Lessons Learned and Successful Root Cause Analysis of Elusive Ethernet Network Failures in Installed Systems

D. DOLEZILEK, J. DEARIEN, and M. VAN RENSBURG
Schweitzer Engineering Laboratories, Inc.
USA
dave_dolezilek@selinc.com

KEYWORDS

IEC 61850 Edition 2, Ethernet, wide-area network (WAN), local-area network (LAN), testing, commissioning, troubleshooting, problem solving.

1 INTRODUCTION

In an effort to improve communications for substations and systems, IEC 61850 Technical Committee 57 has added detail to the communications standard and added references to other companion standards [1]. This paper is a companion to Reference [1] and adds examples of in-service local-area network (LAN) problems and specific details to satisfy IEC 61850 Edition 2 inclusion of IEC/TR 61850-90-4 network engineering guidelines ([2]) and other technical references related to performance tests of Ethernet LANs. Each LAN is a fundamental part of the communications-assisted protection and high-speed automation application that it supports. Therefore, the LAN must be specified, designed, built, and tested with as much rigor as the protection and automation applications. Unfortunately, traditional Ethernet technology does not have the ability to differentiate individual data streams between end devices. Therefore, monitoring and diagnostics are necessary in the Ethernet packet receivers to understand both the signal exchange within the packets and the ability of the LAN to deliver the packets.

This paper summarizes the Ethernet packet performance time definitions in IEC 61850 Edition 1 and the new IEC 61850 Edition 2 time classes, introduces complete and necessary subscribing device features required to confirm proper packet exchange, and illustrates case study examples of in-service LAN failures.

Though not in the scope of this paper, software-defined networking (SDN) technology is a new method being used for operational technology (OT) networks. Although it was first used in reactive IT networks, it is now used for proactive design and configuration of mission-critical OT networks. In addition to fault detection, isolation, and recovery within microseconds, SDN provides very powerful and specific packet monitoring and diagnostics. SDN permits design and real-time performance that better support packet delivery and monitoring.

2 DIGITAL SIGNALING TRANSMISSION, TRANSFER, AND TRANSIT TIME REQUIREMENTS

As summarized in the initial paper on this topic [3], digital signal transmission time describes the time between the detection of signal status change of state in a publisher device, the subsequent publication of this signal in a digital message, and the recognition of that change of state in the logic in the receiver device [3]. The transfer time specified for an application is the time allowed for a signal or

data exchange to travel through a communications system. IEC 61850-5 "Communication Requirements for Functions and Device Models" describes transfer time, shown in Figure 1, as the time between the action of communicating a value from the logic processing of one device to the logic processing within a second device as part of an application [4]. Transfer time includes the transit time and the time it takes to execute the communications-processing algorithm, which encodes the message in the source physical device (PD) and decodes the message in the destination PD. The transit time is the time it takes for the message to travel through the communications network.

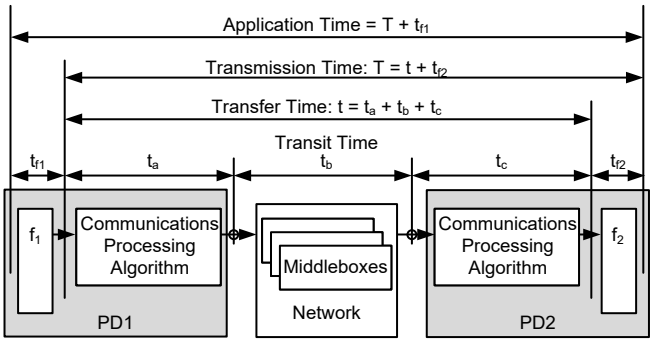

Figure 1: Application, transmission, transfer, and transit time based on IEC 61850-5

The IEC/TR 61850-90-4 network engineering guidelines clarify performance and test requirements and are considered by some to be the most important enhancement among those collectively known as IEC 61850 Edition 2. Of note, they simplify the discussion of transfer time requirements by documenting time classes for different types of messages and their associated transfer times, as shown in Table 1. These guidelines allow network engineers to accurately specify and design LANs to satisfy a transfer time class without needing to understand the underlying protection and automation applications [2].

