WHITE PAPER

Virtualization of protection and control – Evaluation and deployment considerations

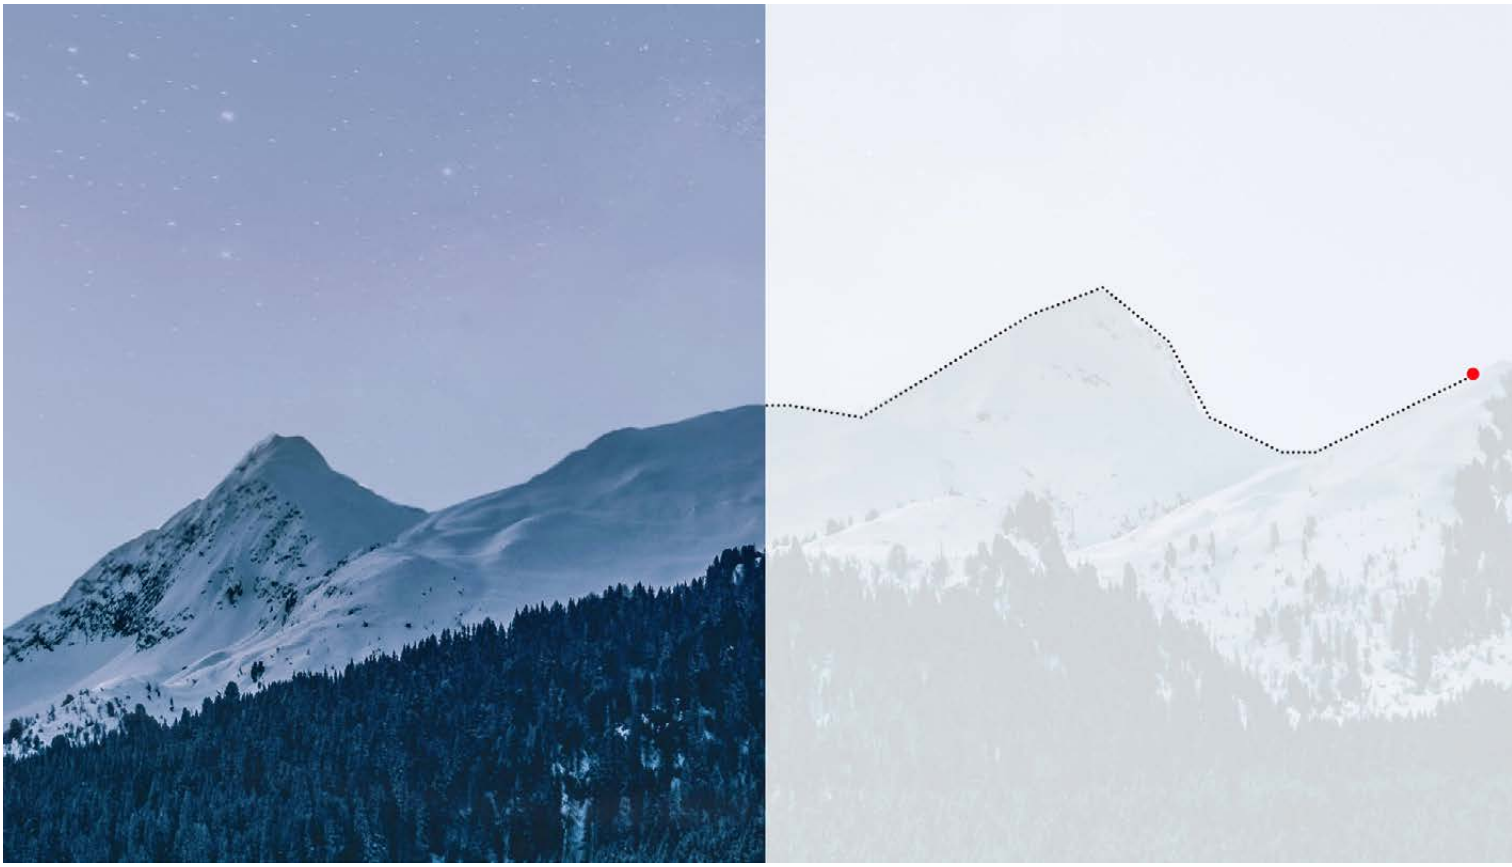

Virtualization of Protection and Control (VPC) allows the use of vendor-independent protection and control (P&C) hardware in IEC 61850 standard-based digital substations.

This paves the way for a novel model for power system protection, which can adapt to the evolving power grid and further offers a pathway to the development of new digital applications.

In this white paper, we address the main system aspects and key technical requirements when implementing VPC systems. We outline virtualization technology and the networking aspects using performance benchmarks laid by IEC 61850 standards. And we present results from the deployment of centralized protection and control (CPC) and VPC, including detailed performance results of protection functions with oscillographic waveforms.

Table of contents

Background – need for reliable flexibility increases

By necessity governmental targets for achieving carbon neutrality are ambitious, with some targets as soon as by 2035. This means that the change needs to happen very fast. Electrical energy will replace other forms of energy, since it is one of most flexible forms of energy that is possible to produce and distribute with a low carbon footprint. Electricity systems are becoming increasingly essential to the functioning of society, and as a result we need to redefine what system resilience means.

