Software-Defined Networking Redefines Performance for Ethernet Control Systems

Mark Hadley
Pacific Northwest National Laboratory

David Nicol
University of Illinois at Urbana-Champaign

Rhett Smith
Schweitzer Engineering Laboratories, Inc.

Revised edition released July 2022

Previously presented at
IEEE ROC&C 2017, November 2017

Previously published in
Sensible Cybersecurity for Power Systems: A Collection of Technical Papers Representing Modern Solutions, 2018,
and *Wide-Area Protection and Control Systems: A Collection of Technical Papers Representing Modern Solutions*, 2017

Originally presented at the
Power and Energy Automation Conference, March 2017

Software-Defined Networking Redefines Performance for Ethernet Control Systems

Mark Hadley, *Pacific Northwest National Laboratory*
David Nicol, *University of Illinois at Urbana-Champaign*
Rhett Smith, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Software-defined networking (SDN) is revolutionizing the data communications networking world by introducing the concept of programmable networking to enable system managers to react and keep up with the ever-changing demands of today's fully connected world. Surprisingly, this same technology brings great advantages to the purpose-engineered industrial control system world that controls the complex systems involved in managing critical infrastructure operations, such as the electric power grid and industrial plant systems.**

This paper outlines the three distinct advantages that SDN brings to Ethernet-based control systems: dramatically improved packet delivery performance under both normal and fault event conditions, greater cybersecurity without added complexity, and centralized situational awareness with disruptionless change control, enabling seamless scalability. The authors compare spanning tree technology with SDN technology in research efforts performed as part of a U.S. Department of Energy-sponsored research and development project that applied SDN technology to the applications and services that run energy control systems. The data show the undeniable advantages SDN has over traditional Ethernet networking.

I. Introduction

Control systems produce, manage, and monitor processes and the production of goods and services for organizations. These organizations are constantly looking for ways to reduce workplace injuries and losses, and increase productivity and system stability. Process improvements and a skilled workforce are key contributors to achieving these goals, but technological advances have contributed to some of the most significant improvements in these areas.

As technology advances, we can make more precise measurements that result in more accurate decisions, and we can do it at faster rates. The amount of data generated as a result of the automation and digital management of control systems is skyrocketing, especially when the devices distributed around the control system coordinate decisions and share data. Systems are collecting direct measurements with synchrophasors and IEC 61850-9-2 Sampled Values, and as applications continue to use these data, Ethernet networks will become even more important.

Distributed generation, for example, is a regular practice that requires significant coordination and communications to keep the system stable. Communications must be as reliable as the critical control systems at the core of the organizations. Ethernet is the world's most interoperable standard for communications and brings all the benefits of switched-packet communications infrastructure.

Critical infrastructure industries are demanding greater operational performance, challenging us to find innovative procedural and technical solutions to break through the limitations that exist today. One example is the power industry, where systems deliver power to the consumer at the speed of light. The devices that control and monitor power systems include nanosecond accuracy and submillisecond execution of contact state changes, so communication between these devices must provide the same level of performance. Traditional Ethernet cannot achieve this level of performance. However, software-defined networking (SDN) is now being applied to operational technology (OT) networks to achieve the desired performance for Ethernet-based control systems running demanding applications.

This paper shares the results of research into using OT SDN for control system local-area networking (LAN). Many aspects of performance were researched, along with the procedural and policy impacts to organizations using OT SDN compared with traditional networking. There are two main categories discussed in technology for LAN and in this paper: the control plane and the data plane. The control plane is responsible for deciding which path packets use to move through the network and reach the desired destination. The data plane is responsible for implementing the control plane plan by identifying and forwarding packets through the individual switches.

II. Today's Ethernet Performance Demands

It is a modern miracle what implementations of today's Ethernet standard have accomplished. Most of the world is connected, and we have large, widespread interoperability across manufacturers. We have proven not only that Ethernet can be used for personal and business use, but also for critical infrastructure control systems. However, the demands for person-to-machine or person-to-person communications are very different than the performance demands for machine-to-machine (M2M) communications. It is M2M Ethernet communications in control systems where we see performance demands surpassing what traditional Ethernet can provide.

Communications-assisted protection systems in the power industry are demanding faster data delivery as the data rate exceeds 100 Mbps, and these systems have become less tolerant of dropped packets. The systems rely on

communications to work the first time to save lives. A good example of such a system is an arc-flash protection system in switchgear. When the light from an arc-flash event is detected, a protective relay generates an Ethernet message and sends it to an upstream protective relay to clear the fault. If the first packet does not make it, the next packet may not be generated for another 4 milliseconds. The longer it takes for the network to successfully deliver the message to the upstream protective relay, the more current flows into the fault, increasing the damage to equipment and danger to people.

