Software-Defined Networking Addresses Control System Requirements

Rakesh Bobba, *University of Illinois at Urbana-Champaign*
Donald R. Borries, Rod Hilburn, and Joyce Sanders, *Ameren Illinois*
Mark Hadley, *Pacific Northwest National Laboratory*
Rhett Smith, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Networking is a central, often essential, function in critical infrastructure today. Unfortunately, most existing networking-related technologies are optimized for corporate or home information technology products and not necessarily for critical infrastructure; the latter requires a different set of use cases and focuses on a different set of priorities. Specifically, critical infrastructure requires reliability, deny-by-default security, latency guarantees, and deterministic transport capabilities. Traditional Ethernet technology is unsuitable for real-time power protection communications. A completely new approach may be the best way to address these gaps. On the other hand, existing technology also provides numerous opportunities for interoperability that we do not want to lose. Hence, we need a way to reconcile these issues.**

This paper discusses the use of software-defined networking (SDN), a new architecture in networking technology, to bridge the gap between interoperability and high-reliability communications transport requirements for the power grid. We present an overview of SDN and the benefits of using this technology, and we address what challenges must be understood before this method can be adopted by the energy sector.

A project sponsored by the U.S. Department of Energy (DOE) called the SDN Project was started in October 2013 to design, develop, and test an SDN-based flow controller for the energy sector. The goal of the SDN Project, built on top of the Watchdog Project also sponsored by the DOE, is to validate whether SDN can play a part in making the energy sector more reliable and economical, as well as safer.

I. Introduction

The network world was born out of the requirement to handle multipurpose computers doing many things over a single physical communications connection. This enabled our multitasking, consistently changing, and dynamic business world to accelerate work completion to never-before-seen levels. In stark contrast, energy sector control systems were purpose-built mechanical or single-purpose-built embedded devices that for the most part had one job to accomplish. Today, these same control systems are built with multipurpose embedded devices with similar demands as corporate networking, but with different performance and priorities. So the obvious question is: can corporate networking technology be used in control system infrastructure? The answer is not without a careful design with well-understood limitations. Power industry professionals are demanding a more scalable, reliable, and easy-to-use network technology framework. Most of the network engineering goals are the same, such as black hole avoidance, loop mitigation, fast convergence speeds, priority control, and support of multiple services all running on a single physical communications channel. However, this still leaves gaps in the capabilities that the engineers designing this critical infrastructure desire that corporate networking technology does not provide. Examples of these gaps include preconfigured primary and failover forwarding paths from end to end, calculated and repeatable latency resulting in managed determinism, and system-wide detailed visualization and monitoring capability, as well as deny-by-default security at all layers of the communications system.

In searching for answers on addressing these gaps, the SDN Project team researched a growing new network architecture called software-defined networking (SDN). Based on this research, we believe that the SDN architecture allows us to keep the parts of the current network technology that work for critical infrastructure and get rid of the parts that do not, replacing them with designed solutions to solve our communications demands.

This paper is intended to highlight the advantages SDN brings to improving the reliability and cybersecurity of control system networks. SDN allows the system owners to design and maintain the network in terms of flows, which are the logical attributes that make up the communications session associated with specific applications. For example, DNP3/IP has a TCP/IP session between a protective relay and the supervisory control and data acquisition (SCADA) master. The packets that travel between the relay and the master make one flow. SDN provides strictly defined forwarding paths for each flow, better scalability, and change control while improving the situational awareness and near real-time monitoring capabilities available to the operators. SDN also allows a deny-by-default cybersecurity model. Combined together, these capabilities make SDN an attractive choice to use in control system infrastructure.

In 2011, the U.S. Department of Energy (DOE) published the *Roadmap to Achieve Energy Delivery Systems Cybersecurity* [1]. This document provides a strategy to address cybersecurity needs in the energy sector and contains the following vision: “By 2020, resilient energy delivery systems are designed, installed, operated, and maintained to survive a cyber incident while sustaining critical functions” [1]. The SDN Project addresses two goal areas within the roadmap. First, next-generation energy delivery system architectures provide “defense in depth” and employ

components that are interoperable, extensible, and able to continue operating in a degraded condition during a cyber incident. Second, collaboration between industry, academia, and government maintains cybersecurity advances.

