ByzSec — A Multi-layered Byzantine Resilient Architecture for Bulk Power System Protective Relays

Christopher Bonebrake*, D. Jonathan Sebastian-Cardenas*, Carl H. Miller, Sahiti Bommareddy*, Yair Amir, Kade Cornelison, Cliff Eyre*, Paul Skare†, Sri Nikhil Gupta Gourisetti, Aditya Ashok, Bev Johnson

* Member, IEEE, †Senior member, IEEE
EED-Pacific Northwest National Laboratory
Richland, WA 99354, USA
Corresponding author: d.sebastiancardenas@pnnl.gov

2024 IEEE Power & Energy Society General Meeting (PESGM) | 979-8-3503-8183-2/24/$31.00 ©2024 IEEE | DOI: 10.1109/PESGM51994.2024.10688803

***Abstract*—Reliability, selectivity, and sensitivity are the fundamental attributes of any protection system, acting as the main drivers in the selection of schemes, and equipment. In high-voltage systems, microprocessor-based relays represent the industry's preferred solution, providing engineers with a vast array of benefits. However, they remain vulnerable to cybersecurity events that may compromise their functionality. To help mitigate against potential cybersecurity risks, this paper presents a fault-tolerant, Byzantine Resilient (BR) architecture that significantly increases the cybersecurity attributes of a protection system while minimizing the amount of performance impacts and integration overheads introduced. The solution relies on an array of independent relays that utilize robust consensus methods (based on Spire [1], [2]) to ensure correct system behavior is achieved even when a relay has been compromised. Furthermore, the solution has been complemented with a custom-built *Situational Awareness* engine that can be used to detect and identify potential threats. The implemented solution has been developed in consultation with three hardware vendors and has been tested to comply with the performance requirements of a 345kV differential protection scheme (87T). The results indicate that the proposed architecture is a comprehensive solution that: supports the strict correctness and performance requirements of the bulk power grid while providing a cost-effective alternative that offers a seamless, long-term solution.**

Keywords—Cybersecure protection architecture, Byzantine fault tolerant systems

I. Introduction

Protection systems are an essential security feature of any power system. A good protection scheme balances the need for continuous system operation while ensuring asset-level protection by selectively and reliably isolating faults. Due to their criticality, protection engineers spend considerable amounts of resources to ensure proposed schemes operate as intended by performing extensive modeling, simulations, and in-field validations. In addition, the host environment may need to be evaluated and hardened in order to attain the expected reliability metrics.

Despite the number of validations, protection misoperations continue to occur, with most of them being attributed to incorrect settings, communication issues, or equipment failures [3]. To help mitigate against equipment-level failures, utilities may implement A) Backup protection schemes — based on time-delayed mechanisms or remote protection schemes; or B) Replica sets — which rely on a redundant set of equipment designed to operate in case the primary fails. However, many of these solutions have an inherent dependency on microprocessor-based relays, which are devices that may be vulnerable to cyber-attacks. Such attacks may cause byzantine faults that could lead to the inhibition of trip actions, a reduction in the selectivity and/or reliability attributes, or enable attackers to initiate trip actions without cause.

Grid-level effects due to a compromised relay (or a set of) are hard to predict, but in general, may lead to A) Loss of load (or generation), B) Equipment damage, C) Loss of synchronism, and D) Cascading failures (if a successful, coordinated attack on protection infrastructure is launched), which may result in wide area disturbances. To help reduce risk, the authors have developed *"ByzSec — A Multi-layered Byzantine Resilient Architecture for Protection Relays*", a composition of technical constructs designed to enable the continuous operation of a protection system despite the presence of compromised nodes (i.e., compromised relays). At its core, the architecture relies on a set of formally proven consensus algorithms, that ensure proposed actions (e.g., a trip command) are the product of a group agreement, providing a natural defense against rogue agents. The platform seeks to offer the following capabilities:

- A highly resilient power system protection architecture based on Byzantine Resilient consensus techniques. The architecture targets protection systems that are deployed in critical asset areas (e.g., high-voltage substations, transmission corridors, or critical loads).
- Aims to reduce adoption barriers by offering a modular design that facilitates integration with existent deployments (i.e., the digital substation environment), while also seeking to minimize operational burdens (e.g., by restricting time delay overheads to less than ¼ cycle).
- Provides system operators with enhanced situational awareness capabilities that can alert of a system's abnormal status, helping to identify threats early on.
- The proposed architecture is intended to offer a solution that can remain functional and secure over a long lifespan, matching the typical, long-term deployment expectations of grid system operators and planners.