Transfer Time Class	Transfer Time	Application Example
TT0	>1,000 ms	Files, events, and log contents
TT1	1,000 ms	Events and alarms
TT2	500 ms	Operator commands
TT3	100 ms	Slow automatic interactions
TT4	20 ms	Fast automatic interactions
TT5	10 ms	Releases and status changes
TT6	3 ms	Trips and blockings

Table 1: IEC 61850 transfer time classes

3 IEC 61850 EDITION 2: VALIDATING TRANSIT TIME AS PATH DELAY OF ETHERNET PACKET SIGNAL MESSAGES BETWEEN ALL SOURCE AND DESTINATION PORT COMBINATIONS

IEC 61850-5 describes transit time, shown in Figure 1, as the time it takes the message to travel through the communications network [4]. Mission-critical signal transfer solutions discussed in IEC 61850 use multicast digital messages made up of individual Ethernet packets for protection and high-speed automation. These signal transfer applications include Generic Object-Oriented Substation Event (GOOSE) (also referred to as Generic Substation Event [GSE]), Sampled Values (SV), and line-current differential (87L) messages within individual Ethernet frames or packets. The only way to accurately know how long it takes a packet to transit through the network is to measure transit time t_b as the difference in time between a packet entering the network as it leaves the source and exiting the network into the destination device. These three message types have features to support monitoring delivery success but not speed.

The IEEE 1588 standard defines Precision Time Protocol (PTP) with a goal of achieving very high precision for time synchronization over a packet-based network, such as Ethernet. PTP takes advantage of special Ethernet hardware for precise time-stamping of the Ethernet frame send and receive times. Using these mechanisms, transit time t_b is measured as the path delay in microseconds of PTP Ethernet packets.

Path delay is first measured for the typical packet path through the LAN based on a correctly operating hardware and cable topology. Path delay must also be calculated for each of the alternate paths that the LAN may create as a result of fault detection and reconfiguration. Reconfiguration may result in a longer path passing through more switches and cables, which results in a longer path delay and transit time. Each path is tested by placing a PTP clock on the source LAN port and a time device capable of synchronizing with a PTP clock on the destination LAN port. Both the clock and the intelligent electronic device (IED) subscribing to the PTP signal must be separately synchronized to the same time reference.

4 IEC 61850 EDITION 2: VALIDATING TRANSMISSION TIME OF ETHERNET PACKET SIGNAL EXCHANGE VIA TIME-DOMAIN LOGIC

IEC 61850-5 describes transmission time, shown in Figure 1, as the time duration between the execution of logic in PD1 and the result of logic it provokes in PD2. For example, when the result of the logic in PD1 is a change of state, it is conveyed as a protection signal to PD2 in a GOOSE message. The transmission time is the time duration between the change of state in PD1 and the subsequent detection of the representative protection signal received in PD2. Reception of this protection signal, in turn, triggers the appropriate communications-assisted action in PD2, such as a breaker trip output or other interlock action. The duration is easily calculated using time-domain devices such as discrete and real-time automation controllers that are synchronized to the same time reference. A timer is started in PD2 at the same instant that the change of state is triggered in PD1. PD2 then calculates the transmission time as the elapsed time between the test start and the reception of the protection signal from PD1 via a GOOSE message. The accuracy of this method is related to the thread time or operating cycle duration and the processing method of each controller.

This is also possible in phasor-based devices, such as protective relays, by using time-domain logic. In both PD1 and PD2, the start of the test to measure transmission time is triggered at the top of the second and top of the minute. This means the start is triggered when the clock reads each minute plus 0 seconds and 0 milliseconds, or x:00:0000. Therefore, PD2 starts a timer each minute plus 0 seconds and 0 milliseconds and PD1 triggers a protection signal change of state. When PD2 processes the receipt of the signal, it stops the timer. The accuracy of this method is related to the operating cycle duration and processing method of each relay. The elapsed time calculated by PD2 provides the transmission time.