Another example is copper reduction efforts. Copper reduction decreases workplace injuries and increases productivity and system stability, but it results in greater reliability on communications infrastructure. Synchrophasors and IEC 61850 protocols [1] drove the development of devices to directly measure the system and turn the values into digitized communications using Ethernet. Looking at the performance demands for Sampled Measured values, one of the applications within IEC 61850, we have an Ethernet packet transmitted every couple hundred microseconds. Depending on the application, the protection signal may go offline and stop protecting the system if the packets stop for as little as half a millisecond. This spotlights the great demand for Ethernet networks that can provide networking healing in less than half a millisecond, even when there is a network event such as link loss or switch failure.

Other performance metrics we need to improve are cybersecurity, situational awareness, and operational management. These control systems manage critical infrastructure and therefore demand the highest level of cybersecurity. Our way of life depends on it, so the stakes cannot be higher. The good news is that control systems have the advantage when it comes to cybersecurity because the system is purpose-engineered (everything on the system is engineered to be there). This establishes the first and most important law in cybersecurity: know the system. It also enables the system owner to safely and accurately deploy allowlisted security controls, which eliminate expensive and error-prone signature and denylisted filter security.

This cybersecurity model requires purpose-engineered Ethernet networking to remove the vulnerable allow-all behavior of plug-and-play features typically found in Ethernet network devices. There are security vulnerabilities in bridged networks that either have no mitigation or complicated mitigation efforts that reduce the stability and usability of the communications infrastructure. One example is bridge priority data unit (BPDU) spoofing, where an attacker sends a malicious BPDU packet to take control of the network and become the root switch. Another example is spoofing a source media access control (MAC) address, convincing a switch to send the attacker another device's packets. There is no known mitigation to BPDU spoofing, and the only way to mitigate MAC spoofing is to deploy complicated IEEE 802.1X [2] infrastructure or perform MAC locks on each port, making it more difficult to replace field devices.

Control systems typically have the advantage of operators monitoring them to ensure they are operating as designed. Because an Ethernet network is now a component of system reliability, operators need visibility into the network to ensure it operates safely and as expected. These operators want to monitor the Ethernet network based on the applications and the service the network is supporting.

For example, when a supervisory control and data acquisition (SCADA) application is running, operators want to know when the data collection is successful at all end points and, failing that, where it is broken. The goal in any Ethernet communications outage is to restore service as quickly and efficiently as possible. The more information the network can provide the operator, the faster the operator can correct the problem. We need to understand where the packets are and what application those packets are associated with. More importantly, we need to move from the network telling the operator when it fails to telling the operator when it is about to fail so remedial action can be taken to avoid a disruption in service or an outage. This identifies a future need for programmable networks so operators can take actions beyond immediate link failure healing.

The performance demands on operational management have elevated to new heights with rapid change management and disruptionless service expectations. To reach organizational goals, new technology deployment occurs more frequently and success for many organizations depends on how efficiently they can scale the size of the control system. Any disruption in the system reduces the organization's revenue potential and, in many cases, puts people in hazardous locations. These performance demands are necessitating options for offline configuration and testing with disruptionless deployment. Adding network switches should not put operational signals at risk of dropping. When events occur and equipment does need to be replaced, there needs to be an easy way to replace it. In addition, the workforce skill level required to replace the equipment needs to be as low as possible, keeping the process economical and the error rate low.

Another operational management performance demand is the reduction of maintenance to the technology deployed on the system. A system's highest probability for misoperation or error occurs when human intervention is required, so the more that can be avoided, the higher the reliability we will achieve. Take, for example, the patch burden on technology. The fewer times we need to patch the device, the less risk of an undesired operation. The best way to reduce the patch burden is to reduce the complexity of the device. In purpose-built control systems, we have the ability to select purpose-built network devices to do the job they need to and no more instead of selecting network devices to be everything for every type of network.

III. Traditional Ethernet Performance

Network engineering today stands on the shoulders of IEEE 802.3 [3] and IEEE 802.1Q [4] to achieve interoperability. Unfortunately, with IEEE 802.1Q, we have the integration of control plane standards mixed with data plane standards, making it difficult to achieve interoperability

when newer control plane technology emerges. This has led to stagnant innovation in the core control plane features of Ethernet networks. The industry has attempted to compensate for this by continuing to layer more encapsulated or additive fields to the existing standards rather than looking at new and more efficient ways to address the root problems to achieve better performance. The performance gap has surfaced because of the real-time requirements of control system applications, and these requirements are not solved by layering more overhead data or technology on top of existing technology. A fresh look at the solution space is needed.

Reliability is critical for control systems, and with networking, that means redundancy and strong cybersecurity such as deny-by-default. Spanning tree algorithms (STA) are the dominant technology used in Ethernet networking for redundancy and to mitigate network loops. STA is a well-documented standard with a proven ability to maintain interoperability. However, in an OT network, STA limits our options for purpose-engineering the network to achieve the primary goals of the organization to build efficient and reliable networks to support all applications safely.

One attribute that limits the OT network engineer is the physical topology design. When using STA, the physical topology directly impacts network performance. How fast the network can converge and heal from failures depends on the network topology and where the network fault occurs. This makes it very difficult to predict reliability under every fault condition. With STA, the OT network engineer is forced to design the physical topology to optimize the algorithm instead of the facility layout, which eliminates topologies that would provide better application priority.

