Case Study: Lessons Learned Using IEC 61850 Network Engineering Guideline Test Procedures to Troubleshoot Faulty Ethernet Network Installations

Marcel van Rensburg, David Dolezilek, and Jason Dearien
Schweitzer Engineering Laboratories, Inc.

Presented at the
Power and Energy Automation Conference
Spokane, Washington
March 21–23, 2017

Case Study: Lessons Learned Using IEC 61850 Network Engineering Guideline Test Procedures to Troubleshoot Faulty Ethernet Network Installations

Marcel van Rensburg, David Dolezilek, and Jason Dearien, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—IEC/TR 61850-90-4 and other technical references describe the need for performance testing to verify design concepts and then confirm performance during factory and site acceptance testing. Even more important is the need to perform testing for diagnosis and troubleshooting as an essential part of identifying and correcting the root cause of errors found during commissioning and operation of in-service systems.**

Ensuring correct setting and installation of Ethernet networks used for communications-assisted protection, safety, and automation is critical. If the local-area network is designed or installed separately from the primary project, it requires a network acceptance test as part of the factory and site acceptance tests. The ability to use shared bandwidth, multipurpose Ethernet connections on microprocessor-based relays, controllers, and other digital devices has simplified physical installations while complicating device settings and configuration. Many of these devices have been designed with internal communications monitoring and diagnostics to provide commissioning tools unavailable in traditional standalone relays. However, installation and commissioning remain complicated when these tools are not used or understood.

This paper provides examples of detecting problems, finding root cause, and correcting communications problems within operational systems around the world. Case study examples illustrate the less well-known consequences of the use of Ethernet technologies, how they negatively impact mission-critical communications, and how to find and resolve the issues prior to a system failure. These examples represent mistakes that are easy to make, hard to diagnose, and difficult to correct after the physical devices are programmed.

Even with better attention to detail during commissioning, communications problems develop over time in active systems. The paper also explains best practices to configure systems to constantly perform self-tests and monitor and record performance statistics. The benefits of analyzing application and communications performance data stored within digital devices make it evident that this is a necessary best practice to reduce communications-related misoperations.

I. Introduction

Ensuring the correct setting and installation of Ethernet networks used for communications-assisted protection, safety, and automation is critical. Local-area network (LAN) topology reconfiguration algorithms must detect internal network faults, enable hot-standby links, and redirect data flows, while remaining secure against delaying or damaging Ethernet packets. The ability to use shared bandwidth and multipurpose Ethernet connections on microprocessor-based relays, controllers, and other digital devices has simplified physical installations, while complicating device settings and configuration. Many of these sensors, controllers, and communications network devices include better commissioning features and standalone tools than those available with older, traditional, standalone relays and devices. However, installation and commissioning remain complicated when the features and tools are not used or understood.

Installation and settings errors continue to be widespread, implying a need for more rigorous commissioning tests. IEC/TR 61850-90-4 network engineering guidelines [1] and other technical references describe the need to perform specific and measurable performance tests of Ethernet networks, such as network acceptance tests, as part of each network’s design verification, factory acceptance test, and site acceptance test.

Even with greater commissioning effort, occasional communications problems develop over time. These problems can best be resolved by analyzing application and communications performance data stored within the digital devices. In the interest of reducing communications misoperations, this paper shares practical lessons learned through experience with troubleshooting, diagnosing, and correcting in-service Ethernet networks. These examples are important for all users and designers of Ethernet communications to review because they represent mistakes that are easy to make, hard to diagnose, and difficult to correct after the fact.

This paper is an extension of [2]. It expands on the use of IEC/TR 61850-90-4 network engineering guidelines [1] and other technical references that describe the need to perform specific and measurable performance tests of Ethernet networks as part of network design verification, factory acceptance tests, and site acceptance.

II. Digital Signaling Transmission, Transfer, and Transit Time Requirements

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 finally the recognition of that change of state in the logic in the receiver device. 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 describes transfer time, shown in Fig. 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 [3]. 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.

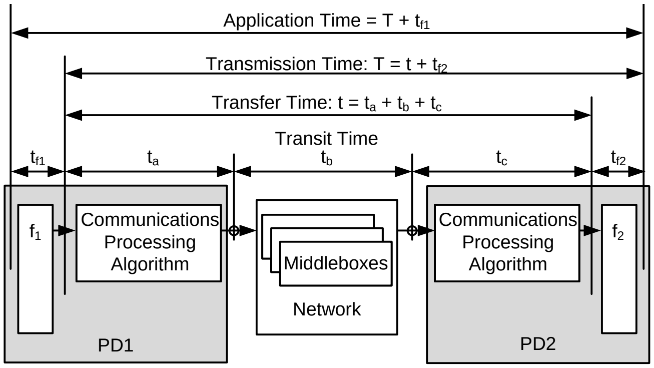
Application Time = T + t_{f1}
Transmission Time: T = t + t_{f2}
Transfer Time: t = t_a + t_b + t_c
Transit Time
t_{f1}
t_a
t_b
t_c
t_{f2}
f_1
Communications Processing Algorithm
Middleboxes
Network
Communications Processing Algorithm
f_2
PD1
PD2

Fig. 1. Application, transmission, transfer, and transit time based on IEC 61850-5

IEC/TR 61850-90-4 network engineering guidelines clarify performance and test requirements. 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 I. These guidelines allow network engineers to accurately specify and design local-area networks (LANs) to satisfy a transfer time class without needing to understand the underlying protection and automation applications [1].