5 IEC 61850 EDITION 2: VALIDATING CORRECT PUBLICATION OF ALL ETHERNET PACKET SIGNAL MESSAGES

Each publisher is unaware of the delivery to the intended subscribers. Therefore, the validation of the message publication can only confirm the behavior of the publisher and the contents of the digital message being published. To validate GOOSE publication, each publishing device must maintain and produce information about the message configuration and real-time performance of the outgoing GOOSE publications. The publisher calculates and stores information for each of the GOOSE messages that it publishes. This information is available in a human-readable format report via an engineering access connection and via a poll-and-response interaction with a data concentrator.

The GOOSE transmit message report contains information that includes the following:

1. Message configuration information.
    a. GOOSE control reference. The IEC 61850 GOOSE control reference information for each message includes the IED name, logical device instance, logical node class, and GSE control block name.
    b. Multicast address media access control (MAC).

c. Priority tag.
d. Virtual LAN (VLAN).
e. Application identifier (AppID).
f. Data set reference. The IEC 61850 GOOSE data set reference includes the IED name, logical node class, and GSE control data set name.

2. Message status.
   a. State number of the most recently published message. This field represents the number of times the data set contents have changed state since the session of the GOOSE publication began and increments each time a state changes.
   b. Sequence number of the most recently published message. This field represents the number of times a message has been published with unchanged data set contents since the last change and increments each time a GOOSE message is published.
   c. Time to live (TTL) value in the most recently published message. This field contains the time remaining in milliseconds before the next message is expected to be published.

6 IEC 61850 EDITION 2: VALIDATING CORRECT RECEPTION OF ALL ETHERNET PACKET SIGNAL MESSAGES

Because protection and automation message packets are often multicast to numerous subscribers, it is necessary to monitor the receipt of each packet at each receiving subscriber. Ethernet packet messages for protection and high-speed automation signal transfer include GOOSE, SV, and line-current differential (87L).

To validate GOOSE subscriptions, each subscribing device maintains and produces information about the message configuration and the real-time performance of the incoming GOOSE subscriptions. The publisher calculates and stores information for each of the GOOSE messages to which it is subscribing. This information is available in a human-readable format report via an engineering access connection and via a poll-and-response interaction with a data concentrator. The subscriber uses the following GOOSE message configuration information to validate that the GOOSE message is from the intended source and matches the engineered subscription design. GOOSE messages that do not match a pre-engineered configuration are discarded.

The GOOSE receipt message report contains information that includes the following:

3. Message configuration information.
   a. GOOSE control reference.
   b. Multicast address (MAC).
   c. AppID.
   d. Data set reference.
4. Message status.
   a. Priority tag. In a GOOSE receive report, this is the priority tag value received in the last message. If the priority tag is not received as part of the GOOSE message and is unknown, then the report will indicate that it was not received as part of the packet header. The report must avoid confusion between the received value of zero and the fact that the tag is nonexistent.
   b. VLAN. In a GOOSE receive report, this is the VLAN value received in the last message. If the VLAN is not received as part of the GOOSE message and is unknown, then the report will indicate that it was not received as part of the packet header. This is done to avoid confusion between the value zero and the fact that the VLAN is nonexistent.
   c. State number: In a GOOSE receive report, this is the state number value received in the last message.
   d. Sequence number. In a GOOSE receive report, this is the sequence number value received in the last message.
   e. TTL. In a GOOSE receive report, this value is updated with the expected remaining TTL in milliseconds, which represents the expected time duration before receipt of the next GOOSE message in this specific subscription.

f. Error code. The report calculates and displays warnings and error conditions defined by IEC 61850, which include the following.
    - GOOSE configuration revision mismatch, meaning the configuration revision number of the incoming GOOSE message does not match the value as configured in the Configured IED description (CID) file.
    - GOOSE commissioning is necessary, meaning the "needs commissioning" flag is set to true in the incoming GOOSE message.
    - GOOSE message received out of sequence, meaning that the consecutively received message state numbers and/or message sequence numbers are not in sequence.
    - GOOSE message received corrupted, meaning that the format of the incoming GOOSE message is not as configured, is incorrectly encoded, or is otherwise corrupted.
    - TTL has expired, meaning that a GOOSE message for this subscription was not received within the expected time interval.

