WHITE PAPER / **SOFTWARE-DEFINED NETWORKING FOR UTILITIES**

PRIVATE UTILITY NETWORKS LOOK FOR SOLUTIONS FOR A MORE FLEXIBLE FUTURE

BY Dan Bayouth, PE, AND **Matt Olson,** PE

Synchronous optical networks (SONET) have been utilities' gold standard for reliability on critical communication applications. Market forces are driving a migration to packet networks, and software-defined networks (SDN) will be next. What changes should utilities expect?

Electric utilities have been building fiber-optic and microwave infrastructure and utilitywide communications infrastructure for 70 years. In the last decade, the rate has accelerated.

Despite different priorities, private utility networks have long followed in the technological footsteps of public carrier networks. Public carriers constantly push to add capacity and capabilities to meet customer needs. Because private networks must be highly reliable, they cannot risk being on the leading edge of technology. They lack the size and financial influence of the much larger public carriers.

Most notably, utilities consider synchronous optical networks (SONET) the gold standard for reliability on critical applications. Dual forces, however, are driving the migration from SONET to packet networks. First, the equipment manufacturing market has largely moved away from SONET, and equipment has gone to end-of-life. Second, grid modernization has produced the need to serve more distributed applications, and SONET does not scale well for distributed applications or any-to-any communication models.

Based on new deployments over the past decade, packet-based technologies like multiprotocol label switching (MPLS) and carrier Ethernet (CE) have become the new utility network standard. While the number of MPLS transport profile (MPLS-TP) deployments in the U.S. is small compared to MPLS and CE deployments, MPLS-TP more closely matches the needs of protective relaying, and CIGRE has selected it as the standard technology for utility protection applications. These technology standards closely mirror the public carriers pushing the technology market forward, but the carriers are now moving beyond traditional packet networks that rely on standards-based protocols to signal and control the network.

Hyperscale cloud data centers developed software-defined networking (SDN) to address the issues they had scaling traditional network designs such as load balancing, traffic engineering and dynamic workload allocations in near real time. This flexibility required a common application interface between the data and control plane on a network for dynamic system control. This, in turn, allows carriers greater flexibility in serving specific needs, which is leading the rapid migration of wide-area networks (WAN) to SDN architecture.

On the surface, SDN sounds like a massive departure from the SONET and packet networks used in past and current-generation networks. In fact, SDN is more of an evolution of the packet networks deployed today and is one of the solutions utility communications infrastructure will need to handle future applications.

COMMON ARCHITECTURE

SONET, MPLS and CE technologies all have a similar fundamental architecture for network control and data forwarding, the two separate “planes” of a network.

The control plane, which is used to signal and set up the network devices (e.g., routers, switches and multiplexers), typically provides protected or secured communication that sets up and changes how the network forwards data. Control plane commands and data go to and from the network node to control network traffic flow. This can take the form of a user changing a device configuration, protocols running on nodes making decisions, or a controller pushing changes out to the network.

The data plane determines how user or application traffic is forwarded through the network. Data plane traffic passes through the network and typically originates and terminates outside the network. Forwarding decisions are made based on the rules provided by the control plane. Data plane traffic is not processed hop-by-hop at the network nodes, but forwarded to the network’s next node or to a destination outside the network. The control plane, in other words, sets up the data plane, and the data plane is responsible for forwarding applications or user traffic.

On a fundamental level, transport networks have not changed since SONET, which required overhead channels for network administration. These channels are used to make cross-connects for a data plane configured by a user or central administration server that automated the task of finding a path through the network. Rather than having a user or server locate and establish paths for the data plane, IP/MPLS uses routing protocols like open shortest path first (OSPF), intermediate system to intermediate system (IS-IS), label distribution protocol (LDP) and resource reservation protocol with traffic engineering (RSVP-TE). MPLS-TP and CE have paths set up by a controller that forms the data plane. SDN also has one or several central controllers that push changes out to network devices that create flow tables that make up the data plane.

The difference is how each technology sets up the data plane. The shift to SDN is a departure from the typical controller or protocol used on these other technologies. Instead of sending configuration changes for protocols, changing cross-connects or pushing MPLS label stacks, SDN uses an application programming interface (API) that makes a whole new level of network provisioning flexibility possible.

SDN FUNDAMENTALS

The basic premise behind SDN is that software controllers push forwarding decisions out to network devices using an API (Figure 1).

For comparison, MPLS uses several standards-based protocols (e.g., OSPF, IS-IS, LDP, RSVP-TE and Ethernet) to manage the control plane. An operator sends commands to create a path to each router using an interface, typically SNMP or SSH. The management plane is the user interface for configuring the system. This plane does not forward data but is the interface for setting up a network node.