TABLE I
IEC 61850 TRANSFER TIME REQUIREMENTS

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

The IEC/TR 61850-90-4 network engineering guidelines technical report defines latency of communication as the delay between the instant that data are ready for transmission and the moment they have been completely received at their destination(s) [1]. IEC 61850-5 describes the traffic recommendations specific to IEC 61850 [3]. It does not identify the other necessary traffic on the IEC 61850 Ethernet network for maintenance, telephony, video surveillance, and so on.

TT0 through TT6 in Table I illustrate time classes that satisfy different types of applications within a multipurpose communications network using protocols that include those within the IEC 61850 standard. The IEC 61850 transfer time requirement for digital signals as part of a communications-assisted protection scheme is TT6 in Table I. IEC 60834 requirements for security, reliability, and dependability are met if the system meets the 3-millisecond transfer time 99.9999 percent of the time and has a delay no longer than 18 milliseconds for the remainder [4].

Questions that must be answered by engineers and technicians during design and commissioning include the following:

1. How do I verify that the Ethernet switches are configured properly for the signal message parameters?
2. How do I validate the time duration between a power system event and a subsequent mitigation reaction in a remote intelligent electronic device (IED)—representing the total signal application time—via an Ethernet signal application?
3. How do I validate the transmission time duration between the detection of an event in one IED and a subsequent mitigation reaction in a second IED?
4. How do I validate the transfer time duration between the publishing of a message in one IED and subsequent message processing in a second IED?
5. How do I validate the transit time duration of message delivery between IEDs?
6. How do I verify the impact of failure and reconfiguration on a hot-standby Ethernet network path for each of the previous questions?
7. Will the signal channel be affected if I expand the network?
8. How do I verify that all published Generic Object-Oriented Substation Event (GOOSE) messages are getting to each destination?

For each of these questions, network, protection, and automation engineers often ask: How would I know during the design phase? How would I know during a factory acceptance test? How would I know during onsite commissioning? How would I know as part of ongoing monitoring? [4] This paper promotes methods to test and diagnose the functionality and performance of devices, LANs, and wide-area networks (WANs) to satisfy the reliability and speed of packet delivery. Many other questions about the IEDs, protocols, and Ethernet message configurations that are equally important to signaling are outside the scope of this paper. Signaling via digital messages requires that specific engineering best practices be used during specification and design. Best practices to deploy Ethernet LANs are discussed in detail in [4] and include fast and efficient spanning tree algorithm processing in switches configured in a ladder topology.

Once these best practices are deployed in the design and construction of networks of IEDs to perform mission-critical applications, it becomes very important to also design methods to test and validate performance.

III. Verifying Correct Local-Area And Wide-Area Virtual LAN (VLAN) Configuration

Before testing network performance, it is necessary to verify that the perimeter and backbone ports are configured correctly. For normal operation and for every failure mode, each perimeter port must demonstrate correct message ingress and egress. This test is performed via a network configuration tester and monitor, as shown in Fig. 2.

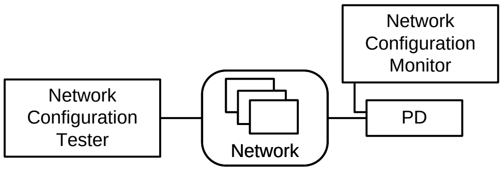

Fig. 2. Network configuration tester and monitor topologies

Best practice is to connect the monitor to the second Ethernet port of the PD and configure the PD to pass through all traffic received on the first port out of the second port. If the PD does not support this capability, disconnect the network cable from the PD and plug in the monitor for the brief duration of the test. Every message configuration combination of the media access code (MAC) address and VLAN is published into the network, and the display shows which messages successfully egress each perimeter port [4]. This answers Question 1 from Section II. For example, when the network tester publishes 150 unique GOOSE signal messages with VLANs from 1 to 150 into the network, the network switch configuration is verified via the network configuration monitor within seconds. The human-machine interface (HMI) on the network configuration monitor illustrates that the network correctly segregates traffic and only permits GOOSE signal messages that match the configuration of the in-service GOOSE and sampled value messages to egress the network to the PD.

Validation of each specific GOOSE exchange is performed by adding an annunciation function in the signal payload. This validation is performed by triggering annunciation at the subscriber to visibly illustrate success. This annunciation element is triggered via a pushbutton on the publisher IED and mapped to an LED or other display on the subscriber front panel. Witnessing the LED on the subscriber change as a result of the pushbutton on the publisher gives immediate confirmation that the signal exchange is configured correctly and that the network is configured to pass the message. This test will succeed even if unwanted traffic is present. The network configuration monitor is necessary to confirm that no unwanted messages are allowed via the network configuration.