II. SDN Defined

SDN is a new approach to the management, configuration, and operation of network systems. This architectural change is revolutionizing the management of large-scale enterprise networks, cloud infrastructures, and data center networks to better support the dynamic changes required many times a day. The reasons SDN has been adopted so much in the corporate world are also why we believe it can have a significant impact in the management of control system networks. SDN allows a programmatic change control platform, which allows the entire network to be managed as a single asset, simplifies the understanding of the network, and enables continuous monitoring in more detail. Control system networks are often more static, while the corporate world is more dynamic. That is, control system flows are more consistent and continuous than the ever-changing nature of a corporate network flow snapshot. This is primarily due to the control system being made up of machine-to-machine communications, while corporate communications are mostly people to machine. This means that the SDN architecture will be applied differently. However, the good news is that SDN architecture is able to optimize for both. The fundamental shift in networking brought by SDN is the decoupling of the systems that decide where the traffic is sent (i.e., the control plane) from the systems that perform the forwarding of the traffic in the network (i.e., the data plane).

The traditional network deployment process begins with designing the topology, configuring the various network devices, and, finally, setting up the required network services. In order to achieve the optimal usage of network resources, the application data must flow in the direction of the routes determined by the routing and switching protocols. In large networks, trying to match the network discovered path with an application desired data path may involve changing configurations in hundreds of devices with a variety of features and configuration parameters. In addition to this, network administrators often need to reconfigure the network to avoid loops, gain route convergence speed, and prioritize a certain class of applications.

This complexity in management arises from the fact that each network device (e.g., a switch or router) has control logic and data forwarding logic integrated together. For example, in a network router, routing protocols such as Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) constitute the control logic that determines how a packet should be forwarded. The paths determined by the routing protocol are encoded in routing tables, which are then used to forward packets. Similarly, in a Layer 2 device such as a network bridge (or network switch), configuration parameters and/or Spanning Tree Algorithm (STA) constitute the control logic that determines the path of the packets. Thus, the control plane in a traditional network is distributed in the switching fabric (network devices), and as a consequence, changing the forwarding behavior of a network involves changing configurations of many (potentially all) network devices.

SDN is a new architecture in networking that simplifies network management by abstracting the control plane from the data forwarding plane. Fig. 1 illustrates the building blocks of SDN, which are discussed in the following subsections.

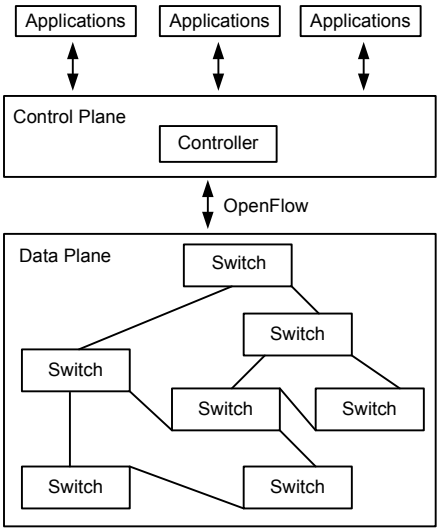

Fig. 1. SDN Architecture Overview

A. Control Plane

At the heart of SDN is a controller that embodies the control plane. Specifically, controller software determines how packets (or frames) should flow (or be forwarded) in the network. The controller communicates this information to the network devices, which constitute the data plane, by setting their forwarding tables. This enables centralized configuration and management of a network. Many open-source controllers such as Floodlight (http://www.projectfloodlight.org/floodlight/), NOX (http://www.noxrepo.org), and Ryu (http://osrg.github.io/ryu/), to name a few, are now readily available.

B. Data Plane

The data plane consists of network devices that replace switches and routers. In SDN, these devices are very simple Ethernet packet forwarding devices with a communications interface to the controller to receive forwarding information. Many vendors today provide packet forwarding devices that are SDN-enabled.

C. Control and Data Plane Interface

SDN requires a communications interface between network devices and the controller, as is evident from the description of control and data planes. A standardized interface between them will allow a controller to interoperate with different types of network devices and vice versa. The OpenFlow protocol is one such standardized interface that is managed by the Open Networking Foundation (ONF) and has been adopted by major switch and router vendors. However, it should be noted that OpenFlow is just a building block in the

SDN architecture and there are other open Internet Engineering Task Force (IETF) standards or vendor-specific standards that are either already available or are being developed.

D. SDN Services

In SDN architecture, the controller can expose an application programming interface (API) that services can use to configure the network. In this scenario, the controller may act just as an interface to the switching fabric while the control logic resides in the services using the controller. Depending on the SDN controller being used, the interfaces may be different. Controllers and their application interfaces can be tailored to meet the needs of an application domain. A controller that is designed and optimized for data centers, for example, may not be suitable for control networks in the electric sector and vice versa. The application domain specific to the industry it is used in will determine the overall system requirements. Tradeoffs between optimizations like single instruction speed or parallel processing determine the best interfaces to use.

