Operational Technology Behavioral Analytics (OTBA) – A Data-Centric Approach for Reducing Cybersecurity Risk

Clifton R. Black[1], *Member, IEEE*, Michael Thorn[1], Adam Carrington[2], Masrur Rahman[3]
Southern Company[1], TEKsystems[2], Securicon[3]
crblack@southernco.com, mithorn@southernco.com, x2ajcarr@southernco.com, masrur.drmc2007@gmail.com

***Abstract*—This paper provides a summary of the methodology, findings, lessons learned, and insights gained from an investigation into the feasibility of the Operational Technology Behavioral Analytics (OTBA) cybersecurity approach. The concept was evaluated with data from the National Carbon Capture Center (NCCC) – a U.S. Department of Energy (DOE) funded facility that is managed and operated by Southern Company at Alabama Power's E. C. Gaston generating power plant in Wilsonville, Alabama. Appropriate data sources for the post-combustion carbon capture system were identified. Infrastructure was deployed to monitor, capture, and archive data for the system. Critical parameters for each subsystem were identified and analyzed. Machine learning algorithms were used to establish and characterize normal operations and subsequently identify anomalies. This effort yielded valuable insights and formed the basis of a data-centric strategy for detecting cyber-attacks along with a coordinated response philosophy. A significant takeaway is that the OTBA cybersecurity approach is quite portable; it can be applied to other critical infrastructure beyond fossil power generation.**

***Index Terms*—Cybersecurity, Behavioral Analytics, Anomalies, Machine Learning, Sensors, and Risk Management.**

I. Introduction

Attacks on critical energy infrastructure, including power generation assets, are increasing at a staggering rate. The need to secure the nation's critical infrastructure is vital as seen by the impacts of successful cyber-attacks around the world. These include but are not limited to the 2015 cyber-attack at a Ukraine power plant and the 2021 cyber-attack on Colonial Pipeline in the US [1], [2]. Current world events have only exacerbated potential security threats. The trend toward digitization for enhanced monitoring, control, and optimization of our critical infrastructures has resulted in the extensive incorporation of information, communication, and technology (ICT) systems. This tight coupling of ICT and physical systems has increased the attack surface for cyber-attacks. Cybersecurity solutions that seek to enable, improve, and protect power systems through the application of applied novel concepts are of urgent importance. This project is driven by a desire to develop approaches to reduce cybersecurity risks associated with the deployment of distributed sensors and advanced controls in fossil power generation environments. The proposed solution integrates data from both the operational technology (OT) infrastructure and the informational technology (IT) infrastructure to characterize normal operations and subsequently identify abnormal conditions that may be indicative of cyber-attacks [3], [4]. This approach incorporates proactive and predictive security measures to increase the possibility of averting cyber-attacks.

The goal of the Operational Technology Behavioral Analytics (OTBA) project was to reduce cyber risks in the production of energy through improved operational intelligence. The basic premise is that the OT infrastructure and the IT infrastructure both contain vital information that when integrated, yield valuable cybersecurity insights greater than the sum of the individual parts. With data from both systems, normal behavior for a system can be characterized, facilitating anomaly identification. Initially anomalies are alert flags. The behavior of other associated parameters during an anomalous event helps to determine if alerts indicate potential attacks. Correlation among anomalies increases confidence that outliers are indicative of a cyber event.

Abnormal system behavior may provide information about impending or ongoing cyber-attacks, and the overall concept facilitates the development of an enhanced knowledge of risks and risk management techniques. It is important to note that the systems to be protected have varying degrees of sophistication and complexity, due to the combination of the intricacy of the innate physical system and communications systems coupled with human behavior. An automated OTBA system needs to be able to learn the way these systems operate over time, as this may change. Machine learning algorithms are used to develop models to capture the normal operation of the system to support identifying abnormal conditions. The incorporation of physics-based concepts as well as insights from subject matter experts (SME) is helpful in distinguishing faults from attacks.

This material is based upon work supported in part by the Department of Energy Award Number DE-FE0031640.