The loop mitigation built into STA technology enables redundancy, but it also blocks ports, reducing the efficiency of the switch and increasing the total cost of ownership. STA is a distributed control plane that requires switches to communicate with each other to operate correctly. This communication occurs continuously, steals bandwidth from the operational data, and potentially adds jitter to the system with priority egressing control plane packets. The switches communicate topology changes through these control plane packets and keep the network stable when network healing occurs.

Network healing performance when using STA varies depending on the topology and the STA configuration, but it is typically 5–100 milliseconds. For the majority of cases, the healing time for any link failure is 10–30 milliseconds. While this performance is adequate for information technology (IT) networks, the most demanding control system applications require network repair in under 5 milliseconds or, in some cases, under half a millisecond. STA does not meet those requirements, and layering standards on top of the existing IEEE 802.1Q will not fix it. We need to use better control plane technology.

When topology changes to an STA-based network occur, there is a risk that the convergence will disrupt existing application communications. This means that as organizations scale the network, they must plan for dropped packets. Attempting to minimize the impact of dropped packets as a result of topology changes increases network asset deployment cost and complexity.

STA also limits OT network engineers in the number of hops away from the root switch they can design for. There is a ring diameter limit of 40 switches. In control systems, applications use Layer 2 protocols to achieve faster speeds. If these applications are used across facilities, we now have the unfortunate likelihood that topology changes in one facility impact the communications in other facilities, possibly exceeding the 40-switch limitation.

Plug-and-play is a phrase used to describe how traditional Ethernet networks work. The primary goal is to deliver a packet to its destination, if possible. This certainly gets a network up and running fast because it eliminates the need for proactive traffic engineering. However, the packet is not inspected in context to determine if it should be there in the first place. This has cybersecurity impacts, but it also has reliability impacts because rogue packets resulting from misconfigured devices or personnel mistakes can impact the authorized packets on the network.

Safeguarding the introduction of unauthorized packets, devices, and applications is key in maintaining North American Electric Reliability Corporation Critical Infrastructure Protection (NERC CIP) compliance and system stability. To find the destination of a packet, traditional switches flood the packet. This duplicates the packet and sends it out to all ports, which expect the real destination to be the only one to reply or use the packet. While this gets the packet to the destination, it does so at the cost of bandwidth and data exposure to any host on the network. System stability has been compromised when this occurs. In one case, the host failed and went offline, but other hosts continued to attempt to communicate with it. After the MAC address aged out in the switches, the switches flooded the packets destined for the failed device. This overwhelmed the channel and saturated the communications, blocking critical control packets from being forwarded through the network.

Packet delivery options in traditional networking are predetermined. All unicast packets are treated the same way, just as all multicast packets are treated the same way. Traditional networking challenges the engineer in implementing selective one-to-many packet delivery functionality. Multicast packets are sent to all destinations on the same broadcast domain, and to change this behavior, the switch configuration needs to partition virtual local-area networks (VLANs) or use multicast MAC filter options. In IEC 61850 Generic Object-Oriented Substation Event (GOOSE) protocol, the host publishes tagged multicast packets and VLANs are used to subscribe to GOOSE messages the device should receive. In scaling to a larger facility, the VLAN management becomes complicated and making changes to the network becomes an expensive task. The performance demands we have for this multipurpose, overlapping application environment—which uses multicast and unicast protocols for the applications—is to provide an easy way to select the source and destination(s) when there is

a point-to-point or point-to-multipoint desired delivery. Another requirement is the ability to add or remove end points or flows without disrupting existing flows.

Priority service performance requirements have become stricter and more important as attempts are made to enable high-demand, real-time applications on a switched-packet network. Traditional networking allows the port on a switch to put priority tags on packets or to act on packets that it receives with priority tags on both the VLAN priority code point (PCP) and the differentiated services Type of Service (ToS) bits.

In today's control system applications, we have many different applications that are untagged but need different priorities. This is not possible in traditional networking because the switch cannot differentiate between untagged packets, for example, to identify the applications involved. All network appliances must be equipped to identify not only the destination and how to forward the packet, but which application that packet belongs to so it can be associated and have priority policy applied to it. This results in critical applications getting higher priority regardless of their location in the network. IT hosts do not publish tagged packets; the switch applies them. OT hosts publish tagged packets and set the PCP values on some applications. Hosts that publish tagged packets allow more granular priority per application but also introduce a new requirement for the network to check the priorities and enforce or overwrite the priorities if they are wrong or the organization wants to make a system-wide change. Traditional networking cannot perform this priority checking, but OT SDN networking can.

Cybersecurity performance requirements are demanding because of the criticality of the OT network engineering systems. Downtimes, mistakes, and exposure can be life-threatening, have significant impact on revenue, and result in federal monetary penalties. Traditional networking in LAN provides limited solutions. The most robust is the application of IEEE 802.1X network access control enforcement. The complication in applying this in an OT network environment is that the switches need to communicate to a policy server to check for and grant access to devices, and if that policy server is offline, it could impact the safety and reliability of the field applications.

