VIRTUAL PROTECTION RELAY: MYTH OR REALITY ON THIS GENERATION?

Jose Ruiz
Doble Engineering

Montie Smith
Dell Technologies

Bryan Gwyn
Doble Engineering

ABSTRACT

It is more than four decades since the introduction of the first microprocessor-based protection relay and more than a century since the invention of protection relaying. Technology changes have made this possible throughout the years, and it has not stopped there.

Many years ago, the idea of a virtual world sounded like science fiction, but the virtualization concept is now a reality in the power industry and, believe it or not, applied to protection relaying.

A new era of protection relays has been born, and the virtual protection relay (VPR) is a reality today.

Digitalization of the copper wires connected to protection relays and standards such as IEC 61850 has opened the door to a new technology change.

Today, computer servers are a hardware platform where IEDs used for protection automation and control (PAC) are installed in a virtual format. With a smaller footprint than traditional protection relays, VPR can reduce the need for space on traditional protection applications.

This paper focuses on introducing the reader, with little or no knowledge of VPR, to this new protection technology. Topics such as introduction to virtual protection relaying, hardware requirements for adopting this technology, lessons learned during a computer server setup, and test results comparison of a protection element with VPR and a digital protection relay will be covered.

1 VIRTUAL PROTECTION RELAYING

1.1 Protection Relaying Synopsis

Since the earliest days of the large-scale power grid, multiple different technologies have been used to provide some form of circuit protection. The need for a reliable and safe method of interrupting circuits was evident even to the pioneers of the grid we have today; as early as1879, Thomas Edison described an early form of circuit breaker in a patent application. While the earliest commercial deployments of power distribution technology used fuses to provide circuit protection, electrotechnical relays (EMRs) soon supplanted fuses as the dominant protection devices after their invention in 1905. Early EMRs were bulky and lacked flexibility and scalability; they required complex wiring for each individual circuit under protection. Early EMRs, as the work "mechanical" in their name implies, relied on the operation of physical switches to break a circuit.

Solid-state relays were later invented, which included electronic components in their design.

The advancement of protection relay technology has accelerated over the past century from the crude devices of the early 20th century. Today, the predominant devices found in substations for circuit protection are microprocessor-based relays (MPRs). These devices have been common in transmission and distribution systems for over 40 years.

As their name suggests, MPRs were the first generation of substation protection devices to primarily utilize onboard digital processing to the protection operation, rather than using purely mechanical means. Early MPRs were designed to essentially emulate the principles of EMRs, to make the transition easier for customers. However, at their core, MPRs utilize software algorithms to detect, evaluate, and trigger protection operations. With this shift in technology came a number of benefits, including faster trip speeds,

event logging, and logic that could be made much more complex than that possible with EMRs. Additionally, MPRs allowed for the protection scheme to consider multiple circuits, evaluating the larger scope of the system rather than an individual circuit.

Since their introduction to the market, MPRs have continued to advance with new capabilities and features beyond the EMR. These include the introduction of more advanced protection schemes and algorithms, as well as deeper integration into DMS/SCADA systems. In the last decade, newer advancements focusing on eliminating analog cabling have grown in popularity, such as the move to the IEC 61850 standard. This standard includes provisions for the digitalization of previously analog signals. These digital signals or sampled value streams allow measurements to move over a digital network rather than analog copper cabling.

MPRs have even evolved into a more simplified concept called centralized protection and control system [1], which allows the integration of multiple protection elements in a single device. This avoids the need for having multiple MPRs for controlling multiple bays. A typical bay consists of a breaker and its associated equipment.

Advancing further into state-of-the-art systems, new advancements have continued to be made in recent times, allowing an even more advanced capability to gain momentum in the industry: the Virtual Protection Relay (VPR).

