



VIRTUAL PROTECTION RELAY

A PARADIGM SHIFT IN POWER SYSTEM PROTECTION

Authors: Intel - Dean Samara-Rubio, Graham McKenzie, Prithpal Khajuria

Kalkitech - Roopesh Ariya, Prasanth Gopalakrishnan, Viju Ravindran

TABLE OF CONTENTS

1	Introduction	
2	VPR Architecture	
2.1	VPR Approach 1 – Direct Porting of discrete IEDs onto a host	
2.2	VPR Approach 2 – Shared Functionality	
2.3	The Role of Software Container Technology	
2.4	Hardware Platform Technologies	
2.4.1	Server Hardware Specification	
2.4.2	Time Coordinated Computing	
2.4.3	Precision Time Protocol	
2.4.4	Next Generation NIC with HSR/PRP, PTP, and SRIOV	
2.5	Hypervisor Selection	
2.5.1	VPR Redundancy and High Availability	
2.5.2	Implementation of Digital twin	
3	VPR Reference System	
3.1	Kalkitech VPR Framework	
3.2	Signal Conditioning within the VPR Reference	
3.3	Laboratory Setup for Testing of VPR	
3.4	Evaluation of VPR Reference Application	
3.4.1	Response Time Dependence on HW Allocation to VM	
3.4.2	Response Time versus Fault Magnitude	
4	The Path Ahead	
4.1	Reference materials	


1. Introduction

As the energy sector transitions to a high penetration of renewable energy sources, many of which will connect to the grid as distributed energy resources (at the medium and low voltage levels), the patterns and levels of power flow are changing significantly. Figure 1 illustrates the extent to which our mix of energy sources is shifting with renewable energy production projected to exceed the sum of both Coal and Nuclear within this decade. The power sector is responding in several ways: First, the relative agility of natural gas plants compared to baseload power stations of coal and nuclear is highly valued to ballast against the variability in renewable energy production and this has accelerated the transition to natural gas (while at the same time reducing carbon footprint compared to coal). Secondly, the adoption of bulk energy storage has begun to further smoothen the peaks and valley of renewable energy supply. Third, investments into increasingly adaptive “intelligent” grid controls is crucial to enable grid operators at all levels to manage voltage levels and supply quality across the grid.

It has been well-documented that one of the key aspects of the grid controls which are put under stress by this penetration of renewables is the circuit protection sub-system. To operate and maintain effective circuit protection with high penetration of renewable resources requires a level of agility and programmability for protection systems not available in today’s mainstream protection relay products. What’s more, when considering that there is a still a large installed base of older generation devices across the grid, utilities are faced with a generational decision on how to re-invent and reform their entire protection infrastructure for the agile and intelligent grid. In this whitepaper we describe the virtual protection relay (VPR) concept – an architecture where software-defined and virtualized platforms are deployed to host the critical circuit protection functions for an advanced and agile grid. We assert that this use of virtualization technologies is needed to unlock the potential of the grid to deliver reliable, available, and resilient electricity supply for our low-carbon economy of the future.

Renewables Projected to Become Largest Electricity Source

Electricity production by source (kilowatt hours)



Source: U.S. Energy Information Administration.

COUNCIL OF
FOREIGN
RELATIONS

Figure 1: Renewables set to become the leading source of electricity. Electricity production by renewable sources is approaching an unprecedented level, set to exceed 2T kWh per year in the USA by 2043.

The first protection relay was developed in the beginning of the 1900's beginning with electromechanical devices that would sense a fault and actuate a mechanical switch (or a series of mechanical switches) to interrupt the source that was feeding the fault with hazardous energies. Power electronics technology later paved way to static or solid-state relays based on electronic components in the 1960's. Then came the development of Microprocessor-based relays in the 1980's including digital signal processing and digital network communications. These relays along with a range of similar devices for monitoring and control of the grid were given a name – "intelligent electronic device", or IED. While the core functionality of power system protection -which is detection of a power system disturbance and isolation of the affected area- has remained relatively unchanged, the successive generations of technology have improved the operation of the grid as it scales in size and complexity.

Each step in the evolution of protection systems has been an effort to standardize the functionality, behavior, and interfaces to enable a robust market of vendors and solutions to serve the global power utility sector. In the early

2000's, the IEC 61850 set of standards brought in a major boost to multi-vendor IED integration and device-to-device data transfer. Standard models of various protection functions were devised for possible interoperability between protection IEDs from any vendor. Standards were prepared for data exchange between devices (station bus) and current/voltage information from field (process bus). Acceptance of the IEC 61850 standards worldwide have resulted in station level and process level communication networks for exchange of digitized raw values (using Sampled Values, or SV, protocol) and processed values/information across the substation devices and beyond the substation to centralized monitoring systems.

Even after the century of advancement, the implementation of protection for even a single medium-voltage substation requires up to 100 individual IEDs each with its own power supply, case, electronic components, firmware, installation and maintenance guide(s), and configuration tool(s). Each vendor of IEDs has their own design paradigm for their hardware and there is a significant engineering cost to swap one make or model of an IED with an essentially equivalent IED from another vendor.

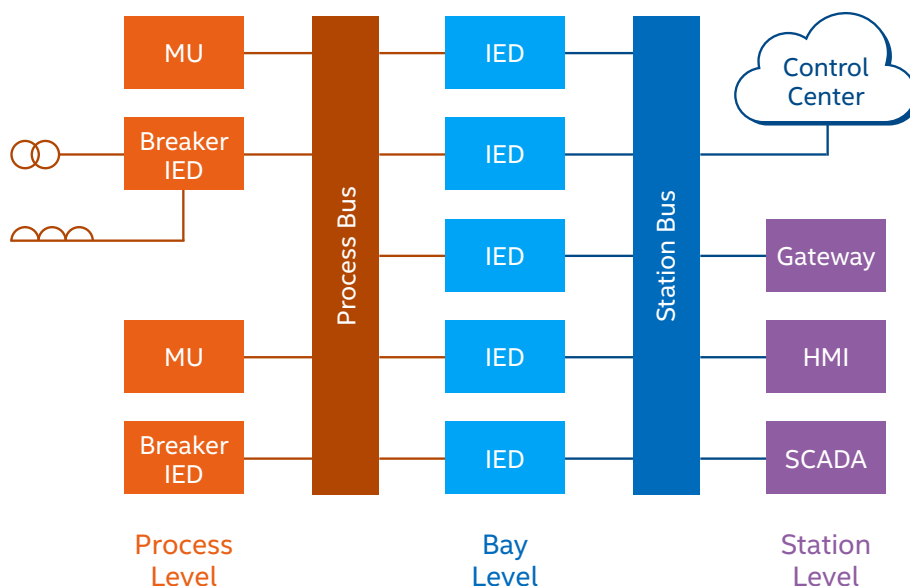


Figure 2: IEC 61850 Digital Substation Network model. Merging Unit(s) (MU), circuit breaker controllers (Breaker IED), relay and control devices (IED), and supervisory control and data acquisition (SCADA) systems are connected across three conceptual levels with communication buses.

VPR eliminates many of these costs., reduces the physical size and complexity of a substation protection and control infrastructure, and opens the market for software-defined innovation where the best-of-breed protection functions and automation algorithms can be deployed with high confidence and reliability but minimum cost.

While these potential benefits of VPR are compelling, there remain questions in the industry as to whether the virtualized hosts and protection functions deployed as software modules can meet the performance and reliability requirements of the modern grid. In this paper, a reference VPR subsystem is described, and the initial performance characteristics are presented. The VPR solution builds upon the preceding century of innovation in the protection industry including the IEC-61850 standard.