IV. Troubleshooting an In-Service System Experiencing Rapid Spanning Tree Protocol (RSTP) Problems

During testing of an in-service network, it was found that when any of the dual-redundant, triple-modular front-end processors (FEPs) on an RSTP ring network 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 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 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 a 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 Fig. 3.

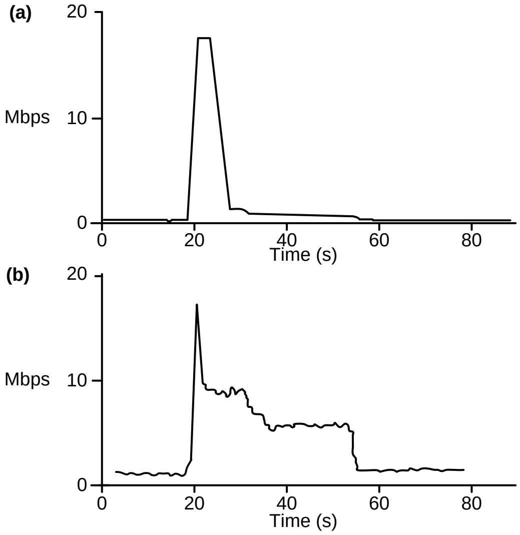

Fig. 3. Spike in WAN link traffic when FEP is power-cycled (a) and gradual return to normal (b)

Fig. 3a 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. Fig. 3b 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.

V. Incorrect LAN Configuration Prevents Failover GOOSE Delivery

During an onsite commissioning process, 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 Fig. 4, 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.

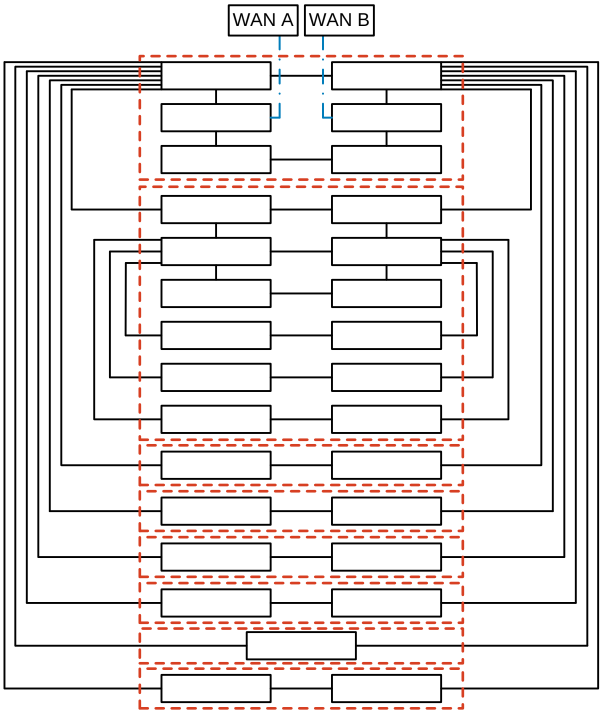

Fig. 4. Physical cabling of a large ladder topology network

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 Fig. 4, 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 Fig. 5, 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.

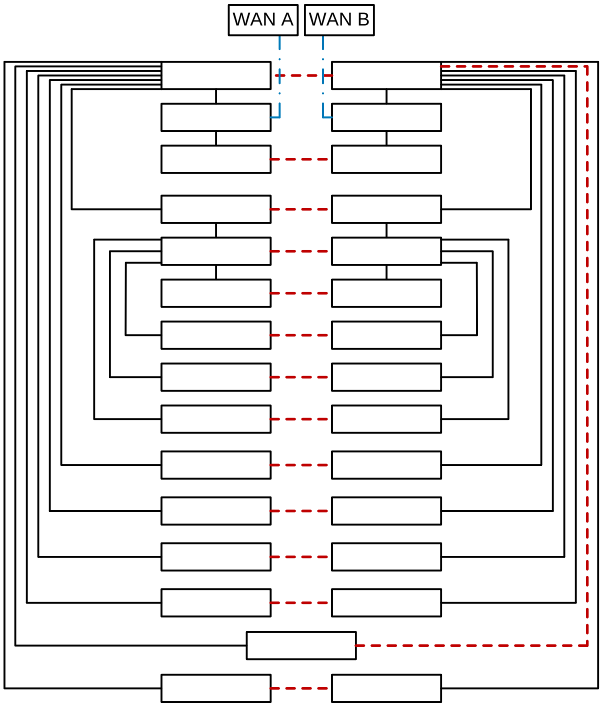

Fig. 5. Illustration of incorrect settings creating unexpected alternate paths in center of network

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 Fig. 6.

