A Substation Automation Solution That Uses Virtualization to Reduce Cost While Ensuring Redundancy and Security Compliance

Sagar Dayabhai
CONCO Energy Solutions

John Prestwich
Schweitzer Engineering Laboratories, Inc.

Presented at the
Power and Energy Automation Conference
Spokane, Washington
March 6–7, 2018

A Substation Automation Solution That Uses Virtualization to Reduce Cost While Ensuring Redundancy and Security Compliance

Sagar Dayabhai, *CONCO Energy Solutions*
John Prestwich, *Schweitzer Engineering Laboratories, Inc.*

***Abstract*—Utilities are currently investigating and implementing various smart grid-enabling technologies and developing strategies that align with the new smart grid paradigm envisioned for their country's electrical infrastructure. Rugged computing platforms with fast processors and reliable storage capacities are essential for utilities to realize the benefits of smart grid technologies in modern substation automation systems. They provide the flexibility to innovate new solutions not provided by off-the-shelf hardware. The number of servers required per substation continues to grow, depending on the application type and level of redundancy required. Regardless of the flexibility computers provide, the skills needed to commission and maintain these systems often discourage a utility from embracing these technologies. Cost is also a factor, and adding the redundancy required to ensure high reliability can cause the cost to compound quickly.**

Virtualization provides an efficient, low-cost method of implementing smart grid-enabled systems and services in a zero-touch deployment approach that offers a mechanism for long-term management of the solution. Virtualization is generally described as creating a virtual device—such as a computer, network router, or switch—via software. This discipline has conventionally been widely adopted in the enterprise and informational technology (IT) environment, and its benefits of reducing cost while adding redundancy and flexibility are slowly being realized in the operational technology (OT) space.

This paper describes a real-world implementation that uses high-performance rugged computers as virtualization platforms in the substation. It also discusses the benefits derived from this application, and the challenges utilities face when implementing these systems. Moreover, this paper discusses the performance criteria of the system, the security regime, and the traffic partitioning architecture developed to facilitate the transfer of mission-critical data from a range of systems and services in a virtualized environment. We also discuss leveraging the common elements within the IT/OT space while maintaining the balance in their convergence and ensuring compliance with cybersecurity requirements.

I. Introduction

With a consistently tight balance between supply and demand, many utilities are experiencing a financial diet, and some are focusing efforts on investing in power-generating assets to meet available demand. Efforts to invest in technologies pertaining to next-generation digital substation for the transmission and distribution networks are often rejected or postponed due to implementation costs. It is, however, prudent not to lose sight of the benefits that can be derived from smart grid-enabled technologies, including demand-side management, efficiency, improved plant availability, and reliability, among other key factors.

Typical substation automation applications found in a digital substation include:

- Intelligent human-machine interfaces (HMIs).
- Phasor concentrator units for wide-area measurement (WAM) systems.
- Localized authentication servers using Lightweight Directory Access Protocol (LDAP) and Remote Authentication Dial-In User Service (RADIUS) protocol.
- Engineering workstations.
- Data historians and warehouses that collect data for analytics and condition-based monitoring systems.
- System logic processors.
- Substation gateways and data concentrators for supervisory control and data acquisition (SCADA) and automation applications.
- Distribution automation applications.
- Large-scale facility SCADA systems for monitoring and control.

These systems and services are designed and built to facilitate the following:

- Long-term operations and maintenance.
- Wide-area protection.
- Improved system availability and reliability.
- Minimized power system disruptions.
- Secure and safe operation.
- Prompt notification of alarms, events, and disturbances.
- Rapid response and diagnosis of disturbances (to minimize restoration times and improve key performance indexes).

The high capital cost with the high operating and maintenance costs over the lifetime of these applications often discourage a utility from embracing these new technologies. Furthermore, in critical transmission substations where redundancy is extremely important, the cost for provisioning redundant hardware to host all of these applications can be excessive.

However, it is possible to reduce the capital hardware costs of modern substation automation solutions by implementing proven concepts such as virtualization in a substation. This technology, if used correctly, provides significant savings and allows the utility to implement an assortment of substation automation applications.

For one high-priority pilot transmission substation in South Africa, virtualization technology was used to provide a reliable and cost-effective method of deploying substation automation applications. This project aimed to alleviate the hardware dependencies typically needed for each substation automation application. The following sections describe the details of the project.

II. Project Overview

Fig. 1 shows the network architecture of a high-priority transmission substation automation system (SAS). Typically, in the operational technology (OT) domain, an industrial computing platform is used for each substation automation application. Significant cost savings can be achieved if these applications are virtualized on reliable industrial hardware designed for the substation environment.