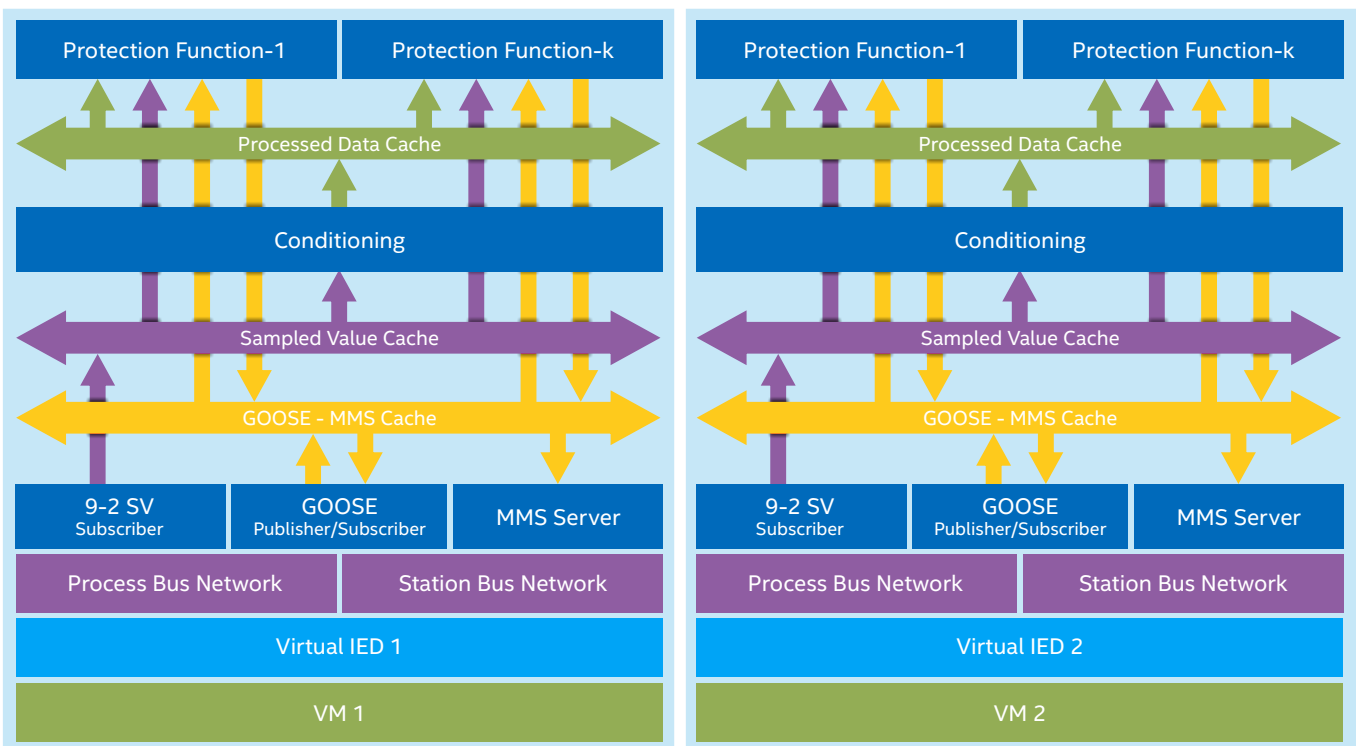
2. VPR Architecture

There are multiple approaches for utilizing the virtualization capability to create a VPR system. Any protection system must meet the basic requirements of Selectivity, Speed, Simplicity, Reliability, Economy. In order to meet these objectives and to maintain a healthy market, additional considerations include the maintainability by the end-customer, the supportability by the technology providers, interoperability between vendor offerings, and the ability for test and certification by independent actors.

2.1 VPR Approach 1 – Direct Porting of discrete IEDs onto a host

A straightforward approach will be to port a complete existing stand-alone IED into a virtual machine (VM) on the host hardware and hypervisor. If this is done for multiple discrete IEDs then a system such as shown in Figure 3 would result. Note that there is no “sharing” of functions and each of the VMs will have its own complete ensemble of capabilities (Sampled Value subscriber, MMS Server, grid event recording, system health monitoring, configuration utilities, etc.)

Figure 3: Direct Porting of Protection relay software
Each Discrete IED is ported to the host in its own virtual machine.
None of the application functionality is shared.



2.2 VPR Approach 2 – Shared Functionality

In an alternative approach the IED functionality is divided so that common modules are consolidated to a separate VM. This can give a level of efficiency when scaling to a larger number of protection functions, but to do so requires an open API to be defined so that each VM can access the shared functions and make data available to all other modules and protection functions across multiple VMs. For example, as shown in Figure 3, the MMS

Server on the Station Bus is consolidated and shared. With the application of advanced network interface virtualization (e.g., the inclusion of the Single-Root Input Output Virtualization, or SR-IOV, technology), and further design optimizations and development of the common API's, it is possible to share more of the functions including the sampled value subscriber, signal conditioning, system monitoring, and other functions of the system to further enhance the scalability of the solution.

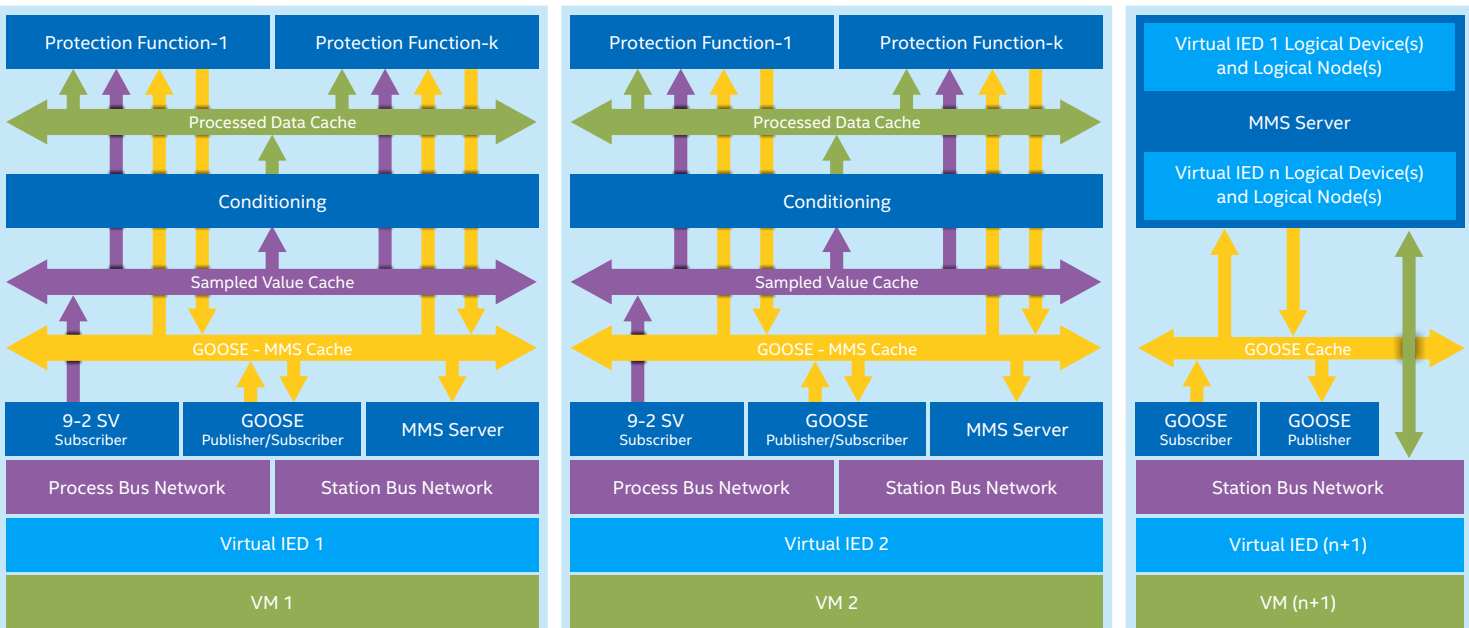


Figure 4: Porting of Relay SW with refactoring to use some shared services. Protection functions can run individually or as a group in one or more VM. Common functions like SV Subscriber, MMS Server, GOOSE publisher, configuration module, system monitoring, etc. can be part of a different VM. To achieve this partitioning and sharing of common functions openly defined API interfaces are required.