With increasing penetration of renewable energy and advent of wide-scale adoption of Battery Energy storage systems (BESS), power generation, and distribution has become increasingly complex and less predictable, the grid must become more flexible and consumer-friendly to facilitate bi-directional flow of power and still meet or exceed system reliability requirements. The required rate of change in the system will be exponentially higher compared to the slow and gradual changes happening today. This will create a challenging combination – how does one simultaneously increase the rate of change and improve robustness and reliability?

The implication of this is that we will see new approaches when it comes to addressing 'system resilience', with digital technologies playing a key role.

Overall system resilience is a key factor in the electric grid, as it will also determine and limit the maximum speed of systemic change. New innovative technologies viz., 5G, virtualized real-time computing etc. can be connected to the grid only if doing so does not reduce the resilience of this critical infrastructure.

With accelerating adoption of EVs, vehicle-to-grid technologies are being piloted or in-phase of deployment. Traditionally resilience has been related to maintaining the power balance and protecting against network faults, resilience of the grid will evolve to a broader definition. During recent years cyber security has also become increasingly important. Furthermore, the aspect of adaptability and updateability of the system will emerge.

Therefore, it is likely that we will see the emergence of new digital platforms that are operating in live mission critical environments, while being constantly updated. The needs are clear, and for the first time, we also have the technologies available for realizing such platforms, including, Machine Learning and Artificial Intelligence (AI) along with 5G, virtualized real-time computing. And a key challenge will be the optimal integration of these technologies.

Protection system evolution

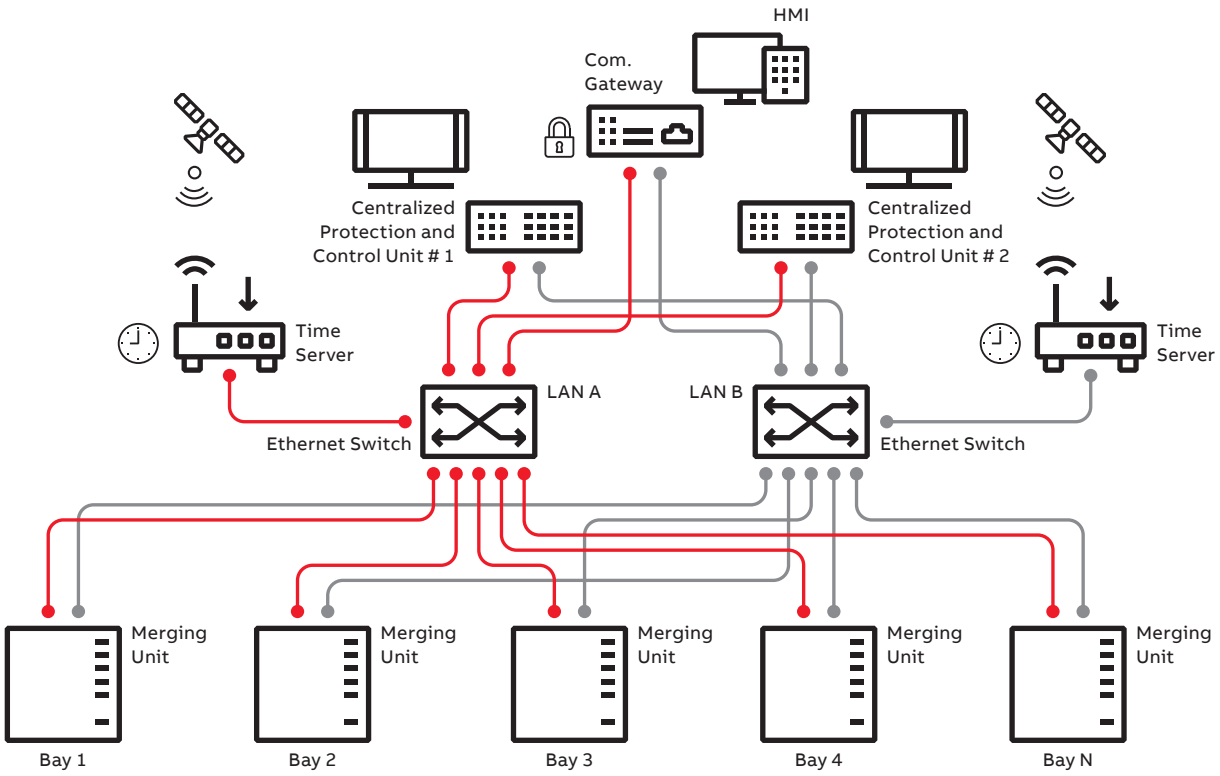

Figure 1. Centralized protection and control system with redundant CPC units [4]

Protection in power systems has been subject to several technological advancements. From electromechanical mechanisms to the microprocessor intelligent electronic device (IED) [1], relaying has been an essential aspect to the continuing development of a more flexible, interconnected, and smart power system that assures dependable and secure delivery of electricity to consumers.

The technologies mentioned above and the standardization offered by IEC 61850 has driven an interest towards centralized protection and control (CPC) systems. With increasing adoption of IEC 61850 and the evolving nature of the power system, products and system engineering tools that support this architecture have already become available on the market.

The main idea of the CPC [2] concept is to move the P&C functionality from multiple bay level devices to a single central processing unit within a substation, leaving only the process interface functionality in the bay level Merging Units (MU). The targeted benefits are related to improved functionality and reduced overall lifecycle costs [3]. An overview of a CPC system, with redundant CPC devices is shown in Figure 1.

