How to Successfully Isolate and Shield In-Service Communications-Assisted Protection From Ethernet Packet Storms

Wesley Roberto, Eduardo Goncalves, Paulo Lima, Thiago Bordim, and David Dolezilek
Schweitzer Engineering Laboratories, Inc.

Presented at the
78th Annual Conference for Protective Relay Engineers at Texas A&M
College Station, Texas
March 31–April 3, 2025

How to Successfully Isolate and Shield In-Service Communications-Assisted Protection From Ethernet Packet Storms

Wesley Roberto, Eduardo Goncalves, Paulo Lima, Thiago Bordim, and David Dolezilek,
Schweitzer Engineering Laboratories, Inc.

***Abstract*—The dependability of a communications-assisted protection system is the ability to operate correctly every time it is required. At the same time, security is the ability to prevent unintended operations and allow the application to continue to function as intended even when the communications system has become degraded due to a data flow problem. In communications-assisted schemes, one of the major design criteria is the high availability of protection signals within the data flow shared among devices throughout the communications network. Message prioritization reduces the likelihood of increased average latency; however, while Ethernet, as defined by IEEE 802.1 standards, has best-effort classifications to prioritize treatment of one packet over another, it does not support the quality of delivery or determinism. It is possible to assign Class of Service (CoS) priority to packets to guarantee preferred treatment within a switch, but without restrictions on traffic generation and bandwidth consumption, priorities cannot prevent frame losses by buffer overflow and bandwidth saturation. Those responsible for the design of IEC 61850 networks face these and other challenges in maintaining performance, resilience, and cybersecurity.**

Considering these definitions, it is important to design systems with dependability and security to maximize performance. This paper discusses the root cause analysis of a real-world Ethernet network storm event and illustrates how a newly installed software-defined network was instrumental in preventing the total collapse of an industrial control system. By addressing methods to dramatically improve data flow availability and determinism as well as message prioritization, the paper highlights the importance of robust network design in ensuring the dependability of protection systems.

I. Introduction

The International Electrotechnical Vocabulary standard, IEC 60050 [1], defines dependability of a protection system as the ability to perform as and when required. In most cases, the term is used as an umbrella term to express its core attributes of reliability, maintainability, supportability, and the resulting availability. The network engineering Technical Report IEC 61850-90-4 [2], in Clause 11.1 about network performance metrics, states that dependability is a term more commonly used in the protection community, while the information technology (IT) network community frequently refers to it as quality of service (QoS). Both terms refer to the collective effect of service performance but not the means of achieving it. Among the techniques employed to better achieve dependability (or QoS) is Class of Service (CoS), which is the process of assigning priority levels to different types of data in a network to improve performance of important data at the expense of less important data.

It is important to note that operational technology (OT) practices support guaranteed data flow performance QoS and the classification of one packet as more important than another to prioritize the processing of the former over the latter. However, the IT network community uses the terms QoS and CoS interchangeably. With respect to mission-critical, time-sensitive data flow, the IT use of QoS to describe IT methods of CoS is incorrect because it does not dedicate resources to guarantee quality or dependability and truly refers to CoS in shared bandwidth systems typically used in substations and offices alike. The IT CoS refers to methods that classify packets containing all or part of a data flow message with various priorities to allocate shared resources when bandwidth is limited. The digital process bus Generic Object-Oriented Substation Event (GOOSE), Sampled Values (SV), and Precision Time Protocol (PTP) messages fit within a single Ethernet frame; for this paper, message and frame are interchangeable. While specialized Ethernet services that provide dedicated bandwidth for data flows exist, such as Ethernet private lines or virtual synchronous networking, they are rarely, if ever, deployed within a substation where shared bandwidth methods are typical. More recently, IT methods of shared bandwidth and topology management via Spanning Tree Algorithms (STAs) supported by constant Rapid Spanning Tree Protocol (RSTP) messages are being replaced by software-defined networking (SDN) used in OT environments (OT SDN).

Unlike OT SDN devices, IT Ethernet network devices are not configurable to guarantee engineered data flow. Even though protection methodology accurately classifies IEEE 802.1Q as CoS, IT personnel refer to it as QoS. Since the terms CoS, QoS, and dependability are widely used in the Ethernet networking community and some discussions of protection applications, this paper will use these three terms with their appropriate definitions.

Section 11.3.3.3 of the communications requirements part of IEC 61850-5 [3] defines the dependability requirement (D) in terms of how the communications network affects protection schemes. This concept refers to dependability in regard to missing commands, such as the loss of trip messages with a protection scheme. Given that P_{mc} is the probability of missing

commands, the dependability is defined as $D = 1 - P_{mc}$. Table I from IEC 61850-5 presents the dependability classes.

TABLE I
DEPENDABILITY CLASSES

Dependability Class: $D = 1 - P_{mc}$		P_{mc}
D1	Low	10^{-2}
D2	Medium	10^{-3}
D3	High	10^{-4}
D4	Very High	10^{-5}

While the dependability classes establish statistical requirements for all types of communications networks, the communications aspects can be divided into individual components to enable further performance metrics to be defined. Latency is another critical metric; if a message is delayed, it loses its usefulness, and being overly late can be even worse for a protection scheme, as it might process it as a valid signal after previously taking action in its absence. In this context, IEC 61850 defines transfer time and its dependence on both intelligent electronic device (IED) processing and network device dwell times. The dwell time is the duration of time a message exists within the confines of a device or application as it is passing through to a destination. Transfer time is the complete transmission delay from the sending application to the receiving application, including all necessary coding, decoding, and media access at both ends, as shown in Fig. 1. Application Function 1 in Physical Device 1 (PD1) transmits data to Application Function 2 in Physical Device 2 (PD2). The timing starts when the sender places the data on its transmission stack and ends when the receiver retrieves it from its reception stack. The total transfer time t comprises the individual times for coding (t_a), decoding (t_c), and the actual network transfer time (t_b), regardless of whether dedicated communications processors are used.

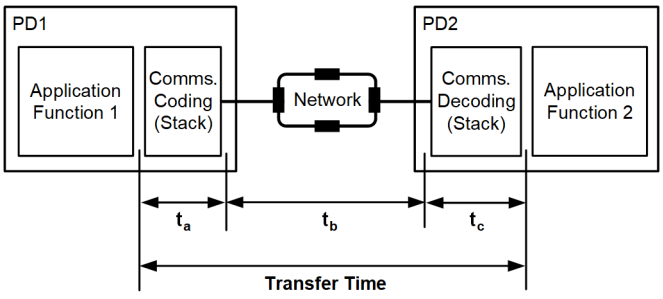

Fig. 1. Transfer time of command message between devices.