While SDN is commonly used for monitoring and programmatically changing network configurations, the centralized nature of SDN is also well suited to meet the security, performance, and operational requirements of control system networks. Control system networks are designed to do specific jobs for many years with as little change as possible. With the help of SDN, operators can take advantage of this knowledge to preconfigure network paths and effectively create virtual circuits on a packet switching network. Power companies can design the virtual circuits they require for communication between certain devices and lock down the communications path. This type of approach can enhance security by reducing the attack surface and provide a clear approved baseline that can be continually monitored to make sure that it is never changed.

III. Challenges Not Met by Traditional Corporate Networking Information Technology Today

This section reviews the gaps in how the current corporate network technology addresses the control system network demands. There are five main categories in which we can organize the networking gaps.

The first is in the planning, design, and testing stages of new projects. The control systems that make up our critical energy infrastructure are purpose-built systems requiring the highest levels of reliability and continuous operation. These systems depend on the network to communicate between the devices doing the monitoring and control as well as between the operators and the control devices. All of these actions are pre-engineered and must strictly follow policy. The networks that carry these critical messages need to match the pre-engineered policy enforcement, high-reliability model. Designers must engineer each communications circuit and failover circuit, prove through professional engineering principles the reliability, and methodically test to make sure the system will perform all desired actions before going live.

The second category deals with change control and scalability of the network after it has been deployed and commissioned. It is desirable for energy sector control systems to minimize the amount of changes required for keeping the system operational. When changes are required, there needs to be a programmatic way to make these changes system-wide at a desired time while having the smallest impact possible to the larger system.

The third category addresses engineering the communications circuits and the required performance, as well as the tools to monitor and guard this performance. The desire is to engineer the complete forwarding path the way we engineer power delivery circuits and their failover circuits, ensuring we do not overload any segment of the circuit. Pre-engineering the forwarding circuits for all communications also brings an expectation that the forwarding path will have the same latency, providing a baseline to calculate the deterministic parameters of the messages and validate they are met for the system. Traditional networking on a switched packet infrastructure takes the approach that more bandwidth and application retries make best-effort delivery good enough. This unknown cloud approach is not acceptable for critical infrastructure. There is also a desire to maximize networking asset utilization, eliminating blocking or other degradation technologies.

The fourth category includes the continuous supervision and visualization of the entire network for operational monitoring and management. Control system operators need to monitor and respond to network conditions like they do power system conditions. To do this, they need to understand the flows on the system and the expected behavior, be alerted when those behaviors change, and have the tools and training to know what to do to get the system back to normal operating conditions.

The fifth critical category is the cybersecurity of the network. Control system networks are unmanned networks that often exist in places that are difficult to access physically. The engineers who design and deploy these systems want the capability to approve all services running on the network and deny all other flows by default. Any new communications flow should be approved before being allowed to connect.

IV. Project Initiation, Design, Deployment, and Testing Process

As we have seen, a successful network design is much more than just the technology, but includes how that technology interacts with the operators and engineers responsible for the care and maintenance of the system. So it is worthwhile taking a look at how these processes merge with the SDN technology architecture. A network project begins with the determination that a new network is required from a particular business segment. A business case is developed that includes a preliminary budget estimate that provides leadership with the ability to approve funding for the initiative.

Once preliminary funding approval has been received, a project engineer or project manager is assigned to the project. The project engineer or manager creates a project scope, which includes the necessary resources (original equipment manufacturers, a network engineer, telephone services, a server team, and so on) to design a successful project.

Once the necessary resources have been identified, the project team comes together to develop a set of requirements that are necessary for implementation of the project. The team begins with a high-level design that provides an overview of the necessary components to create an effective project. Next, an environmental assessment (including possible site visits) is performed in order to determine if sufficient infrastructure is available to support the project scope. Once the high-level design and environmental assessments are complete, a more detailed cost assessment is performed to ensure that sufficient budget dollars are available to move the project forward. The high-level design, including the more detailed cost assessment, is provided to leadership for project approval.