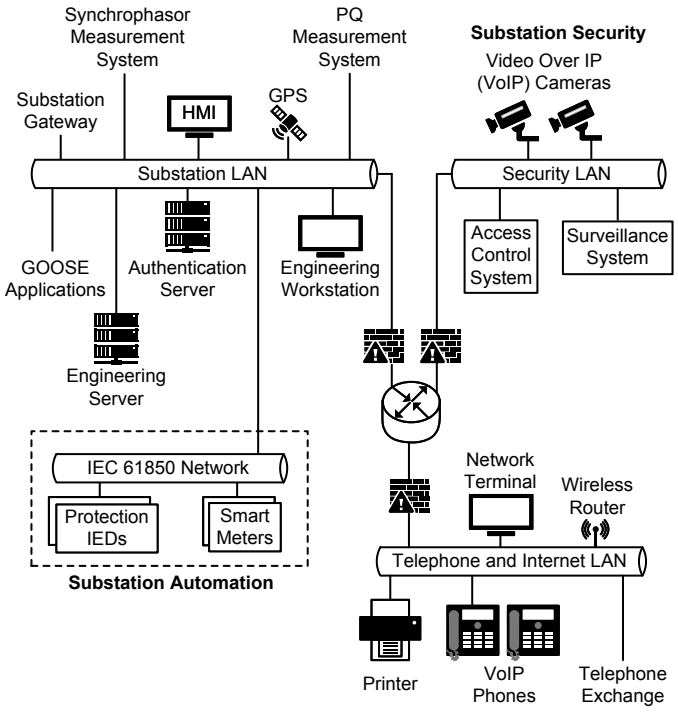

Fig. 1. Network Architecture of a High-Priority Transmission SAS

Fig. 2 shows the virtualized SAS architecture with all the substation automation applications. The project included the design, testing, commissioning, and implementation of a cost-effective substation computing platform capable of autonomously running guest virtual machines (VMs) that host various substation automation applications.

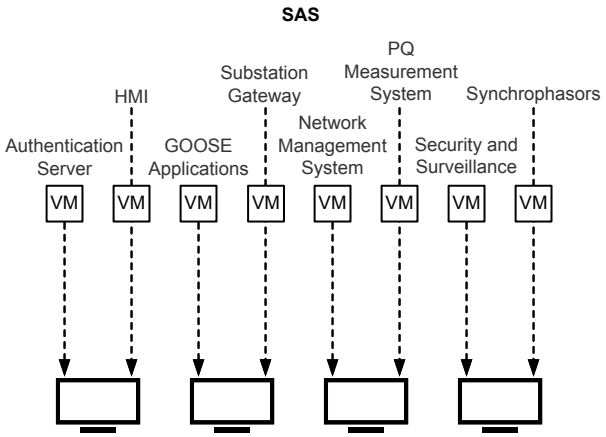

Fig. 2. Example of a Virtualized SAS

Several requirements needed to be satisfied during the implementation phase to accommodate site personnel's lack of experience with virtualization technology. The goal was to ensure that the new technology had a minimal impact on their traditional way of using these applications.

The first requirement was that the system support concurrent sessions. This requirement catered to commissioning and maintenance scenarios where multiple individuals work at the same substation and each person needs to access different applications at the same time to allow for efficient, uninterrupted work flow. Examples of this include metering technicians requiring access to the power quality (PQ) VM while the substation operator needs to access the HMI VM.

The second requirement was that the hypervisor support a user-friendly graphical user interface (GUI) to allow site personnel to seamlessly navigate and access guest VMs without using complex code in a command line shell. Further, the hypervisor solution must present a cost-effective licensing model to accommodate hundreds of substation installations.

Another requirement was that the system use a keyboard video mouse (KVM) switch to allow navigation between multiple computing platforms in a substation through a single keyboard, mouse, and screen. This alleviates the need for site personnel to use a separate terminal in the substation to access each of the automation applications.

The system was also required to be virtual local-area network-capable (VLAN-capable) to support the reception and transmission of tagged Generic Object-Oriented Substation Event (GOOSE) messages during testing and commissioning. Furthermore, it needed to support multiple VLANs to accommodate traffic partitioning for load management and security.

Finally, the system needed to operate autonomously in the substation environment, providing self-supervision of the hardware, host, and management of all VMs. It was also required to provide self-status reporting of all supervised data to the SCADA system.

III. Design Considerations

Virtualization technology requires the use of a hypervisor, which is responsible for running and managing guest VMs that each have their own independent operating system. In the OT environment, this technology is useful to run an assortment of substation automation applications and minimizes the number of required hardware devices. Each automation substation application shown in Fig. 1 operates on a separate guest VM in the substation. These VMs share the virtualized hardware resources that belong to the same high-performance substation computing platform.

During the project design phase, the hypervisor type, redundancy architecture, hardware, network architecture, and system supervision in a substation environment were critically evaluated. These design considerations are discussed in detail in the following subsections.

A. Hypervisor Considerations

Hypervisors are classified as either Type 1 (bare-metal) or Type 2. Type 1 hypervisors run directly on the host hardware and manage multiple guest VM operating systems. These hypervisors operate VMs that are used in enterprise applications. Examples of Type 1 hypervisors include VMware® ESX™/ESXi™, KVM/QEMU, Xen Project™, and Microsoft® Hyper-V™ operating on Microsoft Windows® Server 2008 and later [1]. The Type 2 hypervisor runs above a conventional operating system. An intermediary operating system is required between the hypervisor and the host hardware. Common examples of Type 2 hypervisors include VMware Workstation™, VMware Workstation Player, and Oracle® VM VirtualBox. For this project, the Type 1 hypervisor was used in the transmission substation.