1.2 Transitioning to Virtual Protection Relays

The "virtual" portion of VPR refers to the widespread practice in the information technology (IT) field. Virtualization of software resources is standard practice in enterprise architecture across most industries [2]. Virtualization is the process by which an operating system (OS), and all its associated software functions, is made to run within a virtual machine (VM), rather than a physical machine. The virtual machine itself, alongside potential dozens, or hundreds of other VMs, is then run on a physical computer or server. From within the individual VM, the OS has no awareness that it is not running as the sole OS on a physical server. This architecture allows for many VMs to be consolidated on one or few pieces of hardware, rather than requiring an individual physical server for each OS and associated software workloads [3]. This is not done without benefit; virtualization allows for greater utilization of resources for the same number of workloads. It also allows for the segregation of workloads to aid management and avoid linked failures. For example, rather than run Workload A and Workload B from within the same OS on the same computer, workloads A and B could be run from within VM-A and VM-B, respectively. In this architecture, if a software error causes the operating system of one VM to fail, the other VM would remain unaffected and carry on. However, this can be achieved on a single physical machine, without the need to separate the workloads into multiple pieces of hardware. This can also prevent updates or compatibility issues from adversely affecting multiple workloads on the same machine. Figure 1 below visualizes a comparison between a virtual system architecture and a traditional architecture.

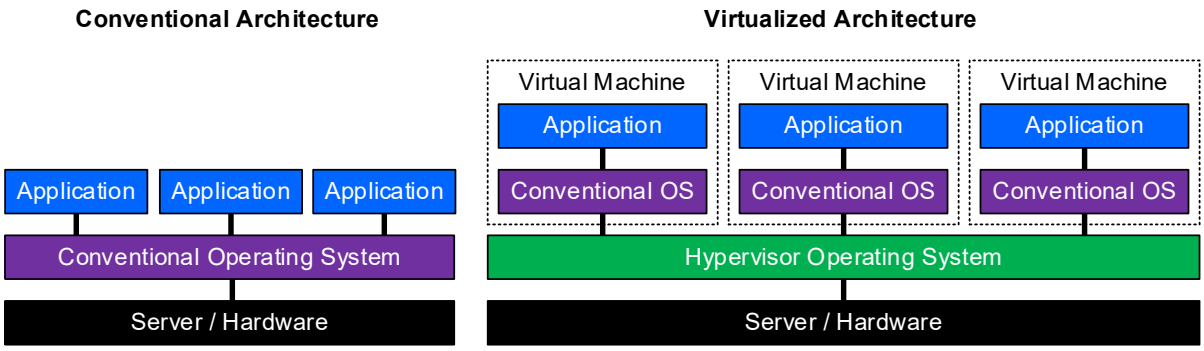

Conventional versus virtual system architecture
Figure 1

To create and manage VMs, a special type of OS called a hypervisor is installed as the base OS on the server. A hypervisor, in contrast to a general-purpose OS like MacOS or Windows, is an OS that has the sole functionality of managing VMs and distribute the underlying computer hardware to those VMs [2]. The hypervisor will allocate a certain number of resources dynamically to the individual VMs and present them as available hardware resources to the VM OS. This includes allocating the central processing unit (CPU) cores, memory, networking, and drive storage.

While in general, virtualization offers several benefits when considering it for IT workloads, the technology has the potential to deliver even greater benefit when applied to operations technologies (OT) workloads; specifically, those found in utility substations. In substations, it is common practice to have many MPRs performing the same function across multiple workloads; it is not uncommon for there to be 10 to 30 identical protection relays for a given substation. In virtualization, there is an opportunity to consolidate and simplify the protection system found in substations from utilizing many relays to a few servers. A single server has vastly more computational capacity than the microprocessors found in even modern MPRs, allowing for the virtualization of 30 protection relays in just 4 computer cores in a modern CPU. This is achieved in a smaller and simpler form factor, with easier upgrades and life cycle management. The number of equipment related outages is also reduced or eliminated.

A centralized protection and control system within a substation allows the integration of multiple protection elements in a single protection device. Nowadays, this protection solution is available on the market from multiple vendor specific solutions. A similar protection concept is integrated into a VPR, but with the advantage of been a vendor agnostic solution.

1.3 Benefits of Virtual Protection

There are multiple benefits of virtual protection, which can be summarized as follows:

- Fewer devices to manage within the protection system (2-3 versus 20-30).
- Less physical space dedicated to protection equipment.
- Fewer outages associated with life cycle management.
- Reduced operations and maintenance costs.
- More efficient software patching.
- Improved security posture.
- Faster protection and better performance.
- General purpose platform, allowing for deployment of future workloads on the same hardware.
- Fewer cables and less wiring complexity.