An evolutionary step in centralized P&C is to make it more software-driven and hardware-agnostic thus resulting in virtualized protection and control (VPC). This means that the protection application is no longer tied to a particular centralized device, but it is a software image that can be freely deployed to versatile industrial server architectures in different environments.

Virtual protection and control overview

The term virtualization broadly describes the separation of a resource or request for a service from the underlying physical delivery of that service [5]. Virtualization, as applied to P&C, is the use of software for the creation of an abstract image of a traditional P&C device, inside a physical host that is ruggedized computing hardware; and this hardware is based on commercial off-the-shelf components.

With a software image providing real-time, deterministic protection & control, the replicability of this image for creating several instances becomes easier. Software images with different applications providing monitoring & diagnostics can reside on the same host machine.

For instance, a set of protection, control and metering functions of several electrical apparatus residing in a software image can provide amongst several elements, line distance, transformer differential, set of directional overcurrent functions, etc.

Addition of an UFLS (underfrequency-based load-shedding scheme) or interconnection specific functions, for IEEE 1547 compliance, would require additional hardware and several different pieces of ancillary equipment.

In a virtualized environment, this software image can be software-updated to provide additional schemes and controls.

Additionally, profiling of load and measurement of system harmonics is no longer restricted to storage and sampling requirements. Utilizing the machine learning system, loading and apparatus anomalies can also be understood more comprehensively.

Dynamic self-regulated control rules can be programmed in separate applications on the same host machine, which can then suggest optimal switching routines. Real-time and deterministic availability of P&C is an obvious challenge that can be addressed by dedicating a specific set of CPU cores for such tasks.

Allocation of cores to ancillary applications and other software images can be performed by virtualization management software called a hypervisor. This machine-level software is responsible for resource allocation of the host machine.

Redundancy can thus be addressed at the hardware level with an active host running in parallel to another host/server and still using the same redundant process[1)] and station data bus[2)] [6]. Station and process data may be combined into the same physical device, but have logical separation by means of VLANs and filter such as MAC-ID, etc.

A VPC system with a virtual switch (explained in “Network virtualization”, page 12) is shown in Figure 2.

Figure 2. Virtual protection and control system overview

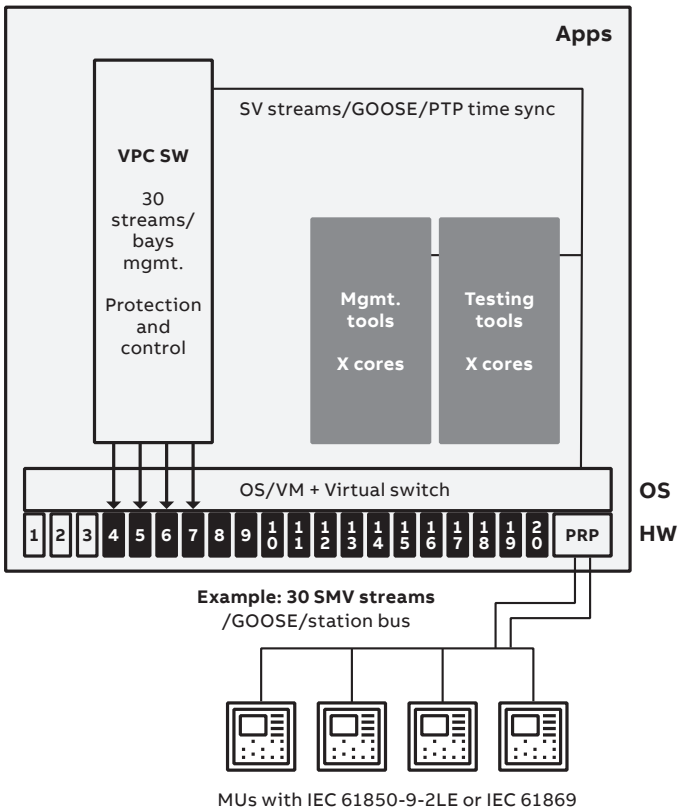

[1)] Process bus connects the primary measurement and control equipment (merging unit) to the VPC. It is expected to provide real-time quality of service.

[2)] Station bus interconnects the whole substation and provides connectivity between central management and the individual bays. It provides only soft real-time quality of service i.e., some jitter in delivery time is acceptable depending on the nature of the traffic.

An example architecture of a substation with VPC is shown below; the choice of virtualization technology between servers may vary and benefits and considerations are further described in “Virtualization solutions and hypervisors”, page 11.

Properly configured merging units provide the measurements to all virtual applications without any additional wiring. This makes the overall solution scalable and reduces the risk of manual errors when adding new applications, which utilize the already existing digital measurement streams.

Further, time synchronization using IEEE 1588v2 over the IEC 62349-3 standard can generate the added level of redundancy in a virtual environment.