2.3 The Role of Software Container Technology

VPR may also utilize the software container technology (container systems include the well-known Docker and Podman projects among others). Container technology is a complement to the VM hypervisor so that, for example, each discrete protection function and supporting microservice can be its own container. The advantage of container technology is the scaling of services with maximum portability and a set of tools that help the software developers in the integration, deployment, and maintenance of their software across the lifecycle of the

systems. As commercial solutions for VPR emerge, we expect that both virtualization hypervisor and software containerization technologies will be utilized together.

2.4 Hardware Platform Technologies

2.4.1 Server Hardware Specification

Today, utilities must handle multiple physical IEDs from different vendors deployed during different periods of time. They must have a survivability plan for maintenance, management, and replacement, for a large ensemble of different product generations, product SKUs, firmware versions, and product options. In the

virtualized environment all the physical IEDs can be accommodated in an Intel Xeon based Server hardware that can withstand the harsh environment in a substation. To achieve the maximum effectiveness of the VPR concept, Intel is working with an ecosystem of computing hardware companies to design and manufacture Intel Xeon-based Substation Servers that are IEC 61850-3 and IEEE 1613 certified for use in a substation and capable of hosting automation, control, and protection workloads. OEMs including Advantech, Crystal Rugged, Dell and others have created powerful rugged servers that meet the needs of the market.

2.4.2 Time Coordinated Computing

Intel-based computing systems have technologies which enable the hosting time-sensitive applications; applications with operations that must execute within defined rules and time constraints. The term “real-time” is used to identify such operations but the term is used in many different contexts and not all “real-time” systems are the same. In practice, response time requirements at the external system level varies from a few microseconds to hundreds of milliseconds, but for every external action (of even an optimized software system for “real-time” computing) there can be a sequence of hundreds or thousands of

internal operations as data is moved between software processes and threads, calculations are performed, etc. In the case of protection and control functions, the timescale for deterministic behavior at the system level is on the order of 5 milliseconds.

In the suggested VPR environment there will be multiple VMs hosting a heterogeneous mix of applications. Ensuring consistent performance and the right priorities for access to hardware resources is a key role of the hypervisor with real-time support. Even in a model where VM's are assigned (pinned) to specific CPU cores and memory is allocated to a VM, there can be shared resources of the system including the network and internal buses, and the last-level cache. Intel Cache Allocation Technology (CAT) address the concerns on the most critical of the shared resource by providing software control of where data is allocated into the last-level cache (LLC), enabling more effective isolation and prioritization of key applications.

Intel Time Coordinated Computing (Intel TCC) including CAT enabled processors deliver optimal compute and time performance for real-time applications. Intel TCC ensure that performance requirements are met in a stand-alone environment and under the presence of concurrency with other workloads.

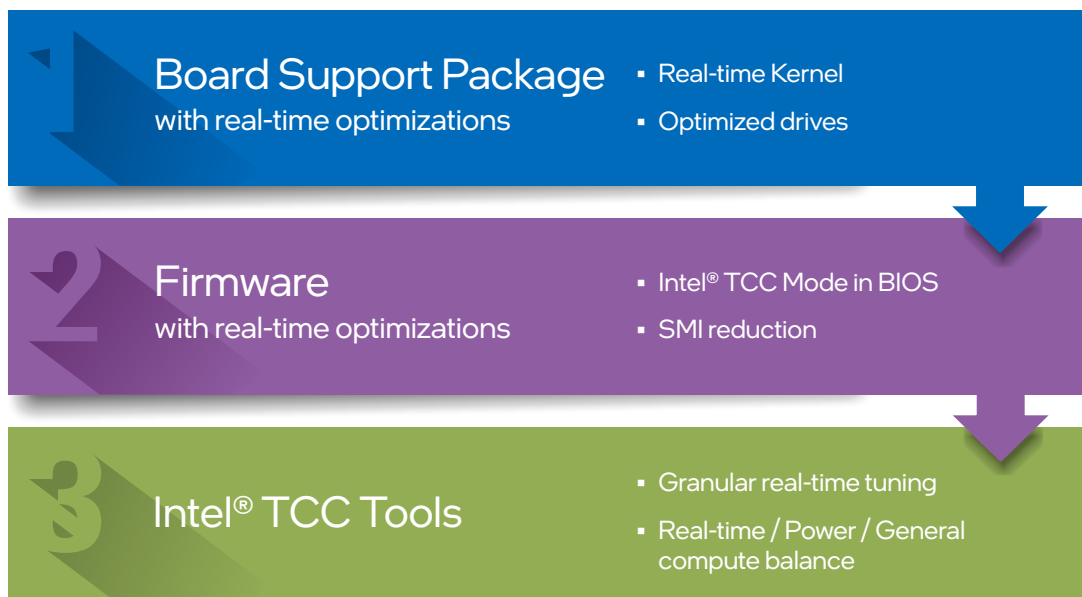


Figure 5: Model to enable Intel Time Coordinated Computing (TCC) for real-time workloads.

2.4.3 Precision Time Protocol

IEDs in the substation operates in real-time where the devices shall be time synchronized to the best possible accuracy. Precision Time Protocol (PTP) provides time synchronization for digital substations with microsecond-level accuracy. PTP, defined in IEEE standard (IEEE1588-2008), is a protocol enabling precise synchronization of device clocks in Ethernet. Devices running PTP are automatically synchronized to the most accurate clock in the network. The protocol supports system wide synchronization accuracy in sub microsecond range with minimal network and local clock computing resources. VM Hypervisor shall support high accurate clock synchronization with dedicated ports in the network interface card (NIC). There are network cards that has HSR/PRP support along with PTP time synchronization feature.

2.4.4 Next Generation NIC with HSR/PRP, PTP, and SRIOV

The network interface is another resource within

the compute platform that needs to be shared between VM's and can be one of the points for performance degradation when hosting time-critical applications. In addition to latency considerations, the VPR system adheres to the IEC61850 standards which dictates zero packet loss for the multiple high speed sampled value streams such that all packets should be received in the virtualized devices.