g. Out-of-sequence count. This is the count messages lost because of both sequence number and state number out-of-sequence errors. It is not recorded for the first message after the device is turned on or reconfigured.

h. Time to live count. This is the count of the number of times a message is not received within the expected time interval, referred to as the TTL.

i. Decode error count. This is the count of the number of messages where enough information is decoded to associate them with a subscription but fails further decoding because of corruption or errors, such as a mismatched data set.

j. Buffer overflow count. This is the count of the number of messages that are discarded because the message receive queue was full. This may occur as a result of time compression in the network that causes two packets from the same subscription to be received within one publication period. The receiving IED should discard the older packets for this subscription and process only the newest one.

k. Message lost count. This is the aggregate count of the estimated number of messages lost because of out-of-sequence errors. For each out-of-sequence error, the number of messages lost is estimated by subtracting the expected state number from the received state number and the expected sequence number from the received sequence number and summing them. This estimate is only made if the state number or sequence number in the received message is greater than expected.

l. Maximum message lost count. This is the maximum estimated number of messages lost for an out-of-sequence error.

m. Total downtime. This is the total time (in seconds) the subscription was in an error state.

n. Maximum downtime. This is the maximum time (in seconds) the subscription was continuously in an error state.

o. Message status history. The GOOSE report maintains statistics for several of the most recent error events, including date of event, time of event, duration of event, and event error code.

7 TROUBLESHOOTING AN IN-SERVICE SYSTEM EXPERIENCING RAPID SPANNING TREE PROTOCOL (RSTP) PROBLEMS

As introduced in [1], an in-service LAN was experiencing problems that could not be diagnosed or fixed using traditional IT methods. The LAN was a ring topology design with packet switching and fault recovery mechanisms based on Spanning Tree Algorithm (STA) via RSTP. During testing of the network, it was found that when any of the dual-redundant, triple-modular front-end processors (FEPs) were power-cycled, communications with remote substations were lost for up to 30 seconds. GOOSE communications from the remote substations were not actually lost, but rather all IP communications among the six FEPs and the human-machine interface (HMI) were disrupted. The amount of time the network lost communications was not consistent, but communications failure was consistent.

An analysis of the logs in the switches in the central substation and the remote substations showed that, on occasion, a LAN switch in a remote substation was attempting to become the root bridge of the extended RSTP network. A Wireshark® protocol analyzer capture of Ethernet traffic revealed that the switch in the remote LAN had sent an RSTP message requesting to become the root bridge. Spanning tree algorithms within switches may make this request before their initial configuration or when they lose communications to the rest of the network for a significant amount of time. The wide-area network (WAN) connection to the remote substations was restricted to 2 Mbps but was adequate for the typical traffic observed with Wireshark. Based on this, the initial theory was that rebooting an FEP caused a spike in network traffic, which in turn saturated the remote connection and caused the switch to attempt to become the root bridge.

The local design was a LAN extended across the WAN so that the local and remote LANs were connected as one distributed RSTP LAN. Wireshark captures were taken at the interconnection between the local LAN and the WAN by monitoring a switch port that was set to mirror the WAN interconnection traffic. The monitoring revealed a surge of WAN traffic when the FEP was power-cycled. The Wireshark bandwidth analyzer function illustrates bandwidth usage in real time or when reviewing a captured file. This function was used to graphically illustrate unexpected and unwanted network behavior, such as the bandwidth consumption graphs in Figure 2.