After the project has received final approval, the project team begins the detailed design activity. This includes identification of what each segment will require to execute the project. This involves not only the necessary equipment and systems that will be incorporated in the project, but also determination of labor resources (internal, external, and hybrid solution) for each segment. The final aspect of the detailed design process includes the development of plans on how to deploy and maintain the proposed project solution. This includes development of a construction schedule and testing plan that address the following:

- Develop implementation and backout plans. With SDN, technology owners can more easily access deployment hardware through central configuration of the controller rather than complex settings or configuration files in each field appliance.
- Perform an operational readiness assessment. SDN provides the metrics required to do thorough analysis to validate all circuits are ready for the new communications load.
- Execute a change management plan. SDN can programmatically track change orders to the person and completeness by user access control and network flow change sets.
- Receive final approval to proceed with execution of the project.
- Go live and provide implementation support. Technology owners can monitor the status of the deployment online and make commissioning changes centrally, eliminating field staff burdens.
- Verify production integrity. SDN can collect the network metrics and diagnostics centrally in near real time so the operators can validate the integrity of the system.
- Conduct post-implementation review to ensure business need has been met.
- Put tools in place to monitor; perform quarterly maintenance as needed. SDN provides continuous monitoring capabilities, and maintenance can happen as needed, rather than having to wait for quarterly work order deployments.

SDN should improve the processes and features of a new network solution by providing the following:

- Easier review and verification process to ensure that correct configurations are deployed in the new devices by using the controller to monitor all communications flows and circuit diagnostics.
- Reduction of link failure recovery times by pre-engineering primary and failover forwarding paths for each communication.
- Easier configuration of networks due to the abstraction of network tagging overhead. Configuration is done through flow paths rather than virtual local-area networks (VLANs), access control lists, Media Access Control (MAC) filters, or route tables.
- Better system-wide visualization because there is a central collection point that is visible to all network appliances.
- Baseline network configuration to verify correct configuration of the network. SDN has a central point in the controller where all forwarding paths of each communication reside, and the overall network is managed as a single asset.
- Centralized operation center that has visibility of all networks (corporate LAN, substation LAN, distribution dispatch, and so on) to more effectively manage these networks.

V. Benefits SDN Brings to Control System Networks to Address the Identified Gaps

Great care and planning go into every new power transmission and distribution circuit installed on any power system. Similarly, communications circuits must be engineered to carry messages to the intended destinations in the time frame expected, in the most reliable way. From the start, the biggest advantage that attracted the team to SDN technology is the ability to engineer all traffic flows on a circuit-by-circuit level, dictating the exact forwarding path the message travels from source to destination. Unlike the data center dynamic requirements driving the SDN revolution, the control system industry greatly benefits from its deny-by-default, circuit-based configuration that can be locked down to a very static topology. Combined with this cloud-evaporating functionality, revealing the exact circuit-based message forwarding paths is a more advanced way for operations to monitor and visually identify what is happening on the network in more real time than ever before.

Designing networks with SDN is now done by simple, physically oriented circuit design principles. This enables power system engineers to do what they are used to doing for transmission lines with communications lines and design the specific path through which they want the electrons to flow. Traffic engineering enables the network owner to have greater control over how the network operates and to maximize the network asset capabilities. No longer is there a need for dynamic negotiation protocols designating or blocking forwarding paths, but all physical ports can be used for forwarding packets. This helps balance bandwidth and segregating services, which maximizes the network asset potential.

Control system networks are deployed in unmanned locations with the desire that staff visit those sites as little as possible. Another major reason SDN holds great potential for the energy sector is the reduction of patch management on network appliances. One of the reasons that work orders are released is to patch or update electronic equipment in the field. The less patch maintenance required, the more savings the power system owner realizes. The SDN architecture reduces the amount of code required in the field network appliance because it is no longer required to manage the forwarding discovery service and control plane features. This code, in turn, resides in the flow controller and not the field network appliance. In theory, the SDN architecture reduces the amount of patch management required in the field. Simply put, the less code deployed, the less patch management required and the more reliable the system is.

Network change control is difficult to manage when dynamic protocols like Rapid Spanning Tree Protocol (RSTP) handle the forwarding decisions for the network based on physical or logical topologies and not the services running on the circuits. SDN, in contrast, allows the forwarding decisions to be based on the applications requiring communications and not the topology. The forwarding circuits are independent of the topology, meaning that no matter how the switches are connected, the configuration of the forwarding path is selected and set based on the desired transport attributes. The more connections between switches, the more options there are for application paths; there should no longer be any empty ports. Specified failover circuit engineering is just as difficult with traditional RSTP technology. SDN now enables the engineer to select the primary forwarding circuit and design the N – 1 failover circuit for a link or switch failure. This then enables the engineer to trust the design and confirm that the design handles the failure cases intended with simpler test procedures. This failover is expected to be faster in SDN due to the elimination of reconvergence times and the network topology discovery demanded by RSTP every time a link or switch fails.