To accommodate this possibility, a policy server can be placed in an unmanned remote facility, but that is not an attractive option because of the need for frequent updates and the deployment of a general-purpose Windows®-based computer in the system. The other choices are to filter a MAC address or VLAN tags. Neither have integrity protection and both can easily be spoofed. When we combine this with the plug-and-play attributes of LAN technology, we are relinquishing our top cybersecurity attribute: known purpose-engineered systems. Cybersecurity performance demands require us to maximize the use of this purpose-engineered aspect of the system and only allow known good traffic based on the packet- and path-level allowlisting.

Situational awareness performance requirements with traditional networking tend to provide topology-level awareness. The operator can see how the network is connected and how the hosts are connected into the network as well as the port status on the switch. With STA, the operator can determine in which path the traffic is flowing because of the singularity of the active paths between the source and destination. However, operators tend to lack the ability to identify which application the packets on the network belong to, so they do not have the ability to see the tiered data they have for the applications. The situational awareness demand for OT network engineering is the ability to see the network topology, hosts on the network, which applications are communicating on each host, and, by application, what the path and network load each introduce. Managed change control is also required to track all electronic actions taken on the devices that make up the network. This is typically done through Syslog and Simple Network Management Protocol (SNMP). This visibility is excellent and should continue regardless of the control plane technology selected.

IV. Software-Defined Networking

SDN is an architectural networking concept that abstracts the control plane out of the switch and centralizes it in software. This central software manages the fleet of switches in its domain. Switches become simpler as the control plane technology is removed. The deployed assets in the field are less complex, resulting in less patch management and likely fewer errors. What was not immediately apparent to the authors was all of the other advantages this simple change in architecture would bring.

The goal of two research projects sponsored by the U.S. Department of Energy (DOE) was to discover what advantages SDN could provide the OT networking industry. These two projects—the Watchdog Project and the SDN Project—brought together industry experts from academia, a national laboratory, and a manufacturer as well as multiple power system owners with the goal of bringing advanced technology to the market. The results, which exceeded the expectations of research participants, are making a significant impact in OT networking performance and security around the world [5].

SDN is an Ethernet technology and continues to stand on the established and proven interoperability of IEEE 802.3. This means that the hosts on the network do not have to change or be altered to work in an SDN network. In fact, the hosts do not know if they are connected to a traditional STA network or an SDN network. It is important to recognize what the implementation of SDN removed. SDN networks no longer use STA, so the dynamic topology discovery and loop mitigation convergence behavior are no longer required. The switches themselves do not have MAC tables, but instead have flow tables that bring that association with the packet to the application at each hop.

OT SDN has not changed the architectural concepts of SDN, nor were the standards and protocols used in SDN systems changed to fit OT systems. However, the way the technology was applied to the OT system is different than how SDN is applied today in IT networks for data centers and carrier industries. Because the standards did not have to be

altered, the interoperability between the different industry SDN solutions remains, which lays the foundation for rapid innovation. Table I shows a summary of how SDN is applied to an OT network versus an IT network.

TABLE I
OT SDN VS. IT SDN

Key Attribute	OT SDN	IT SDN
Network state	Persistent	Dynamic
Network control	Purpose-engineered	Traffic-reactive
Controller purpose post-switch deployment	Monitor	Control
Security	Deny-by-default	Plug-and-play
Fault healing speed	Link detect	Flow setup time
Network management	Traffic-reactive	Fault-reactive

OT SDN proactively engineers the network flows and the redundancy so that all primary and failover paths are planned in advance to achieve the purpose-engineered predictable and repeatable behavior desired for control systems. This proactive traffic engineering drove the difference in how the DOE research team applied SDN to the OT network. For simplicity, the switches used were OpenFlow® 1.3 only and not hybrid STA/SDN switches, which maximizes the performance, minimizes the cost of ownership, and reduces the attack surface of the switch. In control systems, all communications to and from each device are purpose-engineered, making the network flow configuration quick to identify and enter in the switches. The switches store the flow tables, groups, and meters, so the network performance is not dependent on the flow controller being online. This eliminates a potential single point of failure with the flow controller. The rate of change in the OT network is very low, and changes are only needed when devices are added or removed, or when new applications requiring new network delivery requirements are enabled. This plays very well into the proactively traffic-engineered allowlisted model of OT SDN.

The applications described in this paper use the OpenFlow 1.3 protocol to communicate between the flow controller and the SDN switch and use the top-level SDN architecture shown in Fig. 1. Tests were performed on various control systems in which the most demanding Ethernet-based network applications were enabled to capture performance results based on the criteria discussed in this paper.