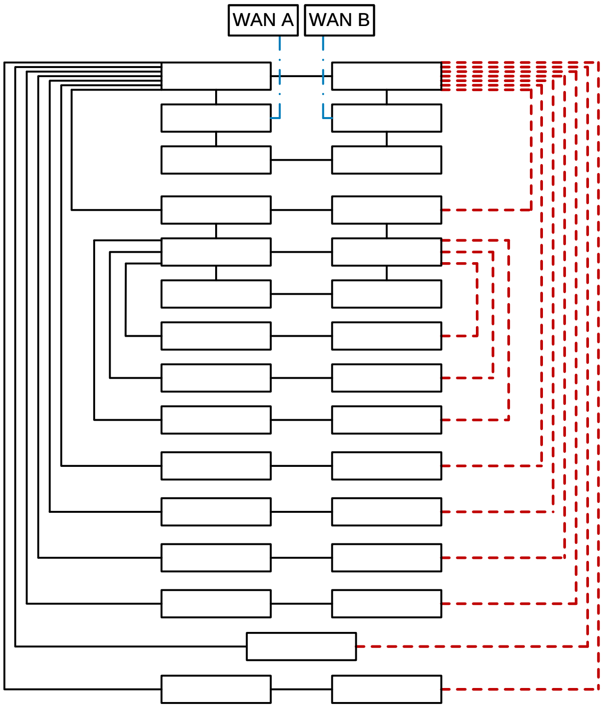

Fig. 6. Illustration of proper network topology with alternate paths on the right

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.

VI. Validating Elapsed Application Time Between Detected Event And Resulting Mitigation Action

The total signal application time duration between a power system event and a subsequent mitigation reaction performed by a remote IED (see Question 2 in Section II) is measured using synchronized-logic IEDs (SLIs). These SLIs are attached to laboratory and in-service systems to simulate power system actions and monitor IED reactions for test purposes. These SLIs have high-accuracy synchronization to an IRIG-B time source, create high-accuracy digital Sequential Events Recorder (SER) reports, and have time-synchronized logic processing. This time-synchronized logic makes time duration calculations and absolute time stamping more accurate than in protection IEDs (PIEDs) synchronized to the power system. The SLIs trigger logic precisely at the top of the second with 1-millisecond accuracy and, when synchronized to the same time source, they start test activities at precisely the same point in time regardless of geographic location [4].

Using synchronized logic, SLI1 in Fig. 7 triggers a simulated power system contingency change of state via a contact output wired to a contact input on PD1 precisely at the top of the second. SLI2 starts a timer at the top of the second. After detecting a contact input, PD1 publishes GOOSE messages with change-of-state data and sends them to PD2, which then closes an output contact as a mitigation reaction. SLI2 detects the PD2 output as a contact input and stops the timer as the total signal application time duration. SLI timers experience error from 0 to 1 millisecond because of a 2-millisecond operating cycle and precision starts. For verification, the SLI1 output contact is also temporarily hardwired to SLI2, and the time duration between the two input contacts on SLI2 is separately measured to confirm the accuracy of the top-of-the-second timer in SLI2.

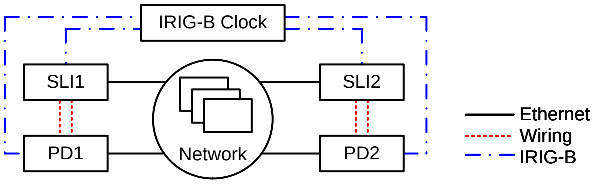

Fig. 7. Test network

This means that multiple SLIs can be distributed over any distance and create precise time measurements via digital messaging alone when synchronized to the same time source. The typical duration of mitigation applications based on digital signal exchange and contact output action is measured to be less than 14 milliseconds, as seen in Table II.

Table II
Measured and Calculated Application Time Duration for a PIED With 2-Millisecond Operating Cycle

Signaling Messages	**LAN Recovery Time**	**Transfer Time**	**Application Time (Digital Input to Digital Output)**
1st (t_0)	No failure	<3 ms	<14 ms
2nd (t_0 + 4 ms)	<3 ms	<8 ms	<18 ms
3rd (t_0 + 8 ms)	<7 ms	<12 ms	<22 ms
4th (t_0 + 16 ms)	<15 ms	<20 ms	<30 ms

This answers Question 2 in Section II. Redundant GOOSE signal messages are published at periods of 4, 8, and 16 milliseconds after the initial GOOSE message resulting from the change of state. If the first published GOOSE signal message after the change of state is not delivered successfully, the redundant publication 4 milliseconds later triggers the mitigation output in less than 18 milliseconds. This burst of redundant signal messages ensures that the transfer is executed within the required 20 milliseconds even if the LAN experiences a failure that lasts 15 milliseconds. Methods to design and validate LANs to recover from failure within the maximum duration of 15 milliseconds are accomplished by way of the ladder topology [4].