Once configured, routers use protocols to negotiate the data plane paths with other network routers. After the new traffic path is signaled and ready, the application or user traffic can be switched or routed in the data plane. This separation is desirable because it separates the control of the network from the user traffic carried by the network. This approach is hampered by the protocols that make it work. To behave predictably, standards-based protocols must be widely adopted and thus are slow to change, rigid by design, and require long lead times for adoption.

Distributed decision-making means the entire network must learn about the network's state. Repair decisions are made on a hop-by-hop basis, which can be relatively slow and can consume a significant amount of network capacity. If nodes have different information, it can lead to incomplete or inaccurate forwarding information. Harsh environmental conditions in small buildings and cabinets, as well as limited battery power found at typical utility locations, limit the processing power of distributed controllers such as routers and switches.

SDN nodes forward data plane traffic based on flows configured by a central controller that looks like SONET

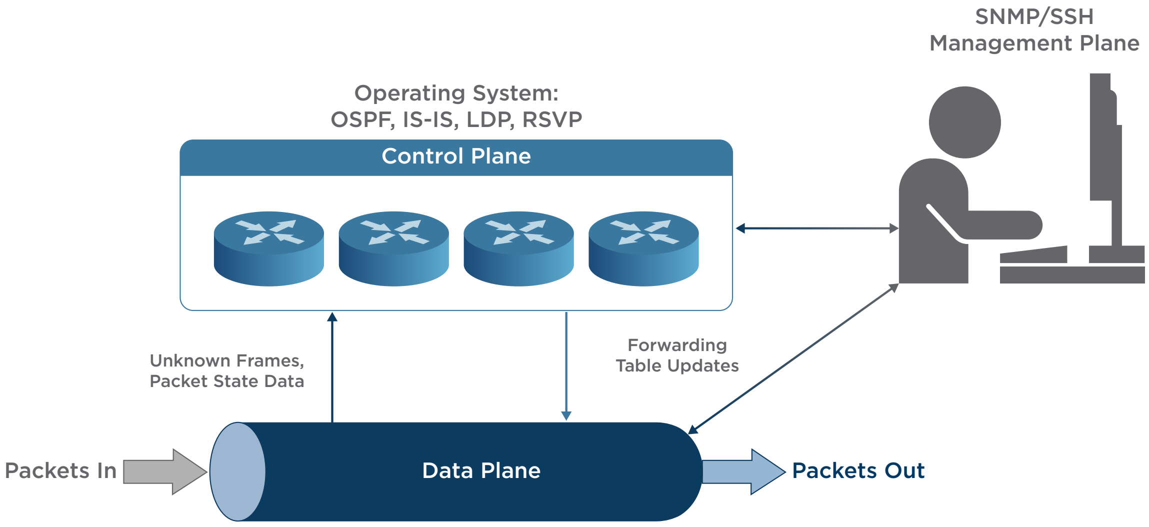

FIGURE 1: *MPLS control and data planes*

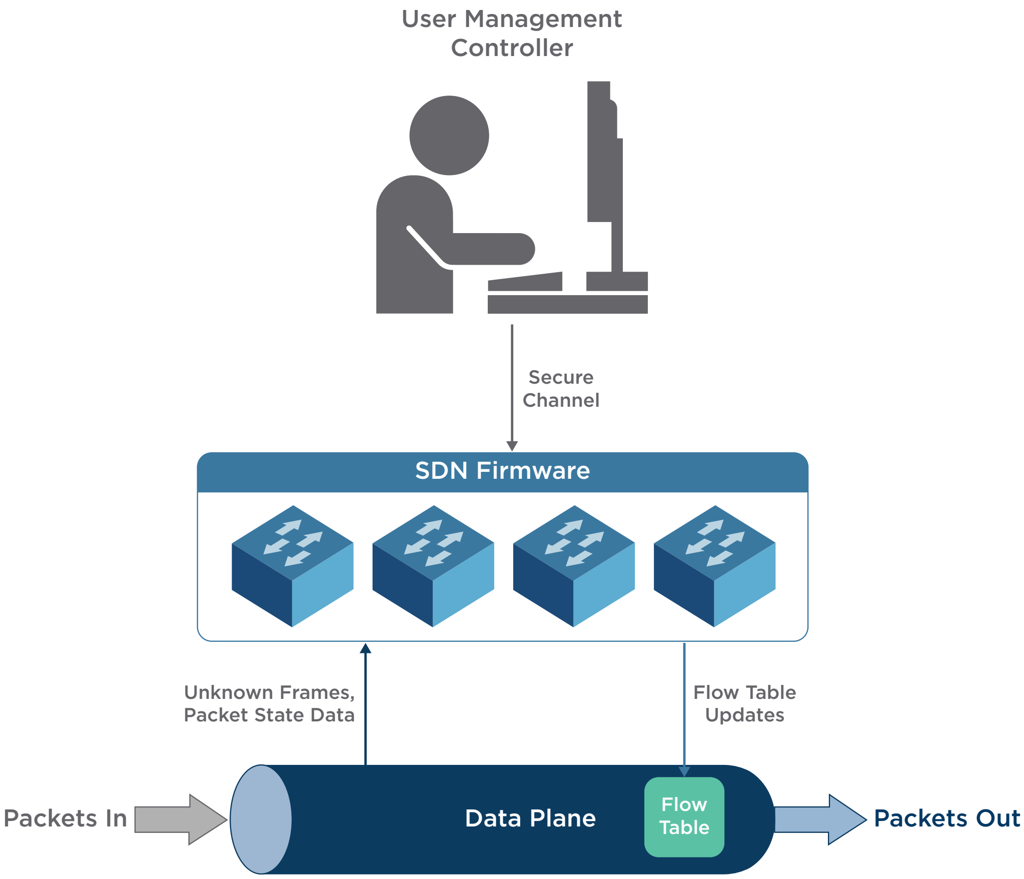

FIGURE 2: *SDN control and data planes*

and MPLS-TP architecture, rather than that of IP/MPLS or CE. Users interact with a controller to establish traffic flow rules. This interaction could be accomplished with a programming language. A utility network, however, would more likely use a set of algorithms built to understand utility application requirements, such as availability or latency for teleprotection and recloser schemes. As packets arrive at a network device, the device checks with the controller for directions on how to handle the flow. If the flow is valid, the device updates its flow table and forwards the traffic (Figure 2). If not, it may drop the traffic or forward it to an appropriate security device for further analysis.

Processing is moved to central locations like control and data centers, where processing power is readily available. Because centralized controllers can take a holistic view of the network, complex operations become fast and flexible. With a programming API, it is relatively simple to, among other things, balance traffic dynamically to reduce congested links.