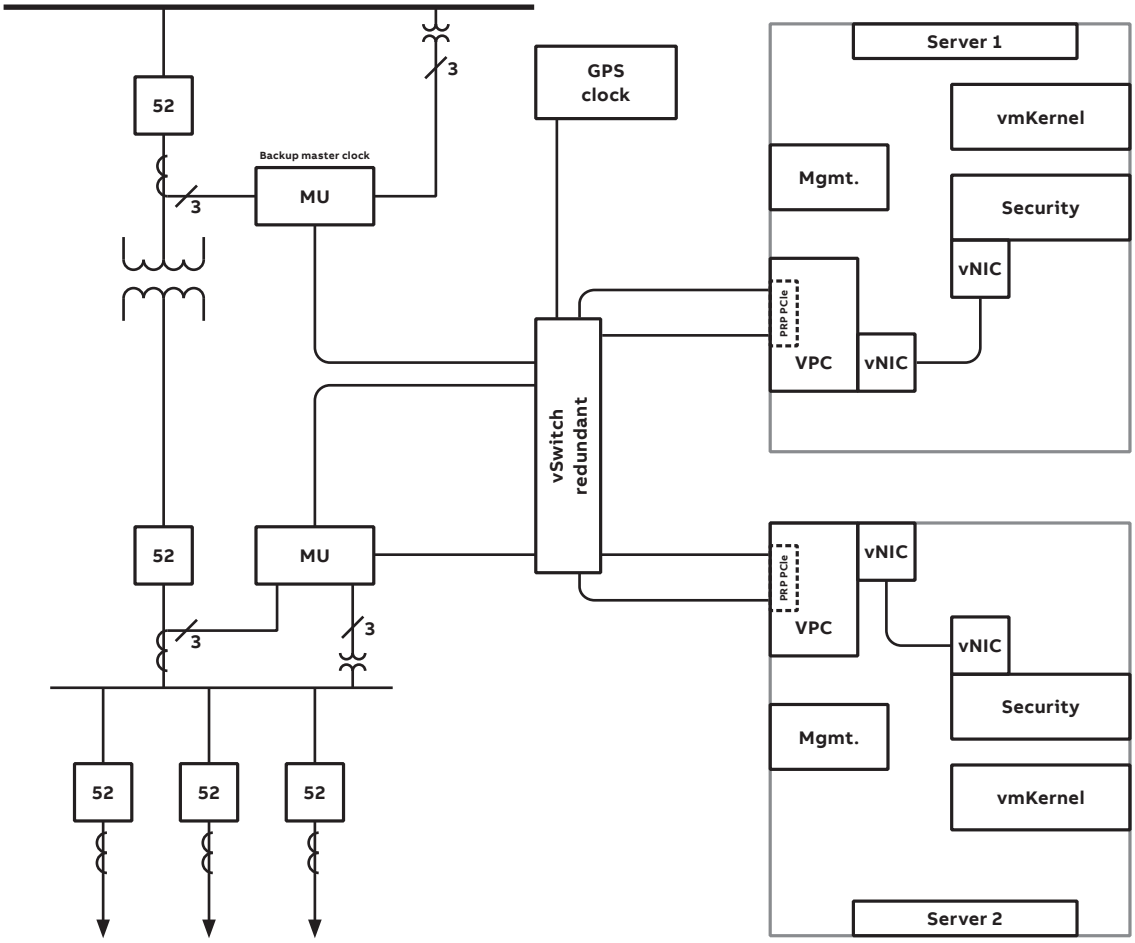

Figure 3: Simplified architecture of redundant VPC in an example substation

Main system aspects

When a VPC solution is being deployed, it is important to think about the overall system aspects first, so that main benefits of the approach can be achieved.

Standardization

One key driver behind the virtualization trend is to increase flexibility and scalability of the utilized solutions and systems. This also indicates the need for an open system that reduces the risks associated with vendor lock-in situations, and in parallel, increases the possibilities to combine (and reuse) the best products/solutions available from different vendors into one solution.

Interoperability and interchangeability of all software (SW) and hardware (HW) components is thus an essential requirement. In practice this means to ensure the solution is based on widely used global standards.

In the area of electricity distribution automation systems, the most important standard is the IEC 61850. Thanks to wide coverage of the standard and increased acceptance if it, it has been the de-facto choice for all modern P&C systems. With Ethernet-based communication and IT-centric modeling concepts it is also the most natural basis for virtualized P&C solutions.

Performance requirements

IEC 61850-5 already defines the application-level performance requirements, seen in Figures 4-6, that can directly be used as performance requirements for virtualized P&C solutions.

Since the requirements are made from the actual P&C application needs point of view, they remain the same regardless of the actual technology that is being used (conventional relay-based protection and control, CPC, or VPC).

—
Figure 4. Application time synchronization classes [7]

Application time synchronization classes for time tagging or sampling		
Time synchronization class	**Accuracy [µs] synchronization error**	**Application**
TL	> 10 000	Low time synchronization accuracy – miscellaneous
T0	10 000	Time tagging of events with an accuracy of 10 ms
T1	1 000	Time tagging of events with an accuracy of 1 ms
T2	100	Time tagging of zero crossings and of data for the distributed synchrocheck. Time tags to support point on wave switching.
T3	25	Miscellaneous
T4	4	Time tagging of samples respectively synchronized sampling
T5	1	High precision time tagging of samples respectively high synchronized sampling

—
Figure 5. Trip message performance requirements [7]

Type 1A "Trip"

The trip is the most important fast message in the substation. Therefore, this message has more demanding requirements compared to all other fast messages. Same performance may be requested for interlocking, intertrips and logic discrimination between protection functions.

Performance class	Requirement description	Transfer time		Typical for interface (IF)
		Class	ms	
P1	The total transmission time shall be below the order of a quarter of a cycle (5 ms for 50 Hz, 4 ms for 60 Hz)	TT6	≤ 3	3, 5, 8
P2	The total transmission time shall be in the order of half a cycle (10 ms for 50 Hz, 8 ms for 60 Hz)	TT5	≤ 10	2, 3, 11

—
Figure 6. Raw data performance requirements [7]

Type 4 – Raw data messages ("Samples")

This message type includes the output data from digitizing transducers and digital instrument transformers independent from the transducer technology (magnetic, optic, etc.). The data will consist of continuous streams of synchronized samples from each IED, interleaved with data from other IEDs.