Often, for very fast and precise applications, it is necessary to remove the ambiguity of the noncontiguous data flow presented in the IEC 61850 illustration in Fig. 1 in which the logic blocks do not connect to each other or the interfaces do not connect to the external world, represented by the edges of the physical device box. Fig. 2 illustrates numerous time durations under consideration in the update of IEEE 1646 [4] as a part of the data flow of a signal entering the sending IED data processing function *F* and resulting in an operation because of the data processing function *F* in the receiver IED. Dwell times of hardware devices in the data path include the time that the message resides within the switches *j* and *k* as well as hardware within the potential fiber or radio wide-area network (WAN). Dwell times of communications processing functions are illustrated as the time durations *a* and *e*. The time duration *a* represents the latency of the sender IED communications processor (CP) encoding and publishing a signal message. The time duration *e* represents the latency of the receiver IED CP subscribing to and decoding a signal message.

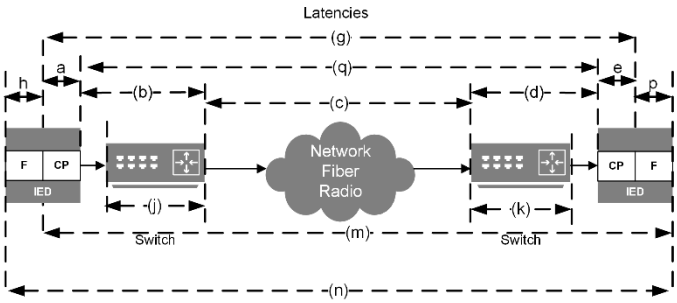

Fig. 2. Time latencies of signal processing and transfer via digital messages between devices under consideration in IEEE 1646.

As an example of the command latency requirements, consider IEC 61850-90-1 [5] Section 6.4.2.1, which outlines types of message performance classes. A Type 1 high-speed message typically contains basic binary information, like a short command or a simple message about the function (e.g., trip, close, reclose order, start, stop, block, unblock, state change, or current state). Type 1 messages characterize critical applications, requiring a receiving IED to act immediately upon receipt. Type 1 messages are further categorized, with the trip messages designated as Type 1A, and recognized as the most important fast binary messages in the substation. For intersubstation communications, the block and release messages might have the same importance due to their association with trip applications. IEC 61850-5 sets the transfer time requirements for Type 1A messages, depending on the application. It specifies that the transfer time, $ta + tb + tc$, of trip messages inside a substation shall not exceed 3 ms, while trip messages between substations shall not exceed 10 ms.

When planning to meet a transfer time of 10 ms, the D4 dependability class mandates that the probability of a command exceeding this transfer time must be less than 10-5, equating to no more than once in every 100,000 commands sent. Ethernet network disturbances or device port hot-standby failover processing may disrupt the command data flow path during reconfiguration or reconnection. Any subsequent loss of messages or delays contributing to a transfer time greater than 10 ms even once within 100,000 command messages sent is a failure to meet the dependability metric.

Traffic prioritization schemes used in Ethernet, which is defined as a best-effort technology, using CoS allow users to prioritize protection messages in an effort to minimize the average latency, but prioritization alone does not guarantee compliance with the transfer time and dependability classes of IEC 61850-5. Clause 6.4.8.2 in IEC 61850-90-4 states that without restriction on the traffic generation rate of all

participating devices in a network, it is impossible to prevent message losses due to buffer overflow. Buffer overflow can occur, for example, when an IED generates large amounts of multicast traffic due to an internal failure, when several IEDs generate large volumes of legitimate traffic in response to an event, when a loop exists in a network and rebroadcasts old traffic, or through a denial-of-service cyber attack. Sophisticated traffic throttling techniques also available in OT SDN can ensure application-specific throughput, maintaining dependability even under peak network load conditions. Low-priority frames are selectively dropped during bandwidth saturation, thereby prioritizing and forwarding high-priority traffic with minimal latency.

Building on the presented network concepts and metrics, this paper examines a real-world Ethernet network storm event that occurred within an industrial control system and showcases how the application of OT SDN successfully isolated and shielded the packet storm, preventing it from spreading and causing the total collapse of the supervision, control, and protection network. IEC 61850-5 presents the concept that message dependability in a specific design is negatively influenced by the probability of a device not receiving commands as the probability of missing commands (P_{mc}). This paper demonstrates how SDN (used interchangeably with OT SDN) minimizes the probability of missing commands (P_{mc}), which, in turn, maximizes the performance to meet the dependability and the transfer time classes for the stringent requirements, especially for large networks. This is achieved through:

- Segregation of Layer 2 networks.
- Granular, multilayer, distributed traffic prioritization and control on each network switch.
- Differentiating services by applying traffic throttling to each data flow, with a focus on IEC 61850 GOOSE traffic.
- Improved cybersecurity through deny-by-default access control policy and support of discard and alarming policies for unexpected traffic on the network.

II. OT SDN Foundations

In traditional Ethernet networks, switches operate on two planes: the data plane and the control plane. The data plane is responsible for receiving and forwarding Ethernet frames. The control plane is responsible for deciding how frames are forwarded through internal management algorithms, such as the media access control (MAC) table, STAs supported by the RSTP, segregation by virtual local-area networks (VLANs), and prioritization, as illustrated in the top diagram of Fig. 3. In traditional IT switches, these algorithms provide convenience and connectivity through automatic network management. The MAC tables determine the forwarding ports, while the RSTP takes care of network convergence in case of path failures. However, the ease of operation comes at the cost of less control, determinism, and lower resiliency against cybersecurity threats, granting it an allow-by-default policy. Settings, such as bridge priority, path cost, and MAC filtering, can be used to influence the network behavior to solve these dynamic algorithms in specific ways and to block traffic from specific segments; however, these are rarely used. These settings require manual configuration and change data flow behavior dynamically, and perhaps unexpectedly, during network disturbances. Additionally, each switch is configured individually, which increases complexity and creates maintainability challenges as the network grows.