It is important to understand the features and limitations of each Type 1 hypervisor model to identify which one satisfies all of the application requirements. Several challenges and limitations with current Type 1 hypervisor products arose during this exercise. This is attributed to the fact that these hypervisors were intended to be mass deployed in large data center environments and were not designed for the OT environment. These limitations and challenges are discussed later in this paper.

Both commercial and open-source hypervisors were evaluated for this project. Commercial hypervisors offer the benefit of direct support, while open-source models can offer a price advantage if the user is capable of self-support. An open-source hypervisor was ultimately chosen for this design because it satisfied a majority of the design criteria and user requirements. Several key factors were considered when selecting this hypervisor.

1) Software Compatibility and Licensing

The licensing model and cost were important considerations. Since the hypervisor is open source, it can be used for free and there is no license management.

Additionally, the supported operating systems needed to be determined [1]. A software audit was performed on all required substation automation applications to ascertain the different operating system requirements needed for the project.

2) Memory Management and Hardware Compatibility

The available memory and storage management mechanisms and how memory and storage usage is dynamically managed between each guest VM were important considerations. Additionally, the hypervisor and hardware driver compatibility, the network interface support and features (including support for multiple VLANs, network interface card [NIC] teaming, and so on), and the serial and USB pass-through capabilities were considered.

3) VM Management Features

The hypervisor live migration feature (or similar capability) was examined. This is the process of moving a running VM or application between different physical machines without disconnecting the application or VM. VM memory, storage, and network availability are transferred from the original guest machine to the destination. This is particularly useful for disaster recovery and for commissioning in the OT environment where availability is critical.

The ability to create snapshots at various instances of the VM lifecycle was also considered. This is a useful feature when the system experiences issues during new installations, software/file corruption, commissioning, maintenance, and so on.

Additionally, the flexibility to create VM templates that can be exported as images for deployment on multiple computing platforms for various substations was a consideration. This prevents the recurring installation of applications, the base configuration, operating systems, and so on for each substation.

Finally, the backup and recovery mechanisms for the VMs and the host were also considered.

4) Local Access to VMs

There needed to be support for a local, user-friendly GUI. Many hypervisors are designed to run headless, meaning they do not have a keyboard, mouse, and monitor connected to them during normal operation. This is primarily because they are designed for use in data centers. The VM manager included with the selected hypervisor allows local access to the VMs using a keyboard, mouse, and monitor. This was a feature not offered by all the hypervisors evaluated.

B. Redundancy Architecture

One of the critical design considerations for the solution was the redundancy architecture, specifically whether the system should be designed for a high-availability architecture or a disaster recovery architecture. Because of the critical nature of a transmission substation, it is common for substation designs to incorporate a high-availability redundancy architecture. This includes redundancy for intelligent electronic devices (IEDs), substation gateways, SCADA infrastructure, networking equipment, cabling, and so on.

For this project, incorporating a high-availability architecture for the transmission substation using the virtualized solution presented in Fig. 2 would have required independence from any hardware failures, protection from planned and unplanned outages [2], automated VM failover, and no single point of failure. These requirements would not have been economically viable for the solution because of the

costs and complexity. For this solution, it was more cost-effective to design for a disaster recovery scenario with adequate redundancy and hardware. Consequently, various hardware considerations were evaluated to ensure maximum reliability and availability.

C. Hardware Considerations

Adequately sizing and specifying the correct hardware platform was essential to support the architecture described in Fig. 2 for a transmission substation and to meet the user requirements mentioned previously. Several factors were considered when evaluating hardware computing platforms for the solution. A high-speed platform with adequate processing capacity and multiple cores per socket was needed to support a Type 1 hypervisor installation with a minimum of three guest VMs. The processor had to accommodate the peak requirements of all three guest VMs. A minimum range of 16 GB to 32 GB of error-correcting code (ECC) random-access memory (RAM) was required. ECC RAM can detect and correct common internal data corruption and was critical for a transmission substation where data corruption cannot be tolerated.

Flash-based nonvolatile storage was preferred over traditional magnetic-based media. This is due to the availability of industrially rated parts and the need for low maintenance (i.e., no mechanical parts to wear out). Different flash-based technologies can be used depending on the longevity needs of the application and the data rate—that is, terabytes written (TBW)—required. Single-level cell-based (SLC-based) flash provides the longest data retention and the highest TBW rating but is more expensive than other flash-based technologies such as multilevel cell (MLC) and industrial multilevel cell (iMLC). Tradeoffs can be made by considering the average data rate for the application and the maximum storage space needed.

The computing platform hardware (central processing unit [CPU], graphics, and I/O) was required to support virtualization to allow multiple guest VM operating systems to share hardware resources and directly use peripheral devices, such as accelerated graphics cards and hard drive controllers. Typical examples of hardware-supported virtualization include the CPU (e.g., Intel® VT-x and AMD-V) and graphics processing unit (e.g., Intel Iris™ Pro) [1].

Facilitating the disaster recovery architecture required that the computing platform, at a minimum, support RAID 1 (redundant array of independent disks). For this application, RAID offered a more cost-effective solution than other high-availability designs because it has minimal downtime in the event of a hard disk failure.