During the rest of this document, a discussion on *ByzSec* architecture will be presented. The paper starts by discussing the fundamental constructs of the architecture (Section II), followed by the implementation-specific details (Section III). Results are presented in Section IV, while Section V summarizes the work's conclusions.

II. ByzSec Foundational Background

The ByzSec architecture is built upon a collection of technical constructs that have been specifically selected and assembled to improve the cybersecurity posture of protection schemes. In this section, a high-level overview of key foundational components will be presented, followed by an overview of the proposed architecture that describes its key operational principles.

979-8-3503-1637-7/24/$31.00 ©2024 IEEE

A. Byzantine Fault Tolerance in Grid Applications.

Fault tolerance refers to the ability of a distributed system to continue to make decisions even when some nodes or links are temporarily down. In distributed systems, the term *safety* is used to describe a system that can maintain a consistent, ordered state (e.g., banking transactions occur in the order that they were executed, and reflect the correct balance). In contrast, the term *liveness* is used to describe a system's ability to eventually reproduce the same states across all correct replicas (e.g., all nodes show the same balance). Attaining perfect *safety* and *liveness* properties when links or nodes fail (known as a partitioned system) is a known impossibility of asynchronous networks [4]. This means that if *safety* is preferred, a network-isolated cash machine would not be able to query up-to-date balances, but it would not halt the network. In contrast, if *liveness* is preferred, the network could become deadlocked due to an unresponsive actor, but all correct nodes would agree on the latest consensus state.

BFT builds upon fault-tolerant principles by requiring a system's safety and liveness properties to be upheld even in the presence of a malicious node, a node that may have the intelligence to create systematic disruptions while at the same time appearing honest. Due to this assumption, and the nature of distributed systems, it would be difficult for well-intended entities to reliably self-detect failing or malicious nodes. Hence, special consensus protocols must be developed to handle incorrect, or conflicting messages (from ill-intended actors). Over the last few decades, multiple algorithms (PBFT, HoneyBadgerBFT, RBFT, Stellar, etc.) have been proposed to achieve such functionality with varying levels of success. Among the proposed approaches, Prime is a leader-based (elected), BFT-capable protocol that seeks to maintain a consistent timing performance (with or without an ill-intended agent) by offering a bounded-delay performance guarantee [4]. The bounded-delay performance guarantee is achieved by assigning fixed-length tasks to elected leaders and monitoring its performance across time — which enables nodes to replace the elected leader if unexpected or malicious delays are detected.

B. Spire

Supervisory Control and Data Acquisition (SCADA) systems are a core feature of modern power systems. Due to their criticality, they often adopt fault-tolerant techniques, which may include the use of redundant communication paths, and mirrored or replicated control systems (mostly via a master-slave architecture). However, the vast majority of SCADA systems do not implement modern voting consensus methods or BFT. To help improve the network resiliency properties of SCADA systems, researchers in [1] have developed *Spire*. *Spire* is a prototypical architecture based on *Prime* that embraces two additional constructs, referred as to diversity and proactive recovery, these can be summarized as follows:

Diversity: BFT systems are often characterized by the number of faulty nodes that can be tolerated (f). However, if a system is homogeneous (e.g., the same devices and the operating system) it becomes feasible for a single vulnerability to serve as the common entry point, potentially compromising more than f nodes. Therefore, *Spire* proposes the use of diverse, heterogeneous systems to mitigate against single points of failure.

Proactive recovery: Even if a system implements diversity if enough time is given, it becomes feasible for attackers to eventually find up to f distinct vulnerabilities. However, if these systems are periodically refreshed (via updates or other ways of creating software diversity), or at least, restored to a trusted state then attackers will have to restart their exploitation efforts to compromise more than f nodes within a limited time window. The original Spire paper also performs a technical evaluation on a variety of potential deployments (by varying the numbers of replicas and assuming a varying number of compromised systems), recommending $3f + 2k + 1$ nodes to support up to f Byzantine nodes, and k systems going under proactive recovery.

