack � Sql Injection	0%	43%	0%	57%
Web Attack � XSS	5%	0%	0%	95%

Evaluation quality metrics

0.99 Overall accuracy

0.534 Mean recall

Per class recall and accuracy

Receiver operating characteristic (ROC) curve

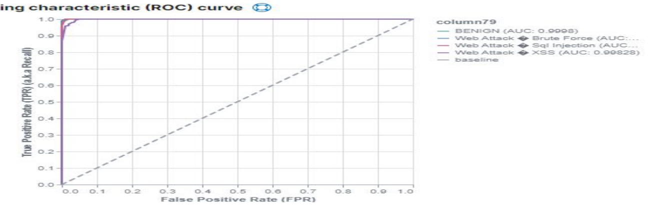

Columns

Actual class

	Benign	modbusqueryflood
Benign	100%	0%
modbusqueryflood	0%	100%

Evaluation quality metrics

1 Overall accuracy

1 Mean recall

Per class recall and accuracy

Class	Accuracy	Recall
Benign	1	1
modbusqueryflood	1	1

Rows per page: 10

1

Receiver operating characteristic (ROC) curve

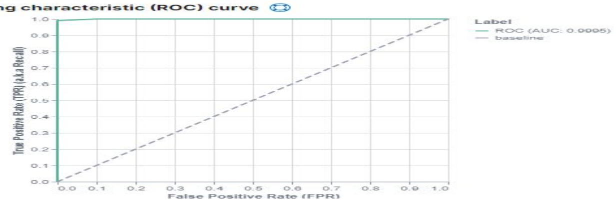

Local: Theoretical Design

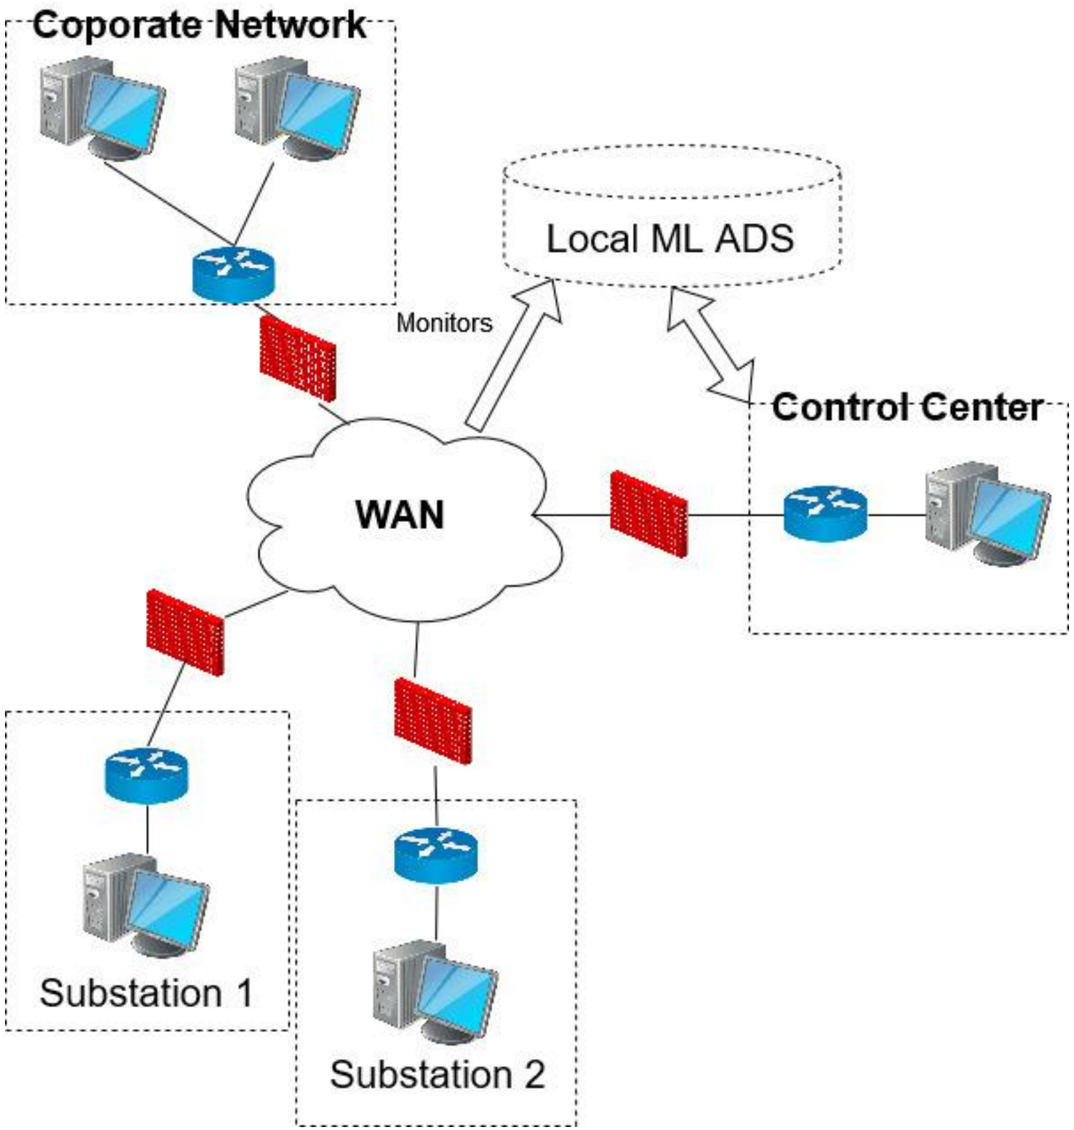

Local: Practical Design

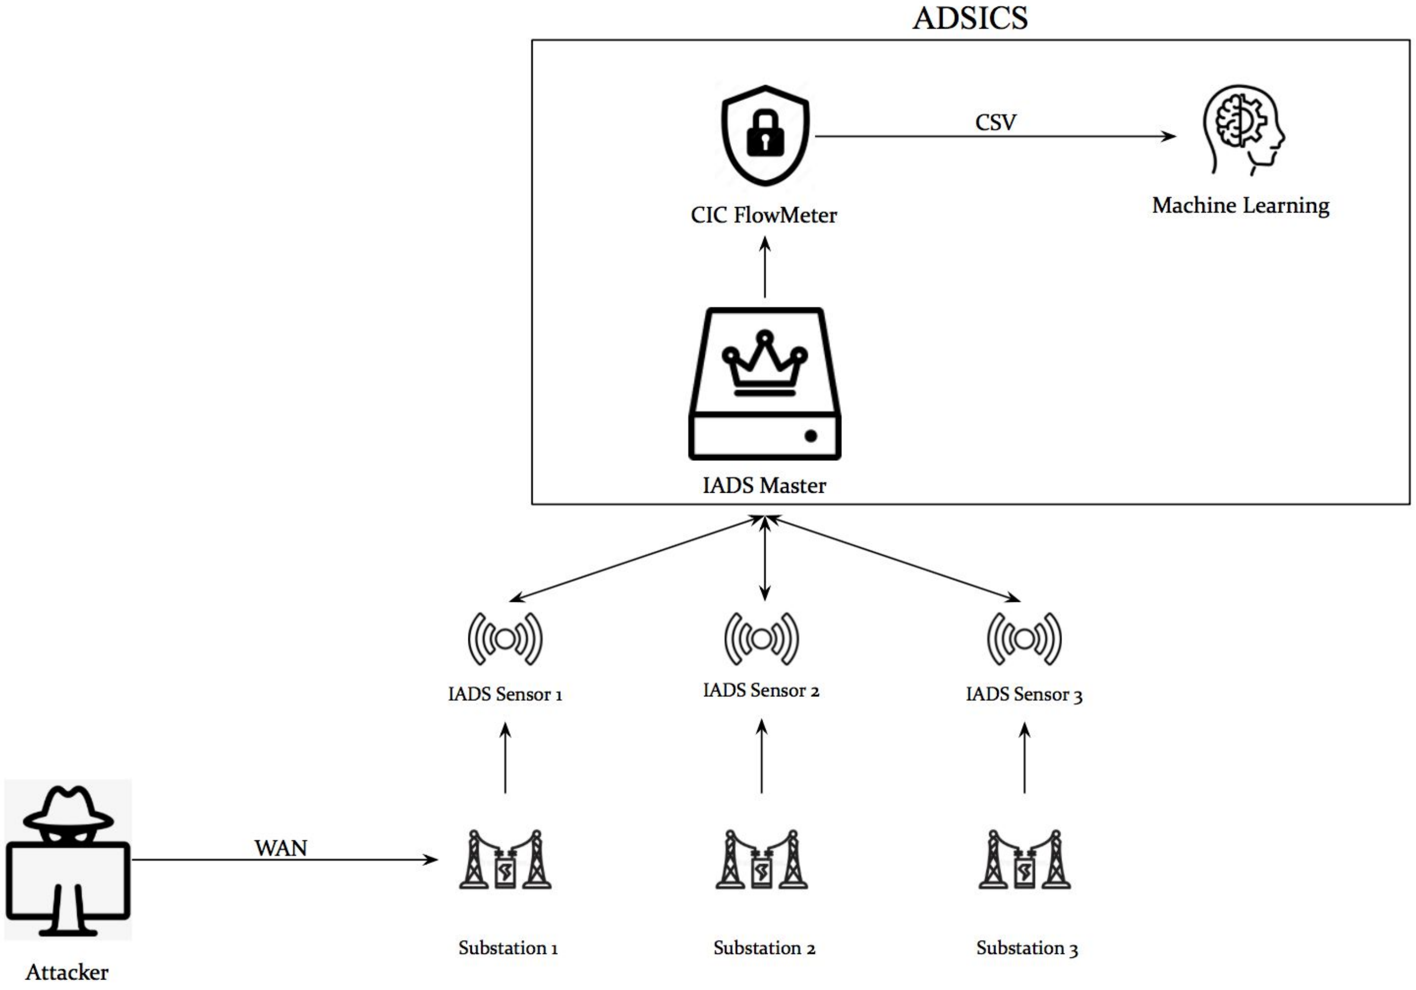

Local: SMOTE/Undersampling

	Benign (Normal)	DoS Slowloris	DoS Slowhttptest	DoS Hulk	Dos GoldenEye
Before SMOTE	440042	5796	5499	231073	10293
Distribution	63.53%	0.84%	0.79%	33.36%	1.49%
After SMOTE	231073	115536	115536	231073	115536
Distribution	28.57%	14.29%	14.29%	28.57%	14.29%

- Imbalanced dataset

Synthetic Minority Oversampling Technique (SMOTE)

Random undersampling

- Undersampling amount = ½ the largest alert count (ie. the BENIGN count).
- Oversampling amount = ½ of the undersampling amount

Local: Cross Validation

- Reduced feature count from 78 to 29 features from Thursday dataset.
- Forward model selection is used to select the features.
- Evaluated by least MSE, applied greedy algorithm to get the best combination of the features
- Set the number of features to 29, so no other criterion is used (such as AIC, BIC, Mallows CP, adjusted R squared)