Transfer time means for the stream of synchronized samples a constant delay resulting in a delay for the functions using the samples e.g. for protection. Therefore, this transfer time shall be so small that no negative impact on application function is experienced.

Performance class	Requirement description	Transfer time		Typical for interface (IF)
		Class	ms	
P7 [a]	Delay acceptable for protection functions using these samples	TT6	≤ 3	4, 8
P8 [b]	Delay acceptable for other functions using these samples	TT5	≤ 10	2, 4, 8

[a] Equivalent to P1
[b] Equivalent to P2

Application architecture

As virtualization enables completely free allocation of functionality, it is important to build the application architecture based on system needs. IEC 61850 modeling provides, also here, good tools for a new kind of application architecture, as each application entity can be modelled as a separate Logical Device (LD).

It is not needed, and often not even beneficial, to repeat old relay-based thinking, where today's bay-level relays would be just designed as bay-level protection applications running in a virtual IED image. Instead, new station (or system) level LDs can be built, where, e.g., low-impedance busbar differential protection, interlocking logics, or directional protection with back-up schemes can be better engineered.

Cyber security

Cyber security (CS) requirements mean both specific technical requirements (e.g., requirements from CS standards like IEC 62443, IEC 62351) and operational requirements (flexible updates in case of CS improvements or vulnerability handling). These requirements relate to both the virtualization environment as well as the individual virtual images. Also in this case, the recommended best practice is to base the solution on widely used standards and to opt for solutions where efficient updating procedures can be achieved.

Operational requirements customized to the substation (Bulk Electric System or otherwise) can be designed per NERC CIP interpretations and applicability [8] [9]. Role Based Access Control (RBAC) configurations can be automatically tied to an enterprise LDAP (Light-Weight Directory Access) server/application. Thus, compliance reporting is simplified without adding additional microprocessors/controllers and networking infrastructure.

Collection of Sequence of Events (SoE) and disturbance records from each electrical bay/switchyard/protection application is centralized into the VPC device and stored at a single location, which can aid in compliance and reporting.

Testing

When a P&C application is virtualized, also new possibilities for testing are achieved, as the testing can be handled in the digital domain. IEC 61850 provides the tools also for this, as it defines specific test modes for LDs and allows all Ethernet traffic to be tagged as ‘simulated’ for testing purposes.

In addition to these new benefits there are also new testing requirements, as the virtualization layer adds a new element to the system that also needs to be tested, and constantly monitored. This means that self-supervision functionality of the system needs to also include new aspects - real-time operation of the virtualized application. This self-supervision functionality may reside in the actual virtual P&C image, or in the hypervisor component, or in both.

Deployment of new firmware

New firmware updates of individual protection applications can be deployed in stand-by mode in a separate virtual image and real-time performance can be tested over a period of time before actual deployment. Incremental updates can be pushed for virtual images and deployed in batches once fully tested. With virtualization technologies it is also possible to take a snapshot of the last known healthy state, then upgrade the application, and roll-back to the snapshot, if something goes wrong during the upgrade.

Main operational aspects

After technical details and system design, it is important to also evaluate the operational side of things – how to operate and maintain the system.

If the organization has an intermediate level of understanding of virtual images and how to deploy them; setting up VPC with the help of guidelines and instructions may be achievable. However, prior experience of IEC 61850 engineering would be needed to ensure that the technology leap and learning curve are not too steep, which might cause execution delays. Selecting this technology can be made in cooperation with qualified system integrators that can contribute advanced expertise.

Key technical considerations

Virtualization technologies have been available for a several years, but so far, the real-time performance for mission critical applications has not been sufficient. Now the situation is changing, and promising solutions are becoming available [10], [11]. Next, we will look at the main technical aspects to be considered.

Virtualization solutions and hypervisors

Virtualization solutions can be divided into Hardware virtualization, Kernel-based virtualization, and OS-level virtualization.

Hardware virtualization is based on abstracting the underlying physical hardware from the guest Operating System (OS), by using specific software components called hypervisors for allocating virtualized system resources and managing the execution of virtualized applications, called Virtual Machines (VM). Hypervisors can be further divided into Type-1 hypervisors that run directly on system hardware, and Type-2 that run on top of the host OS. Type-1 hypervisors are more suitable for providing predictable performance to virtual machines (VMs) than Type-2 hypervisors, as they directly interact with the hardware and have full control of the hardware resources allocated to the VMs. The host OS layer between the hardware and the hypervisor in Type-2 hypervisors can introduce unexpected and unpredictable delays to the hosted hypervisor and then to the VMs on the hypervisor. In addition, any security vulnerabilities and flaws in the host OS could potentially compromise the virtual P&C VM.

Kernel-based virtualization (KVM) means that virtualized applications are executed as part of an operating system. Applications with different safety and security levels can be run on the same hardware, protected from each other by means of software partitioning (also called virtual machines). The partitions can contain different applications with their own operating systems. KVM offers features such a live migration, resource scheduling and control. For virtual P&C this helps in balancing the computing capabilities between two physical host machines and eliminates hardware dependencies. Live migration offered may not be desirable for critical SMV traffic as the failover and recovery times are in the millisecond range.