2 VIRTUAL PROTECTION RELAY HARDWARE

Virtualization inherently aims to be versatile, assuming that the minimum specifications of the specific hypervisor OS are met. When the underlying hardware meets this minimum specification, an individual VM becomes hardware-agnostic, no longer being directly tied with the underlying hardware. All aspects of the interaction between VM and hardware are allocated and managed by the hypervisor.

However, integrating virtualized protection systems into a substation demands additional considerations beyond simply meeting the hardware specification for a specific hypervisor. Stringent industry requirements, such as IEC 61850-3 [4] and IEEE 1613 [5], govern what types of hardware may be deployed in these environments. Substation computer equipment must meet a stricter set of environmental standards compared to those in data centers. Hence, utilities must carefully consider their options and pick hardware that meets both their specific needs and regulatory requirements.

2.1 Processor Specifications

Ensure that the processor selected is compatible with virtualization technology. Indications that a processor is virtualization-ready are the ones that support Intel Virtualization Technology (VT-x) or AMD AMD-V. Settings

to enable these features are found within the server's basic input/output system (BIOS), which enables the division of physical cores for virtualized workloads.

In substations, certain functions such as protection are time-critical, necessitating operation within a set time limit. The processor used for a virtual protection application system should be compatible with time computing features.

While the virtual protection software tested as the basis for this publication is a modest workload, requiring only 4 computer cores, the overall number of CPU cores chosen should be governed by the total planned software workloads intended for deployment to the server.

2.2 Memory Specifications

Random Access Memory (RAM), or simply "memory," is the pool of storage used to store data being used in active processes in a server. For substation servers running critical applications, it is recommended to use a server equipped with Error Correction Code (ECC) memory, which is resistant to data corruption.

Overall, memory requirements are dependent on the nature and number of VMs deployed. The VPR software under evaluation in this paper requires 8 GB of memory. Additional VMs deployed to the substation will have their own memory requirements. A minimum of 64 GB of memory is recommended for a small substation virtualization deployment.

2.3 Networking Specifications

Due to the critical nature of the software applications deployed to substations, it is recommended to incorporate parallel redundancy protocol (PRP) into the substation virtualization architecture. This networking protocol allows for the seamless duplication of network traffic across two separate networks, protecting against the failure of any single network component. In the event of network failure, the traffic will still reach its destination across the secondary network. This is a "hot-hot" topology, as all traffic is continuously sent across both networks. PRP is based on the IEC 62439-3 standard [6]. [7] provides a deeper understanding of redundancy networks.

If a given device does not support PRP natively or through an add-in card, a PRP redundancy box (RedBox) may be used to protect this device's traffic from network failure. A PRP RedBox is a device that acts as an interface between redundant PRP networks and devices which do not support PRP protecting their network traffic.

2.4 Time Synchronization

A common time synchronization source is crucial in digital protection deployment. The multiple merging units (MUs) feeding the VPR must be time synchronized. A MU converts the conventional voltage and current signals from the instrument transformers in the substation to digital signals: sampled value (SV). The server where the VPR will be installed must be capable of supporting any of the available precision time protocol (PTP) standards or profiles: IEEE C37.238 [8], IEC 61850-9-3 [9].

2.5 Environmental Specifications

It is common for substation control buildings to lack air conditioning or other environmental controls. For this reason, any server deployed to a substation should be rated for extreme environments, such as intense heat and cold ambient temperature. The specific ranges required vary by industry standard, but -10 °C to 55 °C is a common operational range. Some manufacturers may have products that operate beyond this range.

Some servers used for rugged applications have been designed to operate without fans. These fanless designs offer several tradeoffs over designs that incorporate fans. While they do have reduced moving parts, they often have lower thermal dissipation thresholds, requiring the use of a lower-power processor. This can limit the maximum capabilities deployable on these systems. CPUs with 4 to 6 cores are common on fanless systems. Other servers that do incorporate fans can incorporate more capable processors due to the additional thermal dissipation abilities provided by forced air cooling. CPUs with 24 cores and beyond are possible in forced-air systems.

2.6 Power Supply Specifications

Power supply redundancy is recommended in the VPR server. No interruption time should be between switching from one power supply to the other one in the event of a power supply failure.

Consider the current type: AC or DC, that will be available on-site as well as its voltage level, when choosing the power supply on the server.

2.7 Hypervisors

When selecting a hypervisor for a virtualized protection deployment, utilities have several options to consider. One factor to consider is which hypervisor is currently used by the organization's IT department. Consulting with the existing IT team can help determine if there is an in-house hypervisor in use, as this can simplify the transfer of institutional knowledge and expertise on hypervisor operation.