C. Differential Protection schemes

At their core, differential protection schemes (DPS) rely on Kirchhoff's law to detect imbalances between the current(s) entering an asset versus the current(s) exiting. Due to their relatively simple design, high sensitivity, and fast operational speeds, DPSs are the preferred solution for providing protection to busbars, generators, and transformers via the 87B, 87G, and 87T protection schemes, respectively. In specific, the 87T DPS is used extensively to protect high-voltage and medium-voltage transformer assets due to their high criticality, high cost, and long manufacturing lead times.

A key design factor of the 87T scheme (as many others) is the system's total clearing time (TCT) requirements, which are dictated by factors such as the system's stability limits, the asset's fault handling capacity, and the fault characteristics. However, in general, high voltage networks expect TCTs to be in the order of 3-5 cycles (for instantaneous trips), requiring relay vendors to implement sub-cycle fault detection algorithms that are able to compensate for the relatively long interruption times of breakers and other data acquisition overheads [5].

D. Digital Substations

Many modern protection deployments are migrating away from the traditional, analog point-to-point connectivity model into a digital, network-based communication infrastructure that simplifies data exchanges while enhancing cross-vendor interoperability. Protocols such as *IEC 61850 8-1 Generic Object-Oriented Substation Event (GOOSE)* are now being used to interact with a multitude of Intelligent Electronic Devices (IEDs), enabling for example, a relay to query or trip a circuit breaker. Furthermore, this digitalization has extended to the process bus, which now supports the real-time transmission of waveform data collected from current and voltage transformers via the *IEC 61850 9-2 Sample Values (SV) protocol.* Timing signal properties (e.g., resolution) have also been improved by protocols such as PTP (Precision Time Protocol). In addition to the benefits of interoperability, a digital substation may enable the deployment of digital situational awareness solutions that enhance an operator's visibility. For example, Machine Learning (ML) techniques may be used to analyze traffic patterns, helping to identify anomalous behaviors early on.

E. ByzSec's Architectural Overview

As outlined during the introduction, the *ByzSec* architecture seeks to incorporate BR mechanisms to enhance the security of the trip decision process. At its core, the solution employs an array of *Byzantine Resilient relay nodes,* that use custom-built consensus protocols (*Arbiter* and *Peer*) to attain a global "state" that represents the relays' agreed action. This "state" can then be used to decide if the breaker needs to trip, close, or remain unchanged. Within ByzSec, consensus is achieved by deploying one of the two consensus protocols that have been specifically designed to operate in time-critical protection environments, enabling consensus to be reached within a ~¼ cycle.

A graphical overview of the proposed solution is presented in *Fig. 1*. The design expands upon the typical relay-breaker arrangement to host the *Byzantine-Resilient relay nodes*, which are a set of traditional relays that have been paired with a *BFT harness.* The *BFT harness* is a software-based adapter that reads and interprets GOOSE packets coming from the attached relay and forwards its intended trip state to the consensus network. At the other end, a destination proxy listens to the consensus network and determines the action that must be executed by the circuit breaker. Due to the modular architecture, different timing and BFT guarantees can be achieved depending on the deployed consensus protocol.

In addition to the BFT-oriented components, the platform implements a situational awareness (SA) tool. The SA tool is an operator-oriented dashboard system that collects data from multiple tools embedded within ByzSec. Among other tools, it includes a real-time state estimator and a transient event monitoring tool that can be used to identify potential power system issues. In addition, an ML-based tool has been developed to provide visibility into the consensus processes. Further details of these tools will be presented in the next section.

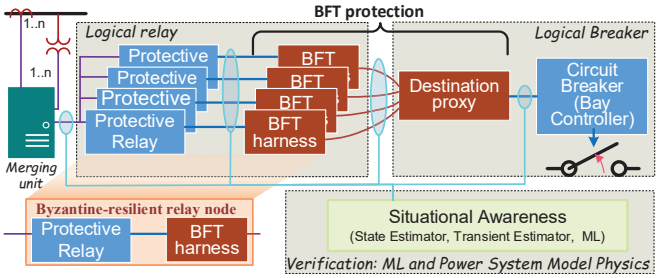

Fig. 1 A high-level overview of the ByzSec architecture. The solution maintains the traditional relay-breaker arrangement while significantly enhancing its cybersecurity posture.