VII. Validating Signal Transmission Time as Part of an Application

IEDs, SLIs, and PDs are computerized products that perform intelligent control and monitoring. PIEDs are devices that track the power system and whose clocks are synchronized to a time source, but their operating cycles are synchronized to the power system. Therefore, input measurements and signals are detected, time-stamped, and recorded as SER reports at some unknowable time within the operating cycle. As a result, time-stamp accuracy varies from being exactly correct to being nearly one full operating cycle late. The end result is that the time stamps of inputs are inaccurate by 0 to 1/2 of an operating cycle. For PIEDs operating every one-eighth power system cycle, input errors vary from 0 to 1.04 millisecond for a 60 Hz system, and from 0 to 1.25 millisecond for a 50 Hz system. For PIEDs operating every one-quarter power system cycle, input errors vary from 0 to 2.08 milliseconds for a 60 Hz system, and from 0 to 2.5 milliseconds for a 50 Hz system. The output time-stamp error for both PIEDs is zero because the operating cycle processes the output and the time-stamp function consecutively in the code. Unfortunately, because the starts of the operating cycles in the source PIED and in the destination PIED are not perfectly synchronized, SER records contain errors relative to absolute time as well as to each other. SLIs that are designed to correctly compensate for the time duration of input measurements and to operate at a 2-millisecond cycle will time-stamp the inputs accurately.

Transmission time duration, as shown in Fig. 1—where the PDs are actually PIEDs—is calculated as the time difference between the time stamp in the SER for the contact input detection in PD1 or SLI1 and the time stamp in the SER of the signal reception in PD2 or SLI2 in Fig. 7. When an external trigger is synchronized to the top of the second, the error in the transmission time duration based on the delta time-stamp method includes a physical input time-stamp processing error on the publisher and a digital message processing time-stamp error on the subscriber. The application time test is performed via an additional application timer test element in the actual signal GOOSE message. Using the actual signal GOOSE message provides vital performance information and acts as a persistent confidence check. The application timer test element provokes an SER in the subscriber and is used to trigger subscriber logic to publish a return GOOSE message that contains a second application timer test element. IEEE refers to this as a ping-pong test, but it uses GOOSE messages rather than a ping command. Ping time is the duration of one direction, and pong time is round-trip.

Although the method of using SER time stamps to calculate transmission time is useful and relatively easy, the error introduced by the asynchronous processing cycles is statistically large compared to the expected values. Therefore, with existing IEDs, it is possible to get an accurate understanding of application times, but it is not possible to get a precise time duration calculation. The time duration calculations within the PIEDs have enough accuracy to confirm when the applications are working correctly and—more importantly—when they are not.

Interestingly, error values are the same for the test case where the transmission time duration is calculated by a timer in SLI2 and PD2. SLI1 and PD1 are triggered by internal logic to publish a test GOOSE signal at the top of the second each minute. As described in the SER test method, this is best done via an additional test element in the actual signal GOOSE message. A timer is started at the top of the second of each minute in SLI2 and PD2 logic and is stopped upon receipt of the test GOOSE signal from SLI1 and PD1. This test case can run permanently and act as a system self-test. These measured values are monitored, and if the transmission time exceeds a threshold value, an alarm is sent to supervisory control and data acquisition (SCADA) and displayed on the IED front panel, and an email is sent to a technician. This is an automatic method to answer Question 3 in Section II. This can only be confirmed in real time and in an ongoing fashion, as required by the network engineering guidelines via self-test mechanisms in the IEDs [1].

For IEDs that are not capable of starting logic based on clock time, such as top-of-second, it is necessary to use a variation of the IEEE ping-pong test to calculate, rather than measure transmission time. In this test, an additional test element in the actual signal GOOSE message or a separate GOOSE message is published, and a timer is started in SLI1 and PD1. SLI2 and PD2 are programmed to immediately publish a pong test signal message of their own in reaction to receipt of a test signal from SLI1 and PD1. Using this method, the round-trip transmission time is the result of the timer stopping in SLI1 and PD1 when they receive the pong test signal from SLI2 and PD2. This round-trip time is referred to as the pong transmission time, and the time divided in half is the ping transmission time, which is an approximation of a single-direction signal transmission time. The ping-pong transmission time calculation is also done automatically in IEDs capable of top-of-second logic execution in order to perform a real-time signal application self-test. For each of these tests, the change of state is controlled automatically for repetitive testing of large numbers of samples or by a front-panel pushbutton on an IED for in-service samples. This method is an automatic and constant way to calculate an answer for Question 3 in Section II.

VIII. Validating Signal Transfer Time as Part of an Application

Transfer time, as seen in Fig. 1, is not directly measurable in IEDs because they do not time-stamp the receipt of messages, but rather their logical reaction to the contents. Therefore, transfer time is actually a calculated value equal to the transmission time minus the duration of the IED processing cycle. This is done manually with the transmission time calculated using the SER method or automatically in the logic of the IEDs performing the ping-pong test. This calculation answers Question 4 in Section II.

Keep in mind that when additional traffic is allowed on the perimeter ports, it also affects processing in the IEDs.