The VPR software evaluated in this paper supports two specific hypervisors: VMware ESXi and Linux KVM. When choosing a hypervisor, it is essential to consider which operating systems are supported or recommended by the particular VPR software you plan to deploy in the substation. This information can guide your decision-making process and ensure compatibility between the hypervisor and the VPR software.

2.8 Security

Physical and cybersecurity implementations should be considered for a VPR.

International standard organizations such as IEEE and IEC have relevant standards that are applicable to a VPR. Refer to IEEE C37.240 [10], IEEE 1686 [11], IEC 62351-3 [12], IEC 62351-5 [13], IEC 62351-7 [14], IEC 62351-8 [15], and IEC 62351-11 [16].

3 SETTING UP A VIRTUAL PROTECTION RELAY

Servers might be new to some people in the power industry, but they are not in other industries. IT teams have been working with servers for years and have plenty of experience with them. Therefore, it might be prudent for protection engineers to continue to focus on what they are good at and work closely with the company’s IT team during the VPR server set up process. At the end of the day, the protection team needs a functional VPR to work with.

3.1 Getting Your Virtual Protection Relay Up and Running

Setting up a VPR is a critical task that requires careful planning and execution. While the specific steps may vary depending on the hardware vendor, there are some key areas you will want to focus on to ensure smooth deployment.

3.2 Prepping the Physical Server

The first order of business is to configure the underlying physical server hardware. This step is crucial because it lays the foundation for installing the hypervisor operating system, which will host the VPR application.

3.3 Managing Remote Access

Many data center servers come equipped with an Intelligent Platform Management Interface (IPMI) port. This handy feature allows you to remotely manage the server, including powering it on/off and installing operating systems. If you plan to use the IPMI port, make sure to isolate it on a dedicated management network for security purposes. If remote management is not needed, it is best to disable the IPMI port altogether.

3.4 Identifying Network Connections

The network interface cards (NICs) are the links that connect your VPR to the power grid. Before installing the VPR software, you will need to properly identify these NICs. If your server has a tool for identifying NICs, you can simply plug in a cable to each port one by one and verify the medium access control (MAC) address displayed by the tool. Once you have the MAC addresses, you can easily allocate the appropriate ports for

the VPR. Ideally, you should have the MAC addresses for each NIC handy when purchasing the server, as these addresses are unique to each card.

3.5 Configuring Disk Redundancy

When it comes to formatting the disks and choosing a redundancy mode, you will need to strike a balance between redundancy and capacity. Since VPR workloads typically do not consume vast amounts of data, but are critical for operations, the recommended approach is to prioritize system redundancy over overall capacity. Redundant disk operating modes like RAID 5, RAID 6, and RAID 10, are good options to consider, as they can protect your VPR server from failing due to single or multiple disk failures.

3.6 Installing the Hypervisor

Setting up a VPR involves installing a hypervisor operating system, which is not a simple task. The level of difficulty can vary greatly depending on your familiarity with hypervisor technologies. Before diving in, it is crucial to weigh your options carefully.

3.6.1 The Open-Source Route: Cost-Effective but Potentially Costly

Open-source hypervisors can be tempting from a cost perspective. However, this route may end up costing you more in the long run if you lack the necessary technical expertise. Without access to dedicated support resources, you might find yourself spending countless hours troubleshooting and resolving issues during the setup process. It is like trying to assemble a complex piece of furniture without instructions.

3.6.2 The Commercial Option: Pay for Peace of Mind

On the other hand, commercially available hypervisor solutions come with a price tag, but they also provide access to a technical support team. Think of it as having a personal assistant to guide you through the process. This can save you valuable time and effort, allowing you to focus on other critical tasks. Additionally, enterprise-grade hypervisors often offer more mature and feature-rich platforms compared to their open-source counterparts, which can be a game-changer for complex deployments.

3.6.3 Making the Right Choice

Ultimately, the decision between an open-source or commercial hypervisor solution will depend on your organization's budget, technical expertise, and the level of support you require. It is like choosing between a do-it-yourself project and hiring a professional contractor – both options have their merits, but the right choice will depend on your specific needs and resources. Carefully weighing the pros and cons of each option will help ensure a successful and efficient VPR deployment, allowing you to focus on keeping the power flowing smoothly.