OS-level virtualization is the newest way of handling virtualization, where the OS components (such as memory access, file system and network access) are abstracted to the guest OS. In this case, the host OS is the same as the one used in each guest OS. Due to so tight coupling, OS-level virtualization imposes low (if any) overhead in terms of resource usage. OS-based services, such as backup and recovery, security management as well as integration with Active Directory (AD) are available, which can help in compliance management for CIP.

Different solutions have their own strengths and weaknesses and these have been extensively elaborated in literature. Hardware Virtualization provides the strongest isolation between different VMs, whereas OS-level virtualized containers provide the lowest overheads. Main aspect to highlight is that with all mentioned solutions it is possible to achieve the required deterministic operation and reliability for P&C purposes.

The case for hardware acceleration in P&C

Hardware acceleration technologies have enabled the idea of the adoption virtual protection as it can provide guaranteed service to SMV and GOOSE traffic for the protection functions. Since P&C is a process-intensive task, using such technology can allow better hardware resource allocation and performance. Type-1 hypervisors use hardware acceleration (enabled through BIOS of the host hardware) and would be mostly preferred for P&C virtualization. Type-2 hypervisors can rely on hardware acceleration through their host OS layer, but will fall back on software emulation in absence of native support.

Network virtualization

A key aspect for VPC solutions is to ensure reliable handling of network virtualization. P&C applications are dependent on SMV and GOOSE data coming from the merging units. Handling this traffic in real time is as important as the execution of the actual P&C application.

Virtualization of the network interface card (NIC) can occur at the media access control (MAC) layer. The physical NIC simply becomes a port forwarding backplane. The process and stations bus traffic from individual bays can be virtualized on the physical host where the VPC is running. With virtualization, it is possible to utilize the guaranteed bandwidth feature, which allows virtual NICs (vNICs) to reserve a hardware traffic lane [12] for SMV traffic, which is a stream of 80 samples/cycle (approx... 5..6 Mbit/s) constant load on the network. [3)]

Different network virtualization technologies are evaluated in more detail in a Linux environment in [13]. The simplest solution is to use direct host access, but it provides no isolation and can lead to issues in both network performance and management. A physical Ethernet switch can be emulated as a bridge, or as a MACvlan (version of software bridge in Linux). Single root I/O virtualization (SR-IOV) is an extension to the PCI Express specification, and allows a device, such as the NIC, to virtualize its resources. Results in [13] show that different network virtualization technologies are indeed sensitive to system and network load situations, and can fail in stress tests, if not engineered properly.

Nevertheless, the results also show that with proper design (e.g., hardware assisted SR-IOV solutions) a robust operation of the system can be guaranteed. Internal delays such as any store-and-forward delay, vSwitch (virtual switch providing logical connection between different applications and/or VMs) latency and packet queue latency is minimized compared to a conventional system and the delay settings for SMV traffic can be internally set for each MU.

[3)] 5 Mbit/s (50Hz) or 6 Mbit/s (60Hz) per one MU streaming three phase current and voltage and ground/neutral current and voltage. Actual size depends on SMV ID, OptFlds as well as length of control block and data set name.

Loops between redundant LAN A and LAN B virtualized network(s) in both the physical domain and the virtual domain can create broadcast storms that will impact protection system reliability. To mitigate, network load monitoring, balancing and rate limiting algorithms need to be embedded into the vSwitch. To mitigate the effects of broadcast storms, at minimum the MUs and the network switch would need to have rate limiting functionality.

Time synchronization

With Ethernet-based technology it is possible to achieve software-based time synchronization with an accuracy of 1 ms quite easily, and without any help from hardware. This is also what the IEC 61850 standard refers to as the basic time synchronization accuracy class (T1).

An older and more common protocol is the SNTP (Simple Network Time Protocol), which is suitable for local substation synchronization in relatively small systems. However, if the SNTP server is behind multiple Ethernet nodes, the latency increases, which reduces the accuracy of the time synchronization. Therefore, SNTP is not an ideal solution for system-wide implementation.

Normally a Global Positioning System (GPS) or equivalent time synchronization resource is required in every substation. IEEE 1588v2 and IEC 61850-9-3 deal with these issues and make it possible to achieve a time synchronization accuracy of 1 µs. This is required if an IEC 61850-9-2 process bus [6] is used. In case of a loss in GPS clock; one of the merging units can take over as a master clock per Best Master Clock Algorithm defined in IEEE 1588v2 and synchronize the servers and other MUs. PTP pass-through between different VMs and applications in a server can provide a highly accurate level of time synchronization. VMs and Apps can be transparent clock themselves, and publish processing time in the stack as ‘ΔT’ in the correction field. For non-time sensitive ancillary applications, SNTP can be utilized.

Redundancy

Two crucial parameters for architectures utilizing a virtual P&C system are high availability and high reliability of the communication network. The IEC 61850 standard identifies this need, and specifically defines in IEC 61850-5 the tolerated delay for application recovery and the required communication recovery times for different applications and services.

The tolerated application recovery time ranges from 800 ms for SCADA, to 40 µs for sampled values. The required communication recovery time ranges from 400 ms for SCADA, to 0 for sampled values. To address such time critical need for zero recovery time networks, IEC 61850 standard mandates the use of IEC 62439-3 [14] standard wherein clause 4 of the standard defines Parallel Redundancy Protocol (PRP) and Clause 5 defines High-Availability Seamless Redundancy (HSR).