The computing platform was required to reliably operate, withstanding extreme temperatures, electromagnetic radiation, electrostatic discharge, and extreme vibration/shocks, as dictated in different sections of the IEEE 1613, IEC 61850-3, IEC 61000-4, and IEC 60255 standards.

The computing platform required NICs capable of VLAN tagging and trunking that could be equipped with multiple independent Ethernet ports. This would support the segregation of different systems and services per port.

The support period and services of the original equipment manufacturer (OEM) were carefully considered to ensure the sustainability of the computing platform and the solution.

D. VLAN Architecture

Networks in an SAS require careful planning and design for performance and redundancy. Mission-critical processes require data within strict time constraints, but other important data flows also need to be accommodated on the same networks. Traffic partitioning using VLANs can address data separation while traffic prioritization ensures network performance criteria are met when multiple data flows must share common trunks. Modern networks are expected to handle the transfer of data from equipment—ranging from IEDs to IP cameras and telephones—while meeting the demands of each category of data.

Fig. 3 describes the VLAN architecture implemented for this application. Access to the system is provided through the engineering VLAN. The substation router/firewall illustrated in Fig. 1 manages the inter-VLAN routing from the engineering VLAN and the security functions and policies. A separate physical Ethernet interface is used for the management VLAN, which is used for out-of-band management and host system access. All authentication and authorization is managed via the authentication server. This VM is not linked to any physical interface but is linked through an internal private virtual switch (vSwitch). The vSwitch is the first entry point for all VMs sending traffic to the physical Ethernet ports and on the network. It acts as a bridge between the logical Ethernet ports and the physical Ethernet ports.

Most hypervisors come with some form of vSwitch technology built in. For this particular project, the built-in vSwitch of the selected hypervisor did not offer all the features needed. Fortunately, the hypervisor did provide the ability to integrate other vSwitch packages with the design, and this allowed for the inclusion of an open-source, feature-rich, multilayer vSwitch that did meet the requirements.

Fig. 3 shows how the synthetic or logical Ethernet ports connect to the vSwitches.

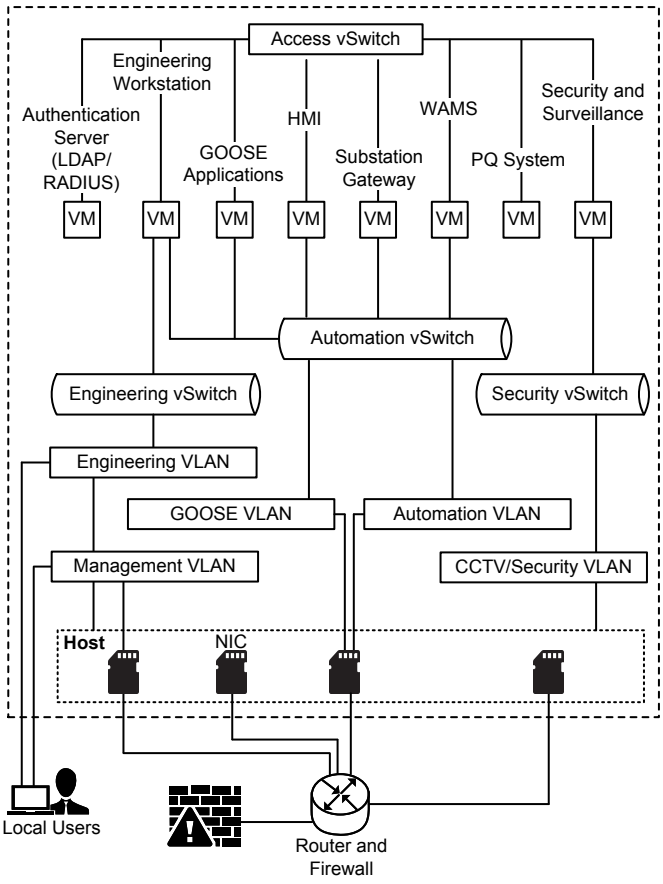

Fig. 3. VLAN Architecture

With the new vSwitch package, it became possible to configure multiple VLANs on a single vSwitch, allow GOOSE-tagged traffic through the vSwitch, and have enforced quality of service (QoS) policies on the vSwitch for traffic prioritization. Fig. 4 shows a high-level overview of the features and functions of the vSwitch [3].

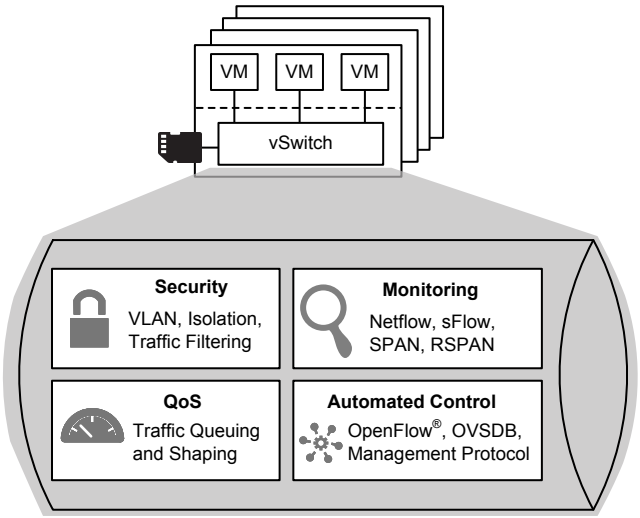

Fig. 4. vSwitch Architecture [3]

The vSwitch can also be used in substations with software-defined networking (SDN) technology and can be controlled as a node using third-party controllers and managers. However, these applications are outside the scope of this project and paper.