3.6.4 Installation Process

When it comes to the installation process itself, commercial hypervisor vendors typically provide step-by-step instructions for setting up a complete system. It is like having a detailed roadmap to follow, which can be a significant advantage, especially for those with limited experience in hypervisor deployments.

In contrast, open-source hypervisor solutions may require a more in-depth understanding of the platform to define and execute the necessary steps for deploying the hypervisor on a server. Without proper guidance or support, this process can be more challenging and time-consuming.

4 TESTING COMPARISON

4.1 Test Bed Set Up

The positive- and zero-sequence components of a 69 kV feeder and nominal frequency of 50 Hz were used for this testing comparison. 50 Hz was used as the nominal frequency because it is used by most European countries, and because of the main audience for this paper is based in Europe.

With the assistance of a power system fault calculator, Figure 2, phase-to-ground (AG), phase-to-phase (BC), and three-phase (3-Ph) fault types were simulated. Each fault type was simulated at 25%, 50%, 75%, and 90% from the protective relay.

The fault calculator automatically calculated the required voltage and current signals during the pre-fault, fault, and post-fault state.

The calculated fault values were used by a secondary power system simulator test set to inject conventional voltage and current values to a MU.

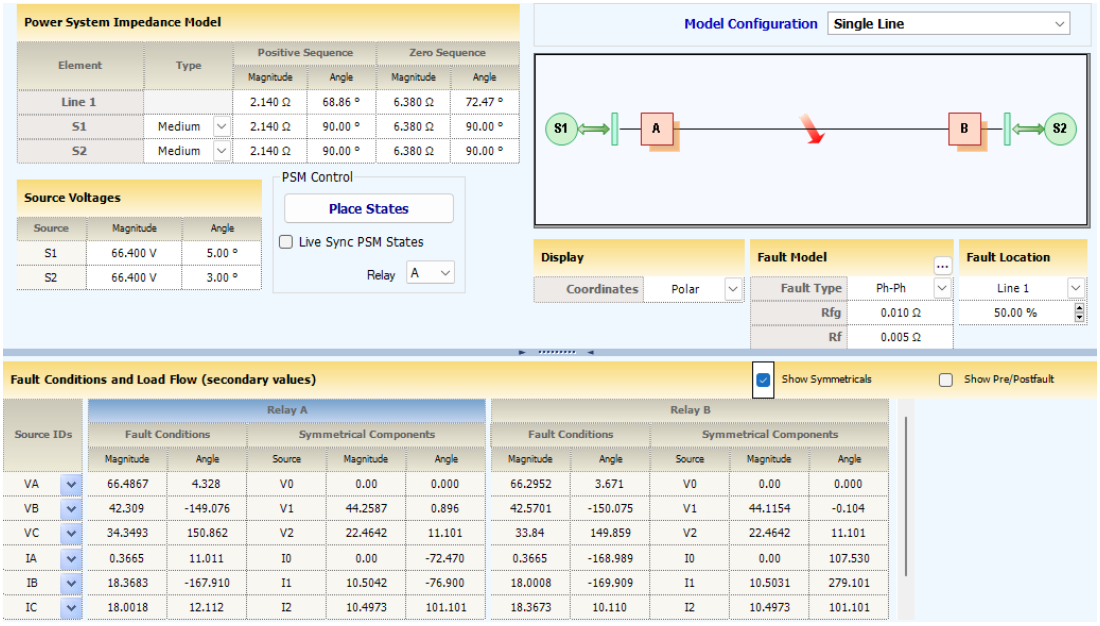

Power system fault calculation in secondary values
Figure 2

Two protection relays were used for this comparison: an MPR and a VPR. The MPR only supports SVs and Generic Object-Oriented System Event (GOOSE). A GOOSE message with the trip signal (PTRC) was configured on each relay.

The VPR device is an IEC 61850-3 compliant server with an Intel Xeon Gold 6312U CPU processor running at 2.4 GHz. It has a RAM of 256 GB. The CPU supports virtualization technology, and the VPR software is running as a VM within a type-1 hypervisor. The server meets all other specifications laid out in this paper for a recommended VPR system.

The Ethernet switch is managed with support for IEEE 1588 and it has gigabit ports.