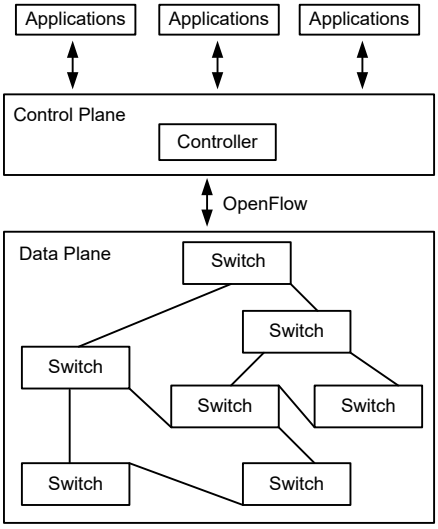

Fig. 1. SDN Architectural Overview

V. SDN PERFORMANCE

This section looks at SDN performance in the OT network applications described in this paper and compares them with traditional networking results.

A. Network Healing

SDN uses multilayer packet matching and programmable instruction sets to forward packets and avoid loops through traffic engineering rather than STAs. This allows topology-independent network healing performance. Unlike STA-based healing, SDN healing performance is uniform across all failure events regardless of where failures occur in the network or how big the network is. This removes the limitations of STA, where topologies must be optimized to achieve healing times. Now, network topologies can be designed to match what makes sense for the facility, not the STA topology.

Because the SDN switches do not have a convergence time, the switches heal on any link or switch failure in the time it takes to detect the link loss. For the switches tested in this research, healing occurred in less than 100 microseconds. This is a significant improvement in network healing time over STA-based technology, which is typically more than 10 milliseconds. This performance enables the control system owner to use even the most demanding applications where signal loss can occur if network outages extend past half a millisecond, such as Sampled Values.

In addition, this performance significantly increases safety in applications like arc-flash protection, where the worst-case scenario is that the network has a link failure at the same time an arc is detected. To understand the impact to applications when a packet is dropped, the team researched arc-flash protection with GOOSE communications. In this example, the first packet could be lost because the link could fail when the packet is traveling across it, but with SDN the network heals in less than 100 microseconds and has plenty of margin in that healing time before the next GOOSE message needs to be transmitted 4 milliseconds later. So, even if the first control packet is lost, the signal will not be lost because the second packet will be forwarded correctly down the alternate path.

As shown in Fig. 2, personnel in an arc-flash event are at an elevated risk if the network does not heal in less than 16 milliseconds. SDN allows network engineers to determine how many redundant paths they want to plan for because the fast failover groups can include many action buckets. When links come up, STA converges again to a new topology, resulting in a chance to drop packets. SDN, by design, does not drop packets because only engineered forwarding occurs. Network healing performance clearly favors SDN.

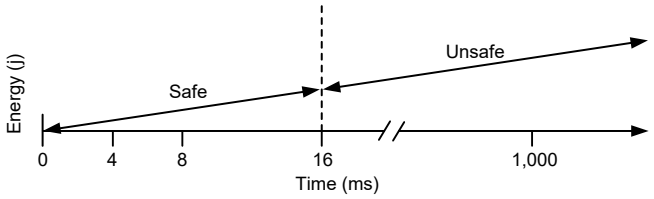

Fig. 2. GOOSE Transmit Intervals

B. Network Hop Latency

When comparing the hop latency between STA- and SDN-based technology, the STA switch hop latency measured 5 to 7 microseconds on average, and the SDN switch measured 9 to 15 microseconds. The variation in the STA-based technology was primarily the result of using different manufacturer devices, while the variants in the SDN-based technology were the result of how many flow tables the packet had to travel through before being forwarded. For every flow table the packet went through, the hop latency was extended. Network hop latency performance favors STA.

C. Packet Delivery

Packet delivery performance is measured by how fast and how accurate the technology is in determining how to forward packets to the destination. By design, SDN technology has the advantage in two areas: first, it includes multilayer matching; and second, the SDN switch did not have to learn the path to reach the destination because it was proactively traffic-engineered. The multilayer matching allowed the network engineer to identify the packets that belong to the application flow of packets by using more than just the MAC address or VLANs.

This is important because in today's multiapplication devices, we need the ability to prioritize packets ingressing and egressing each hop to best support the application requirements. Without higher layer packet matching, this is impossible. SDN allows matching on network Layers 1 through 4, which enables us to identify the application session uniquely. With proactive path planning, the switch knows how to forward the packet and does not have to learn by floods. This removes the extra packets on the network and burden on all of the network end devices.

The research team discovered an unexpected performance benefit in the ability to configure where the packet goes regardless of the packet cast type. So, regardless of whether the packet is unicast or multicast, the delivery of the packet is traffic-engineered the same way. This also benefits the performance because different paths can be used for different applications, removing potential congestion. This removes the complication of VLAN management or multicast MAC filters and the unwanted extra packets multicast generates. Packets are only delivered to the destinations that the network engineer designates. Similarly, if it is desired that unicast applications such as User Datagram Protocol (UDP) packets travel to multiple destinations, the network engineer can, with SDN, treat unicast packets as multicast and send them to many destinations. This solves the problem of engineering one-to-one and one-to-many packet delivery without replacing existing equipment or specifying expensive application upgrades. Packet delivery performance favors SDN.