III. ByzSec Implementation Overview

A. *Byzantine resilience for Time-Critical Protection Applications*

BFT protocols such as Prime are often referred to as classical state machine replication (SMR) models. SMR-based approaches are an ideal solution for systems whose current state is dependent on a series of past system states (e.g., an account balance is a product of past deposits and withdrawals). However, due to implicit requirements of maintaining a time-ordered sequence of events, SMRs typically introduce large processing delays. On the other hand, protection systems can only have two, mutually exclusive states (i.e., tripped or closed) leading to only two state transitions (i.e., trip, close). Based on this observation, the ByzSec architecture offers two distinct consensus protocols that are implemented using finite-state machine models (instead of SMR), due to their simplified architecture timing delays have been reduced significantly. Extensive discussions into these protocols (named the *peer* and *arbiter protocol*) can be found in [6]. However, these can be summarized as follows:

The Arbiter protocol: In the arbiter protocol, a breaker node collects and validates signed messages emitted by independent relay nodes to obtain a global decision based on a threshold count. Each signed message contains the relay's intended action, as well as a time-based nonce to prevent replay attacks. As part of the verification, the breaker node asserts the liveness of the messages (by enforcing a 1 ms validity window). High-resolution timing comparisons are possible due to PTP.

The Peer protocol: In contrast to the arbiter protocol which relies on a single breaker node to determine the breaker state (which may introduce a single point of failure), the peer protocol implements a fully distributed consensus approach based on *Spire* that offers stringent BFT guarantees. The BFT network requires the presence of $2f + k + 1$, independent actors to operate, where f, represents the number of compromised devices (relay nodes), and k represents the maximum number of devices undergoing an active recovery process. Each node implements a custom-built state machine that dictates its behavior (see Fig. 2). Transitions within the state machine are executed once a predefined threshold is reached (e.g., $f + 1$ messages agree upon the next state). To ensure the security of the threshold evaluation mechanism, the solution relies on threshold cryptography to validate and prove agreement among the distinct messages.

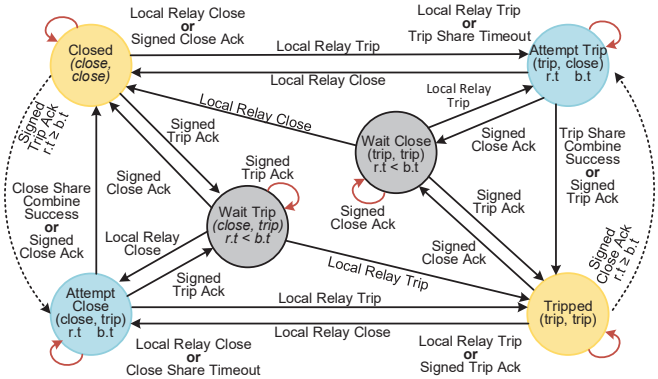

Fig. 2 The peer's protocol state machine, adapted from [6]

B. *Protection relays*

As defined above, the peer protocol requires a minimum of $2f + k + 1$ independent nodes to operate. Hence if a system must tolerate $f = 1$ (a compromised relay node), and $k = 1$ (a relay node undergoing proactive recovery), then a minimum of 4 independent relay nodes are required to attain BFT guarantees. To help organizations manage costs, ByzSec can use a mixture of *physical* and *software-based* relays to support the BR requirements (refer to Fig. 3, which shows the "as built" topology). Relay components can be summarized as follows:

Software-based Relay: A custom-built, software-based relay has been developed by the PNNL team. The codebase executes over a preemptive, real-time Linux-based kernel. The system has been extensively tested to ensure its reliability and to ensure it detects and operates under a ¼ cycle during an instantaneous trip condition. The relay was guided by the IEEE C37.91-2008 recommendations, implementing a dual-slope protection curve with optional third, and fifth-harmonic cross-blocking mechanisms (which prevent false trips during transformer inrush events). The relay settings can be configured via a user-provided JSON file, and can communicate via IEC 61850 9-2, IEC 61850 8-1 through the LibIEC61850 library [7].

Hardware relays: Three distinct, hardware-based relays from different vendors have been integrated into the platform. The relays are unmodified, commercially available units that have been configured to operate under the 87T scheme, using standard protection settings. A network-isolated, timer-based power supply unit has been installed at the power receptacle to provide a periodic, out-of-band disconnect-reconnect capability (to trigger a full system restart and enable proactive recovery).