In the network between the powerline sensing components and the VPR server exists the process bus (as defined for example in the IEC 61850-9-2 standard). The process bus is typically implemented as a redundant network making use of high-availability seamless redundancy (HSR) and parallel redundancy protocol (PRP) equipment. One approach to VPR is to terminate the redundant network at the ingress to the VPR server on the network interface card (NIC) in hardware, thus eliminating the software from having to process redundant packets and reduce latency. The NIC would use single root I/O virtualization (SR-IOV) technology to share the network interface

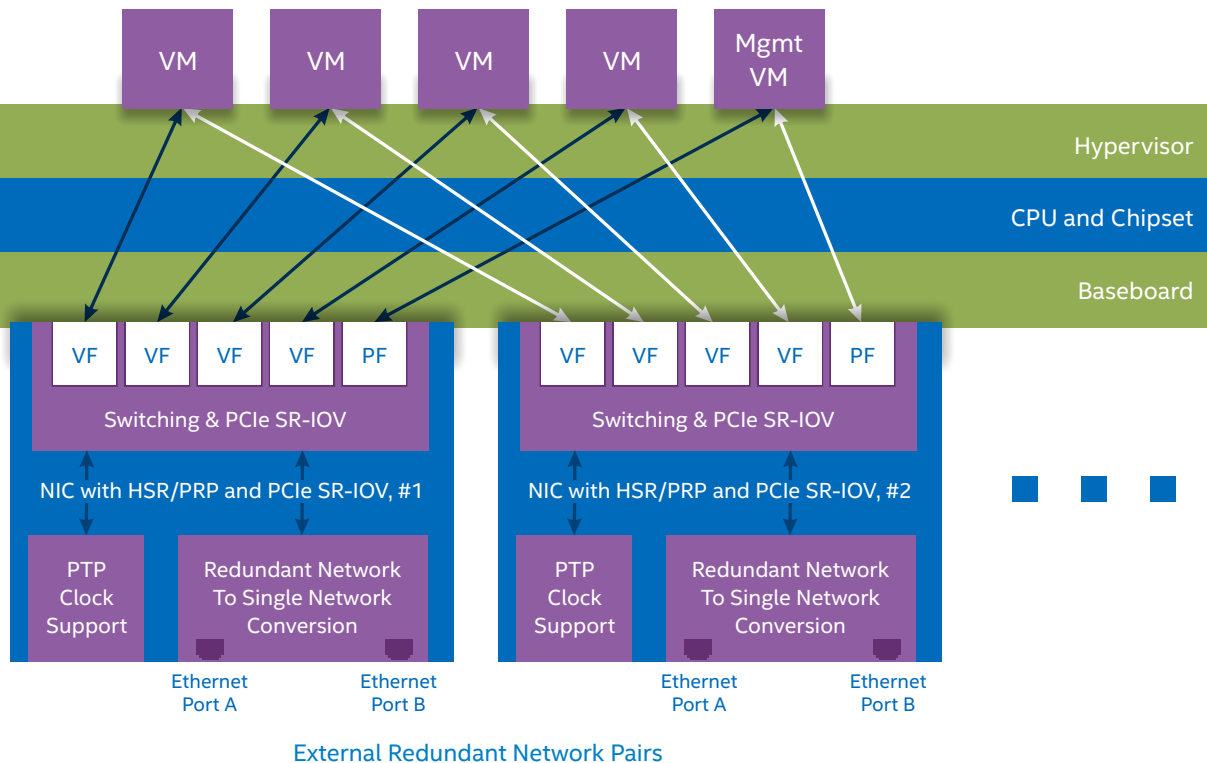


Figure 6: Server with NIC Cards that support HSR/PRP networking, PTP clock sync, and SR-IOV. Multiple such NIC cards can connect to PCIe card slots of server to support redundant networking requirements and virtualized network interfaces for efficient access to networks from multiple VMs. The combination of HSR/PRP, PTP, and SR-IOV technologies are expected to achieve a high level of both performance and scalability for VPR.

efficiently and reliably to the VMs within the VPR server. The NIC would thus implement the combination of HSR/PRP and SR-IOV capabilities as shown conceptually in Figure 6. This approach routes network packets directly to the required VM using PCIe SR-IOV virtual functions to meet the reliability and latency requirements for VPR.

2.5 Hypervisor Selection

Hypervisors are the essential ingredient of the virtualized system as they are the software layer that receives and conveys requests between the physical and virtual resources. Modern hypervisors have modes and schedulers that provide various options for supporting the virtualization of real-time operating systems and real-time applications.

When a hypervisor is installed directly on the hardware of a physical machine, it is called a bare metal hypervisor. Bare metal hypervisors may also be embedded into the firmware enabling better hardware handling. Other hypervisors approaches are possible including host-based hypervisors which actually run inside a host operating system of the physical machine. An example for a bare metal hypervisor is VMWare ESXi. An example of a hypervisor installed with an operating system is the TTTech Nerve platform which offers a Linux KVM hypervisor.

2.5.1 VPR Redundancy and High Availability

Each protection relay shall meet the dependency definition basically the ability to operate or trip for a fault within its protective zone. In many cases the protection engineer will specify a scheme for a substation or a powerline that has multiple relays which operate in a way that they have overlapping responsibility such that if any one relay fails to trip on a fault, another relay in the system will detect that fault and trip. In some cases, the utility protection engineer may deploy two complete protection relays in parallel with the same algorithm on the same zone. Thus, a protection engineer uses her knowledge of the grid to avoid the chance that a failure of a relay could result in a fault condition that damages equipment or creates a persistent safety hazard to utility personnel or the public. A second case where redundancy in the protection scheme is required is for the rare but sometimes necessary maintenance or replacement of protection system components.

The introduction of VPR offers new ways to economically add redundancy to the protection scheme. Modern hypervisors allow deployment of VMs in high availability cluster mode. The host server or VM management console will handle the redundancy and failover of the virtual machines with negligible or even zero downtime. Hence a failure of a single VPR system will not affect the protection system operation.




Figure 7: VM Hypervisor redundancy

Two VPR servers, each with VM Hypervisors, can be coordinated in a side-by-side redundant topology to add resiliency to a VPR-based protection system. The OEMs can use the fundamental capabilities of virtualization in this way or with other topologies for enhanced performance, availability, and resilience.

2.5.2 Implementation of Digital twin

More utilities are now fascinated with the idea of implementing digital twins in the substation for design reviews, testing, managing, and upgrading of the system. VPR infrastructure allows easy implementation of digital twin for the protection relays. Digital twin implementation allows the simulation of the protection relays and real-time updates. With VPR, the simulated relay will have the same configuration and receive the same data from the field as the original relay. Thus, with VPR the digital twin concept can be leveraged for fault trouble shooting and replay based on IEC 61850 model and Disturbance data from the field.

3. VPR Reference System

Kalkitech has created a set of a reference protection applications to pair with VPR Server hardware, hypervisors, and other components to create a complete VPR system to serve multiple roles in the market adoption: as a proof of concept, a demonstration system, and to help the industry to evaluate the capabilities, trade offs, and challenges posed by the VPR architectural options. A key objective is to prove that multiple protection relays can work in a virtualized environment meeting or even exceeding the

performance and time requirements of the modern protection IEDs while adding the redundancy, resiliency, digital twin, ease of deployment, and other features mentioned earlier in the paper.

The Kalkitech VPR reference application has been designed, developed, and tested in virtualized platform running in Intel Xeon based Server hardware. The reference application is available for the utility engineers and system integrators to test and compare the performance with the contemporary protection systems. Figure 8 provides a map of a protection sub-system including the powerline, current and voltage sensors (CT and PT, respectively), the sampling and digitizing unit for the current and voltage signals (merging unit, or MU), the breaker control unit (BCU), the process bus network, the VPR node with internal details exposed, the station bus. The set of software components needed to configure and monitor the protection system including the human-machine interfaces (HMI), and a disturbance recorder (DR) is shown to the right side of the diagram. The VPR can be interfaced to the utility's supervisory control and data acquisition (SCADA) system as well.