IX. Validating Signal Transit Time as Part of an Application

Transit time, as seen in Fig. 1, is not directly measurable in IEDs because IEDs do not time-stamp when messages enter and leave the LAN. There are sophisticated and expensive test tools that are used to measure transit time. However, with knowledge of Ethernet switching methods, transit time can be easily calculated to answer Question 5 in Section II [4]. Be aware, however, that analyzer tools based on nondeterministic operating systems, such as Microsoft® Windows®, capture data but the time-stamp accuracy varies widely regardless of the apparent resolution, and the tools are not useful for latency and duration measurements.

It is important, however, to understand the change in transit time, if any, as a result of LAN failure and recovery. Hundreds of thousands of failure scenarios have been tested that provide enough data to answer Question 6 in Section II [4].

When using a ladder topology, the longest path that a GOOSE message travels includes two perimeter cables at 100 Mbps, three switches, and two backbone cables at 1 Gbps. Meanwhile, other Ethernet traffic is segregated so that it does not have an effect. Transit time through a correctly operating ladder topology LAN is 30 microseconds. The time to recover from failure modes varies from 1 millisecond to less than 15 milliseconds, depending on the type and location of the failure. Therefore, transit time is calculated to vary from 30 microseconds to 15 milliseconds.

However, if the LAN is based on any other design, such as a ring, transit time cannot be calculated because of the influence of other Ethernet traffic, and it needs to be tested. Reconfiguration of any other topology is much longer than 15 milliseconds for every type and location of failure. To know the reconfiguration time, it is necessary to test each possible failure scenario after installation. For non-ladder LAN topologies, onsite testing after each topology change is necessary to answer Questions 6 and 7 in Section II. However, the elegant design of the ladder topology creates an answer for both questions that does not change as the network grows and changes.

X. Validating Correct Delivery of All GOOSE Signal Messages

The only accurate way to monitor the correct delivery of GOOSE signal messages is to keep track at the receiver. The construction of the GOOSE message includes sequence numbers and state numbers to communicate when data changes and to uniquely identify each consecutive message. Each subscriber IED must monitor these parameters and record any abnormalities in signal message delivery. Fig. 8 illustrates subsets of internal IED diagnostic reports that provide information on the subscription activity to a GOOSE signal application as well as an 87L application. These internal diagnostics provide the answer to Question 8 in Section II.

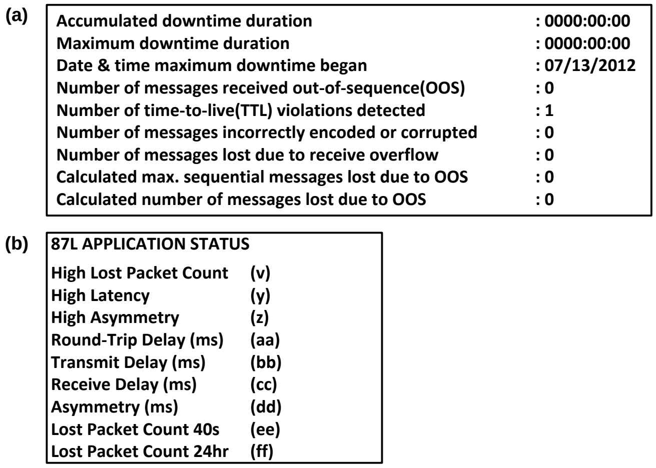

(a)

```
Accumulated downtime duration                              : 0000:00:00
Maximum downtime duration                                  : 0000:00:00
Date & time maximum downtime began                         : 07/13/2012
Number of messages received out-of-sequence(OOS)           : 0
Number of time-to-live(TTL) violations detected            : 1
Number of messages incorrectly encoded or corrupted        : 0
Number of messages lost due to receive overflow            : 0
Calculated max. sequential messages lost due to OOS        : 0
Calculated number of messages lost due to OOS              : 0
```

(b)

```
87L APPLICATION STATUS
High Lost Packet Count     (v)
High Latency               (y)
High Asymmetry             (z)
Round-Trip Delay (ms)      (aa)
Transmit Delay (ms)        (bb)
Receive Delay (ms)         (cc)
Asymmetry (ms)             (dd)
Lost Packet Count 40s      (ee)
Lost Packet Count 24hr     (ff)
```

Fig. 8. Internal IED GOOSE reception diagnostics (a) and 87L packet exchange diagnostics reports (b)

XI. Troubleshooting an In-Service System Experiencing GOOSE Problems

A. *Understanding the Symptoms*

While commissioning an upgrade to an in-service network, it was found that wide-area-distributed communications-assisted remedial action schemes (RASs) were operating less quickly than they had previously. These mission-critical applications were performed via GOOSE messages traveling from a detection device to a mitigation device over fiber-optic channels among substations hundreds of kilometers apart. The system upgrade included adding time-division multiplexers at each substation plus an additional centralized RAS.