C. *Situational Awareness components*

Steady-State State Estimator: The platform implements a positive-, real-time state estimator that relies on the SV data stream to compute a physics-driven representation of the power system. In addition to the local substation measurements, the state estimator is assumed to have access to nearby nodes and generator states (e.g., the network equivalents). The state estimator continually asserts the system's health status by comparing the estimated state versus the locally observed state.

Transient state health monitor: To complement some of the limitations associated with steady-state estimators during transient events, an ML-based transient event classifier has been implemented. The transient classifier relies on a convolutional neural network that constantly analyzes waveform-level data to identify and pinpoint a fault location. The design is based on the work previously published in [8], and can be used to verify the presence and location of a fault (due to its complexity results are delayed by 4-8 seconds).

ML-Based Situational awareness: In addition to the physics-oriented supervisors, the ML-based component analyzes the network traffic present in the substation. It has been specifically designed to detect anomalies within the consensus protocol by inspecting and comparing 1) The GOOSE-based relay signals, 2) The encrypted consensus network packets; 3) The GOOSE message sent to the circuit breaker; and 4) The breaker output.

IV. ByzSec Prototype Environment and Testing

To help validate the proposed approach, a full-scale, self-contained, mobile environment was built. The environment enables the execution of a wide variety of network and physics-based tests that have allowed for extensive functional and red-team evaluations. The design (see Fig. 3) follows a modular approach that is intended to simplify future adoption, but also to offer a multi-layered defense system that 1) Prevent attacks, by adopting basic network security practices as well as specialized techniques such as diversity, 2) Provides operators with situational awareness tools that can detect cyber-physical anomalies in real-time, and 3) Use BFT methods to ensure the system operates as intended even if the defenses fail. To help better illustrate the capabilities of the architecture, this section summarizes key tests carried out during the validation phase.

D. *Grid model*

To help validate the performance of the proposed solution a high-fidelity, electromagnetic transient (EMT) simulation based on the reference model reported in [9] was implemented within the Opal-RT simulator. Following an extensive validation procedure (based on the outputs given by [9]), the model was scaled to a 345kV base (from the original 230kV system) and a 345/34,5 kV $\Delta - Y$ transformer that supplies a 100+25i MV load was added. The resulting 5-node, 3-machine model (shown in Fig. 4) was then used as the reference model for all subsequent grid tests.

E. *87T protection operational tests*

Based on the Opal-RT grid model, a series of tests were performed to ensure the protective relays and situational awareness modules are in accordance with the simulated grid conditions (e.g., relay actions match the fault type and locations). Waveform data was transmitted into the platform by using a *SV protocol* adapter that was developed by Opal-RT. The adapter transmits three-phase current and voltage signals that are sampled at 4800 Hz. In addition to the functional tests, fault signals were recorded and can be re-played by using a custom-built SV-player that allows the platform to operate without the Opal-RT simulator (useful for mobile demonstrations).

F. *Consensus*

As part of the evaluation phase, the team performed a series of performance tests to ensure the consensus layer operated as intended and within the allotted time bounds. The tests aim to capture the system's performance during worst-case conditions (e.g., when a node is compromised, and proactive recovery is active). The results indicate that no performance violations were detected within the 1-million sample size, for each operational condition (see Table I).