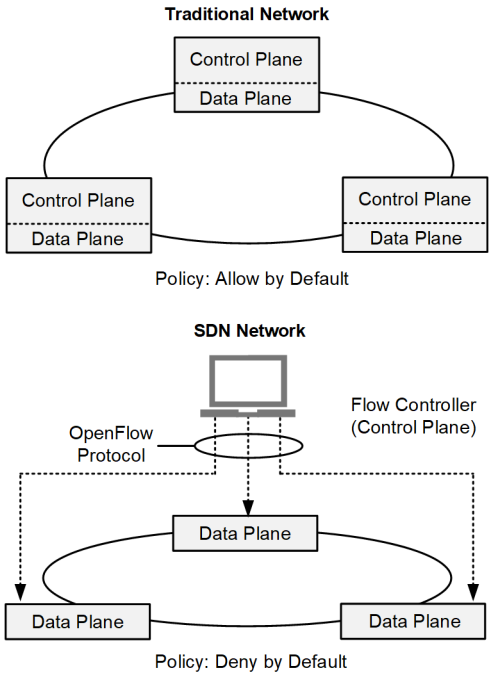

Fig. 3. Traditional versus SDN network architecture.

In contrast, SDN switches apply only the specific data flow rules designed for and implemented in the data plane. The control plane for all switches is centralized in a software application flow controller, as shown in the diagram on the bottom in Fig. 3. The interface between the flow controller and the SDN switches occurs through the OpenFlow protocol, standardized by the Open Networking Foundation (ONF) [6] and now managed by the Linux Foundation. In the controller, the control logic is defined to create rules to inform switches when and how the Ethernet frames should be forwarded. Flows are programmed, which consist of rules applied at ingress to every packet on every port of every switch, a set of forwarding instructions, and other optional actions. The rules are multilayered, corresponding to the main fields of the protocol's implementation layers: ingress port, Layer 1; Ethernet header, Layer 2; IP header, Layer 3; and Transmission Control Protocol (TCP)/User Datagram Protocol (UDP) header, Layer 4. The proactive approach of traffic engineering flows makes the flow controller optional in the network once the SDN switches are configured, allowing it to be removed to further enhance cybersecurity by eliminating the attack surface of the centralized management.

For a better understanding of the pre-engineered and precise principle of operation of SDN networks, consider the example with the GOOSE messages in Fig. 4. The controller has configured the switch with the flow rules illustrated in the flow

table. When the GOOSE 1 message is received on Port 1, it is inspected and its contents compared against a lookup table of rules, or flow table. The entry finds a match with Rule 1, and the switch forwards it to Port 3, as instructed in the output port. When the GOOSE 2 message arrives on Port 2, it is inspected and compared against the same lookup table. Although, in this case, the message contents find no exact match with Rule 1, but they do match with Rule 2, a generic and empty rule, preconfigured to discard unplanned network traffic. Additional rules can be applied so that instead of simply discarding this unexpected traffic, it can be sent to an intrusion detection system (IDS) for further analysis.

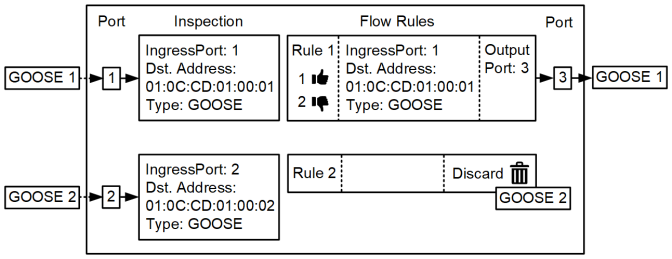

Fig. 4. Example of SDN switch operation.

Though Ethernet, by design and definition, is inherently nondeterministic, pre-engineered data flow rules in OT SDN provide more discrete control of the data flows. On an OT SDN workflow, all network traffic is preconfigured in a centralized controller and all unplanned traffic is discarded from the network or sent to an IDS, which grants the SDN a more predictable, controlled, and secure approach, giving it a deny-by-default policy. Additionally, the centralized controller enables the flow configuration to focus on the application rather than on each individual switch, which simplifies the configuration process, and it is much less impacted by the network size.

III. Case Study

The system under analysis is a large facility of a multinational company in Brazil. It has protection and control IEDs; controllers; and a local supervision, control, and engineering access system. The architecture is abstracted in this paper to keep the security of information related to the installation.

Fig. 5 illustrates the legacy network architecture of the local supervision and control system. The complete network of the system is made of several substations (SS1–SS10) interconnected in two large rings. The RSTP manages the entire topology, forming a large Layer 2 domain. Some of these substations still had unmanaged switches. Each substation makes an additional ring with an IED panel, also managed by the RSTP. This network presented management and performance challenges due to its dimensions and complexity, and for this reason, the company invested in the retrofit of the entire network.

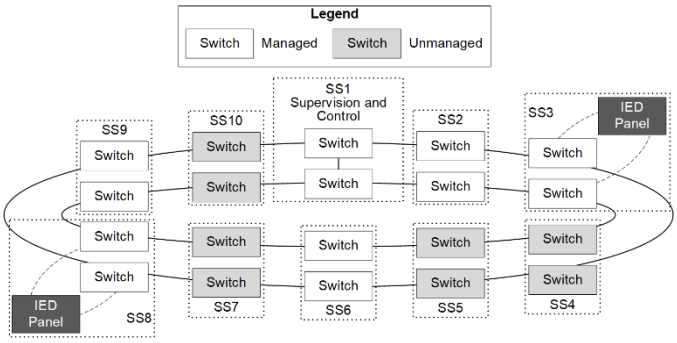

Fig. 5. Legacy network architecture.

A. Best-Known Methods (BKMs) Versus the Existing Design

Like most Ethernet networks in use today, the existing protection and control Ethernet switches were restricted to the IT mechanisms of traffic management. IT switch filtering of Layer 2 GOOSE messages is limited to MAC filtering on the switch egress and VLAN filtering on the switch ingress. IEC 61850 and other technologies do not require appropriate configuration of either of these parameters. The BKMs for configuring mission-critical Layer 2 messages like GOOSE based on IEC and IEEE standards are as follows, as quoted from [7]:

- Assign each GOOSE message a unique VLAN based on IEEE 802.1Q, referred to as a QVLAN. [When this is not possible, carefully group alike GOOSE messages into a single VLAN.]
- Assign each GOOSE message a unique [destination] MAC address.
- Assign each GOOSE message a unique application identifier [(AppID)].
- Assign a [short but] descriptive GOOSE identifier (GOOSE ID) rather than generic IDs in the IED to improve documentation and troubleshooting.
- Label GOOSE message payload contents with [short] descriptive names, rather than generic names, in the IED to improve documentation and troubleshooting.
- Carefully design payload size and contents to facilitate appropriate GOOSE application processing [(GAP)]—*mind the GAP*.
- Carefully choose IEDs that process incoming GOOSE messages appropriately fast for protection-class applications—*mind the GAP*.
- Do not publish multicast messages on the network without QVLAN tags.
- Disable all unused … [device] communications ports.
- Monitor GOOSE message attributes to derive the quality of the message [reception at each subscriber].
- Use the GOOSE attributes of sequence number and state number to determine if all wanted messages reach the receiver.
- Monitor, record, and alarm failed GOOSE message receptions.
- Provide GOOSE reports with configuration, status information, and statistics pertaining to GOOSE messages being published and subscribed by the IED.