B. *Diagnosing Traffic Among Mitigation Devices*

During the network upgrade, the GOOSE reception diagnostics report (as shown in Fig. 8) immediately indicated that not all expected GOOSE messages were being received. This was made evident from error codes, on time to live expired, and messages received out of sequence, which meant some messages were not being received by the mitigation devices. At this point, it was necessary to determine if the GOOSE messages were not being published by the source IED, not being delivered by the network, or not being received by the destination IED.

Because the source IED and destination IED had not changed, and because communications had been normal previously, the new network was investigated first. By using Wireshark, it was quickly discovered that new and unexpected traffic patterns existed at each substation. Previously unseen distributed RAS GOOSE messages from neighboring stations were now visible in addition to new centralized RAS GOOSE messages from the control center. Best engineering practices require that the last octet of the MAC address and the VLAN identifier match and be unique from any other GOOSE message and that they be used to prevent these messages from entering LANs and LAN segments where they do not belong. It was determined that this additional traffic was interfering with the distributed RAS GOOSE messages.

The unwanted distributed RAS messages had been segregated from the substation networks with the previous wide-area communications, but they were now being allowed on local network segments where they did not belong. The unwanted centralized RAS messages had been newly added to the system, but they were being incorrectly delivered to local network segments where they were not needed.

Other centralized RAS messages were being correctly delivered to the local network segments, as identified by their MAC addresses and VLAN identifiers. However, the timing and frequency of these messages were unexpected.

C. Troubleshooting the Network to Identify Root Cause of Unexpected GOOSE Traffic

The Wireshark captures revealed that the previously designed VLAN segregation was no longer working, which pointed to the new WAN multiplexers that were found to have incorrect settings for VLAN management. Once corrected, all of the unexpected distributed RAS GOOSE messages were correctly blocked, but the unneeded new centralized RAS GOOSE messages were still present on the network. Careful examination of the messages using Wireshark illustrated that these centralized GOOSE messages had been incorrectly configured to use the same VLAN tags as other system GOOSE messages. This caused the WAN to deliver needed distributed RAS GOOSE messages and unneeded centralized RAS GOOSE messages to a substation LAN because they each had the same VLAN configuration. Once these centralized RAS GOOSE messages were corrected so that all GOOSE messages in the system had unique VLANs, messages were correctly segregated and only delivered to the LANs where they were needed.

Next, by reviewing Wireshark, it was observed that after the VLAN management was corrected, new centralized RAS GOOSE messages were being correctly delivered to perform infrequent low-speed analog set point changes. However, the messages had inadvertently been configured to be sent in a very rapid burst after a set point was changed and to be repeated often. This was unnecessary and actually saturated the WAN GOOSE links because the messages were so large. Once the publication schedule was engineered to match the type of data being delivered, the WAN link saturation was corrected.

Finally, a review of the bandwidth provisioning of the WAN GOOSE links revealed that the links were too small to meet the speed criteria for GOOSE delivery. Bandwidth is often mistakenly provisioned based on throughput when networks are designed for information technology (IT) purposes. Throughput provisioning is typical and adequate for business information and often for slow SCADA systems as well. However, the throughput provisioning method calculates bandwidth by considering the total number of bits in all the messages that need to be delivered each second as bits per second. Using this method, IT staff often incorrectly provision bandwidth to be only large enough to pass the number of bits in a GOOSE message within a second, considering this as bits per second. The flaw in this method is that it creates bandwidth that may take up to a full second when delivering a GOOSE message. Operational technology (OT) methods instead calculate bandwidth based on the required speed as the number of bits in the GOOSE message divided by the required transit time. The required protective GOOSE transit time is typically 1 millisecond, which means that the bandwidth is calculated by dividing the number of bits in a GOOSE message by 1 millisecond.

In this case, once it was correctly configured, the WAN time-division multiplexing system correctly and quickly delivered all of the distributed and centralized RAS GOOSE messages in addition to all of the other substation communications.

XII. Troubleshooting GOOSE Problems During Commissioning

A. Understanding the Symptoms

During commissioning of a substation system previously staged in the factory, the system began experiencing GOOSE message quality failure. By definition, the message quality of GOOSE subscriptions is set to failed if GOOSE messages are lost, late, corrupted, in test mode, or if the configuration is changed. It was suspected that many types of Ethernet packets were being lost in the network, but only the GOOSE packet loss was being detected. These losses were only being detected by IEDs with correctly functioning message quality monitors and alarms that alerted the technicians.

The system had been tested in the factory with an Ethernet network configured by the application design team. However, the customer had contracted a separate IT group to provide and configure the substation Ethernet network. The application design OT engineers were asked to install and commission the substation IEDs, controllers, and computers by using the IT-installed Ethernet network. Because the IT Ethernet network providers did not fully understand IEC 61850 messaging, IEEE 802.1p packet priority, or IEEE 802.1Q VLAN segregation, the network was incorrectly and incompletely configured.