Making changes to the configuration of the network when scaling to larger systems or downsizing can be tedious because a network engineer must take into consideration every network appliance that could be impacted by the change. This has traditionally been done by expensive automation software packages or homemade scripts that interface with the command line interface of each network appliance. With the abstraction of the control plane to a central location, changes to the network are simpler programmatic alterations where the changes are entered into the flow controller and, in turn, the flow controller updates all the forwarding tables in each network appliance. The key is that the controller already understands the associations between the network appliances, and it will capture all the impacted appliances and update them accordingly to support the new change. These changes can be tested ahead of time, entered, confirmed, and scheduled to be applied system-wide with little concern about the order in which the configuration changes are applied.

A huge advantage the team sees with better change control in SDN is the ability for the system owners to control when changes happen. No longer will network disruptions happen any time a cable is plugged in or unplugged at will (for example, when a technician accidentally plugs a cable into an unused port between switches), but disruptions will only happen when configuration changes are committed to the flow controller. Unused ports are off because the network is not programmed to forward any packets out of the port and any new communications attempts that show up on that port have to be allowed by the flow controller or preprogrammed for the port. RSTP will disrupt the network any time the topology changes (cables plugged in or unplugged), impacting the entire system. With SDN, only the circuits on which the link failure or change happens will be impacted. With RSTP today, negotiation and discovery happen when designated ports have failures or changes made to them; this is called the convergence process. During the time this convergence is happening, there could be communications disruptions happening on other circuits, not just on the circuit that the change was made to. This improves the reliability of the overall message delivery system.

In power systems, it is critical to understand the real-time or near real-time state of the system. This is typically done through SCADA or other state measurements like synchrophasors. These measure the state of the power flow through the system. The communications infrastructure connecting all the assets that make up the power system is just as important to monitor and control in order to keep it stable and healthy. Today, this is very difficult due to the distributed control plane architecture. Each individual network asset has its own view of the world and its neighboring network connections. Trying to piece together all of these small windows to provide the overall system state is challenging.

Vendor-centric solutions do exist to configure and monitor network switches across the enterprise. One such example is the Cisco® Network Assistant product. These tools partially meet the needs of the network administrator to visualize and configure networking equipment. Depending upon the capabilities of the software and the network switches being managed, an administrator could draw a network topology, monitor resource utilization, enable or disable switch ports, push a saved configuration to a switch, push a firmware update, or back up the configuration of a switch. While all of these features sound wonderful, they do not contain a full

feature set. For example, vendor software typically only functions with network switch equipment from the same vendor. Another shortcoming is the need to individually configure switches; there is no standard configuration for all switches. With SDN, the vendor differences are unified under the forwarding table rule syntax governed by the protocol communicating between the network appliance and the flow controller. For example, if multiple vendors support OpenFlow 1.3, the various vendor products can all be programmed by the same controller. The operator entering the configurations for the network does not need to know if the network appliances in the field are all from the same vendor or from various vendors. This improves supply chain security and allows the system owners to always use the best hardware technology at the time of purchase.

With SDN, there is a single point of control for the forwarding across all network appliances and the system owners have a global view of the entire network and can monitor it as a single asset. This visualization advantage provides system-wide operational views of what communications are allowed, where they are, and what path they are taking to get to the destination. This takes the complex nature of the interconnected networks and provides a method for structuring and maintaining order. SDN also provides several advanced monitoring and troubleshooting capabilities that were either not possible or difficult to achieve with traditional networks. These advanced features include the following:

- Mirroring any selected flow rather than the whole port.
- Alarming on bandwidth when it gets close to saturation.
- Providing many metrics for each flow. In SDN, these metrics are counters and meters. Selecting just a few meters, for example, provides functionality such as quality of service or rate limiting, and counters track counters like packet counts, errors, drops, or overruns.
- Allowing the operators monitoring the communications infrastructure to think of applications instead of VLANs or MAC addresses.

It is much easier for people to ask "where are all my DNP3 flows?" instead of "where are all my VLAN 100 ports?" Any time we remove a potential translation error, the system becomes more reliable. What this means is the visualization of the network is not limited to Layer 2 or Layer 3; it is not limited to layers at all. Operators will have the ability to engineer all the virtual circuits that every communications flow travels on, preconfigure response actions to events, monitor communications flows, and react to undesired behavior to keep the critical systems operational. The technology will provide operators a quick visual representation of what happened, which communications are impacted, and how they are impacted. This ranges from which wire was cut to which network segment is experiencing a denial-of-service (DoS) attack.