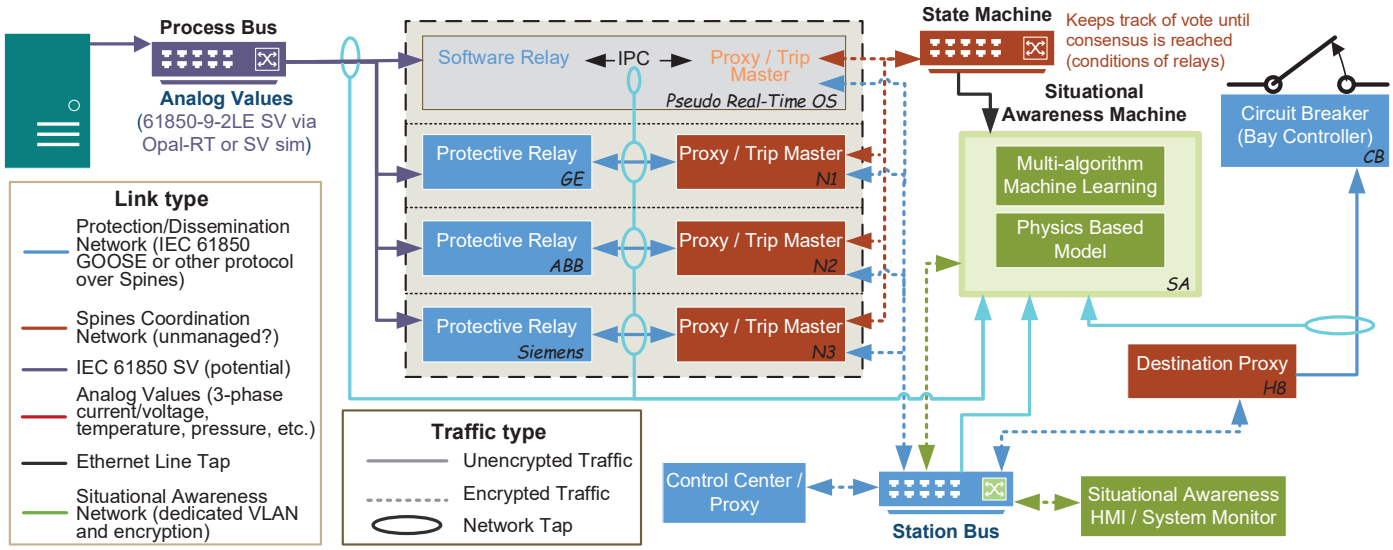

Fig. 3 The ByzSec architecture, an intrusion-tolerant design that uses modern fault-tolerance mechanisms to provide microprocessor-based protection systems with BFT-capabilities. It successfully demonstrates how a network of BFT-enabled relays can be used to ensure the reliability and selectivity of the protection system even when a malicious actor is present within the system.

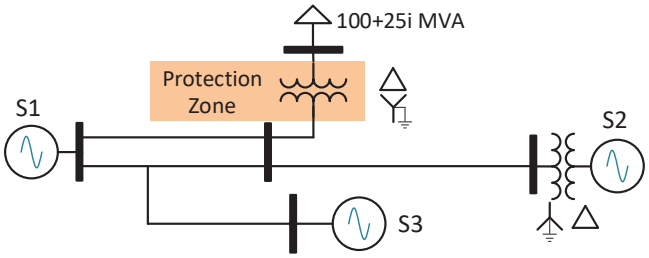

Fig. 4 A simplified view of the grid model used in this paper, electrical parameters derived from [9].

Table I. Statistics on the processing delay in the peer protocol *[2]*.

Operating Condition	Real-Time Kernel (microseconds)		
	Minimum	Average	Maximum
Fault-Free (Normal)	1637	1950	3596
Fail-Stop Fault or Proactive Recovery	1608	1976	3726
Fail-Stop Fault and Proactive Recovery	1750	2015	3996
Byzantine Fault	1665	1984	4002
Byzantine Fault and Proactive Recovery	1767	2019	4101

G. *Purple testing*

As mentioned in the introduction, ByzSec aims to achieve a long-term, functionally secure environment for protection systems. In accordance with this principle, an independent purple-team evaluation was performed by researchers at Sandia National Lab. Purple testing is an evolution of red-team testing in which developers and the evaluation team work together to facilitate vulnerability discovery and exploitation. Full results are intended to be published in a separate report, however, some relevant findings include: A) The consensus layer performs as originally designed (it satisfies the fault tolerance guarantees); B) Rogue instances could join the consensus network but no visible impact was detected; C) Replay, DoS, and DDoS attacks were not successful; D) Attempts to manipulate encrypted packet were unsuccessful; E) Encrypted UDP packets can be dropped; F) In general Cybersecurity best practices are implemented and software is functioning as expected

H. *Situational awareness*

ByzSec offers a web-based dashboard that can be used to monitor the system state. One of its core features is providing visibility into each of the Byzantine-enabled relays, enabling systems operators to quickly identify subsystems that are failing or in disagreement with other related systems (see Fig. 5).