Both methods of network recovery provide 'zero recovery time' with no packet loss in case of a single network failure. Handling of redundancy within the server, using a dedicated PRP PCIe interface is the preferred option, when large number of bays are to be connected. Using HSR topology for connecting MUs with VPC can restrict the number of MUs, but can minimize the amount of needed hardware.[4)]

Redundancy can be strictly provided for station and process bus traffic, but other apps, such as management and security, may use virtual red-box (provides Singly Attached Nodes (SAN) access to redundant LAN A and LAN B).

Hardware

The goal with virtualization is to provide as much flexibility as possible for HW selection. However, main environmental (e.g., Electromagnetic Compatibility, EMC) requirements should be met, ensuring that the hardware reliability and robustness is at suitable level for electrical substations. These requirements are further elaborated in IEC 61850 [7] and IEEE 1613.

It is assumed that most ruggedized servers will not be installed in an outdoor environment, and as such outdoor operating temperature ranges would not be applicable.

[4)] To comply with sampling rate of 80 samples/cycle; a limited amount of devices are supported at 100 Mbit/s; for increased number of devices 1 Gbit/sec interface may be considered and evaluated against the cost of using a PRP-based topology and associated hardware.

Practical experience

VPC has been explored by ABB in a live substation, with the target to compare performance between an existing CPC solution available on the market, and a virtualized version. The following over-current incident was successfully detected by the physical CPC unit and the virtual algorithm at exactly the same time. The directional phase overcurrent, low stage element (NSI. 67/51P) function operated at the same time in bay J07. Oscillographic (COMTRADE) indicate both J01 and J07 observed a fault involving Phase B, C at instant of pickup.

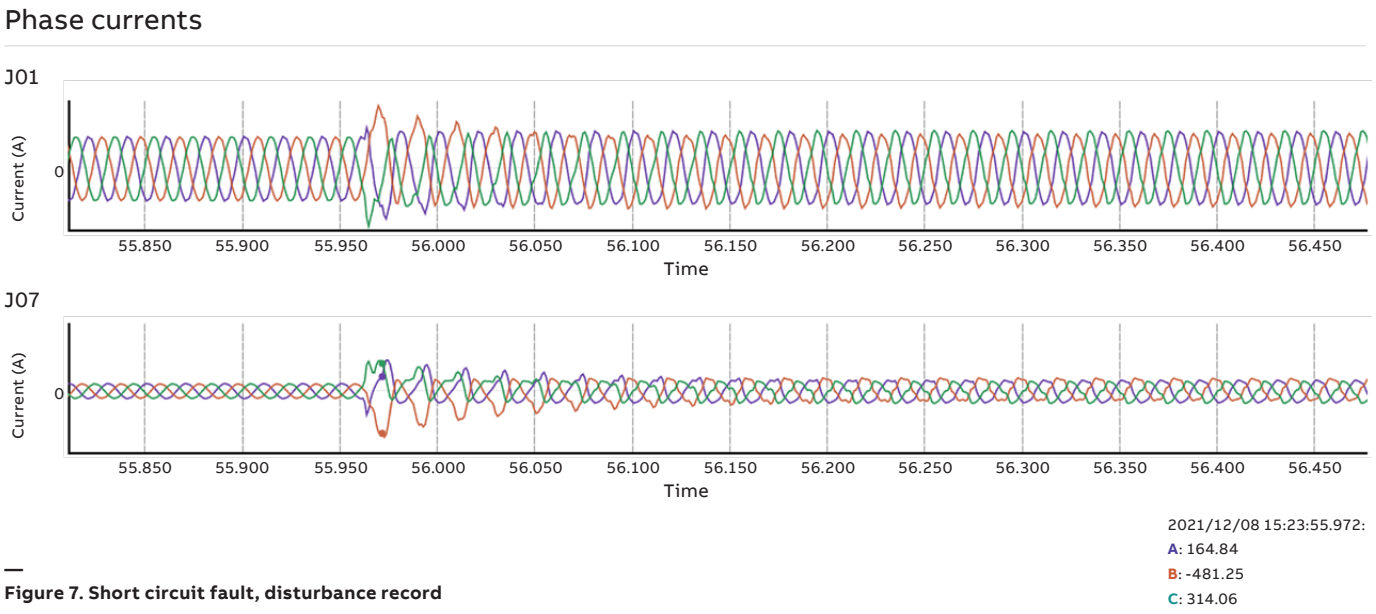

Figure 7. Short circuit fault, disturbance record

TIMESTAMP ▾	SUBSTATION	BAY	FUNCTION	DESCRIPTION	VALUE	IDENTIFIER
8.12.2021 15:23:55.990	VPC1	J07		Start	True	LDO.DPHLPTOC15
8.12.2021 15:23:55.990	VPC2	J07		Start	True	LDO.DPHLPTOC15
8.12.2021 15:23:55.990	CPC	J07		Start	True	LDO.DPHLPTOC15

Figure 8. Short circuit fault analyzed by different applications (1 CPC device, 2 VPC images)

Conclusions

Virtualization has a lot of potential, and we're now at the stage that the technology is ready for wider usage. There are benefits that can be achieved with virtualization, but there are many factors that need to be taken into account to ensure the result is reliable, robust and secure.

The best way forward is to partner with good technology vendors and as a first step to run pilot installations in order to fully determine the best solution for different cases.

Centralized protection and control solutions are already readily available on the market, and they bring several benefits. In the available solutions, software-orientation and the overall concept is the same as with virtual protection, just with pre-defined hardware.