- Record and alarm failed quality of GOOSE messages for use in local and remote applications.
- Display status of GOOSE subscriptions and alert operators of failure.
- Configure each switch port to block the ingress of unwanted and allow wanted multicast messages via VLAN and MAC filtering. This reduces the multicast traffic through the network to only that which is required.
- Configure each switch port to block the egress of unwanted and [only] allow wanted multicast messages via VLAN and MAC filtering. This prevents unwanted messages from reaching the IEDs.
- Use switches designed for rugged environments and Layer 2 multicast among … IEDs in a fixed address network.
- Do not allow dynamic [IED data model and reporting] reconfiguration; this leads to systems different than commissioned.
- Use switches that provide real-time status of traffic behavior and network configuration [7].

Careful consideration must be used when building a network based on IT switches, which are categorized for use in conditioned environments without ruggedized power supplies and with spanning tree algorithm resolution too slow for a mission-critical protection signal transfer. The switch interconnection topology is most important, and influencing the spanning tree algorithm is done by setting the bridge priority and root path cost in each switch correctly. However, the switch network designers are often unaware of the need for message segregation among the IEDs. OT SDN Ethernet networks are simpler because they use fast static instructions in lookup tables based on a pre-engineered data flow. This allows more granular message filtering. Even more important to the network behavior is the performance of the fault mitigation of the dynamic spanning tree decisions [8].

A big advantage in this case is that OT SDN traffic engineering is done in advance and remains in the as-commissioned state unless someone intentionally uses the SDN controller to change rules. It is not possible to bypass the flow controller software application, log into a switch, and use a web interface to accidentally or intentionally change settings and, therefore, behavior. Similarly, it is not possible to accidentally remotely log into the incorrect switch. Also, the deny-by-default of unwanted traffic prevents the propagation of unanticipated messages. As the previous list of BKMs suggests, it is possible to create settings for every IED and every IT switch, configure and test them in the field to safely manage Layer 2 command and protection messages, and then supervise that no unintended changes are made in switch settings. However, with a network of IT switches, this monumental task also requires the knowledge and desire to do so.

IT traffic engineering shortcuts, like grouping many or all GOOSE messages on the same VLAN, may appear to work when commissioning testing; however, these practices will often put constant bandwidth management stress on the network and devices, which will not be evident without detailed scrutiny. This constant stress on the system may go undetected for lengthy periods of time until even the smallest disturbance can create a network communications failure.

BKMs are suggestions, not requirements, because the relevant standards and technical reports do not clarify specific parameter values and are often ambiguous or silent on the use of settings and their values. The multilayer inspection of packets for OT SDN match rules provides more detailed differentiation between packets and their intended destination. This not only aids the design of a system from scratch but also provides superior protection to an existing network when changes are being made.

OT SDN matches are made with any of the contents from Layer 1 to Layer 4, whereas in this case study, the existing network was built with switches using traditional IT methods. The topology appeared to be chosen due to geographic convenience and was not set to optimize bridge priority or bridge path. More importantly, the GOOSE messages were not given unique VLANs or even grouped into logical VLANs based on their purpose. All GOOSE messages had the same VLAN setting, which would force every IT switch to allow them to propagate all over the network and prevent any traffic throttling or bandwidth management; however, in this case, the situation was made worse by the selection of VLAN-unaware designation within the switches and the value of zero for the VLAN identifier in each GOOSE message.

Based on the IEEE 802.1 standards, setting the VLAN identifier within a message to zero causes the VLAN to be ignored by IT Ethernet switches. Setting switches to be VLAN-unaware accomplishes the same behavior, and both situations prevent any segregation. Essentially, every GOOSE message was being delivered to every IED and network segment in the system, creating constant stress on the system. These two settings created a situation that was not visible where every GOOSE message was being sent everywhere all at once. This stress remained persistent and undetected until an IED issue created a message storm that caused multiple failures. Further, the choice of a VLAN identifier of zero meant that no IT switches could be added or modified to provide segregation.

B. Proposed Architecture Using OT SDN

Fig. 6 illustrates the proposed architecture for the system retrofit. All unmanaged switches were replaced with managed switches, adding traffic control and prioritization. The architecture was changed from a double ring to a pseudo-ladder made to match the preferred ladder topology as closely as possible, which optimizes network performance [8]. The term ladder comes from the similarity with a real ladder. The pseudo assignment is because the SDN switch transparently forwards the RSTP Bridge Protocol Data Unit messages received in one ring direction to the other and vice versa, not actively participating in convergence, as occurs in ladder architectures. What was a single Layer 2 domain becomes smaller domains for each substation interconnected by the SDN switch. Each SDN switch enables traffic segregation of GOOSE messages via rules based on Ethernet packet parameters other than the VLAN identifier.

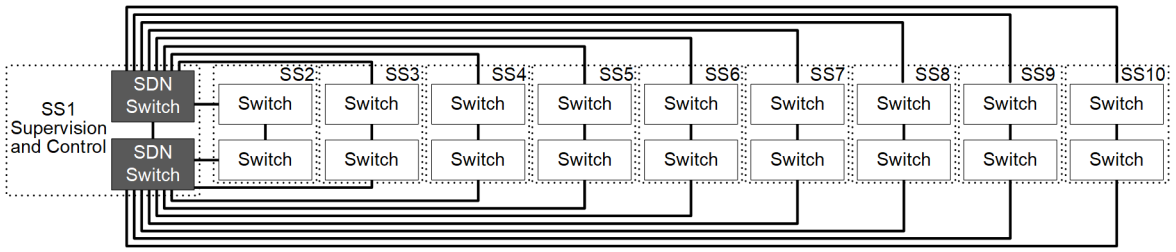

Fig. 6. Proposed architecture.