The IT Ethernet network provider installed and tested Layer 3 addressing, ping command message exchange, and spanning tree reconfiguration. However, the network was not configured for the pre-engineered OT IEEE 802.1Q VLAN management. Ping command messages are unique and not used in the substation systems, so testing with them is not useful and provides false confidence of performance. The spanning tree reconfiguration needed to be tested using true GOOSE messages to confirm failover times for protection speeds, which is also often misunderstood by IT Ethernet designers. After the OT application engineers correctly configured the IEEE parameters for priority and VLANs, the PIEDs still showed failed GOOSE message quality.

B. Diagnosing the Network

Similar to the situation described in Section XI, GOOSE reports immediately indicated that not all of the expected GOOSE messages were being received. Again, in this system it was necessary to determine if the GOOSE messages were not being published by the source IED, not being delivered by the network, or not being received by the destination IED. Because the source IED and destination IED had not changed, and because communications had been normal during factory testing, the new network was investigated first.

C. Troubleshooting the Network

A typical GOOSE exchange publisher and subscriber pair of IEDs that were experiencing failures were chosen. A test IED was also configured to subscribe to the same GOOSE messages being published. The application subscriber IED was put into pass-through mode so that all traffic received on the primary Ethernet port would pass through the second port, which was cabled to the test IED. The test IED showed the same missing packet behavior. Next, the test IED was directly connected to the publisher relay and it was discovered that no GOOSE messages were reported missing.

The GOOSE packets passed through four consecutive Ethernet switches, with Switch 1 connected to the publisher IED, Switch 4 connected to the subscriber IED, and the two others located between them. The test IED was moved and connected to the link between Switches 3 and 4, and a GOOSE reception diagnostics report, like the one shown in Fig. 8, revealed dropped packets. This was repeated for the links between Switches 2 and 3 and between Switches 1 and 2 with the same results.

This troubleshooting method revealed that the messages were being correctly published to a directly connected subscriber IED, but some were being dropped if a single IT-provisioned Ethernet switch was located between the two IEDs. Careful comparison of the switch port settings revealed that the IT Ethernet designers had disabled autonegotiation on the IED ports. Though the autonegotiation setting exists, it should never be used. Autonegotiation not only checks for speed settings but also duplex and crossover settings. Because autonegotiation in the IT Ethernet was disabled, the IED was unsuccessful in performing autonegotiation to the switch, but it was successful when directly connected to the test IED. Per Clause 28 of the IEEE 802.3 standard, if autonegotiation is not performed or if it fails, the IED port defaults to half-duplex. A common performance issue on 10/100 Mb Ethernet links occurs when one port on the link operates at half-duplex while the other port operates at full-duplex and packets are dropped because of the mismatch. Both sides of a link should have autonegotiation enabled to be compliant with IEEE 802.3u.

Once autonegotiation was enabled on the Ethernet switches, the subscriber GOOSE reception diagnostics report verified 100 percent GOOSE packet delivery.

XIII. Conclusion

Communications problems caused by intermittent network saturation are tricky to find and difficult to fix. Finding them requires an understanding of what is likely to be happening and the ability to decide where to perform network analysis. An understanding of what causes these problems helps ensure that all testing is performed and that all appropriate measures have been taken to prevent problems caused by high-traffic events.

Simple tools, application and test IEDs, and very specific network test devices play an important role in Ethernet network performance testing. IED features should be deployed for acceptance testing and ongoing monitoring of application behavior. However, Ethernet network reconfiguration testing requires new special-purpose test devices to verify configuration and performance. These devices must be configurable to use enough resolution and accuracy to measure true performance and automatically trigger link loss and bridge failure to collect statistically meaningful results. Also, they must use appropriate technology to verify network behavior for the specific signal message types, such as multicast GOOSE messages [4].

XIV. References

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

[2] 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.

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

[4] 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.

XV. Biographies

Marcel van Rensburg received his National Diploma in Electrical Engineering from Central University of Technology, Bloemfontein, South Africa, in 2012. After graduation, Marcel took a position at Schweitzer Engineering Laboratories, Inc. in South Africa, where he is in the position of Associate Automation Application Engineer. Marcel has experience in integration, automation, communications, and electric power protection.

David Dolezilek received his B.S.E.E. from Montana State University and is the international technical director at Schweitzer Engineering Laboratories, Inc. He has experience in electric power protection, integration, automation, communication, control, SCADA, and EMS. He has authored numerous technical papers and continues to research innovative technology affecting the industry. David is a patented inventor and participates in numerous working groups and technical committees. He is a member of the IEEE, the IEEE Reliability Society, CIGRE working groups, and two International Electrotechnical Commission (IEC) technical committees tasked with the global standardization and security of communications networks and systems in substations.

Jason Dearien received his B.S. from the University of Idaho in 1993. After graduation, he was a founding member of a small startup software contracting business. Later, he was involved in ASIC development at a fabless semiconductor company, working on compression and error correction technologies. In his 15 years at Schweitzer Engineering Laboratories, Inc., he has led various product development projects and is presently a senior application engineer in the R&D Systems Engineering department, focusing on network communications and security products.

20170127 • TP6789-01