E. System Supervision

A critical requirement for the project was system supervision. The solution needed to operate autonomously in the substation environment and provide self-supervision of the hardware, host, and VMs. Furthermore, this supervisory information had to be reported to the substation SCADA system using an open standard telecontrol protocol.

Several third-party applications were considered for the system monitor, but none met all of the design requirements. Because of this, a custom system monitor was designed using the Python™ programming language and development environment to satisfy all the requirements for autonomous system supervision. Fig. 5 provides a high-level overview of the developed system monitor.

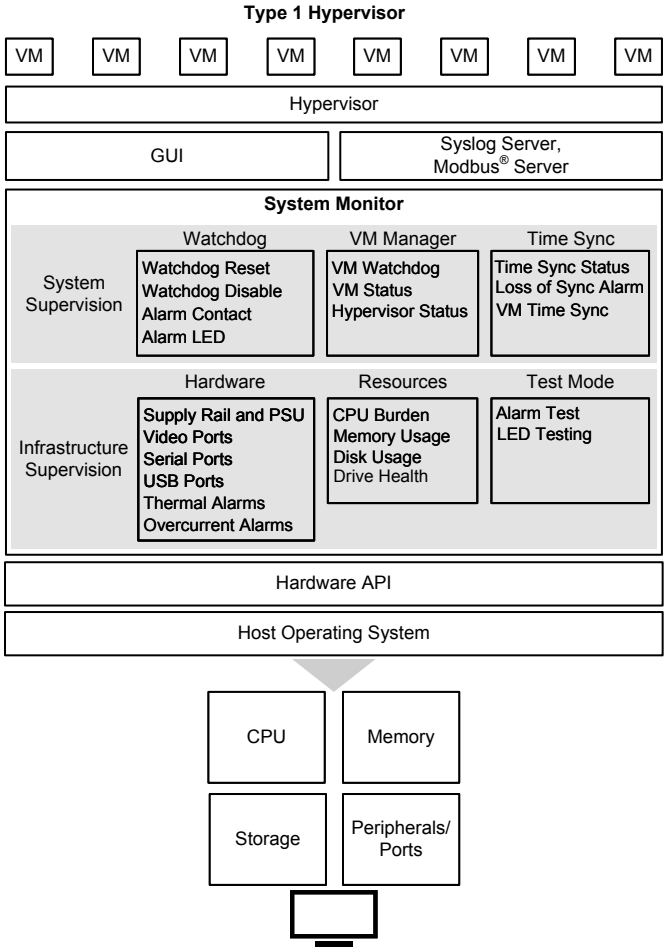

Fig. 5. System Supervision Architecture

1) Hardware Supervision

The key responsibilities of the system monitor are to configure, monitor, and report hardware health information. The hardware computing platform used for the project provided a hardware application programming interface (API) that can be used to retrieve the following real-time statuses, alarms, and measurements:

- Overcurrent alarms on peripheral devices, such as serial, USB, and video ports.
- Main power supply and voltage of all supply rails on the motherboard.
- Thermal alarms from various components of the motherboard.
- Health alarms using Self-Monitoring, Analyzing, and Reporting Technology (S.M.A.R.T) to monitor the health of all installed drives.
- Overall system loading, including CPU burden, free memory, and free disk space.
- Auxiliary light-emitting diodes (LEDs), alarm output contact, and so on.
- Hardware asset information.

Using the information retrieved from the hardware API, the system monitor GUI provides the user with the ability to configure alarm thresholds and operating ranges. It also allows the enabling and disabling of individual alarms for reporting to a variety of interfaces. Furthermore, using the hardware API, the system monitor manages the hardware watchdog system.

2) Watchdog

The system monitor is responsible for servicing the watchdog periodically to prevent it from restarting the host. If the system monitor stops or is unable to reset the watchdog due to a host operating system failure, the hardware watchdog automatically resets the hardware.

A test mode feature (shown in Fig. 5) was developed to allow site personnel to test the alarm functions, auxiliary LEDs, and watchdog functions, and to activate the alarm output contact during commissioning.

3) VM Supervision

To provide supervision of all of the VMs, thin client applications were developed that operated as a software service in each VM. These clients were responsible for communicating periodic heartbeats to the system monitor to prevent it from restarting a VM in an attempt to automatically return the VM to a good operating state.

To provide complete VM supervision, loading information including CPU burden, free memory, and free disk space is also available on a per-VM basis via the hypervisor.

4) Time Synchronization

The system monitor is also responsible for managing the system input time. Several time-synchronization options were developed to accommodate a multitude of substation automation applications. The system monitor can be configured to accept IRIG-B time for the host or the VM (for time-sensitive applications) or to accept Network Time Protocol (NTP) time. The time system is constantly supervised, and a loss of synchronization alarm is generated if the time source fails. This information can be reported via a variety of interfaces, discussed in the following subsection.