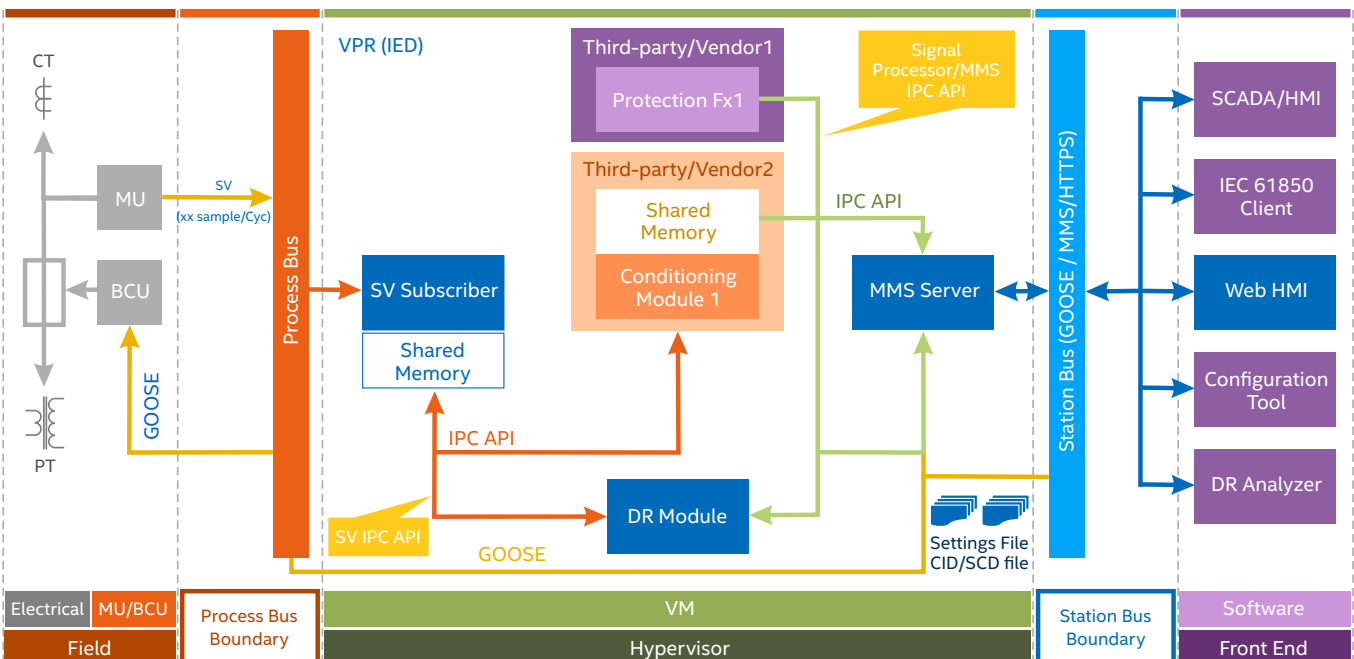


Figure 8: VPR Reference Application within the chain of components for VPR. Some details of the SW modules in the VPR Server, and the connections between those modules, are shown. One MU/BCU pair is shown in the diagram, but it is expected that once the system is mature, a single VPR server could serve more than 100 MU/BCU pairs.

Within the VPR node, the details shown in Figure 8 include the SV subscriber to collect the streams of digitized current and voltage signals from the process bus, conditioning module(s) for digital signal processing to extract key parameters from the data streams, one or more protection functions to detect abnormal power system condition(s) as per their algorithm and settings, and a disturbance recorder (DR) module for capturing detailed event data.

In Figure 8, only one MU and BCU are shown, but for a VPR server in the field it is expected that one VPR server (once mature) can handle data from 100 or more MU's and BCU's. Similarly, within the VPR node, a multiplicity of protection functions and conditioning modules would be needed to support multiple MU's and BCU's.

The VPR reference application is developed using IEC 61850 VPR infrastructure provided by Kalkitech. It provides basic infrastructure backbone for implementing (Virtual) Protection Relays based on IEC 61850 for Digital

Substations. Reference protection functions available for evaluation includes Distance Protection (21), Transformer Differential Protection (87T), Time Delayed Phase Overcurrent Protection (51), Time Delayed Derived Neutral Overcurrent Protection (51N), Instantaneous Phase Overcurrent Protection (50), Instantaneous Derived Neutral Overcurrent Protection (50N) and Breaker Failure Protection (50BF).

3.1 Kalkitech VPR Framework

The Kalkitech Framework provides a set of base ingredients which OEMs and utilities can use to deploy an IEC 61850 VPR infrastructure to evaluate migration of their capabilities to the virtualized environment using the APIs provided. As such, the term "Third Party / Vendor" in Figure 8 is used in the areas where domain experts and developers from multiple companies can insert their conditioning and protection function components. Figure 9 shows a few models of how the Kalkitech VPR

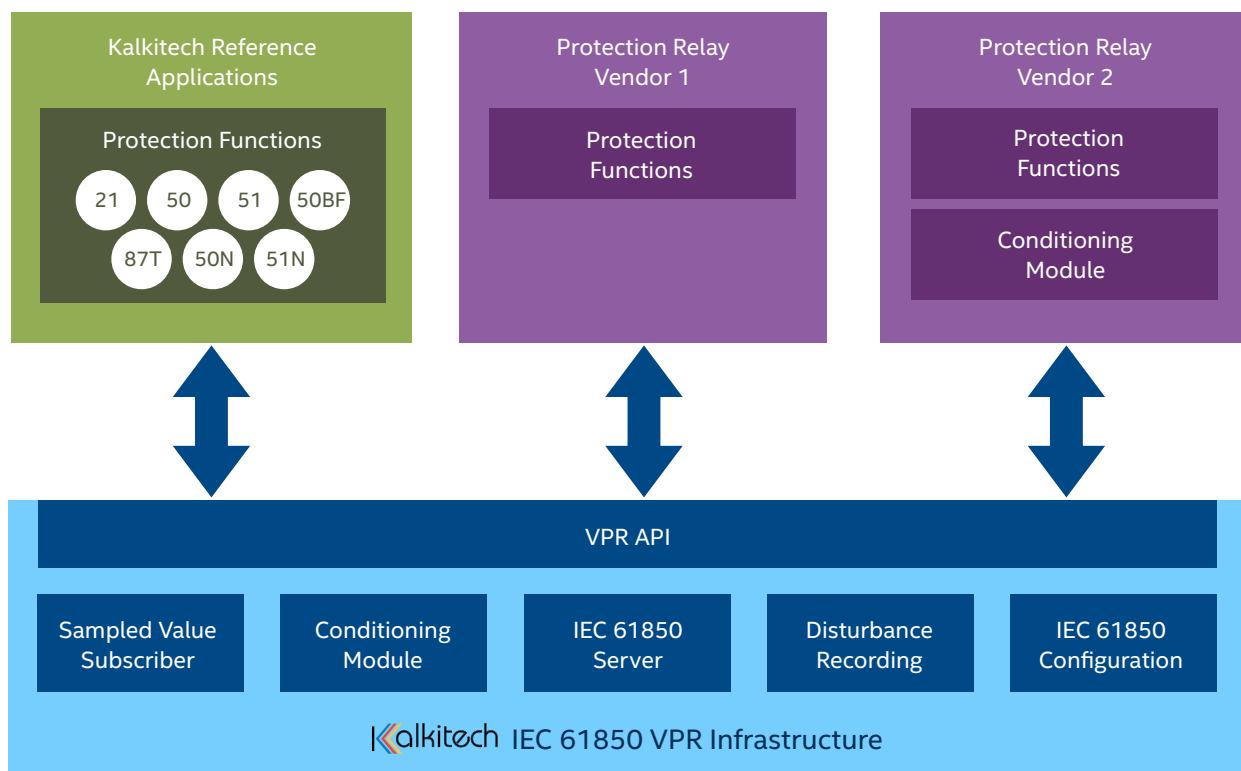


Figure 9: Kalkitech IEC 61850 VPR Framework. The VPR Infrastructure and Reference Applications for demonstration and verification of other components of a VPR system. The Kalkitech VPR Infrastructure can also support OEMs' algorithms and signal processing functions to create a complete VPR server. The protection functions are labeled as per ANSI / IEEE Standard C37.2 Standard for Electrical Power System Device Function Numbers, Acronyms, and Contact Designations.

infrastructure and reference applications can be used and co-exist with third party components.