	Feature Name	Description
1	Destination Port	Destination Port
2	Total Length of Bwd Packets	Total size of packet in backward direction
3	Fwd Packet Length Min	Minimum size of packet in forward direction
4	Bwd Packet Length Min	Minimum size of packet in backward direction
5	Bwd Packet Length Mean	Mean size of packet in backward direction
6	Bwd Packet Length Std	Standard deviation size of packet in backward direction
7	Bwd PSH Flags	Number of times the PSH flag was set in packets travelling in the backward direction (0 for UDP)
8	Fwd Header Length	Total bytes used for headers in the forward direction
9	Max Packet Length	Maximum length of a packet
10	Packet Length Mean	Mean length of a packet
11	Packet Length Std	Standard deviation length of a packet
12	FIN Flag Count	Number of packets with FIN
13	PSH Flag Count	Number of packets with PSH
14	ACK Flag Count	Number of packets with ACK
15	URG Flag Count	Number of packets with URG
16	CWE Flag Count	Number of packets with CWE
17	ECE Flag Count	Number of packets with ECE
18	Down/Up Ratio	Download and upload ratio
19	Avg Fwd Segment Size	Average size observed in the forward direction
20	Fwd Avg Bytes/Bulk	Average number of bytes bulk rate in the forward direction
21	Fwd Avg Packets/Bulk	Average number of packets bulk rate in the forward direction
22	Fwd Avg Bulk Rate	Average number of bulk rate in the forward direction
23	Bwd Avg Bytes/Bulk	Average number of bytes bulk rate in the backward direction
24	Bwd Avg Packes/Bulk	Average number of packets bulk rate in the backward direction
25	Bwd Avg Bulk Rate	Average number of bulk rate in the backward direction
26	Subflow Fwd Bytes	The average number of packets in a sub flow in the forward direction
27	Subflow Bwd Bytes	The average number of packets in a sub flow in the backward direction
28	Init_Win_bytes_backward	The total number of bytes sent in initial window in the backward direction
29	min_seg_size_forward	Minimum segment size obeserved in the forward direction

Local: Decision Tree

```
DECISION TREE WEDNESDAY TEST
[[17380     8   428    26]
 [   10  8902     0    16]
 [   10     1 17760    14]
 [    3     0  2120  6787]]
              precision    recall  f1-score   support

         0.0       1.00      0.97      0.99     17842
         1.0       1.00      1.00      1.00      8928
         2.0       0.87      1.00      0.93     17785
         3.0       0.99      0.76      0.86      8910

    accuracy                           0.95     53465
   macro avg       0.97      0.93      0.94     53465
weighted avg       0.96      0.95      0.95     53465
```

```
DECISION TREE THURSDAY
[[167541     48    533     64]
 [     2    220   1213     72]
 [    22      0    629      1]
 [     0      0      0     21]]
              precision    recall  f1-score   support

           0       1.00      1.00      1.00    168186
           1       0.82      0.15      0.25      1507
           2       0.26      0.96      0.42       652
           3       0.13      1.00      0.23        21

    accuracy                           0.99    170366
   macro avg       0.55      0.78      0.47    170366
weighted avg       1.00      0.99      0.99    170366
```

Local: Random Forest

```
RANDOM FOREST WEDNESDAY TEST
[[17400     2   421    19]
 [    9  8903     0    16]
 [    9     1 17764    11]
 [    3     0  2120  6787]]
              precision    recall  f1-score   support

         0.0       1.00      0.98      0.99     17842
         1.0       1.00      1.00      1.00      8928
         2.0       0.87      1.00      0.93     17785
         3.0       0.99      0.76      0.86      8910

    accuracy                           0.95     53465
   macro avg       0.97      0.93      0.95     53465
weighted avg       0.96      0.95      0.95     53465
```

```
RANDOM FOREST THURSDAY
[[167627      6    496     57]
 [     3    219   1213     72]
 [    21      0    630      1]
 [     1      0      0     20]]
              precision    recall  f1-score   support

           0       1.00      1.00      1.00    168186
           1       0.97      0.15      0.25      1507
           2       0.27      0.97      0.42       652
           3       0.13      0.95      0.23        21

    accuracy                           0.99    170366
   macro avg       0.59      0.77      0.48    170366
weighted avg       1.00      0.99      0.99    170366
```

Local: Modbus

- reference power grid data

Modbus_Combined_Train			
Three different DDoS. Modbus query flood, Ping flood, tcpSYN flood			
Network Traffic lasts 6,12,12 hours			
	IP		
Attacker	172.27.224.50	172.27.224.70	172.27.224.80
Victim	172.27.224.250		
Type of Attack		**Description**	
Modbus Query Flood		30 minute duration	
Ping Flood		30 minute duration	
TCP SYN Flood		30 minute duration	