For this project, the system was configured to time-synchronize the VM using IRIG-B time received directly from the substation Global Positioning System (GPS) clock. The VM, in turn, serves as an NTP server for all the other VMs and for the host.

5) Self-Status Reporting

The system monitor includes a self-status reporting engine that can be configured to report alarms, measurements, and statuses with a variety of interfaces. These include:

- The system monitor GUI, which provides a graphical view of the system supervision statuses, alarms, measurements, and so on.
- The hardware platform ALARM output, which is configurable via the system monitor. This can be hardwired to the substation remote terminal unit (RTU) to detect alarm conditions.
- The auxiliary LEDs for visual indication for site personnel. These are configurable in the system monitor.
- A SCADA interface. The system monitor was developed to include a Modbus server with access to all the data acquired by the system monitor. This allows all the data collected by the system monitor to be reported to a SCADA system or any of the VMs using the Modbus Transmission Control Protocol (TCP), as shown in Fig. 5.

IV. Implementation Challenges and Remedial Actions

The project challenges and limitations are detailed in the following subsections, in addition to the remedial actions that were implemented to overcome these technical issues.

A. Serial Pass-Through Connections

Several IEDs in the substation use serial communications for command line interface (CLI) access, firmware upgrades, engineering and configuration, and data retrieval. Normally, Type 1 hypervisors are unable to natively perform permanent serial and USB pass-through connections due to the way the hypervisor presents these types of devices to the VMs. The selected hypervisor did allow serial port pass-through connections from the onboard serial ports of the hardware platform. Due to the limited number of onboard serial ports, an external IP-based serial device server was employed to accommodate additional serial ports. This resulted in a couple technical challenges.

First, not all serial device servers can support multiple client access (i.e., two VMs accessing the same serial device server over the same Ethernet connection but different serial ports). For example, if one guest VM required the use of Serial Port 1 for purposes of configuration and engineering and the other guest VM required the use of Serial Port 2 for event file retrieval, it became challenging to manage multiple client access to a single serial device server. Several performance tests

were conducted to evaluate and find a serial device server from an OEM that could accommodate multiple client IP connections.

Second, in cases where the application in the VM did not support Ethernet-tunneled serial communications and access to the device from the application can only be achieved using serial communications, it became necessary to use serial port servers that supported serial port emulation. This scenario was mostly for legacy applications and IEDs that could only support serial communications. These port servers emulate physical serial ports, allowing applications on an operating system direct access to the serial ports on a IP-based port server.

B. USB Pass-Through Connections

Many OEMs use USB dongles to license applications. For this reason, it was necessary to have USB support in multiple VMs to accommodate the licensing dongles. The selected hypervisor offered the feature of redirecting certain hardware peripheral devices from the host computer (e.g., the USB port) to each VM. Two methods were available on the hypervisor for passing the USB port to the VM: using the USB redirector or using the USB pass-through mechanisms. Both methods had some limitations. The USB redirector required manual mounting of the USB dongle each time the VM was restarted, making it impractical for an unmanned substation. The USB pass-through method did not have this issue, autonomously mounting the USB at startup instead. If, however, the USB device was not available or connected, the VM would not start. The second method was used for this project because it provided the desired functionality with the least availability and reliability risk.

IP-based USB devices were also considered. However, analogous to the serial port server issue, a similar issue exists for IP-based USB device servers relating to multiple client access. In a scenario where different USB ports belonging to the same device server are used by different VMs for the purpose of assessing a USB licensing dongle, it becomes necessary to evaluate the USB device servers from several manufacturers to identify a product that can accommodate multiple client access.

C. Clipboard Functionality

During the testing phase, it was discovered that the selected hypervisor did not provide any clipboard functionality. Files and folders could not be transferred between the host and the VM. USB redirection was used to transfer files to the VM during commissioning. While other hypervisors may support this functionality, it was a manageable limitation considering the overall requirements were met.

D. Graphics Acceleration and Video Port

While the selected hypervisor presented a user-friendly GUI, several graphics-related challenges were experienced during the solution implementation. One was that video port detection was not enabled at startup by the hardware platform after the hypervisor was installed. This required a BIOS firmware upgrade from the OEM to resolve the issue.

The use of a KVM switch was a user requirement, as indicated previously in this paper. Even though no issues were anticipated, there were incompatibility issues between the KVM switch software and drivers and the hypervisor. The use of this switch also compromised the calibration of the touchscreen monitor, which is managed through a USB port. It was not possible to navigate the built-in menu of the touchscreen when the screen was connected through the KVM switch.

Regardless of the resolution capabilities of the video card on the hardware platform, the resolution of the system was governed by the KVM switch. The KVM switch used for this implementation had a lower resolution than the video accelerator available on the hardware platform, which reduced the overall resolution capabilities of the solution.

One limitation that exists with hypervisors is the inability to dedicate a guest VM to a specific video port in the computing platform autonomously at startup. This is useful when a substation HMI is used on one of the guest VMs and needs to always be visible for the substation operators.