D. Size and Scaling

IEEE 802.1Q limits STA to a 40-switch ring diameter. This means that Rapid Spanning Tree Protocol (RSTP) does not work past 40 hops from the root switch. In the research discussed in this paper, the network healing times slowed down the larger the network and further away from root the failure occurred. In SDN, there is no limitation on network size and no negative impact to network healing the larger the network gets. There are also no root switches and therefore no ring diameter limits. The research team discovered that because the traffic is both packet and path traffic-engineered, the larger the network got, the more path choices there were. This made configuration easier without any risk of broadcast storms because the team also had control over broadcast and multicast forwarding application by application. Network size performance favors SDN.

E. Priority Services

STA-based technology uses the following three methods to manage priority at each hop:

1) PCP in the VLAN.
2) Default port priority.
3) ToS in the differentiated services (DiffServ).

IT hosts typically do not serve VLAN-tagged packets. However, there are applications in control systems that do, so OT hosts can serve application-based tagged packets, allowing us to take advantage of using the tagged traffic direct from the host for priority services. Routers are typically the network devices that set the ToS field, but the switches have the ability to be ToS aware and read this value, mapping it to an egress queue. Then, the STA-based switch can have a default priority per port if the packet does not have a PCP or a ToS tag. When the STA technology has multiple applications using untagged

packets, they all have to be set to the same priority, so there is no way to distinguish between the applications.

SDN programmability enables network engineers to craft the priority on their network with maximum flexibility. SDN is able to work in the same way the STA technology works, but can also have default priority per flow, not per port. This means the complication of priority tagging is removed in the network and the network engineer is able to declare the priority to use for each hop and each flow directly without the extra overhead of more tags being applied to the packet. If the network engineer desires to add tags to the packet, SDN allows additional VLAN tags, the PCP value is fully configurable, and the DiffServ ToS values can be set hop by hop. This flexible programming allows the network engineer to drill into each application and confidently engineer the priority services without any burden of adding packet crafting to the end hosts, simplifying the overall system. Priority service performance heavily favors SDN.

F. Baselining

Control system industries need to design for specific applications and depend on the overall performance of those combined applications on the network to validate operations. These industries are requesting more streamlined methods to establish this baseline and confirm the network continues to perform at the approved baseline. NERC CIP requires critical facilities to have a documented baseline of network traffic and system owners that monitor the approved baseline to ensure it matches the operations.

STA-based technology relies on additional logging and monitoring technology such as SNMP, Link Layer Discovery Protocol (LLDP), and Syslog to collect data in software that computes the current state. The visibility is at the device and port level, and without additional network appliances to do deeper packet inspection, we lose the ability to identify the application flows. SDN has the central monitoring architecture already in place with the abstraction of the control plane and has counters as part of the OpenFlow standard reporting packet and byte counts for every flow, port, table, meter, and group used. This gives unprecedented situational awareness of the near real-time state of the network, allowing a designed baseline to be recorded and monitored continuously from that point forward. Baselining performance favors SDN.

G. Change Control

The research team looked at how flexible both technologies were in making changes and if they could be preconditioned to coordinate with work order processes and safeguard against impacting already operational applications when changes are applied. STA, because of is topology-dependent performance, limits the scalability of the network. When new STA-based technology is introduced into an existing network, a convergence event occurs. This event can disrupt operational applications, so system owners are burdened with extensive planning tasks and expenses every time there is a change. SDN, with its topology independence, left the scalability options open to the network engineer to do what makes the most sense to the organization. Even better, the research team discovered that adding SDN devices to scale the network out did not risk a disruption in the existing operational applications. There is no convergence of STA, so there is no control plane disruption in SDN. This reduces risk and expenses and increases safety and reliability.

SDN also provides the ability to engineer the network virtually to validate path planning. Once work crews arrive on site and install the hardware in the facility, the network engineer adopts the hardware with the virtual configuration node, turning the virtual network into a physical network. Change management performance favors SDN.

H. Performance Summary

Table II shows the top-level results of the observed performance advantages of SDN over STA. SDN technology overcomes the limitations critical infrastructure industries have had to engineer around and delivers the performance required for the most demanding applications, improving control system safety and reliability.

TABLE II
OT SDN VS. OT STA

Performance	OT SDN	OT STA
Network healing	X	
Network hop latency		X
Packet delivery	X	
Size and scaling	X	
Priority services	X	
Baselining	X	
Change control	X	

VI. CYBERSECURITY

There are two ways the team researched the performance of the technology in the area of cybersecurity: by reviewing known vulnerabilities and evaluating their impact on the focus technology, and by reviewing the security controls that exist in the technology and performing a threat modeling exercise.

Looking at STA, two vulnerabilities came to mind quickly: MAC table poisoning and BPDU spoofing. In MAC table poisoning, attackers convince the switch that they are a MAC address that they are not and receive traffic that should be sent to another device. This vulnerability still exists in the STA switch, and the only way to mitigate it is to use MAC locked port technology. The drawback in using this is that when devices need to be replaced for maintenance or failure, it costs more money because the network teams need to coordinate with the technicians to unlock and relock the MAC addresses.