979-8-3315-8138-1

II. PROJECT DATA

Since the OTBA cyber security concept is predicated upon the characterization of normal operations and the subsequent identification and alerting of anomalies at the site, from both the IT and OT vantage points, data from both domains were collected and archived. The aim was to collect as much raw data as possible to establish normal baseline network traffic behavior for the two-year period, January 2019 to December 2021. Several factors impacted data collection, including planned plant maintenance outages, COVID-19 related shutdowns as well as issues with the data collection infrastructure. To account for contiguity inconsistencies, the initial project timeline was extended via no-cost time extension to facilitate the collection of adequate raw data. Ample datasets are needed for both training and testing the data models.

Network visibility was achieved by configuring switched port analyzers or span ports from critical physical interconnection points then consolidating all traffic streams to a single aggregation device. Span ports, also called mirror ports, allow a copy of the traffic to be sent to another interface for monitoring or analysis. Priority was placed on the network data across the operating plant which monitors the command and control of low-level field devices and associated managing software applications. The OT data are transmitted using operational device protocols such as Distributed Network Protocol (DNP3) and Open Platform Communications (OPC) that convey command traffic to program the field devices. The IT data are transmitted using network traffic protocols like Domain Name System (DNS) and Address Resolution Protocol (ARP) [5], [6], [7].

The OT data are continuously archived in a relational database. These points include but are not limited to timestamped measurements from temperature and pressure monitors, flow level indicators, motors, and valves. Overall, the project's dataset was collected over a two-year period to ensure that adequate data were available to characterize the normal plant operations and capture as many potential anomalies as possible. The IT and OT datasets were both analyzed separately and then together.

Once the project data were captured and ingested into the analytic engine, data collector, forwarder, and indexer requirements were evaluated to ensure efficient data normalization and analysis. Analytics SMEs utilized data discovery and capture requirements to engineer detailed hardware and software specifications to ensure successful collection and data integrity.

III. BASELINING AND ANOMALY DETECTION

Figure 1 encapsulates the tools and functions used to develop baselines for normal system operation as well as that used to identify and alert on anomalous activities. The raw events include items from both the IT and OT arenas. The OT data was transferred from a relational database to the data analytics tool [8]. The IT data from the various observation ports were moved to a data aggregator. An industrial control system (ICS) traffic analysis tool subsequently processed the aggregated raw network traffic and produced metadata that fully characterized the data in terms of network device categories, network protocols, and communication alerts. This network metadata was then streamed to the data analytics tool for analysis, including correlation with the OT data.

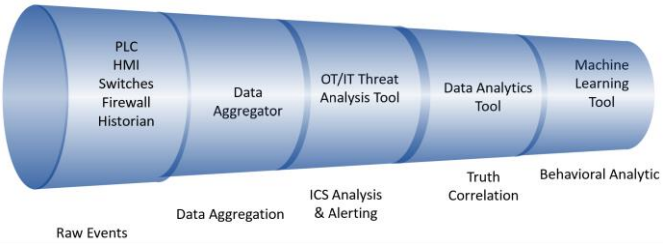

Fig. 1. OTBA Overall System

The detailed characterization of the network data along with the analysis of the data across various axes facilitated the development of a baseline of normal activity. The axes of analysis included:

- Time – Data were represented in a time-series format such that patterns over different periods (days, weeks, months, and years) could be ascertained. The patterns included fluctuations in data volume or other cyclic variations.
- Threat category – This is one of the categorizations generated from the network data. The make-up of this feature when no known security events were present in the data was referenced as the baseline.
- Severity – This is another categorization generated from the network data. The make-up of this feature when no known security events were present in the data was referenced as the baseline.
- Network Protocol – At any given time, the set of expected protocols on the communications network is precisely known by IT SMEs. This information was used as a baseline.
- IP address - At any given time, the set of expected source and destination IP addresses as well as the schedule of their communications based on their defined functions is precisely known by IT SMEs. This information was used as a baseline.