The framework provides mandatory communication modules for VPR implementation. Protection developer or vendor can integrate the modules using the API library provided by the framework. The framework already has the IEC 61850 implementations like IEC 61850 9-2 LE Sampled Value Subscriber, IEC 61850 /MMS Server, GOOSE Publisher module, Disturbance recorder and API library to use the functions. It also has a signal processing module that serves data to the reference protection functions. The data from this signal processing module can also be used by other protection modules or applications.

3.2 Signal Conditioning within the VPR Reference

This section of the paper describes additional details of the implementation of VPR that has been tested. The SV receiver module extracts instantaneous values of voltages and currents are shared with the signal conditioning module which interpolates and down samples the data to adjust for frequency deviations between the nominal frequency and the actual grid frequency. These down-sampled signals will have 16 samples per cycle for each data stream compared to the 80 samples per cycle from

the MU over the process bus (i.e., 960 samples per second for each current and each voltage signal inside the VPR as compared to 4800 sample per second for each data stream from the MU, in a 60Hz AC grid system). These 16 sample per cycle data are used for phasor estimations on each stream: recursive Discrete Fourier Transform (DFT) algorithms are used to calculate the fundamental and up to the 5th harmonic. Sequence components are used for the estimation of signal frequency.

Figure 10 shows conceptually how the signal conditioning recursively updates its estimations on each sample (e.g. 960 times per second in the case of a 60Hz grid frequency). This approach provides both signal continuity and rapid detection for large instantaneous changes to the signals.

3.3 Laboratory Setup for Testing of VPR

The complete reference VPR application has been tested and evaluated using a Protection Relay testing kit.

The VM Hypervisor shall handle multiple network cards to support process bus and station bus. Protection Relay test kit is used for generating sampled values which serves the purpose of Merging Unit. It also measures the operating times based on receipt of respective GOOSE messages. VPR configuration and monitoring is from another VM. This is achieved over the station bus or by using VM virtual network.

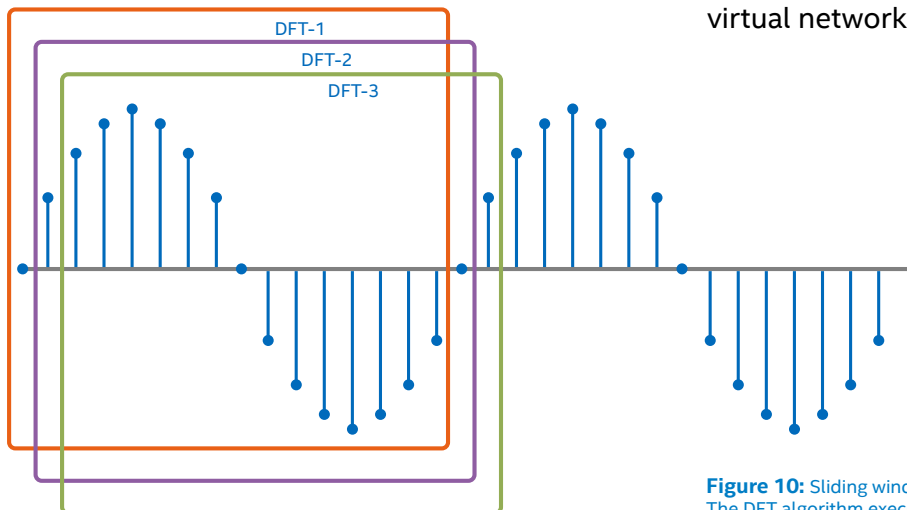
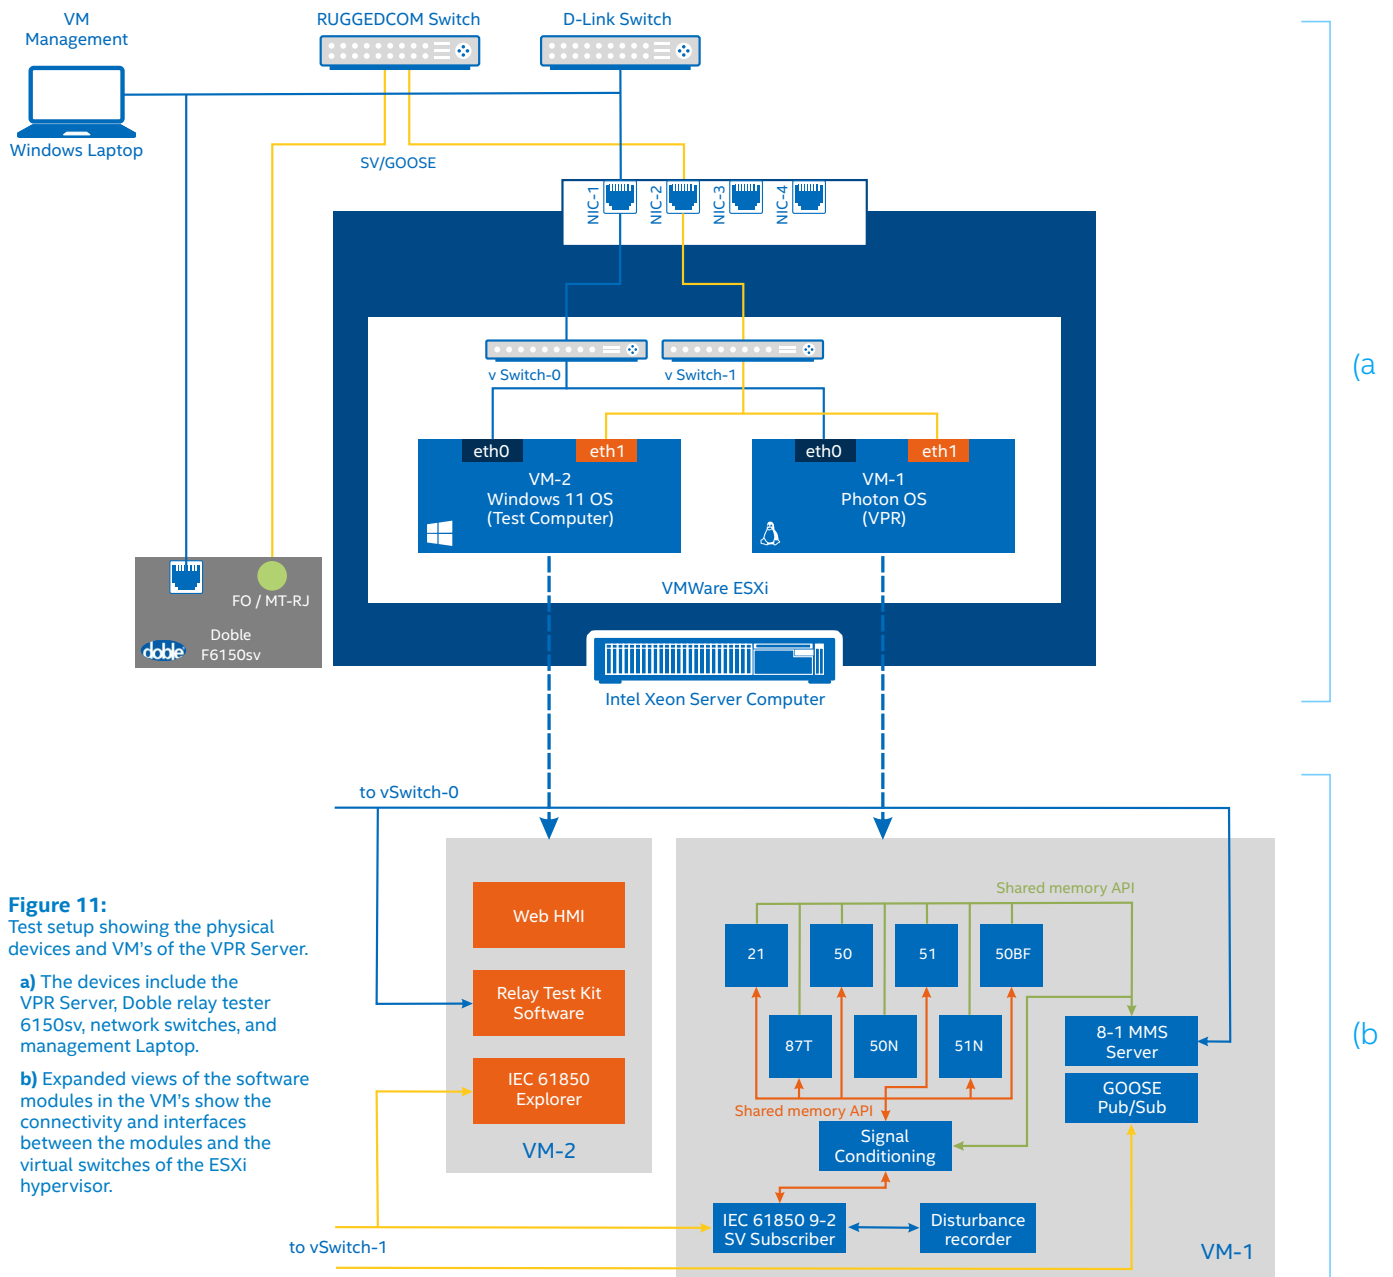