Abstraction is not new to the power industry; take, for example, Sampled Measured Values (SMVs). This is the abstraction of the analog-to-digital conversion from the applications and services that use these data. The goal is to shorten the deployment time for new services to be applied to the power system. The adoption of new services is slow in the energy sector because the threat of unintended consequences negatively impacting the system is too great to take any chances. However, if the new services could be applied in such a way that they would, by design, not impact the live system, these new services could more quickly be applied. SDN is similar to the SMV abstraction, where any new service that runs on network metric data can be applied to the flow controller and can harvest the data without the threat of the new service impacting the live communications. This also eliminates the need to upgrade the firmware of all the field-deployed network devices to realize the new service.

It is important to address the impact to the network if there is a controller failure or the controller is unable to communicate with the network appliance. In this case, the network appliance continues to operate normally and reliably forwards all approved communications. The only impact to the system is that when unconfigured new communications start, the communications will not be forwarded. Typically, new applications should only appear when new devices are added to the network. This is a very controlled deployment in the energy sector and should be planned in advance.

Cybersecurity is another reason the team is excited about the positive impact SDN will have on energy sector networks. The system owners will finally have a deny-by-default network access control solution for flows of traffic, not just MAC addresses and ports. Once again, the SDN technology is not limited to network Layer 2 or Layer 3 of security controls and is established more by thinking about communications flows between hosts and what type of flows are allowed based on the many attributes of that flow. Any flow the switch has not seen before is sent to the controller for approval before being allowed to be forwarded. Not only does this safeguard the system from rogue flows, but it enables the system operators to see when devices are plugged into the system, where they are, and what they are trying to do. Flows can be dropped, altered, or recorded. SDN provides the ability for the communications network to be baselined and monitored. Response to cyber intrusions can be predefined to keep critical systems operational. The biggest cybersecurity advantage for the power industry is that it is very static in nature, allowing the asset owners to baseline known good states and monitor these states to ensure that they do not change.

VI. Cybersecurity With SDN

This paper has discussed many ways SDN provides advanced cybersecurity benefits over traditional networking technology. For the most part, the power of this cybersecurity is in the fact that the engineering and operations team can configure exactly what communications flows should be on the network and what path they take and deny all other flows.

This follows the security practice to know the system well, baseline the known good state, and watch for changes.

The SDN architecture allows for a dominant whitelist security model but also supports blacklisting, as the team has discovered. With integration of tools like Snort®, this is not only possible but easy. This whitelist and blacklist approach is even more simplified by the ability SDN gives the end user to manage communications by flow and not packets, making it easier to understand and manage long term. SDN does provide the ability to make changes to egress packets, enabling the system operators to predetermine response actions to take on certain intrusions or reliability events. There are two methods that can achieve this. Method 1 in Fig. 2 shows the Snort deep packet inspection (DPI) engine connected to the interface of the flow controller. The process for Method 1 is as follows:

- The switch identifies DNP3/IP and sends all packets to the Snort server.
- Snort examines DNP3/IP for approved use.
- Snort informs the flow controller regarding how to handle the DNP3/IP flow (e.g., drop).
- The flow controller pushes the action to the switch.
- The switch performs the action (e.g., drops the traffic).

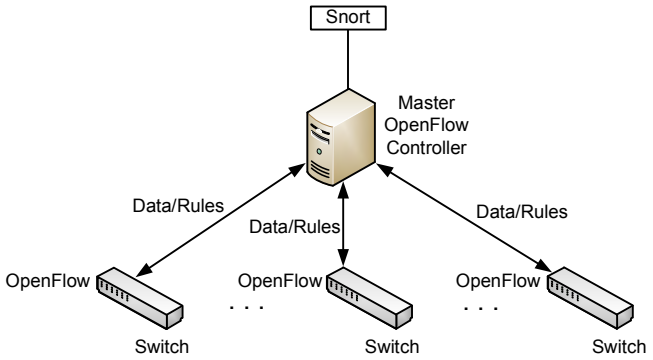

Fig. 2. Centralized Deep Packet Inspection Engine

Method 2 shows the Snort DPI engine local to the network appliance and the flows to it controlled by the flow controller. We believe this will improve the latency and throughput performance of the DPI functionality. Method 2 is shown in Fig. 3 and works as follows:

- The switch identifies DNP3/IP and sends all packets to the local Snort server.
- Snort examines DNP3/IP for approved use.
- Snort informs the local OpenFlow controller regarding how to handle the DNP3/IP flow (e.g., drop).
- The local OpenFlow controller pushes the action to the switch.
- The switch performs the action (e.g., drops the traffic).

Note that for Method 2, the local OpenFlow controller is configured as a redundant flow controller and contains a copy of the master flow controller configuration. The local OpenFlow controller is the failover for the master of the local site in the event of a loss of communication. The two methods may be used together, where the rules pushed to the network appliance are traps for specific vulnerabilities and the central DPI engine controls the whitelisted communications.