Finally, it was required that the hardware platform be connected to a screen at startup for the video card to function. Connecting the screen after startup would not render any video capabilities and would require a system restart. While this was a limitation, it was deemed manageable given the fixed substation environment in which the system would be used.

E. Software Licensing

The use of virtualization in a substation introduced a few challenges in the context of software-based licensing. A finite number of operating systems are supported by each hypervisor. Consequently, the number of operating systems required for guest VMs was carefully managed. During this project, it was discovered that the latest Microsoft Windows 10 operating system was not supported. Fortunately, none of the guest VMs required this operating system. This did not impede project progress, and it is anticipated that this operating system will be supported in the near future.

In the context of licenses for substation automation applications, OEMs were extremely cautious about provisioning licenses for a virtualized environment because of concerns about license duplication, software piracy, and so on. For this reason, many OEMs adopt the USB licensing paradigm to mitigate this risk. However, with the advent of substation automation and the next-generation smart grid network, the increase in the number of available substation automation applications will result in an increase in the number of USB dongles required.

Consequently, it is critical for software OEMs to consider how their applications can operate in a virtualized environment. Considerations when developing a licensing model for this purpose include catering backup and restoration functions to VMs and importing and exporting images and live migration scenarios.

F. Hardware Support

Additional limitations from a hardware perspective were identified during the project. One was the need for substation-class computing platforms to support the Parallel Redundancy Protocol (PRP). IEC 61850-9-2 Process Bus recommends that systems based on this standard use PRP in the substation network. IEDs, substation gateways, meters, and more are now being developed to support PRP. Computing platforms that support PRP are recommended to prevent the use of redundancy boxes (red boxes) for each computing platform that exists in a substation.

Some effort was required to find an acceptable method of sharing high-accuracy IRIG-B time between the host level and guest VM. This is particularly important for GOOSE and synchrophasor applications. With the help of the hardware manufacturer, solutions were found that accommodated the VM time-synchronization needs of this project. More work is needed to improve the virtualization of this hardware resource.

Another challenge was the limited processing capacity and cores found in substation-class servers to manage multiple guest VMs. Heat dissipation is the main limiting factor of processor capacity in a substation. This is due to the extended temperature ranges required by substation equipment and the desire to use low-maintenance cooling solutions (i.e., no moving parts).

As processor technology progresses, the number of processing cores per processor increases due to the effects described by Moore's law. Historically, this increase in the number of processing cores has come with a decrease in the power requirements of each core, giving each generation a proportionate increase in capacity over time while not increasing the power budget. This boost in performance is likely to hold true at least for a few more processor generations.

Finally, faster and higher core count processors currently do exist, but the demand for such computing capacity in the substation has not historically existed. If the demand grows due to applications such as virtualization, solutions to accommodate the increased heat dissipation requirements of these higher power processors will likely be developed to meet the demand.

G. Skills

It was realized during the project that there is a dire shortage of skills in the OT domain related to virtualization and basic networking IT domain concepts that are typically used in the corporate environment. This skills shortage impedes project progression and is a challenge that needs to be addressed to fully capitalize on the associated benefits [4].

H. IT/OT Convergence

Traditionally, the IT domains in a power utility were responsible for antivirus, software, and patch management as well as handling operating system licenses and security updates. However, with the advent of virtualization in substations, traditional methods needed to be reconsidered and a strategy to manage the cross-boundary technology deployment between IT and OT needed to be devised.

The technology and platforms used in IT and OT are becoming increasingly similar. The reality is that many modern OT systems are underpinned by platforms, software, security and communications standards that have traditionally been associated with the IT world. The implementation of this project required that the profiles and skills of employees working in the IT and OT disciplines be converged.

However, many organizations insist on strict separation of the domains from a management and operational perspective. This is largely attributed to the nature and focus of the OT environment on continuous (24/7) real-time operations and on equipment reliability. This makes the OT domain a conservative adopter of technology, which must be well-proven before it is even considered for implementation. IT, on the other hand, focuses on cost reduction, standardization, and resource optimization and can often get away with an 8 hours a day, 5 days a week approach to operations. IT can therefore afford to be more aggressive in its approach of adopting new technology.

However, during this project, several common elements were discovered that could be leveraged by both IT and OT to become proficient in technologies and save costs. Using these common elements, IT and OT applications could do the following:

- Develop common standards and policies (as applicable).
- Leverage common infrastructure and share common technology platforms.
- Share enterprise software license agreements and support contracts, such as operating system licenses.
- Use corporate contracts for hardware and software procurement (where applicable).
- Develop common governance procedures.
- Manage IP address allocations for wide-area network integration.

I. Security

The North American Electric Reliability Corporation (NERC) developed reliability standards that include a set of nine Critical Infrastructure Protection (CIP) standards. These standards address cyberasset security to ensure reliable electric grid operation. While compliance with these standards is mandatory for North American utilities, most utilities worldwide also subscribe to these standards to varying extents because they present an industry-accepted, best-practice approach to most issues governing cybersecurity.

Within the context of cybersecurity, it is essential to treat each VM and each host server with its own entry on the cyberasset register as a bulk electric system (BES) cyberasset. The security model of the solution is currently being designed and reviewed before the solution is accepted for mass deployment. As discussed previously, while there are many common elements between the IT and OT domains, the security model that needs to be adopted for this solution will follow the OT domain.