Figure 10: Sliding window concept used for phasor calculations. The DFT algorithm executes at 960 times per second in a 60 Hz grid system.

3.4 Evaluation of VPR Reference Application

Kalkitech has evaluated the reference applications for functionality and end-to-end latency in their own lab. In these tests, a Dell PowerEdge R740 server with Intel Xeon Gold 6248 R (24 Core) processor and 192 GB RAM hosted a VMWare ESXi Hypervisor and two virtual machines, as shown in Figure 11. IEC 61850-9-2 LE compliant SV streams are generated by Doble F6150sv test kit with a user-defined set of test cases representing the grid current and voltage. The test cases are created

according to industry practices to cause the protection function under test to operate. The Doble test kit and the VPR server are networked using a RuggedCom switch. The VPR reference applications receive and process the sampled value data to generate a trip decision periodically (approximately 960 times per second for most of the protection functions). When a fault condition is detected, the VPR reference applications generate and publish a GOOSE message which is received by the Doble test kit to determine the operating threshold and time.



The Breaker Failure function is a multi-step function meant to take effect if an initial overcurrent protection function responds to an overcurrent condition, but the overcurrent condition is not eliminated (which can be a sign that the circuit breaker IED or the circuit breaker itself are not operating as expected). The Breaker Fault function has re-trip time and then a backup trip time parameters which govern the speed at which the backup protection message will be generated and published to the process bus.

The Transformer Differential Protection function requires two SV data streams to be compared to test whether essentially the amount of energy entering the transformer and the amount of energy exiting the transformer are approximately equal. Because the comparison is required, the algorithm requires additional time to respond.

The response time is measured by the test kit as the delta time from the start of the fault until the receipt of the associated GOOSE message. The response time quoted below include the network transit time(s), which is measured separately to be in the order of 1 millisecond round trip.

Table 1 gives the protection function response time when the VPR reference application is tested in a Photon 4.0, 64-bit operating system VM. The IEC-61850 sampled values have been generated using the Doble kit at a sampling rate of 4800 samples/sec for a 60Hz frequency power system. The test pattern for generating the trip signal are step functions starting at zero RMS current level and rising to a level of 10 A RMS which is 200% of the 5A RMS trip threshold.

3.4.1 Response Time Dependence on HW Allocation to VM

The ESXi hypervisor provides configuration settings to tailor the portions of the physical hardware that is allocated to each VM on the system. For a robust design, one criterion is to ensure that the behavior of the protection functions did not depend significantly on the resource allocation to the VM over a wide range of configurations. As an initial test of this, a Kalkitech reference application VM (Photon OS) was provisioned with between 256MB and 8GB of RAM and from 1 core to 5 cores of the CPU. In this range of configurations, the 50, 50N, 51, and 51N protection functions were confirmed to have no significant variation in response time. Below 256 MB of memory the VM functioning can degrade. A set of this data is presented in Table 2 where a Photon OS VM with 4 CPU cores and varying memory allocation is presented. Additional in-depth testing of these and other ESXi hypervisor configuration settings combinations is ongoing.

Protection functions response time (milliseconds)						
No. of SV streams	50	51	50N	51N	21	87T
1	14	15	10	13	14	--
2	15	14	9	11	13	23
3	15	14	8	12	14	22

Table 1: Responses Time for Protection with 1 to 3 SV streams processed by one VPR system. Protection Functions designed as per ANSI Standard codes: 50 - Instantaneous Phase Overcurrent Protection, 51 - Time Delayed Phase Overcurrent Protection, 50N - Instantaneous Derived Neutral Overcurrent, 51N - Time Delayed Derived Neutral Overcurrent, 21 - Distance Protection, 87T - Transformer Differential Protection. In this figure, each SV stream is a simulated test pattern of current and voltage timeseries telemetry data streams of a three-phase power system node (80 samples per cycle).

Protection functions response time (milliseconds)				
RAM	50	51	50N	51N
256 MB	11.6	12.0	11.7	12.1
512 MB	11.5	11.9	11.6	12.0
1 GB	11.5	11.6	11.6	11.7
2 GB	9.7	11.0	9.7	11.1
4 GB	10.5	11.2	10.7	11.3
8 GB	11.1	11.9	11.2	12.0

Table 2: Response time of protective functions with different memory level allocations. The test pattern is for RMS current is a step function from 0 to 200% of threshold.

3.4.2 Response Time versus Fault Magnitude

The following tests were conducted to understand the dependency of the protection algorithm on the fault magnitude. These tests were carried out for the instantaneous overcurrent protection (50) in a virtual machine running Photon Linux with 1CPU core and 4GB RAM.

The protection function design with iterative moving window phasor estimation at 16 samples per cycle are expected to respond faster when the magnitude of the fault current is increased. The data shown in Table 3 and Figure 12 confirms this expectation.

Shot No.	0-10A	4-10a	4-15A	4-20A	4-25A	4-50A
1	17.29	8.13	5.73	4.79	4.69	5.21
2	13.44	8.44	10.1	5	6.77	6.77
3	9.69	10.31	8.23	7.6	5.42	5.21
4	14.38	10.94	6.77	6.77	5.42	3.65
5	10.42	6.15	7.6	5.63	4.69	4.06
6	11.98	8.13	6.15	4.79	7.29	5.1
7	13.75	6.25	5.73	8.96	6.88	4.69
8	10.1	6.77	5.63	5.42	4.69	5.21
9	9.58	10.1	6.77	7.19	7.19	3.13
10	10.42	7.92	6.15	4.69	5.52	5.73
Mean	12.105	8.314	6.886	6.084	5.856	4.876
STD-DEV	2.55	1.69	1.41	1.47	1.07	1.05
STD-DVE/Mean	21%	20%	21%	24%	18%	22%

Table 3: Response times for varying fault magnitude (trip threshold is 5A). Each scenario was tested 10 times with the values, mean, and standard deviation reported.