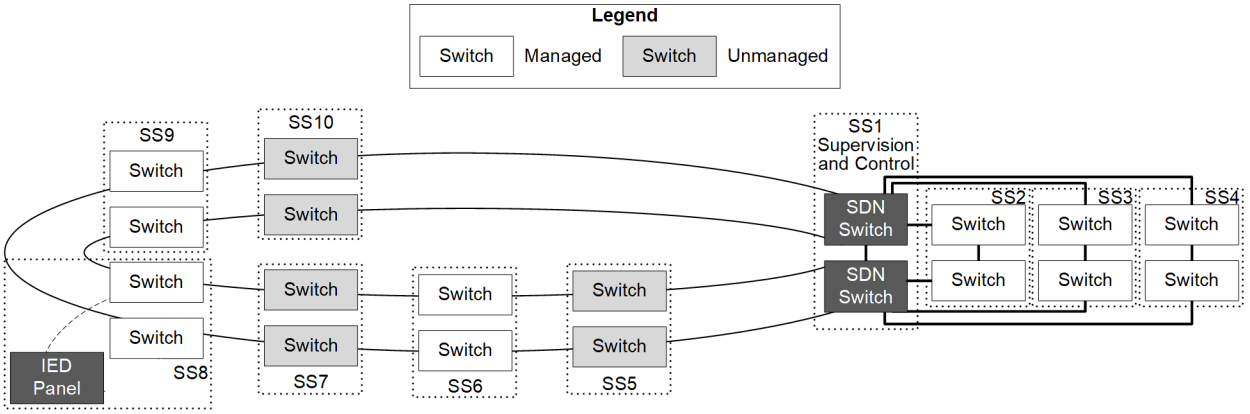

Fig. 7. Architecture at the time of the occurrence.

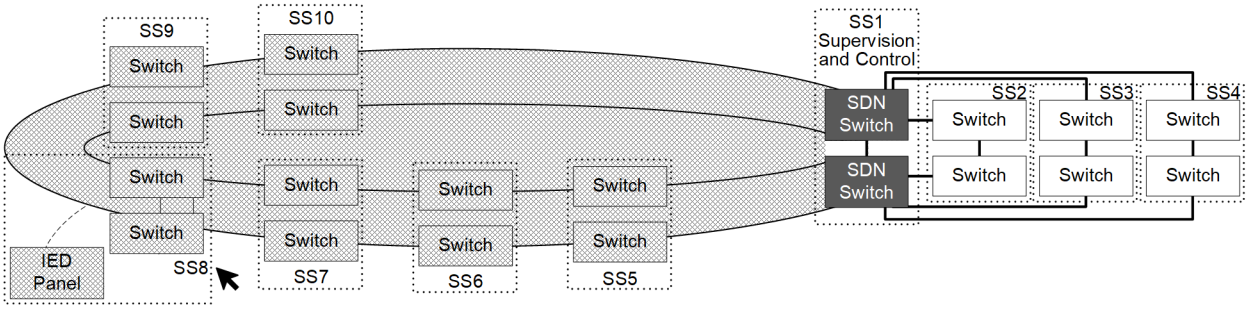

Fig. 8. Location of the occurrence and affected substations.

C. *Architecture at the Time of Occurrence*

As the system was in service, the transition between the legacy architecture of Fig. 5 and the proposed architecture in Fig. 6 was implemented in parts. Fig. 7 represents the transition status at the time of the occurrence of the failures, where approximately 30 percent of the substations had been migrated from the double-ring architecture and single Layer 2 domain to the proposed architecture, with smaller Layer 2 domains interconnected through the SDN switches. Subsystems, along with the network of IEDs installed as a panel in SS8, included protective relay IEDs with the internal Ethernet switch capability supporting RSTP. Therefore, the SS8 relay internal switches and two standalone Ethernet switches were all performing STAs to manage the physical network ring. Given the fact that all switches were VLAN-unaware and GOOSE messages had the VLAN identifier defeated with the value of zero, it was essential and challenging for the devices to perform correctly and prevent a loop during a spanning tree change or failure.

IV. Analysis of the Occurrence

During the time that the OT SDN upgrade was being installed, an unrelated fault occurred in a remote part of the system not yet being upgraded, SS8. SCADA personnel were made aware of multiple communications alarms for multiple IEDs. Following this, it was not possible to communicate with the IEDs to begin diagnostics. The system owner requested urgent local support from the commissioning team. The first analysis performed was in the supervisory system sequence of events, where it was observed that the problem started in Substation SS8, indicated in Fig. 8, with the failure of communications over a fiber-optic connection. The supervisory system also recorded network failures that spread over several substations and their IEDs, as highlighted in Fig. 8. Switches and IEDs all had their applications in the network nonoperational, which, in the context of dependability, translates into a 100 percent probability of missing commands and trips (P_{mc}). The team noted, however, that the event did not affect the substations that had already been migrated to the SDN switches.

The field analysis revealed that one of the switches in Substation SS8 was saturated, making it impossible to access it to collect diagnostic data. By disconnecting the cable from the IED panel to this switch, access to the switch was restored. This raised the suspicion that one of the IEDs in this panel was affecting the communications. The initial review detected symptoms indicative of a device flooding the network with excessive messages leading to the saturation of the interconnection switch. The next step in the analysis was to perform a network capture. Since it was not possible to access and mirror the switch ports with the IEDs panel connected, the network capture was conducted by directly connecting the engineering access machine to the IED panel.

A. Network Capture Analysis

This section analyzes the network capture performed with Wireshark [9]. The first factor derives from the properties of the capture file. The capture lasted approximately 3 minutes, during which about 2.3 million packets were recorded, averaging 12.7 thousand packets per second. Each packet had an average size of 362 bytes, resulting in a total of 810 million bytes and an average bandwidth consumption of 36 Mbps. Fig. 9 presents the protocol hierarchy statistics, which helps to visualize the predominant protocols in the capture. The GOOSE protocol stands out, accounting for nearly 100 percent of the total bytes in the capture. These statistics separate Layer 2 headers and trailers (3.9 percent) from the GOOSE payload (96.1 percent). Given that the byte count of all the other protocols is negligible, the analysis in this paper focuses on the GOOSE protocol.

Protocol	Percent Bytes	Bytes
Frame	100.0	809913009
Ethernet	3.9	31340871
GOOSE	96.1	778543358
Address Resolution Protocol	0.0	12434
Internet Protocol Version 6	0.0	8920
Logical-Link Control	0.0	3276
Internet Protocol Version 4	0.0	560
Link Layer Discovery Protocol	0.0	263

Fig. 9. Per-protocol utilization statistics.

Fig. 10 shows the GOOSE publications present in the capture. Although the number of packets of the publication with the destination MAC address 01:0C:CD:01:00:0B (referred to as GOOSE 0B) stands out, other publications were analyzed first for a baseline.