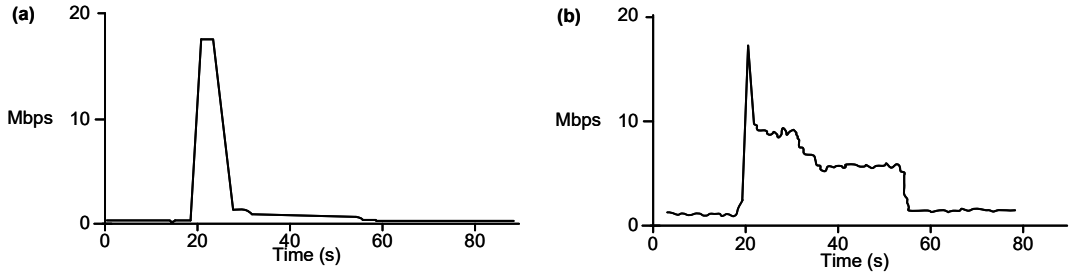

Figure 2: Spike in WAN link traffic when FEP is power-cycled (a) and gradual return to normal (b)

Figure 2a illustrates a spike to almost 20 Mbps of traffic attempting to traverse the WAN link provisioned for 2 Mbps. Repeated tests revealed that this occurred consistently each time a FEP was power-cycled. This confirmed that there was enough traffic to cause bandwidth saturation for a period long enough to cause the remote switch to detect RSTP failure and attempt to become the root bridge.

Normally, when a switch receives a packet for a destination MAC address found in its MAC address lookup table, it sends that packet to the appropriate destination port. When a switch receives a packet for a destination MAC address that is not in its MAC tables, it floods that packet to all of its other ports. This is why it is essential to not publish messages to network addresses that do not belong to any device or to a device that has been removed. By design, when a link is lost on a switch port the switch flushes the MAC table entries for that specific port. When the port is a backbone port, messages are redirected based on the other MAC table entries.

However, when the port is the only perimeter connection to a singly attached end device, like an FEP, that switch loses knowledge of where to send packets destined for that FEP and floods. When this happens, switches flood the message out all ports, except the one it was received on, until the network once again discovers the destination port for that specific MAC or the source stops sending messages.

Using features in Wireshark, it was determined that most of the traffic was destined for the FEPs from redundant FEPs and the HMI. The redundant FEPs normally publish User Datagram Protocol (UDP) messages to the other FEPs every few milliseconds, which aggregated to nearly 20 Mbps per FEP. For some reason, once the FEP was power-cycled, the network became flooded with messages from the redundant FEPs and HMI. Figure 2b shows the high network bandwidth utilization and how it drops as traffic associated with each of the redundant FEPs changes from the unexpected LAN flooding behavior to low data transmission volumes. This decreased traffic is the result of the source FEP suspending the transmission of data to the now unknown destination.

Study of the Wireshark captures indicated that during power cycling the switch flushed its MAC table and flooded messages destined for the FEP. The switch began flooding each message destined for that FEP and only stopped flooding as the sending devices stopped sending them. It was found that the sending devices were timing out within 30 seconds after the FEP was turned off. This happened when their address resolution table flushed the FEP MAC entry because that MAC was not responding to a MAC refresh process.

Normally, messages would be sent to the switch port connected to the FEP. However, when the switch flooded the messages out every other port, they also attempted to exit the WAN port. This flood lasted up to 30 seconds, with messages attempting to egress a switch port connected to a 2 Mbps WAN link. This link became saturated, and the remote LAN segment perceived that it had been disconnected. At this point, one switch in the remote LAN segment attempted to heal the LAN by declaring itself the root bridge. Once communications were normalized, the spanning tree algorithms in the local substation rejected requests from the remote substation to change the root bridge. The resulting RSTP exchange created messages like those recorded in the first Wireshark capture, which were a byproduct of the switch flooding that resulted from the lost FEP MAC table entry.

As a solution, MAC security management was used to whitelist which source MAC addresses are allowed to connect to the WAN link port. Using this method, the flood messages are prohibited from exiting the WAN link because the source destination MAC in the UDP messages is not on the list. This was accomplished by adding a new managed switch between the local substation LAN and the WAN connection. This additional switch was necessary because MAC security management is an ingress function and could not be configured on the existing egress link to the WAN.

8 INCORRECT LAN CONFIGURATION PREVENTS FAILOVER GOOSE DELIVERY

During an onsite commissioning process, as introduced in [1], it was recognized that some IEDs were not properly communicating with each other. IEDs at the site were meant to be sharing GOOSE messages, but some devices were reporting failures receiving the messages. It was not obvious why these messages were not being received while messages from other neighboring IEDs connected to the same LAN were being received without errors. The IEDs sending the messages were not reporting any errors and had network connectivity, but the messages they were sending were not reaching their destination. To diagnose the problem, the IEDs were rebooted one at a time. As these devices were rebooted, the message failures recovered except for one message from one IED.