Data scientists engaged with plant partners to create a system component data baseline using plant historian(s), traditional IT systems, and individual ICS component data. For example, OT operators assisted in identifying relevant programmable logic controller (PLC) devices for air flow valves, pressure gauges, and tank volume. Across the performance period, the project team maintained continuous monitoring and analytics on the collected data to detect anomalies. The project team also included events that occurred on an irregular basis to the baseline as each was discovered and investigated.

At the NCCC, technology developers operate intermittently, evaluating various carbon capture systems and their components. In support of technology testing and

development, the site leverages multiple balance of plant (BOP) systems including instrument air, nitrogen, cooling tower system, waste neutralization system, and other ancillary systems. Data from these systems constitute the operational, or OT, project data.

With help from the engineering SMEs, the relevant parameters characterizing the operation of each BOP system were identified. The data for each of these parameters are archived with specific associated identifiers called tags in the relational database. While some parameters are numeric measurements that change over time, others are operational status indicators reflecting whether a valve is opened or closed or whether a pump is running, stopped or shutdown.

The data analytics tool was used to perform the baseline and anomaly detection on the OT data. Machine learning was used to create outlier models for each of the tags in each system [9]. Specifically, the numeric standard deviation function was used to create machine learning models to generate baselines and identify anomalies for each parameter.

A description of and baselining analysis for the Instrument Air BOP system follows.

A. *Instrument Air*

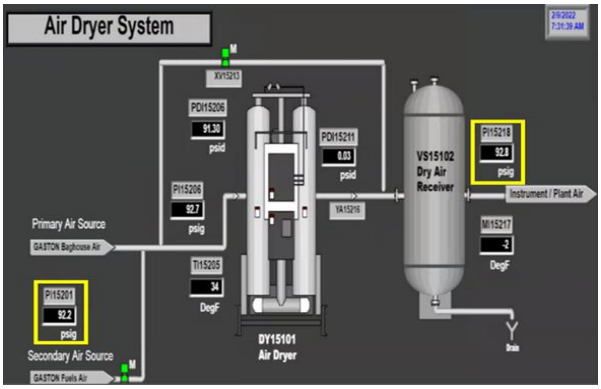

Fig. 2. Instrument Air BOP System

The Instrument Air system, shown in figure 2, supplies the entire site with instrument-quality air that is critical to the operation of hundreds of control valves, purging of gas analysis equipment, and various other operational needs for the site. This is a critical system for the NCCC; all operation would halt within an hour of its loss. The air is supplied by Plant E.C. Gaston through its air compressor systems from a primary and a secondary source.

The system performance is monitored primarily through pressure indications and a moisture analyzer. The moisture analyzer is a performance metric not a system health metric. Two tags which provide key insights to the system performance are PI15218 (Instrument Air header pressure after the air receiver tank) and PI15201 (back-up supply for Instrument Air). Each of these tags indicate the pressure measurement at a certain point in the process. A pressure loss at one or both locations would indicate a significant system upset or performance deviation.

From a total of 25,561 measurements, 607 outliers were detected for the data available from year 2019 to 2021. Values were in the normal range 98 percent of the time. The average operational range was determined by the algorithm. It should be noted that this average range is not a static window. As measurements are ingested over time, the average range is updated to reflect the changing operating state. The OTBA system is "learning" the operational characteristics of the BOP subsystem. The same baselining analysis was performed for the Nitrogen, Waste Neutralization, and Cooling Towers BOP systems using relevant parameters.

IV. Correlating Network and Operations Data to Identify Anomalies

IT and OT data were both analyzed individually and then together, to understand what occurred during a given period of operation. Anomalies in the operational data from the air instrument system pumps were identified along with their frequency of occurrence. The probable cause for their occurrences were investigated. Secondly, associated IT data were analyzed to note outliers. Thirdly, IT data were filtered to the same timeframe in which the OT anomalies were identified in the operational data and examined for the existence of correlations between the two datasets.