Fig. 11 shows a portion of the communications traffic for the second GOOSE publication listed in Fig. 10. The *Time delta from previously displayed frame* column represents the time interval between the current message and the previous one. The stNum (state number) and sqNum (sequence number) fields were added to verify if the behavior aligns with the GOOSE publication mechanism. Two anomalies were found. The first anomaly was messages with the same StNum in a short interval, indicating a possible loop in the network. The second anomaly was messages out of sequence; for example, the packets with SqNum ending in 78, 80, 81, 82, and 84 were missing, indicating network saturation and the probability of missing commands (P_{mc}). Similar behavior was observed in all other GOOSE publications.

Ethernet · 47 | IPv4 · 4 | IPv6 · 23 | TCP | UDP ·

Address A	Address B	Packets	Bytes
[redacted]:c4:ea:f5	01:0c:cd:01:00:0b	2,236,053	809 MB
[redacted]:c5:05:20	01:0c:cd:01:00:11	190	49 kB
[redacted]:c5:21:df	01:0c:cd:01:00:19	174	46 kB
[redacted]:c5:23:77	01:0c:cd:01:00:0e	169	45 kB
[redacted]:c5:21:e8	01:0c:cd:01:00:16	156	41 kB
[redacted]:c5:1f:66	01:0c:cd:01:00:18	136	36 kB
[redacted]:c4:e6:2a	01:0c:cd:01:00:12	126	34 kB
[redacted]:c5:1f:6c	01:0c:cd:01:00:13	123	33 kB
[redacted]:c4:ea:ec	01:0c:cd:01:00:17	118	31 kB
[redacted]:c4:ce:06	01:0c:cd:01:00:0f	112	30 kB
[redacted]:c4:d2:7a	01:0c:cd:01:00:10	99	26 kB
[redacted]:c5:1f:60	01:0c:cd:01:00:14	96	25 kB
[redacted]:c4:e6:27	01:0c:cd:01:00:0d	86	23 kB

Fig. 10. Per-protocol utilization statistics.

No.	Time delta from previously displayed frame	Destination	stNum	sqNum	Protocol
14241	0.000000	IecTc57_01:00:11	28	10967277	GOOSE
14242	0.000000	IecTc57_01:00:11	28	10967277	GOOSE
14245	0.000000	IecTc57_01:00:11	28	10967277	GOOSE
14334	0.007001	IecTc57_01:00:11	28	10967277	GOOSE
38200	4.163221	IecTc57_01:00:11	28	10967279	GOOSE
92874	8.112173	IecTc57_01:00:11	28	10967283	GOOSE
92942	0.002888	IecTc57_01:00:11	28	10967283	GOOSE
135267	4.079525	IecTc57_01:00:11	28	10967285	GOOSE
135268	0.000000	IecTc57_01:00:11	28	10967285	GOOSE
135489	0.050193	IecTc57_01:00:11	28	10967285	GOOSE
135573	0.011950	IecTc57_01:00:11	28	10967285	GOOSE
156791	2.006735	IecTc57_01:00:11	28	10967286	GOOSE
156801	0.000697	IecTc57_01:00:11	28	10967286	GOOSE
156807	0.002105	IecTc57_01:00:11	28	10967286	GOOSE
156867	0.005526	IecTc57_01:00:11	28	10967286	GOOSE
156886	0.001602	IecTc57_01:00:11	28	10967286	GOOSE
156932	0.003353	IecTc57_01:00:11	28	10967286	GOOSE
157059	0.020462	IecTc57_01:00:11	28	10967286	GOOSE

Fig. 11. Anomalies in the GOOSE publication.

The same analysis of the GOOSE publication mechanism for GOOSE 0B was performed. The stNum remained constant at 615 throughout the entire capture, as Fig. 12 shows. This number is significantly higher than the maximum value of 28 found in other publications for stNum, suggesting a high number of events, possibly caused by excessive variation in some point mapped to the message data set. The same two anomalies found in the segment of this analysis repeat throughout the entire capture. Among the approximately 2.3 million packets of this publication, only 47 different SqNum were found.

No.	Time delta from previously displayed frame	Destination	stNum	sqNum	Protocol
196493	0.000694	IecTc57_01:00:0b	615	130294	GOOSE
196494	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196495	0.000000	IecTc57_01:00:0b	615	130433	GOOSE
196496	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196497	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196498	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196499	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196500	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196501	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196502	0.000000	IecTc57_01:00:0b	615	130433	GOOSE
196503	0.001046	IecTc57_01:00:0b	615	130433	GOOSE
196504	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196505	0.000000	IecTc57_01:00:0b	615	130294	GOOSE
196506	0.000000	IecTc57_01:00:0b	615	130294	GOOSE

Fig. 12. Anomalies in the GOOSE 0B.

The graph in packets per second in Fig. 13 shows the high publication rate of the GOOSE 0B, reaching a peak of approximately 25 thousand packets per second.

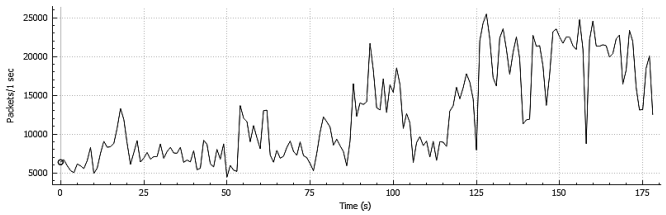

Fig. 13. Packet publication rate of GOOSE 0B.

In the graph in bits per second in Fig. 14, it is possible to observe the high bandwidth consumption of the same message, reaching a peak of approximately 74 Mbps.

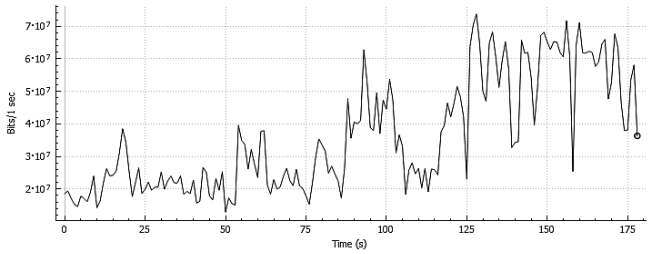

Fig. 14. Bandwidth consumption of GOOSE 0B.

Comparatively, Fig. 15 shows a graph in packets per second of a second network capture performed in Substation SS3, which was already transferred to the proposed architecture and not affected by the network storm. This substation was chosen because it has the most IEDs connected. The capture mirrored all traffic, yet the highest peak observed was approximately 125 packets per second.