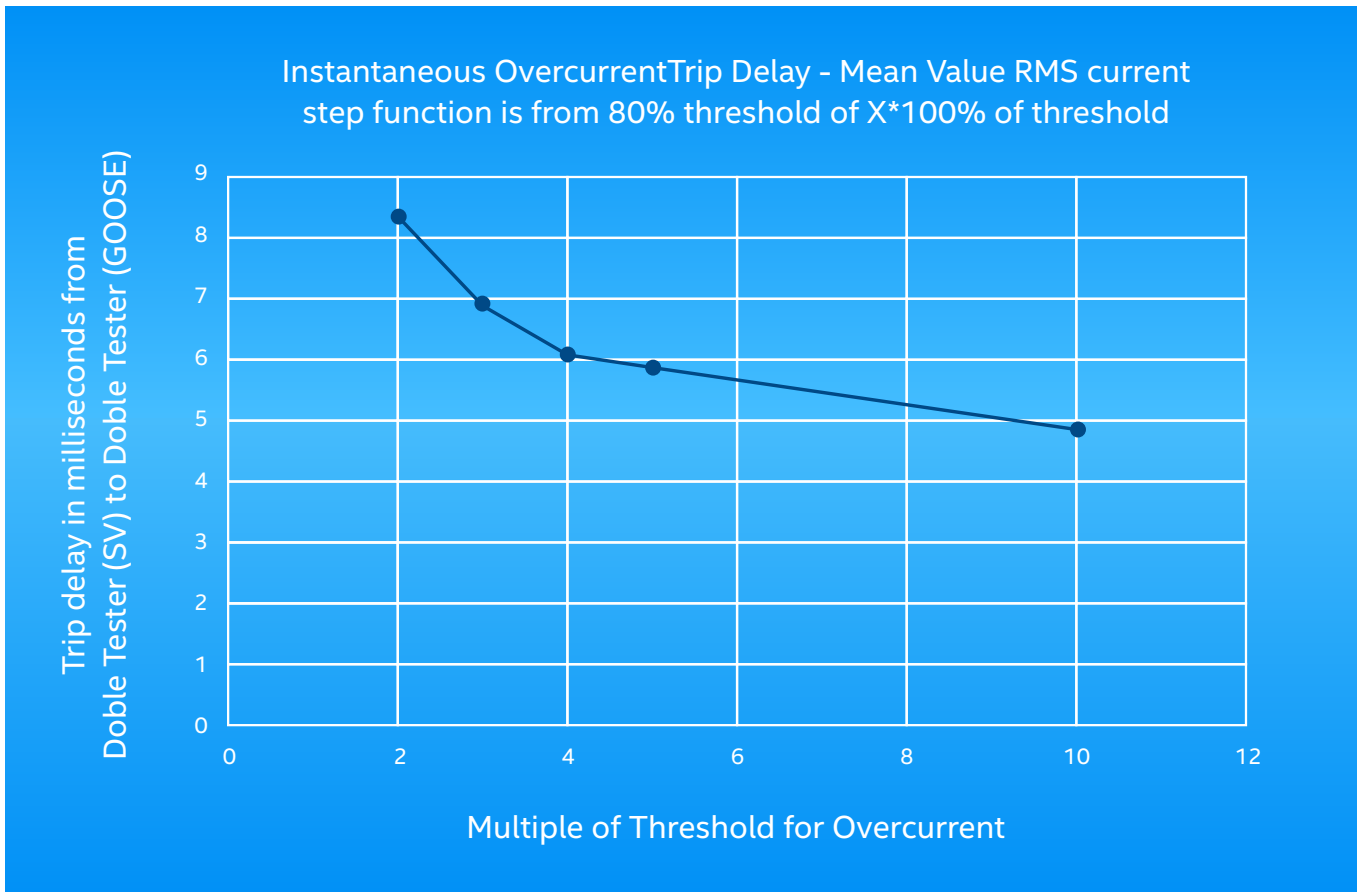


Figure 12: Mean response time with respect to amplitude change in current. The trip threshold is 5A in all conditions; each condition was tested 10 times and the standard deviation of the measurements is approximately 0.2 of the mean value. The data shows the expected trend of reduced delay to pick up the fault as the fault current level increases.

The absolute values of delay approaching ¼ of a 60Hz power cycle, using widely adopted relay test equipment, and with a realistic VPR reference software stack, and using relatively conservative CPU and memory resources from the server platform under test, is intended to be a first step for advancing the VPR concept towards definition, trial use, standardization, and eventually broad commercialization by the industry.

4. The Path Ahead

Any protection system must meet the basic requirements of Selectivity, Speed, Simplicity, Reliability, Economy. These qualities are best achieved by using proven algorithms and methods on top of the best technology available at a given period in time. The advent of IEC 61850 has standardized the modeling of protection functions, the way data is shared both at process level and station level and how data can be shared between protection modules over a communication channel. This

has brought in the much-needed modularity that makes possible the consideration to build software-defined systems that can take advantage of the maturing virtualization technologies applied to real-time and industrial applications.

At the same time, the industry is facing the rapid increase in renewable energy resources on the grid (and especially distributed renewable energy resources); which bring variability and altered power flow patterns not seen before at-scale. These renewable resources are projected to become not only the single largest category of bulk energy produced, but at unprecedented levels in absolute terms. The virtualized and software-defined industrial systems paradigm presents an opportunity for the power utility sector to re-imagine their protection and control systems' implementations and to use these technologies to maintain the ubiquitous power quality and availability that sustains our society and economy.

Hardware vendors such as Dell, Advantech, Crystal Rugged, Moxa, and others have built commercial hardware devices that are candidates for the VPR solution. These devices adhere to IEC61850 and IEEE 1613 specifications and are compatible with advanced networking interface cards for process bus and station bus implementations.

The adoption of VPR has the promise to consolidate the large number of different purpose, make/model, size, and age of discrete devices (hundreds of them in many substations) to a small stack of redundant servers which are standardized and interoperable. Minimization and standardization of hardware can lead to more efficient and effective management and maintenance not to mention ease of replacement and upgrade as well as dramatic space savings. The hardware devices can fit within the management responsibility of the enterprise IT teams of the utility, for example, leading to better total lifecycle cost. Cybersecurity and planned software upgrades are among the services for which the established IT practices can be applied.

Work continues for advanced time synchronization and new techniques for network virtualization. Industry collaboration is essential for standards to configure and manage protection functions in a virtualized environment. Ultimately industry collaboration and standards for performance and interoperability certification will be created in sharing and processing data across multiple modules, processes, and applications.

4.1 Learn More about Intel Technologies

1. Hardware and Software for Real-Time Computing <https://www.intel.com/content/www/us/en/developer/articles/technical/real-time-computing.html>

2. Introduction to Cache Allocation Technology in the Intel® Xeon® Processor E5 v4 Family <https://www.intel.com/content/www/us/en/developer/articles/technical/introduction-to-cache-allocation-technology.html?wapkw=cache%20allocation>

Notices & Disclaimers

Performance varies by use, configuration and other factors.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0522/TB/VC/PDF