The physical network configuration at this location is a typical ladder topology, as shown in Figure 3a, where the A side of the network is on the left and the B side of the network is on the right. The network switches have RSTP enabled to provide loop mitigation and provide redundancy.

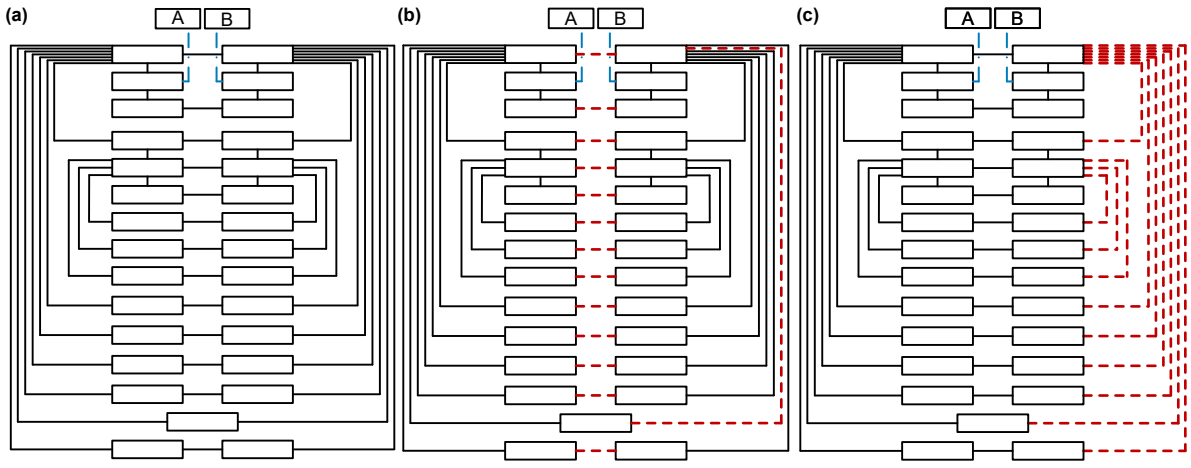

Figure 3: Physical cabling of a large ladder topology network (a), illustration of incorrect settings creating unexpected alternate paths illustrated as dashed lines in center of network (b), and illustration of proper network topology with alternate paths illustrated as dashed lines on the right (c)

The top of the ladder in this installation is connected to a larger network through a WAN device. This WAN device is not a Layer 3 boundary device and instead connects this local Layer 2 network to another similar Layer 2 network. This creates a very large Layer 2 network that spans multiple locations as a single RSTP domain with the RSTP root switch at a different location.

The IEDs in the network are configured in failover mode and have one port connected to a switch on the A side and the other port connected to a different switch on the B side to ensure full redundancy and avoid any single point of failure.

GOOSE messages are multicast messages and will traverse everywhere in a Layer 2 network unless controlled. The IEDs use GOOSE for their protection signaling and the network configuration uses VLANs to manage the propagation of the GOOSE messages. The GOOSE messages in one LAN are never required in any other LAN, so they are blocked from entering the WAN with VLAN filtering at the WAN/LAN interconnections.

Upon further investigation, it was determined that the A port on the device had a failure such that the device was always linking to the B network. The IEDs involved in this situation have a failover mode with redundant ports where the first port that gains a link becomes the active port and the other port becomes the backup port. The active port remains active until a loss of link is detected, and then the device fails over to the backup port. This port remains active until it fails, even if the first port becomes active again. This IED failover behavior means that at any time any IED could be communicating on the A or the B side of the LAN. Which side of the network is used is determined by the order in which the device links became active, which can be determined by the order the switches were powered on.