BPDU spoofing is a vulnerability with no known mitigation. This is the act of sending an unauthorized BPDU into the network and taking control of how packets are forwarded on the network. It can cause complete or selective disruption in the network. Running the threat modeling exercise, the research team found the cybersecurity risk to be very dependent on the features of the STA switch because the access control is all done in the switch. The additional

technology required for that switch causes more overhead administration costs to configure, patch, and maintain the technology on every switch. For example, traditional switches integrate everything local to support a command line interface or web access. The switches forward all packets by default, so it is up to the owners to selectively choose which packets to watch for and drop. This puts the burden of security on the quality of the filter configuration.

Looking at SDN, the research team could not identify any similar control plane vulnerabilities when using the mutually authenticated and encrypted options in OpenFlow. MAC table poisoning and BPDU spoofing do not work on SDN because there are no MAC tables in the switch to spoof and the switches do not use BPDU, so any unauthorized packet is dropped. In fact, the deny-by-default architecture of SDN technology drops any unauthorized packet and only forwards packets that match proactive traffic-engineered flows.

Control plane communications in STA are unprotected. In SDN, they are encrypted and authenticated. Configuration communications in SDN occur on the control plane, and the configuration is through encrypted and authenticated OpenFlow communications, reducing the attack surface and protecting all configuration packets. This simplifies the deployed devices in the field, which reduces the total cost of ownership and increases cybersecurity greatly.

SDN provided even greater cybersecurity than expected because every packet on every hop was inspected at multiple layers—not just at the MAC address and VLAN—and was allowlisted throughout the network both at the packet and path levels. This means the packet and the path it was on had to be authorized, enforcing multiple layers of security across the logical and physical domains.

One use case the team discovered is the deployment of intrusion detection systems (IDSs). Historically, IDS platforms are hard to deploy in control systems because they demand frequent signature updates and accelerated hardware platforms to keep up with the packet load they have to process. This is because they typically hang on a span port of the switch, and all packets are sent to the IDS. Hardware that performs this level of processing is not industrial environmental-rated hardware. With the multilayer packet matching and flow associations created using SDN, we can design the deployment of an optimized OT IDS and only send packets that do not match approved and authorized flow entries. Simply put, only the packets that should not be there get deep packet inspection, eliminating false positives. Because of this, the IDS can be deployed on hardware that is industrial-rated because the packet load is manageable, which saves money.

This approach also permits the enforcement of OT protocol behavior to protect against mistakes and insider threats. If a packet is sent to the IDS, we want to know what it is because the purpose-engineered network was not engineered to have that packet on the network. There can only be two results for this: either the network engineer needs to add a flow for a packet that should be authorized or there is an unauthorized packet on the network and we need to know how it got there and how to permanently turn it off. It is very clear that SDN provides tremendous advantages in cybersecurity over STA, and its core attributes match the control system's core attributes, complementing each other as purpose-engineered and predictable allowlisted technologies.

VII. Plug-and-Play vs. Purpose-Engineered

Plug-and-play is a term used to signify ease of use and simplicity. What plug-and-play really means is that the choices for packet delivery have been made for network owners and changing the network behavior can be quite difficult if they do not like the choices the technology made for them. A good example of this is how STA mitigates loops and trunks all applications over a single link, forcing all applications to share a common path. Changing this takes additional layered technology and considerable configuration. In a plug-and-play network, the packets are not filtered but delivery is always attempted. This has obvious cybersecurity impacts, but it also has significant reliability impacts when unintended mistakes are not stopped. If they are allowed to travel to the destination or are broadcast to all hosts on the network, the applications running on the critical systems may be impacted.

An alternative is to engineer the network the same way we engineer the critical control systems in the first place: purpose-engineer it. Purpose-engineering enables network owners to configure a network to deliver the packets they want on the path they design. This is extended to event conditions as well so we can purpose-engineer the network to react to network events like link failures or switch failures in the manner we engineer them to instead of in the singular forced method STA performs based on path costs and bridge priority.

Plug-and-play does make it easy to get initial traffic running, but it also makes it more complicated to engineer specific behavior in the entire system if network engineers are not satisfied with the limited choices they are offered or want more ability to analyze traffic based on observed unwanted behavior. Purpose-engineered networks can be perceived as more difficult to configure because they are deny-by-default, dictating that the engineer know all flows of traffic on the network and design how to deliver those packets to meet the application requirements. For critical control systems, though, this is exactly what we must engineer or we will not know if the application requirements will be met in all states of the network.

To achieve the high reliability and system stability we demand for these systems, we purpose-engineer the system and the applications running on the system. Now we have the technology that enables us to apply the same professional engineering disciplines to network engineering. The result is an end-to-end purpose-engineered system allowlisted to safeguard the operations to perform the task it was engineered to do and nothing else.