Having the ability to visualize the entire network as a single asset is a huge cybersecurity safeguard. This allows the operators to monitor and accurately react to disruptions or changes. SDN allows the right subject matter expert to consume the information. Getting the right data in the hands of the right person so the right decisions can be made is critical. For example, SDN allows operational health data to be fed to control room operators so that when there is a link failure or SCADA event, they are the ones to determine the correct trouble ticket to issue. However, if there is a DoS attack or new devices appear on the network, information technology staff consume that data so they can take defensive countermeasures to contain the compromised segments.

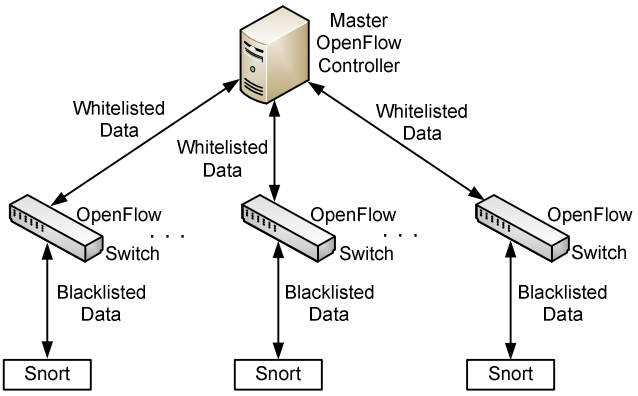

Fig. 3. Distributed Deep Packet Inspection Engine

While SDN provides many benefits over traditional networking, some of the vulnerabilities of traditional networks persist in SDN while giving rise to new problems unique to this domain. In particular, the very same features of SDN that are desirable (centralized configuration and management, for example) become attack targets. Specifically, some of the traditional network attacks can result in more significant damage to SDN because of the centralized nature of the control plane. This section discusses each of the topology blocks that make up the SDN architecture, their security challenges, and possible mitigations.

The SDN controller becomes an attractive attack target because bringing the controller down or gaining control over the controller can have significant impact on the control network. In particular, gaining control of a controller may allow an adversary to learn sensitive information about the network, make direct changes to network appliances, interfere with critical flows, and so on. Similarly, breaking the controller can impact the continued operation of services dependent on the network, especially if forwarding rules have time-out periods. Redundancy for controllers or distributed controllers can mitigate this risk, and the research community is actively working on addressing this issue.

Applications running on top of the controller would be responsible for most of the control plane functionality, such as generating critical alarms in case of a component failure or calculating backup routes. Compromise of these applications can be equally damaging if the controller does not limit the authorization privileges for such applications and provides unchecked access to the API for the network. Multiple

applications running on top of the controller can also interfere with each other and issue conflicting rules. Appropriate access control for applications can mitigate this threat, and solutions like FortNOX are available.

Transport security of the traffic between the controller and the switches can be compromised if proper precautions are not taken. An attacker can masquerade as a controller and direct switches to carry out any action (i.e., essentially take over the network). An attacker can also pretend to be a switch and send a burst of fake packets toward the controller to launch a DoS attack. Even if Transport Layer Security (TLS) or Secure Sockets Layer (SSL) is used as proposed to secure OpenFlow, many problems exist such as management of public key infrastructure, use of legacy devices in SCADA systems, or use of a vulnerable TLS or SSL implementation. The issue for secure communication becomes even more complicated if some middlebox is inserted between the controller and the switch (such as FlowVisor) for network slicing. Careful key or certificate management and the use of only cryptographically secured communications between the controller and all network appliances are the best ways to mitigate this risk.

The network appliance gained a big boost in its defensive profile simply because of the abstraction of the control plane computations. The result is less code in the device and should minimize the number of protocols it has to support for engineering access or configuration. By funneling all input configurations through only a couple of interfaces, the defensive protection and monitoring technology can be focused. As identified previously, removing the control plane functionality from the field boxes and centralizing it in the flow controller reduce patch management for those field devices and the risk of unintended changes during upgrade processes.

VII. Conclusion

Any advancement in the energy sector must start with the business need and either support the existing reliability and safety standards or improve them. Networking is no different. The network needs to be engineered for the application-specific communication that the control systems in the energy sector demand while maintaining the highest levels of reliability. The business need for increasing safety and reliability while driving operation costs down requires a more centralized technology and more informed workforce. SDN is a promising technology in that it supplies both central change management and visualization while allowing the workforce to configure, test, and maintain the network with a service-oriented, not packet-oriented, mentality. It also blends the worlds of engineering power lines or pipelines with networks. Engineers can apply the same principles of design and validation to flows of electric power, oil, or communications. Moving electrons or packets from Point A to Point B becomes an engineering solution that can be designed and tested to N – 1 or N – 2 conditions and that can have performance metrics measured before being applied to the live system. It is easier for operators to monitor and react to services and flows alerting them to applications and direct root cause rather than to more abstract events like RSTP convergences or link bounces.