The observation that the remaining failed device was using the B side of the network led to the understanding that the messages from the IEDs that were not successfully being delivered were in fact being dropped at the LAN/WAN interconnect by the VLAN filtering settings that were meant to keep local GOOSE messages contained in the local LAN. When the IEDs in question were rebooted, their active port reset to the A network because both switches were available at boot time, except in the case of the failed device that had a failed A port.

A properly configured ladder topology is wired as shown in Figure 3a, but cabling alone is not enough to create the proper configuration. Proper settings are required to make the topology respond quickly to network failures, and those same settings force the traffic on the network to prefer the A side of the network whenever possible.

The cabling of the network was correct, but the settings were defaults. This LAN was part of a much larger network and part of a single large RSTP domain. The root bridge for the RSTP domain is outside of this local LAN, and the shortest path to the root (with default settings) for the B side switches is up the B side and out to the WAN. Without proper settings, the RSTP network converged as shown in Figure 3b, which effectively splits the network down the middle. Any device linked on the A side of the ladder must go through the WAN/LAN interconnect to communicate to a device linked on the B side.

The result of the network converging this way, and having VLAN filtering at the WAN/LAN interconnect, means that any devices linked on the B side of the network are unable to send GOOSE messages to devices linked on the A side of the network. There is an even more subtle problem than that. The GOOSE messages are filtered at the WAN/LAN interface and are therefore dropped. Errors can be identified, but there is traffic that is not filtered by the VLAN filtering that does get from the A side of the network to the B side by going out to the WAN to make the journey. This means that there may be unwanted traffic on the WAN network that should never be there.

The solution to this problem is to properly configure the RSTP settings for the switches in the LANs to keep local LAN traffic in the local LAN. By setting the path cost on the B switch WAN/LAN interface ports very high, the network will move traffic to the A side of the network whenever possible. This change causes the local LAN traffic to stay within the LAN.

It is also important, for performance reasons, to set the path cost between the local LAN roots on the B switch appropriately high to cause all the rungs of the ladder to immediately prefer the A side.

Applying the proper settings to the network switches should result in the network convergence depicted in Figure 3c.

If it were not for the failed port on the one IED, the misconfiguration of the network may have gone unnoticed until the failure caused a misoperation.

Proper settings and network configuration are critical, as is proper and complete testing.

9 CONCLUSION

The IEC 61850 Edition 2 reference to the IEC/TR 61850-90-4 network engineering guidelines clearly indicates that digital message publishers and subscribers require internal purpose-built monitoring, diagnostic, and reporting capabilities. These reports need to contain observations and calculations that document the health and performance of digital messages used for signal exchange. Every characteristic of these sophisticated messages needs to be monitored because these data are essential to recognizing and diagnosing communications problems. Testing, commissioning, and ongoing maintenance require that IEDs and network packet switching devices be capable of monitoring and recording packet publication and subscription configurations, behavior, and reliability [5].

Clearly, this work demonstrates the need to specify, design, build, test, and monitor the Ethernet LAN for OT behavior within mission-critical applications. That Ethernet switches will automatically perform packet transfer has confused many engineers into thinking that default settings are adequate. However, it is necessary to understand the true nature of Ethernet, spanning tree algorithms, and recovery mechanisms to understand the failure modes and engineer mitigation factors into the system.

The authors recognize and recommend the use of SDN to create more precise packet delivery and better fault detection, isolation, packet flow reconfiguration, and reestablished packet delivery.

REFERENCES

[1] M. Van Rensburg, D. Dolezilek, and J. Dearien, “Case Study: Lessons Learned Using IEC 61850 Network Engineering Guideline Test Procedures to Troubleshoot Faulty Ethernet Network Installations,” proceedings of the Power and Energy Automation Conference, Spokane, WA, March 2017.

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

[3] D. Dolezilek and J. Dearien, “Lessons Learned Through Commissioning and Analyzing Data From Ethernet Network Installations,” proceedings of the 5th International Scientific and Technical Conference, Sochi, Russia, June 2015.

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

[5] S. Chelluri, D. Dolezilek, J. Dearien, and A. Kalra, “Understanding and Validating Ethernet Networks for Mission-Critical Protection, Automation, and Control Applications,” March 2014. Available: https://selinc.com.