A phase instantaneous overcurrent (50P), a phase inverse time overcurrent (51P), a residual instantaneous overcurrent (50N), and a residual inverse time overcurrent (51N) protection element were set up on each protection relay. All four protection elements were mapped to the PTRC signal on each relay.

The relay settings in secondary values for each protection element are summarized in Table 1.

Table 1
Protection elements settings in secondary values

Protection Element	Value	Description
50P	15.3 A	Pickup
51P	4 A	Pickup
51PC	IEC Normal Inverse	Time-current operating characteristic
51PTD	0.1	Time dial
50N	10.5 A	Pickup
51N	3 A	Pickup
51NC	IEC Normal Inverse	Time-current operating characteristic
51NTD	0.1	Time dial

A MU provided the required SV for both protection relays under test. The MU has high-speed (HS) contact outputs. The GOOSE messages published by the MPR and VPR containing the PTRC signal were sensed by the MU. Each relay PTRC signal was mapped to an individual HS contact output in the MU.

Time synchronization using PTP was used to time synchronize the MPR, VPR, and MU.

The secondary power system simulator, upon injecting the voltage and current values to the MU, started two timers during the fault state. The timers were used to measure each relay operating time, which was given by each of the HS contact outputs from the MU.

Figure 3 shows the test set up for the relay operating time comparison.

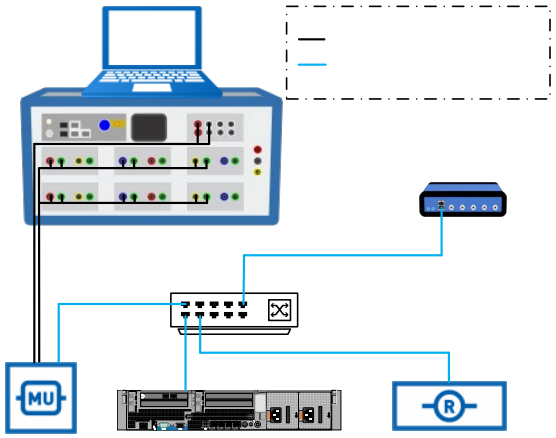

Test set up
Figure 3

The purpose for the multiple fault simulation scenarios was to check the VPR operation response time, when compared to a known relay technology: an MPR.

4.2 Test Results

One of the most curious points about this new technology is how quickly a VPR could operate, when compared to a known MPR with SV inputs and GOOSE messages.

There is no doubt that the processing power on the VPR is higher than the one found on the MPR, which will help to process the measurement values, voltage and current, in a faster way. But the most important aspect of it is that the VPR would make the right decision at the time of tripping. This is achieved with known

protection algorithms that have been in service for many years in MPRs because the same manufacturer has a strong reputation in the protection relay industry.

The comparison was not made with an MPR from the same manufacturer, but rather with a different one available on the market with a similar trajectory in the power industry.

Table 2 summarizes the test results for the 12 faults applied to the MPR and VPR.

It can be observed that the VPR's instantaneous overcurrent element operates faster than the one in the MPR. In the case of the ground fault, the VPR's inverse time overcurrent protection element operates a few milliseconds slower than the MPR. The result is similar for the phase element, i.e. 530.1 ms.

Table 2
Trip operation time comparison between MPR and VPR

Fault	MPR (ms)	VPR (ms)
AG at 25 %	24.5	12.0
AG at 50 %	26.2	11.2
AG at 75 %	41.5	14.0
AG at 90 %	568.98	576.48
BC at 25 %	22.6	15.1
BC at 50 %	26.4	16.4
BC at 75 %	46.3	18.8
BC at 90 %	530.1	530.1
3-Ph at 25 %	19.7	9.7
3-Ph at 50 %	23.4	13.4
3-Ph at 75 %	25.2	15.2
3-Ph at 90 %	34.1	16.6

CONCLUSIONS

Like any other new technology, VPR has two major changes when compared to classic MPRs: it uses a centralized protection approach, and it runs on a server. Usually, VPRs lack visual indications or target LEDs, that are common in MPRs. Some hardware specific platforms include target LEDs, but getting one of these hardware options could make the VPR lose its feature of been a hardware agnostic application. To compensate for the lack of target LEDs, a permanent HMI is needed to access the VPR and monitor its virtual alarms.