V. Conclusions and Future Work

In this work, a highly resilient power system protection architecture based on modern Byzantine resilient (BR) consensus techniques was presented. The proposed approach is able to meet the operational requirements of modern protection systems while offering a design that minimizes potential adoption barriers. The design has been rigorously tested using a variety of subject-specific evaluations to ensure the protection engineer's expected reliability and selectivity characteristics are maintained at all times. In addition to fulfilling the design goals, this work may also serve as a turning point in the adoption of BR-based protocols, which are usually considered unfit for grid applications due to their potential overheads. However, solution developers must still employ validation and benchmarking to justify their decisions (i.e., replicate the work presented in [10]). In conclusion, to the best of our knowledge, this paper represents one of the most mature applications of BR in the context of high-speed grid applications, with most other published work remaining theoretical in nature or with limited testing [11].

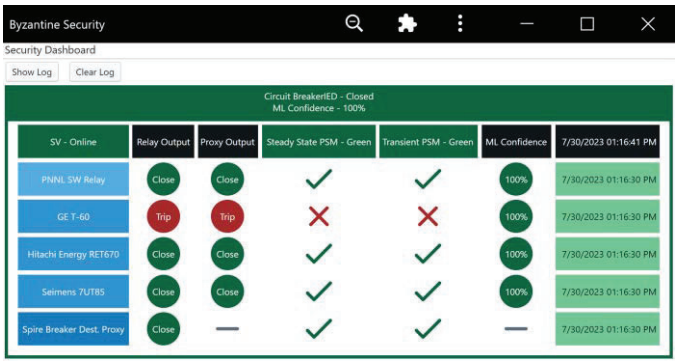

Fig. 5 ByzSec's situational awareness dashboard, in this case, the first physical relay is in direct contradiction with other systems.

Acknowledgments: The results of this project have been developed over a period of 3+ years, with participation from engineering teams from Hitachi-ABB, GE, and Siemens. Support has also been received from researchers at Sandia National Labs (red team testing), Lawrence Berkeley National Labs (BFT evaluations), and current and former staff at PNNL (R&D). Special thanks to the Distributed Systems and Networks lab at the Department of Computer Science at JHU. This work has been supported by the Department of Energy, under the Energy's Grid Modernization Initiative.

VI. References

[1] A. Babay, J. Schultz, T. Tantillo, S. Beckley, E. Jordan, K. Ruddell, K. Jordan and Y. Amir, "Deploying Intrusion-Tolerant SCADA for the Power Grid," in *49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, 2019.

[2] S. Bommareddy, M. Khan, D. J. Sebastian Cardenas, C. Miller, C. Bonebrake, Y. Amir and A. Babay, "Real-Time Byzantine Resilient Power Grid Infrastructure: Evaluation and Trade-offs," in *43rd IEEE Real-Time Systems Symposium*, 2022.

[3] North American Electric Reliability Corporation, "Technical Assessment of 2022 Bulk Power System Performance," 2023.

[4] Y. Amir, B. Coan, J. Kirsch and J. Lane, "Prime: Byzantine Replication under Attack," *Transactions on Dependable and Secure Computing,* 2010.

[5] S. Meier, T. Werner and C. Popescu-Cirstucescu, "Performance considerations in digital substations," in *13th International Conference on Development in Power System Protection*, 2016.

[6] S. Bommareddy, D. Qian, C. Bonebrake, P. Skare and Y. Amir, "Real-time byzantine resilience for power grid substations," in *41st International Symposium on Reliable Distributed Systems (SRDS)*, 2022.

[7] Michael Zillgith, MZ Automation GmbH, "libIEC61850 open source library for IEC 61850," [Online]. [Accessed Nov 2020].

[8] T. McDermott, N. Shepard, M. AP, R. M, J. Doty and J. Kolln, "Protection of Distribution Circuits with High Penetration of Solar PV," PNNL-32230, 2021.

[9] IEEE PES Power System Relaying and Control Committe , "EMTP Reference Models for Transmission Line Relay Testing," 2005.

[10] J. R. Clavin, Y. Huang, X. Wang, P. M. Prakash, S. Duan, J. Wang and S. Peisert, "A Framework for Evaluating BFT," in *IEEE 27th International Conference on Parallel and Distributed Systems (ICPADS)*, 2021 .

[11] N. Jacobs, A. Summers, S. Hossain-McKenzie, D. Calzada, H. Li, Z. Mao, C. Goes, K. Davis and K. Shetye, "Next-Generation Relay Voting Scheme Design Leveraging Consensus Algorithms," in *2021 IEEE Power and Energy Conference at Illinois (PECI)*, 2021.