Looking into the future, business demands are accelerating with demand response and remedial action automation, so technology must enable scalability. SDN provides the links needed to plug in more application services to harvest the network metrics or cloned data as they are invented while abstracting the impact to the live system. This enables the system owners to more aggressively apply new software tools without the threat of live system impacts. This may lead to predictive automation preventing communications outages or failing over to alternate paths before the primary path is down (resulting in zero packet loss), taking us to a level of reliability never seen before.

The DOE cost-sharing project led by Schweitzer Engineering Laboratories, Inc. in partnership with Ameren Illinois, Pacific Northwest National Laboratory, and the University of Illinois at Urbana-Champaign is working to integrate energy sector-specific demands into an SDN flow controller that will be commercially available by 2016. This project builds on the developments of a previous DOE cost sharing project called the Watchdog Project, which is focused on research and development of an energy sector SDN-enabled Ethernet switch. The Watchdog switch is environmentally hardened and will have the SDN interfaces the SDN Project will communicate with to provide a complete SDN solution for the energy sector. The SDN and Watchdog Projects help meet the DOE *Roadmap to Achieve Energy Delivery Systems Cybersecurity* goals to have resilient energy delivery systems designed, installed, and operational to survive a cyber incident while sustaining critical functions by 2020.

VIII. Reference

[1] U.S. Department of Energy, Energy Sector Control Systems Working Group, *Roadmap to Achieve Energy Delivery Systems Cybersecurity*, September 2011. Available: http://energy.gov/oe/downloads/roadmap-achieve-energy-delivery-systems-cybersecurity-2011.

IX. Biographies

Rakesh Bobba is a Research Assistant Professor in the Information Trust Institute (www.iti.illinois.edu) at the University of Illinois at Urbana-Champaign, with joint appointments in the Electrical and Computer Engineering and Computer Science departments. He obtained his Ph.D. in 2009 from the University of Maryland at College Park. His research interests are in the design of secure and trustworthy networked and distributed computer systems, with a current focus on cyber-physical critical infrastructure. He is a member of IEEE, the IEEE Computer Society, and the IEEE Power and Energy Society.

Donald R. Borries is the Supervising Engineer at the Technology Applications Center for Ameren Illinois, located adjacent to the University of Illinois in Champaign, Illinois. His current responsibilities include the analysis of new smart grid devices and the evaluation of equipment ranging from cybersecurity to 69 kV electrical apparatus. During his career with Ameren, Donald has worked extensively in power generation, relaying protection, and substation maintenance. He received his B.S. degree in Electrical Engineering from the University of Illinois and has served with the U.S. Coast Guard as an Electronics Engineering Chief Warrant Officer 4.

Rod Hilburn is the Manager for the Technology Applications Center for Ameren Illinois. He has 28 years of utility industry experience, where he has held positions in distribution system engineering, substation design, and substation construction and maintenance. In 1985, he received his B.S. degree in Electrical Engineering from the Missouri University of Science & Technology.

Joyce Sanders received her engineering degree from the University of Missouri – Columbia in 1985. She has worked for Ameren for 29 years. She began her career in the nuclear industry supporting the Callaway Energy Center. She moved on to work in information technology in 1997 to support the email messaging environment. In 2007, she was promoted to a position in the IT Project Management Office (PMO) and received her Project Management Professional (PMP) certification in 2008. She was promoted to Supervisor of Cyber Security in 2012. She currently supervises seven Security Analysts and a Consulting Engineer, who are primarily responsible for Control Systems Security. This includes cybersecurity for various aspects of electric and gas transmission and distribution, energy center generation (including nuclear), and the smart grid advanced metering infrastructure (AMI) and meter data management (MDM) systems. She received her GIAC Security Leadership Certification (GLSC) in 2012.

Mark Hadley is a Cybersecurity Researcher in the Secure Cyber Systems group at Pacific Northwest National Laboratory. In 1987, he received his B.S. degree in Computer Science and Mathematics from the University of Puget Sound. Mark has over 25 years of experience in application development, network engineering, and cybersecurity research for critical infrastructure.

Rhett Smith is a development manager in the local-area network and security solutions research and development group at Schweitzer Engineering Laboratories, Inc. (SEL). 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. Rhett is a Certified Information Systems Security Professional (CISSP).