Unlike SDN, traditional IP networks are designed to be robust during fault conditions and to have no central point of failure, which is why these networks used distributed control protocols. SDN controllers, on the other hand, compute a primary and at least one backup flow. Bidirectional forwarding detection (BFD) can be used to monitor traffic flow across the data plane. In the event of a forwarding failure, a backup path can be used. One example: a unidirectional path switched rings (UPSR) SONET, where a path selector at the receiving location selects the incoming data flow. To conserve capacity, most packet networks switch flow direction, rather than send multiple flows. This approach prevents them from sending duplicate copies of information on the network.

Some packet network manufacturers, however, have still implemented redundant flows with a path selector specifically for hitless teleprotection applications.

Because IP/MPLS — rather than MPLS-TP — is currently the dominant technology in modern utility networks, it deserves a special look. At a high level, MPLS data plane traffic is labeled when it comes into the network. Labels are then used to switch packets across the network, allowing a single traffic inspection at ingress and greatly simplified forwarding on subsequent nodes. Using control plane protocols like RSVP-TE and LDP, labels are established as the data plane is set up. The introduction of SDN doesn't make existing infrastructure obsolete. Segment routing (SR) makes it possible for an MPLS network to have the flexibility of SDN, but technically SR is not a pure SDN implementation.

Given the flexibility it provides for managing traffic flows on MPLS networks, RSVP-TE made a departure from traditional IP-routed networks possible. SR can do the same without the forklift upgrade required to retire SONET networks.

By replacing RSVP (or LDP) as the mechanism for establishing labels and forwarding traffic in an MPLS network, SR can look a lot like SDN for the WAN. Using a stateful path computation element (PCE), SR allows labels to be distributed by a controller instead of protocols. This effectively reduces, but does not eliminate, reliance on protocols in favor of an API-based controller that can be added to an existing MPLS node.

A stateful PCE requires "strict synchronization between the PCE and not only the network states (in term of topology and resource information), but also the set of computed paths and reserved resources in use in the network," according to IETF RFC 4655. IP/MPLS that uses RSVP-TE would have a traffic engineering database (TED) stored on the routers, with each node having its own database. A stateful PCE would have a powerful TED that reflects the entire network. This central view could be used to look for congested utilized links and move eligible traffic to possibly higher-latency, less utilized paths to free up resources. With such a TED, the network has the potential to look for specific impairments or conditions and automatically adjust to perform better within a defined rule set. For example, the rules might prevent teleprotection channel paths from being moved without approval, while other applications could.