VIII. Conclusion

It is an exciting time to be a network engineer. We have technology that removes the restrictions of STA, enabling us

to purpose-engineer networks and achieve performance that redefines what is possible. We have the ability to establish the programmable network infrastructure so that we can create new best known methods to deliver information between applications and services. It is rare to come across technology that improves the system performance in technical, procedural, and policy aspects without breaking existing interoperability with Ethernet. The team did successfully interoperate an STA and SDN switch after configuring the SDN switch to handle BPDUs appropriately. The team also researched the application of SDN network technology in the Parallel Redundancy Protocol (PRP) architecture, and it performed successfully. All of the listed performance advantages in this paper can be applied to PRP deployments. However, a question did arise: "With SDN achieving network healing times better than 100 microseconds, is PRP still needed?" After all, it is about maintaining the signal, not the packet.

SDN is not a lossless technology. A packet can be lost in transit if it is on the link or in the switch as it fails, or if a port buffer is overrun. By design, OT SDN is a next ingress packet-healing technology. This means that the very next packet that ingresses the switch will be forwarded correctly. No longer are there extended wait times for the network to converge before the switch knows how to forward packets again.

SDN technology allows OT network engineers to purpose-engineer their networks to support even the most demanding applications used to operate, control, and monitor our critical infrastructure. It allows the system owners to centrally monitor and deploy managed change control services without the risk of application disruption. NERC CIP compliance efforts can be supported to provide near real-time centralized reporting on the ports and services running on any network, potentially removing the time and cost of deploying crews to manually check this information.

The cybersecurity advances even force us to rethink what a subnet is and how packets should be filtered through each hop. The problems we have with flat networks today using STA are not present with SDN, opening even more possibilities and calling into question our network architecture's considered best known methods for routers and subnetting. There are certainly many more research opportunities to explore how SDN can bring additional benefits to OT network engineering, but the significant performance increases discovered in this research are already changing what is possible in OT network engineering.

IX. References

[1] IEC 61850, Communications Networks and Systems in Substations.

[2] IEEE 802.1X, IEEE Standard for Local and Metropolitan Area Networks: Port-Based Network Access Control.

[3] IEEE Standard 802.3, IEEE Standard for Ethernet.

[4] IEEE 802.1Q, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks.

[5] R. Bobba, D. R. Borries, R. Hilburn, J. Sanders, M. Hadley, and R. Smith, "Software-Defined Networking Addresses Control System Requirements," April 2014. Available: https://selinc.com.

X. Biographies

Mark Hadley received his B.S. in computer science and mathematics from the University of Puget Sound (1987). He has been a senior research scientist at Pacific Northwest National Laboratory (PNNL) since 2001. His research focus the past 12 years has been critical infrastructure protection with the Secure Cyber Systems group. Hadley has a proven track record working on collaborative projects with industry in the energy sector. He led the PNNL team transferring the Secure SCADA Communication Technology to Schweitzer Engineering Laboratories, Inc. under the Hallmark Project. He also led the PNNL team providing cybersecurity expertise and assessment experience to the SDN and Watchdog Projects sponsored by the U.S. Department of Energy (DOE). Hadley is a cybersecurity analyst providing SCADA and industrial control system knowledge to the Cyber Security Risk Information Sharing Program (CRISP). He is also involved with collaborative activities relating to the standardization of the Secure SCADA Communications Protocol (SSCP). His work in critical infrastructure protection extends to other DOE, Department of Homeland Security, Department of Defense, and private industry clients, where he provides cybersecurity expertise, training, and assessments of products and systems. Prior to joining PNNL, Hadley was a network architect for the Washington State Department of Personnel for 11 years. In total, he has 30 years of application development, network security, and critical infrastructure protection experience.

David Nicol (Fellow) received his B.A. in mathematics from Carleton College in 1979, and his M.S. (1983) and Ph.D. (1985) in computer science from the University of Virginia. Before joining the Department of Electrical and Computer Engineering (ECE) at the University of Illinois at Urbana-Champaign (UIUC) in 2003, where he is now the Franklin W. Woeltge Professor of ECE, he served on the computer science faculties at William and Mary (1987–1996) and Dartmouth College (1996–2003). At UIUC, he is director of the Information Trust Institute and also leads the U.S. Department of Energy-funded Cyber Resilient Energy Delivery Consortium (CREDC) and the Department of Homeland Security-funded Critical Infrastructure Resilience Institute. He is cofounder of the company Network Perception and is the inaugural recipient of the ACM SIGSIM Distinguished Contributions award. He was elected Fellow of IEEE in 2003 and Fellow of the ACM in 2005.

Rhett Smith is the senior product manager for the wired networks department in research and development at Schweitzer Engineering Laboratories, Inc. (SEL). He was the principle investigator and project director for the Watchdog and SDN projects sponsored by the U.S. Department of Energy. In 2000, he received his B.S. degree in electronics engineering technology, graduating with honors. Before joining SEL, he was an application engineer with AKM Semiconductor. Smith is a Certified Information Systems Security Professional (CISSP).

20220720 • TP6791-01