NIST SP 800-82, Guide to Industrial Control Systems (ICS) Security, states: "ICS security objectives typically follow the priority of availability, integrity and confidentiality, in that order" [5]. For an ICS to continue performing the task it was designed to perform, it must prioritize data by availability [6].

This is because ensuring that the process remains within normal operating parameters is paramount. Data availability is followed by data integrity and confidentiality. In the world of IT, the data priorities are reversed because information security is paramount.

The cybersecurity design currently includes (but is not limited to) the following key components:

- Defense-in-depth with firewall layering and a demilitarized zone (DMZ).
- Role-based access management.
- Encryption of remote access connections.
- Secure authentication and authorization.
- End-point protection with antivirus and malware protection on the host.
- Patch management and security updates.
- Operating system application whitelisting.

V. Conclusion

Modern substations continue to require increasing amounts of computing power to manage communications, automation, and HMI needs. Computing assets that meet the stringent standards required to ensure reliable operation in the substation environment are more expensive than those of their conventional counterparts, making it desirable to maximize the use of each asset. Virtualization is a well-established technology developed by the IT community that is designed to help maximize computing asset usage in data centers, but it is only beginning to be explored in the power industry.

The project described in this paper is a practical solution using virtualization that could be adopted across a diverse set of substations. It was designed to be easily deployable by exploiting virtualization features such as VM templates and snapshots. Depending on the applications required, it could produce up to a three-times reduction in the computing hardware needed for a substation. The design also enabled the system to maintain local access to applications running in VMs, allowing personnel to use them while visiting the substation to perform maintenance and troubleshooting.

During the project, there were several technical hurdles to overcome. Most hypervisor technology virtualizes the underlying hardware to provide multiple VMs access to each hardware asset (network ports, hard drives, and so on). This makes it possible to maximize the use of each asset. This project found that several hardware assets used often by substation applications (serial and USB ports, for example) did not virtualize well, requiring the development of workarounds to address these concerns. Also, the unique networking requirements of a substation required special consideration in terms of the vSwitch technology used. Fortunately, there are many choices in hypervisors and their supporting software packages, and this provided many options for working around the most critical issues.

There are many features that could be improved in both the hypervisor and the substation applications to make virtualization of these assets more effective and efficient. Among these features are:

- Better virtualization of serial and USB ports.
- Creation of more virtualization-friendly substation software licensing by providing allowances for virtualization features such as templating and snapshots.
- Adoption of licensing management systems that do not require USB dongles or other mechanisms that are hard to support in a virtualized environment.

VI. Acknowledgment

The authors gratefully acknowledge the contributions of Peter Diamandis, Bruce Mackay, and Wesley Phillips for their work on the project that inspired this paper.

VII. References

[1] S. Lowe, "Hyper-V™ vs. vSphere™: Understanding the Differences," 2012. Available: http://www.apmdigest.com/sites/default/files/images/VMvSphereHyperV_Whitepaper.pdf.

[2] B. Berger, *Hyper-V Best Practices*. Packt Publishing, Birmingham, UK, 2014.

[3] Linux Foundation Collaboration Projects, "Production in Quality, Multilayer Open Portal Virtual Switch," 2016. Available: http://openvswitch.org/.

[4] Garter Newsroom, "Gartner Says the Worlds of IT and Operational Technology Are Converging," March 2011. Available: http://www.gartner.com/newsroom/id/1590814.

[5] K. Stouffer, V. Pillitteri, S. Lightman, M. Abrams, and A. Hahn, *Guide to Industrial Control Systems (ICS) Security*, May 2015. Available: http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r2.pdf.

[6] IEC 62351, Power Systems Management and Associated Information Exchange – Data and Communications Security – Parts 1–13.

VIII. Biographies

Sagar Dayabhai is currently pursuing his PhD at the University of Witwatersrand (Wits), South Africa in the field of smart grids. He received his BSc Eng. (electrical) degree from Wits in 2009 and earned his MSc Eng. (electrical) degree in 2014. He is a registered professional engineer with the Engineering Council of South Africa (ECSA). Sagar was one of the IEC Young Professionals elected for South Africa in 2016 and is currently a member of IEC TC57 WG10 and WG15. After working at Eskom in the field of telecommunications and SCADA, he moved to Consolidated Power Projects (CONCO) as a Senior SCADA/Automation Engineer. Sagar now holds the position of System Control Manager at CONCO Energy Solutions.

John Prestwich received his BS in Electronic Engineering from Utah State University in 1994 and an MS in Computer Science from Boise State University in 2009. He has a diverse background, with experience in embedded system design, computer design, computer applications, and information technology. Upon graduating, he worked for Bently, Nevada providing custom eddy current proximity sensor designs and also worked with their SCADA collection systems. After Bently, Nevada, he worked in the automotive and integrated circuit industries designing and testing embedded computing systems. In 2009, John joined Schweitzer Engineering Laboratories, Inc. in the computing systems group. He now serves this group as a lead application engineer. His interests include robust computer system design, the leveraging of open-source software for the control industry, and computer security. John is a current member of IEEE.