THE FUTURE OF UTILITY NETWORKS

For electric utilities, what happens next will be similar to what happened when packet-based technologies displaced SONET in carrier networks. Public carriers continue to drive the market, and their networks are evolving. Utilities with large networks depend on products that can scale to large deployments. That means using the same technology as public carriers. API-based network control, like SR or SDN, is in utilities' future.

Utilities can benefit from the flexibility SDN and API interfaces allow. It makes it possible to stretch capacity to allow higher resolution security video, a major driver for higher speed utility networks. Private networks are designed to be very predictable, with utility-specific applications like teleprotection and SCADA being extremely predictable. Because central controllers allow better visibility into traffic all over the network, they enhance security by detecting anomalies at the source. Nonconforming traffic can be flagged immediately and sent for analysis, rather than distributed across the network. Analysis can be performed at central locations where processing power is available — something that historically has been difficult to do at scale in harsh environments.

Compared to carrier networks that serve ever-changing customer demands and skyrocketing data consumption, utility applications and needs are relatively simple. Above all, utilities demand reliability and predictability — which API-based control can provide.

Interest is growing in private LTE infrastructure as the next-generation technology for internal radio systems, mobile workforce data, distribution automation and distributed generation. As customers migrate to electric vehicles (EV) and as energy storage demands grow, the need for more and better distribution system control is inevitable. These applications will drive the need for greater communication network control and flexibility.

Another way API-based control is an evolution — rather than a replacement — of older systems is its ability to work in unison with existing MPLS or routed networks, where a central controller influences existing forwarding protocol metrics. A controller could signal new switching paths that are load-balanced, reverting upon failure to distributed control protocol for repairs, with the controller coming through to optimize the network again. In an IP-routed network, a similar approach is used wherein route metrics are influenced by the central controller via OpenFlow, while the distributed protection control remains in operation.

MOVING FORWARD

The good news is, network hardware sold today isn't like older routers that had a purpose-built processor and chipset. Rather, new hardware is commonly FPGA-based design. This means it can be repurposed without a rip-and-replace of major components. A modern MPLS router or CE switch can be loaded with SDN firmware and controlled from the same central system, or an IP/MPLS router can be updated with SR and PCE support. Current packet network equipment manufacturers have included SDN controller capabilities in some product lines. While they may not be branded as such, anything advertised as providing advanced load balancing or SD-WAN capabilities makes use of these techniques.

Another forklift SONET-to-packet conversion is not likely in the near term, nor do utilities face a pressing need to make concrete plans for the next few years. New network purchases or vendor decisions, however, should be guided by answering some basic questions, including:

- What is the SDN road map on this product line?
- Will API-based control be implemented into the existing network management software or as a new product?
- Will utility-specific application requirements like teleprotection (relay) and SCADA be built into the controller to simplify deployment, or could they reasonably be written by the utility?
- Will utility staff need programming language training to use or customize the SDN platform?

The answers to these questions can help guide decision-making. And they can help lead your utility private network into the future.

BIOGRAPHIES

DAN BAYOUTH, PE, is a project manager at Burns & McDonnell specializing in telecommunications and network engineering. His experience includes Multi-Protocol Label Switching (MPLS) WAN, substation IEC 61850, and DWDM network design construction and integration. He has served as project manager and lead architect for the design and construction of multiple utilitywide network installations that encompass NERC CIP physical and cybersecurity, protective relay transport, DC power evaluation and design, volt/VAR control, recloser automation, and automated metering infrastructure. He is a WAN design specialist with experience in network topology, quality of service (QoS) design, and network testing and acceptance.

MATT OLSON, PE, is a projects director and electrical engineer in the Networks, Integration & Automation department at Burns & McDonnell with more than 15 years of experience in wireless telecommunications systems for critical infrastructure networks, enterprise and carrier network design, converged network architecture and network management systems. He is experienced in all aspects of project engineering and management, including scope specification, consulting, design, contract administration and operational acceptance.

ABOUT BURNS & McDONNELL

Burns & McDonnell is a family of companies bringing together an unmatched team of engineers, construction professionals, architects, planners, technologists and scientists to design and build our critical infrastructure. With an integrated construction and design mindset, we offer full-service capabilities with offices, globally. Founded in 1898, Burns & McDonnell is a 100% employee-owned company and proud to be on *Fortune*'s list of 100 Best Companies to Work For. For more information, visit **burnsmcd.com**.

08616-SDN-0820