Protection algorithms do not change because it is a VPR, but the processing power is higher when compared to MPRs. This is an advantage on VPRs for processing the voltage and current signals faster and making the decision to operate or not.

Testing this technology like any other new technology in the power industry will be key for gaining confidence in it and moving forward with this technology. Testing a VPR is not much different than testing a fully digital MPR as presented on this paper.

Migrating from conventional MPRs might be smooth if the adopter has already moved to complete digital technology. Otherwise, a stepper learning curve will be in the process.

The test results presented in this paper intend to demonstrate that the trip operation time is not compromised when compared to the known technology available in MPRs.

New skills will be needed for the initial deployment of this technology such as the ones in the IT and OT domains, which are new to the power industry, but they are not in other industries such as banking.

REFERENCES

[1] I. P. W. K15, "Report: Centralized Substation Protection and Control," 2015

[2] IBM, "What is virtualization?," [Online]. Available: https://www.ibm.com/topics/virtualization. [Accessed 10 05 2024]

[3] R. Hat, "What is virtualization?," [Online]. Available: https://www.redhat.com/en/topics/virtualization/what-is-virtualization. [Accessed 10 05 2024]

[4] "Communication networks and systems for power utility automation - Part 3: General requirements," IEC 61850-3, 2013

[5] "IEEE Standard Environmental and Testing Requirements for Communications Networking Devices in Electric Power Substations," IEEE 1613, 2003.

[6] "Industrial communication networks - High availability automation networks - Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR)," IEC 62439-3, 2021

[7] "Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines," IEC/TR 61850-90-4, 2013

[8] "IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications," IEEE C37.238, 2017

[9] "IEC/IEEE International Standard - Communication networks and systems for power utility automation - Part 9-3: Precision time protocol profile for power utility automation," IEEE/IEC 61850-9-3, 2016

[10] "IEEE Standard Cybersecurity Requirements for Substation Automation, Protection, and Control Systems," IEEE C37.240, 2014

[11] "IEEE Standard for Intelligent Electronic Devices Cybersecurity Capabilities," IEEE 1686, 2022

[12] "Power systems management and associated information exchange - Data and communications security - Part 3: Communication network and system security - Profiles including TCP/IP," IEC 62351-3, 2023

[13] "Power systems management and associated information exchange - Data and communications security - Part 5: Security for IEC 60870-5 and derivatives," IEC 62351-5, 2023

[14] "Power systems management and associated information exchange - Data and communications security - Part 7: Network and System Management (NSM) data object models," IEC 62351-7, 2017

[15] "Power systems management and associated information exchange - Data and communications security - Part 8: Role-based access control for power system management," IEC 62351-8, 2020

[16] "Power systems management and associated information exchange - Data and communications security - Part 11: Security for XML documents," IEC 62351-11, 2016

BIOGRAPHY

Jose Ruiz received his master's degree in sciences (Power System concentration) from the University of Tennessee at Chattanooga – USA - in 2012. He worked as a protection application engineer and later as a team leader in ABB Inc. from 2011 until 2016. In 2016 he joined Doble Engineering as an application engineer for the protection testing solutions of Doble. Nowadays, He works with the protection solutions team on the development of new hardware and software solutions.

Jose is an active IEEE PSRC and PSCCC senior member and enjoys sharing his knowledge on IEC 61850 standard-based protection application solutions.

Montie Smith is an Energy Field Director with Dell Technologies, where he works to develop energy focused solutions and position go to market efforts focused on the utility industry. He has particular interest in the convergence of IT/OT systems, including virtualization and its applicability in digital substations. He graduated from the University of Tennessee with a Bachelors and Master's in Electrical Engineering with a focus on power electronics and microgrids.

Bryan Gwyn is a seasoned Power Utility and Consulting Professional with over 30 years of international experience in electric utility Transmission, Generation, and Distribution. His expertise spans protection, control, and telecommunication engineering, as well as operations and management. Currently, he leads the development of Protection, Asset Management, Monitoring, and Security solutions at Doble Engineering. Bryan holds a BEng (Hons) in Electrical and Electronic Engineering and a PhD from City University, London, UK. His distinguished career includes roles at National Grid in both the UK and US, and Quanta Technology. Bryan is a Chartered Engineer and a Senior Member of IEEE, having led global teams of subject matter experts throughout his career. He resides near Boston, Massachusetts, in the United States.