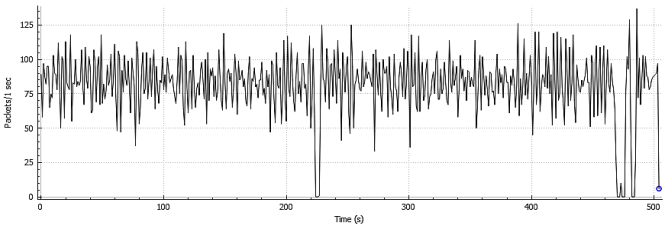

Fig. 15. Packet publication rate in Substation SS3.

The analysis shows that the network storm was caused by messages generated from a single IED with the possibility that this IED also created the loop in the network of the IED panel. Legitimate and repeated traffic was being injected into the interconnection switch of Substation SS8, which spread to the others, as shown in Fig. 8, causing the failures in switches and IEDs indicated in the supervisory system.

Another interesting point noted in the capture is the difference of 4 days between the capture date and the date present in the GOOSE messages. Possibly, there is an indication that the problem was already present 2 days before the signaling in the supervisory system.

B. Why Were Only the Double-Ring Substations Affected?

The network collapse only affected the substations interconnected through traditional switches in the double-ring architecture, making up a single Layer 2 domain managed by the RSTP. The problem spread due to the lack of traffic control configuration in the IEDs and switches and the absence of traffic control features in the unmanaged switches. VLAN identifiers, VLAN-aware switches, and multicast MAC filters would have helped to contain part of the collapse and reduce P_{mc}. This event did not spread to the rest of the network because the other substations were already isolated through the SDN switches, which performed important traffic filtering.

The supervision and control in Substation SS1 and Substations SS2, SS3, and SS4 were not affected by the network storm and did not collapse. None of the IEDs in these already isolated substations subscribe to GOOSE 0B, meaning that there were no flows configured in the SDN switches to forward the message. The SDN switches received the flooded messages and inspected its ingress port and Layer 2 content; after not finding a match in the flow table, they applied the deny-by-default policy, discarding the messages. All applications running in the isolated substations did not receive this traffic, remaining healthy and operational and, according to the confirmation of the network manager, did not present any missing commands or trips (P_{mc}).

The event showed the importance of continuing the transition from the legacy to the SDN architecture, separating the network into smaller Layer 2 domains so that localized problems, like the one analyzed, have less impact on the dependability of the overall system.

V. Network Traffic Throttling

OT SDN would have prevented cascading of the GOOSE storm, even if IEDs in the already isolated substations needed to subscribe to GOOSE 0B. Cascading of the GOOSE storm would have also been prevented if a GOOSE message originating from another IED panel and with one or more destinations coincided with those of GOOSE 0B. In these cases, traffic control and prioritization become crucial. According to IEC 61850-90-4, Clause 6.4.8.2 indicates that while prioritization contributes to the QoS, it is not sufficient on its own. Without restricting the generation rate of all participating IEDs, frame losses due to buffer overflow cannot be prevented.

If required, the OT SDN isolated system can be designed to allow the egress of GOOSE 0B messages to devices on other subnetworks, which are configured to subscribe to GOOSE 0B. In these non-OT SDN subnetworks, traffic control and prioritization by themselves may still allow equipment to be rendered unavailable due to the large number of GOOSE 0B messages, increasing the probability of missing commands (P_{mc}). To address this, OT SDN supports another feature to shape essential network traffic to prevent saturation of network segments. This section introduces a feature available in SDN architectures for network traffic throttling to allow wanted data flow to reduce the probability of missing commands (P_{mc}) while simultaneously preventing poor traffic management from unnecessarily repeating correct data flow messages such that they become incorrectly delivered and unwanted. IEC 61850-5 presents the concept that message security is negatively influenced by the probability of a device receiving unwanted

commands (P_{uc}). Therefore, throttling and traffic shaping are useful to increase both security and dependability of communications-assisted protection and control schemes. Fig. 16 represents the physical connections on the top and the data flow diagram on the bottom.

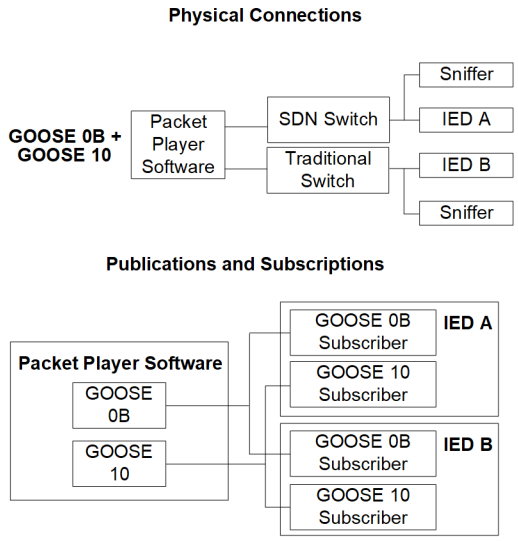

Fig. 16. Setup and diagram for traffic throttling simulation.

A packet replay tool replicated the GOOSE 0B messages from the occurrence, along with an added GOOSE 10 message. The total traffic is represented as GOOSE 0B + GOOSE 10 in Fig. 16. The purpose is to verify the probability of GOOSE delivery problems contributing to the system missing trip actions in two traffic restriction scenarios. The first scenario uses per-port traffic control available on traditional switches, and the second uses per-flow rule traffic control available on SDN switches.

The first scenario applies a traffic restriction of 1 Mbps on the port of a traditional switch, where traffic from GOOSE 0B and GOOSE 10 ingress. The graph in Fig. 17 derives from a capture mirroring the traffic received at IED B, presenting a logarithmic scale for ease of visualization. Note that at the time of replaying GOOSE 0B, GOOSE 10 begins to experience message losses and returns to normal after completing the replay. The two separated blue lines in the figure show that, rather than consistent delivery, a few GOOSE 10 messages made it through during the congestion scenario, but it is entirely random, demonstrating the lack of determinism in this solution.

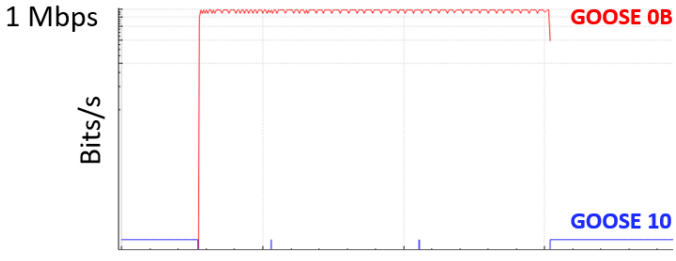

Fig. 17. Traffic restriction per port.