Local: Modbus-Combined

```
[[60811   787]
 [ 5635 92568]]
```

		precision	recall	f1-score	support
Category	0	0.92	0.99	0.95	61598
	1	0.99	0.94	0.97	98203
	accuracy			0.96	159801
	macro avg	0.95	0.96	0.96	159801
	weighted avg	0.96	0.96	0.96	159801

Local: Modbus-Separate

explain why we want to have this for anomaly detection

```
[[60809     0   789     0]
 [    0 26333    21     0]
 [ 5635 14882  5765 14988]
 [    2 14896   649 15032]]
```

		precision	recall	f1-score	support
BENIGN	0	0.92	0.99	0.95	61598
MODBUS Query Flood	1	0.47	1.00	0.64	26354
Ping Flood	2	0.80	0.14	0.24	41270
TCP SYNFlood	3	0.50	0.49	0.50	30579
	accuracy			0.68	159801
	macro avg	0.67	0.65	0.58	159801
	weighted avg	0.73	0.68	0.63	159801

Comparison

	Elasticsearch		Local ML			
	Wednesday	Thursday	Wednesday	Thursday	Modbus	
					Combined	Separate
Max Precision	100.0%	100.0%	100.0%	100.0%	99.0%	95.0%
Min Precision	99.0%	0.0%	87.0%	23.0%	92.0%	24.0%
Accuracy	99.8%	99.0%	95.0%	99.0%	96.0%	68.0%
Macro Average	99.4%	53.4%	95.0%	48.0%	96.0%	67.0%

Conclusion

- Elasticsearch worked well with IT attacks
- Local ML was able to capture Modbus attacks
- Refining ML model is needed to increase accuracy and provide additional information to the client
- *Elasticsearch may provide OT support later on? <- what do you think*

REFERENCE POWER GRID

Recommendation:

Elasticsearch system could be used to detect IT attacks to the control system, however, it lacked OT protocol support. For OT network traffic, Local ML Anomaly detection can be used to supplement a company's anomaly detection system until this functionality is added to Elasticsearch.

Thank You For Your Time!
Any Questions?

References

Khraisat, A., Gondal, I., Vamplew, P. et al. Survey of intrusion detection systems: techniques, datasets and challenges. Cybersecur 2, 20 (2019). https://doi.org/10.1186/s42400-019-0038-7

https://www.linkedin.com/company/sqrrl

Intrusion Detection Evaluation Dataset (CIC-IDS2017). University of New Brunswick. https://www.unb.ca/cic/datasets/ids-2017.html

Applications - CICFlowMeter (formerly ISCXFlowMeter). University of New Brunswick. https://www.unb.ca/cic/research/applications.html#CICFlowMeter

ahlashkari/CICFlowMeter. https://github.com/ahlashkari/CICFlowMeter

Frazão, I. and Pedro Henriques Abreu and Tiago Cruz and Araújo, H. and Simões, P. , "Denial of Service Attacks: Detecting the frailties of machine learning algorithms in the Classication Process", in 13th International Conference on Critical Information Infrastructures Security (CRITIS 2018), ed. Springer, Kaunas, Lithuania, September 24-26, 2018, Springer series on Security and Cryptology , 2018. DOI: 10.1007/978-3-030-05849-4_19

modbus TCP SCADA #1. https://github.com/tjcruz-dei/ICS_PCAPS/releases\

Network datasets, Overview of Network Forensic Tools and Datasets. Github. https://martinazembjakova.github.io/Network-forensic-tools-taxonomy/network-datasets.html

References

Extremists

https://www.utilitydive.com/news/extremists-have-developed-credible-specific-plans-to-attack-the-us-power/617741/

Global Threats

https://semiengineering.com/power-grids-under-attack/