Bibliography

[1] B. Lundqvist, 100 years of relay protection, the Swedish ABB relay history, Technical report, ABB Automation Products, Sweden, 2010.

[2] R. Das et al., Centralized Substation Protection and Control, WG K15 Report, IEEE Power System Relaying Committee, 2015.

[3] J. Valtari, Centralized Architecture of the Electricity Distribution Substation Automation – Benefits and Possibilities, Ph.D. Thesis, Tampere University of Technology, Finland, 2013.

[4] ABB, Centralized protection and control – Enhancing reliability, availability, flexibility and improving operating cost-efficiency of distribution substations, ABB White Paper, 2022, [ONLINE]. Available: https://search.abb.com/library/Download.aspx?DocumentID=2NGA001420&LanguageCode=en&DocumentPartId=&Action=Launch

[5] VMWare, Virtualization Overview, White Paper, 2022, [ONLINE]. Available: https://www.vmware.com/pdf/virtualization.pdf

[6] IEC, Communication networks and systems for power utility automation – Part 90-4: Network engineering guidelines. Technical Report, The International Electrotechnical Commission, 2020.

[7] IEC, Standard for Communication Networks and Systems for Power Utility Automation, IEC Standard 61850-5 Edition 2. The International Electrotechnical Commission, 2013.

[8] "ERO Enterprise CMEP Practice Guide," NERC, [Online]. Available: https://www.nerc.com/pa/comp/guidance/CMEPPracticeGuidesDL/CMEP%20Practice%20Guide%20%20Virtual%20Systems.pdf

[9] "Compliance Guide," NERC, [Online]. Available: https://www.nerc.com/pa/comp/guidance/Pages/default.aspx?View={873894d8-c688-4340-9c23-7cece8f0af5c}&SortField=Standards_x0020_Grouping&SortDir=Asc

[10] A. Moga, T. Sivanthi and C. Franke, Os-level virtualization for industrial automation systems: are we there yet?, SAC '16: Proceedings of the 31st Annual ACM Symposium on Applied Computing, p. 1838-1843, 2016.

[11] L. Abeini, A. Balsini and T. Cucinotta, Container-Based Real-Time Scheduling in the Linux Kernel, SIGBED Rev., vol. 16, no. 3, p. 33-38, 2019.

[12] Anjing Wang et al., Network Virtualization: Technologies, Perspectives, and Frontiers, Journal of Lightwave Technology, 2013.

[13] G. Albanese, G. Giannopoulos, T. Sivanthi, R. Birke and T. Sivanthi, Evaluation of Networking Options for Containerized Deployment of Real-Time Applications, 26th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), 2021.

[14] IEC, Industrial communication networks – High availability automation networks – Part 3: Parallel Redundancy Protocol (PRP) and High-availability Seamless Redundancy (HSR) Edition 4, The International Electrotechnical Commission, 2021.

List of abbreviations

AD	Active Directory
AI	Artificial Intelligence
BESS	Battery Energy Storage Systems
CIP	Critical Infrastructure Protection
CPC	Centralized Protection and Control
CPU	Central Processing Unit
CS	Cyber Security
EMC	Electromagnetic Compatibility
EV	Electrical Vehicle
GOOSE	Generic Object Oriented Substation Event
GPS	Global Positioning System
HSR	High-Availability Seamless Redundancy
HW	Hardware
IED	Intelligent Electronic Device
KVM	Kernel-based Virtual Machine
LDAP	Light-Weight Directory Access
LD	Logical Device
MAC	Media Access Control
MS	Millisecond
MU	Merging Unit
MFA	Multi Factor Authentication
NERC CIP	North American Electric Reliability Corporation Critical Infrastructure Protection
NIC	Network Interface Card
OS	Operating System
P&C	Protection and Control
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect express
PRP	Parallel Redundancy Protocol
PTP	Precision Time Protocol
RBAC	Role Based Access Controls
SAN	Singly Attached Nodes
SCADA	Substation Automation and Data Acquisition
SoE	Sequence of Events
SMV	Sampled Measured Values
SNTP	Simple Network Time Protocol
SW	Software
UFLS	Underfrequency-based Load-shedding Scheme
VM	Virtual Machine
vNIC	Virtual Network Interface Card
VLAN	Virtual Local Area Network
VPC	Virtualization of Protection and Control

Acknowledgements and trademarks

This white paper is based on the conference paper first published at the CIGRE US National Committee 2022 Grid of the Future Symposium.
Authors: Jani Valtari and Dinesh Baradi.

ABB and Relion are registered trademarks of the ABB Group. All other brand or product names mentioned in this document may be trademarks or registered trademarks of their respective holders. ABB is a Titanium IoT Solutions member of the Intel Partner Alliance.

Additional information
We reserve the right to make technical changes or modify the contents of this document without prior notice. With regard to purchase orders, the agreed particulars shall prevail. ABB does not accept any responsibility whatsoever for potential errors or possible lack of information in this document.

We reserve all rights in this document and in the subject matter and illustrations contained therein. Any reproduction, disclosure to third parties or utilization of its contents – in whole or in parts – is forbidden without prior written consent of ABB.

For more information, please contact your local ABB representative or visit the centralized protection campaign page
solutions.abb/centralizedprotection

© Copyright 2023 ABB. All rights reserved.
Specifications subject to change without notice.

2NGA001642 REV A EN 3.3.2023