A man-in-the-middle (MITM) type event was detected. It was deemed "of interest" as it occurred the same day and time as the operational anomalies. An investigation determined that no actual malicious MITM attack on the system occurred. However, the network behavior as observed had similar characteristics to a MITM attack. A MITM attack is a type of eavesdropping attack in which the attacker intercepts an existing data transfer or conversation [10]. This discovery would have more than likely been missed through normal alerting and monitoring practices. This hypothesis paves the way for more advanced detections to be created by combining and centralizing IT and OT data and monitoring these data together in an automated approach versus a manual approach, which will allow for further observations to discover stronger correlated events that may not be immediately discovered without the use of machine learning.

V. Data Centric Anomaly Detection Strategy Architecture

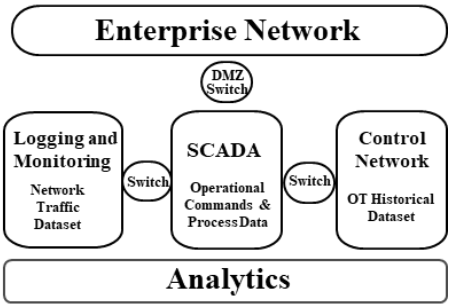

Fig. 3. Anomaly Detection Strategy Architecture

A. *Machine Learning Automation*

The use of Machine learning in the project environment currently required manual intervention through the entire lifecycle. Manual data cleaning requires a significant amount

of time and collaborative effort among staff. It also takes time to discover what data impacts the machine learning models most significantly. Feature selections were based on SME's expert opinion and statistical measurements to validate hypothesis to understand if an observed occurrence appeared by chance. Different models must be investigated to determine which is best be suited to address project objectives, whether geared towards classification or regression. With a variety of models that can address the problem, it is time consuming to select, run, and adjust parameters for each model. This is where machine learning automation can provide great benefits. With high volumes of data generated each day from both IT and OT, it takes significant effort to manually determine the key features for both domains independently. Alternatively, creating an automated pipeline that prepares the data, only allowing for the most important features and/or statistically significant features to be selected for model building, reduces the overall time and effort. Figure 3 shows the architecture for the developed anomaly detection strategy.

Significant efficiencies can be gained when the ML processes are automated [11], [12]. The main aspects of the process are shown in figure 4 below.

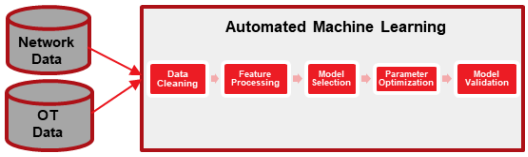

Fig. 4. Automated Machine Learning

Utilizing Automated Machine learning unlocks the potential to scale machine learning algorithms, allowing a brute force approach that will run multiple or all algorithms that have been automatically selected to solve a given problem such as regression, classification, ranking and more. This approach will aid SMEs in determining and identifying patterns or anomalies not otherwise detected by conventional methods.

B. *Anomaly Monitoring and Alerting*

Alerting and monitoring capabilities can be enhanced by machine learning in conjunction with an IT SME monitoring all operations [13], [14]. The process is to combine both datasets. Combining both datasets enables analysis and determination regarding whether an abnormal condition has been observed and if it is correlated. This finding will be provided to the SME for further investigation and validation. This approach keeps the human in the loop, allowing the SME to be a key part of the retraining and continuous improvement of the machine learning models. Keeping the SME in the life cycle also provides a backup to the machine learning model as there could be instances where the model has not been exposed to or trained on a particular scenario. The SME can provide direct feedback to the model and issue a retraining event so that the model can be updated and retrained on the new data points. The SME acts as a fact checker or redundancy plan, so that if the ML driven alerts fail or are unavailable, a SME will always be available to ensure that alerts are not missed and can push the alerts to the appropriate stakeholders. The overall monitoring and alerting process is depicted in figure 5 below.