The second scenario applied a traffic restriction of 1 Mbps to the GOOSE 0B on the SDN switch. The graph in Fig. 18 derives from a capture mirroring the traffic received on the port of IED A, also presented on a logarithmic scale. Note that at the time of replaying GOOSE 0B, GOOSE 10 does not experience any message losses, as shown by the uninterrupted blue shape, reducing the probability of missing trips (P_{mc}).

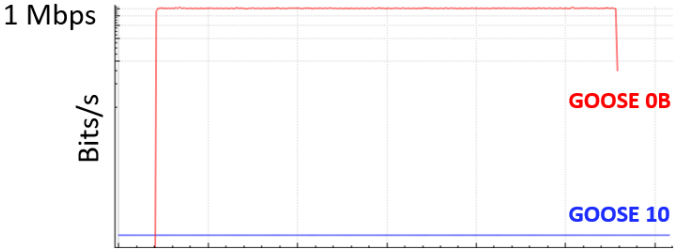

Fig. 18. Throttling of traffic generation per-flow rule in SDN.

VI. Conclusion

While traffic control and prioritization mechanisms help mitigate the risk of missing trips, they are insufficient to guarantee QoS. Message loss can still occur, even when using port-based traffic restrictions, as seen with traditional switches. An overflow situation, such as the one observed with GOOSE 0B, can consume all available bandwidth and saturate the network, disrupting the exchange of healthy GOOSE messages between IEDs. This results in message loss, increased probability of missing commands (P_{mc}), and, hence, reduced dependability.

SDN technology offers an excellent solution for interconnecting Layer 2 domains within a substation or even for interconnecting networks of different substations or agents [10]. SDN provides precise traffic control and prioritization with granular, multilayered rules that better confine network issues and enhance domain isolation. It inherently offers a stronger cybersecurity posture with its deny-by-default policy. The combination of per-flow traffic throttling and traffic prioritization ensures a reduction in the probability of missing trips, thereby maximizing performance and meeting the most stringent dependability classes required for critical applications. These multilayered traffic management OT SDN features provide confidence and safety when modifying an in-service network, even one built without BKMs. In fact, it is nearly impossible to learn how an in-service IT-based network is configured, and using OT SDN avoids the need to know the existing network design and overcome incorrect descriptions. Instead, OT SDN uses flow control rules based on the communications needed to satisfy the applications to pre-engineer the control plane and subsequent flow control rules. This paper has demonstrated with a real-world example how industrial control systems can benefit from SDN technology in preventing network events from causing a total network collapse.

VII. References

[1] IEC 60050-192-01-22, *International Electrotechnical Vocabulary (IEV) – Part 192: Dependability, 2024.* Available: electropedia.org/iev/iev.nsf/display?openform&ievref=192-01-22.

[2] IEC TR 61850-90-4, *Communication Networks and Systems for Power Utility Automation* – Part 90-4: Network Engineering Guidelines, 2020.

[3] IEC 61850-5, *Communication Networks and Systems for Power Utility Automation* – Part 5: Communication Requirements for Functions and Device Models, 2013.

[4] IEEE Std 1646, *IEEE Standard Communication Delivery Time Performance Requirements for Electric Power Substation Automation*.

[5] IEC TR 61850-90-1, *Communication Networks and Systems for Power Utility Automation* – Part 90-1: Use of IEC 61850 for the Communication Between Substations, 2010.

[6] Open Networking Foundation. Available: opennetworking.org/.

[7] D. Dolezilek, "Ethernet Design for Teleprotection and Automation Requires a Return to First Principles to Improve First Response," proceedings of the 14th Annual Western Power Delivery Automation Conference, Spokane, WA, March 2012.

[8] D. Dolezilek, J. Dearien, A. Kalra, and J. Needs, "Appropriate Testing Reveals New Best-in-Class Topology for Ethernet Networks," proceedings of the 13th International Conference on Developments in Power System Protection, Edinburgh, United Kingdom, March 2016.

[9] Wireshark. Available: wireshark.org/.

[10] E. Goncalves, A. Guglielmi, et al., "Proposal for Secure Integration of Networks for Interface Between Transmission Agents," proceedings of the XVI STPC, Rio de Janeiro, Brazil, October 2022.

VIII. Biographies

Wesley Roberto received his bachelor's degree in automation and control engineering from the Federal University of Itajubá in 2020. He was certified as a Global Industrial Cyber Security Professional by GIAC in 2023. He has worked at Schweitzer Engineering Laboratories, Inc. (SEL) in Brazil since 2019 as an automation application engineer. He develops technical applications for power system automation and is an instructor at SEL University.

Eduardo Goncalves earned his BSEE in electrical engineering from the Federal University of Itajubá in 2014. That same year, he joined Schweitzer Engineering Laboratories, Inc. (SEL), where he has worked in numerous roles as a project engineer, application engineer, and SEL University instructor. He earned a specialization degree in power systems automation in 2021 from the National Institute of Telecommunications in Brazil. In 2023, he transferred to Pullman, Washington, to join the Research and Development division, where he works as a development lead engineer.

Paulo Lima received his BSEE in electrical engineering from the Federal University of Itajubá, Brazil, in 2012. In 2013, he joined Schweitzer Engineering Laboratories, Inc. (SEL) as a protection application engineer in Brazil. In 2018, he became coordinator of the application engineering group. Since 2020, he has been the regional technical manager for Brazil. His experience includes application, training, integration, and testing of digital protective relays. He writes technical papers and white papers and provides training associated with SEL products and is a qualified SEL University instructor.

Thiago Bordim graduated in Industrial Automation Technology from the Federal University of Technology – Paraná in 2007 and received his bachelor's degree in electrical engineering from the University of São Francisco in 2014. In 2014, he joined Schweitzer Engineering Laboratories, Inc. (SEL) as a commissioning engineer in Brazil. In 2021, he became a technical leader of protection and automation projects and coordinator of the commissioning team.

David Dolezilek is a fellow engineer at Schweitzer Engineering Laboratories, Inc. (SEL) and has three decades of experience in electric power protection, automation, communication, and control. He develops and implements innovative solutions to intricate power system challenges and teaches numerous topics as adjunct faculty. David is a patented inventor, has authored dozens of technical papers, and continues to research first principles of mission-critical technologies. Through his work, he helped coin the term operational technology to explain the difference in performance and security requirements of Ethernet for mission-critical applications versus IT applications. David is a founding member of the DNP3 Technical Committee (IEEE 1815), a founding member of Utility Communications Architecture 2 (UCA2), and a founding member of both IEC 61850 Technical Committee 57 and IEC 62351 for security. He is a member of the IEEE, the IEEE Reliability Society, and several CIGRE working groups.