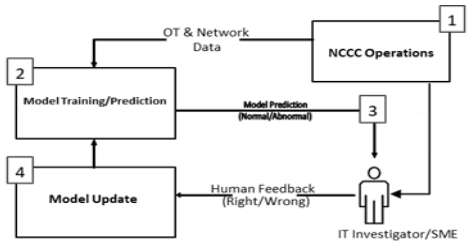

Fig. 5. OTBA Monitoring and Alerting Process

C. *Software*

The monitoring and logging collection architecture stack depicted in figure 6 was designed to use vendor supported tools to allow quick implementation at the plant location. Components for raw data aggregation, ICS analysis, and truth correlation were chosen as logical functional sub-boundaries for the dataset collection stack.

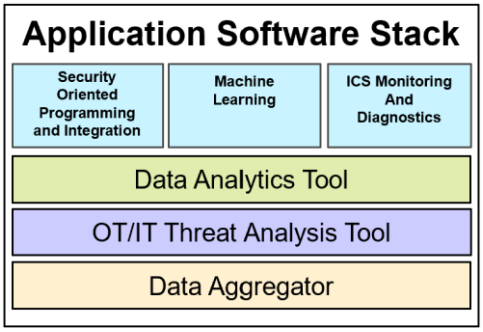

Fig. 6. Application Software Stack

A raw data aggregation appliance was used to route network traffic occurring in the Supervisory Control and Data Acquisition (SCADA) and control network. An ICS traffic analysis tool for OT/IT data, was selected as the categorization engine used to normalize the aggregated raw network traffic.

Datasets were normalized over the collection period to allow advanced data analysis to be executed. In addition to the core enterprise software, the data analytics tool also has the option to utilize data specific add-on applications to narrowly focus on individual data types or compute precise reporting results. This project primarily utilized Machine learning, however separate add-on applications for OT/IT Threat Analysis Tool and ICS Monitoring and Diagnostics were reviewed. In figure 6, refers to an application add-on used to parse and visualize data from the threat analysis tool.

VI. LESSONS LEARNED

A. *Scalability*

Accounting for significant amounts of data needed to accomplish analytic initiatives should be addressed early to ensure a solution can easily scale without the need for reconfiguration, significant additional cost, or time delays.

B. Deep Learning Can Provide Greater Insights

While statistical analysis and machine learning provide valuable information, deep learning can provide even greater insights. Characteristics and behaviors that are not easily detected can be brought to light. Deep learning techniques explore hidden relationships, allow automated feature learning, and moves us closer to true artificial intelligence.

C. Manual Processes Create Botthenecks

Due to various project constraints, several manual processes were used to interact with the data. These processes created bottlenecks which led to delays. Automation would help eliminate such dependencies as data can be scheduled, refreshed, and prepared for faster and more seamless analysis.

VII. Conclusion

The cybersecurity risks associated with the deployment of distributed sensors and advanced controls can be reduced with the proposed OTBA approach. A pilot system for IT/OT anomaly identification and alerting system was developed and yielded insightful information regarding the overall feasibility of the system. Previously unknown activities were identified, including a scenario where nontraditional characters were included in DNS traffic requests to the control network. While this event was not a security threat, it facilitated the development of additional monitoring and analysis to improve the situational awareness of the plant.

The incorporation of physics-based approaches along with SME insights is helpful in distinguishing faults from attacks. This is important as frequent false positives degrades confidence in the system results. The accuracy and speed of OTBA will be critical to its usefulness and acceptance. Physical systems sometimes operate in abnormal modes. The monitoring systems need to account for these operational variations. These items will be addressed in future work.

All systems have limitations. It is important to understand where “blind spots” exist, so that additional processes and procedures can be employed to minimize security risks. Finally, the OTBA cybersecurity approach is quite portable and can be applied to a range of infrastructures to reduce cybersecurity risks.

Disclaimer

"This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof."

References

[1] University of Washington, The Henry M. Jackson School of International Studies. (2017, Oct. 11). *Cyberattack on Critical Infrastructure: Russia and the Ukrainian Power Grid Attacks*. [Online]. Available: https://jsis.washington.edu/news/cyberattack-critical-infrastructure-russia-ukrainian-power-grid-attacks/

[2] D. E. Whitehead, K. Owens, D. Gammel and J. Smith, "Ukraine cyber-induced power outage: Analysis and practical mitigation strategies," 2017 70th Annual Conference for Protective Relay Engineers (CPRE), College Station, TX, USA, 2017, pp. 1-8, doi: 10.1109/CPRE.2017.8090056.

[3] W. A. Conklin, "IT vs. OT Security: A Time to Consider a Change in CIA to Include Resilienc," 2016 49th Hawaii International Conference on System Sciences (HICSS), Koloa, HI, USA, 2016, pp. 2642-2647, doi: 10.1109/HICSS.2016.331.

[4] P. K. Garimella, "IT-OT Integration Challenges in Utilities," 2018 IEEE 3rd International Conference on Computing, Communication and Security (ICCCS), Kathmandu, Nepal, 2018, pp. 199-204, doi: 10.1109/CCCS.2018.8586807.

[5] R. Amoah, S. Camtepe and E. Foo, "Securing DNP3 Broadcast Communications in SCADA Systems," in IEEE Transactions on Industrial Informatics, vol. 12, no. 4, pp. 1474-1485, Aug. 2016, doi: 10.1109/TII.2016.2587883.

[6] M. H. Jalalzai, W. B. Shahid and M. M. W. Iqbal, "DNS security challenges and best practices to deploy secure DNS with digital signatures," 2015 12th International Bhurban Conference on Applied Sciences and Technology (IBCAST), Islamabad, Pakistan, 2015, pp. 280-285, doi: 10.1109/IBCAST.2015.7058517.

[7] H. Y. Ibrahim, P. M. Ismael, A. A. Albabawat and A. B. Al-Khalil, "A Secure Mechanism to Prevent ARP Spoofing and ARP Broadcasting in SDN," 2020 International Conference on Computer Science and Software Engineering (CSASE), Duhok, Iraq, 2020, pp. 13-19, doi: 10.1109/CSASE48920.2020.9142092.

[8] Dr. Nadine Shillingford, Data Analytics Using Splunk 9.x: A practical guide to implementing Splunk’s features for performing data analysis at scale, Packt Publishing, 2023.

[9] M. Cersosimo and A. Lara, "Detecting Malicious Domains using the Splunk Machine Learning Toolkit," NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium, Budapest, Hungary, 2022, pp. 1-6, doi: 10.1109/NOMS54207.2022.9789899.

[10] A. Prasad and S. Chandra, "Defending ARP Spoofing-based MitM Attack using Machine Learning and Device Profiling," 2022 International Conference on Computing, Communication, and Intelligent Systems (ICCCIS), Greater Noida, India, 2022, pp. 978-982, doi: 10.1109/ICCCIS56430.2022.10037723.

[11] T. Peng, "Edge Computing Service Framework for IT/OT Converged Networks for Cyber-Resiliency: Machine Learning-Assisted SDN," 2023 IEEE 3rd International Conference on Electronic Communications, Internet of Things and Big Data (ICEIB), Taichung, Taiwan, 2023, pp. 230-234, doi: 10.1109/ICEIB57887.2023.10170421.

[12] C. Feng, S. Wu and N. Liu, "A user-centric machine learning framework for cyber security operations center," 2017 IEEE International Conference on Intelligence and Security Informatics (ISI), Beijing, China, 2017, pp. 173-175, doi: 10.1109/ISI.2017.8004902.

[13] V. W. Samawi, S. A. Yousif and N. M. G. Al-Saidi, "Intrusion Detection System: An Automatic Machine Learning Algorithms Using Auto-WEKA," 2022 IEEE 13th Control and System Graduate Research Colloquium (ICSGRC), Shah Alam, Malaysia, 2022, pp. 42-46, doi: 10.1109/ICSGRC55096.2022.9845166.

[14] S. Ndichu, T. Ban, T. Takahashi and D. Inoue, "A Machine Learning Approach to Detection of Critical Alerts from Imbalanced Multi-Appliance Threat Alert Logs," 2021 IEEE International Conference on Big Data (Big Data), Orlando, FL, USA, 2021, pp. 2119-2127, doi: 10.1109/BigData52589.2021.9671956.