GE Vernova's
Electrification Segment

VALIDATING VPAC

PERFORMANCE, RELIABILITY, AND SCALABILITY IN VIRTUALIZED PROTECTION SYSTEMS

GE VERNOVA

INTRODUCTION

Virtualized Protection, Automation, and Control (VPAC) offers greater flexibility, scalability, and easier maintenance than traditional hardware-based IEDs. By separating software from hardware, VPAC enables remote deployment, faster upgrades, and optimized resource use, helping utilities reduce costs and improve system efficiency.

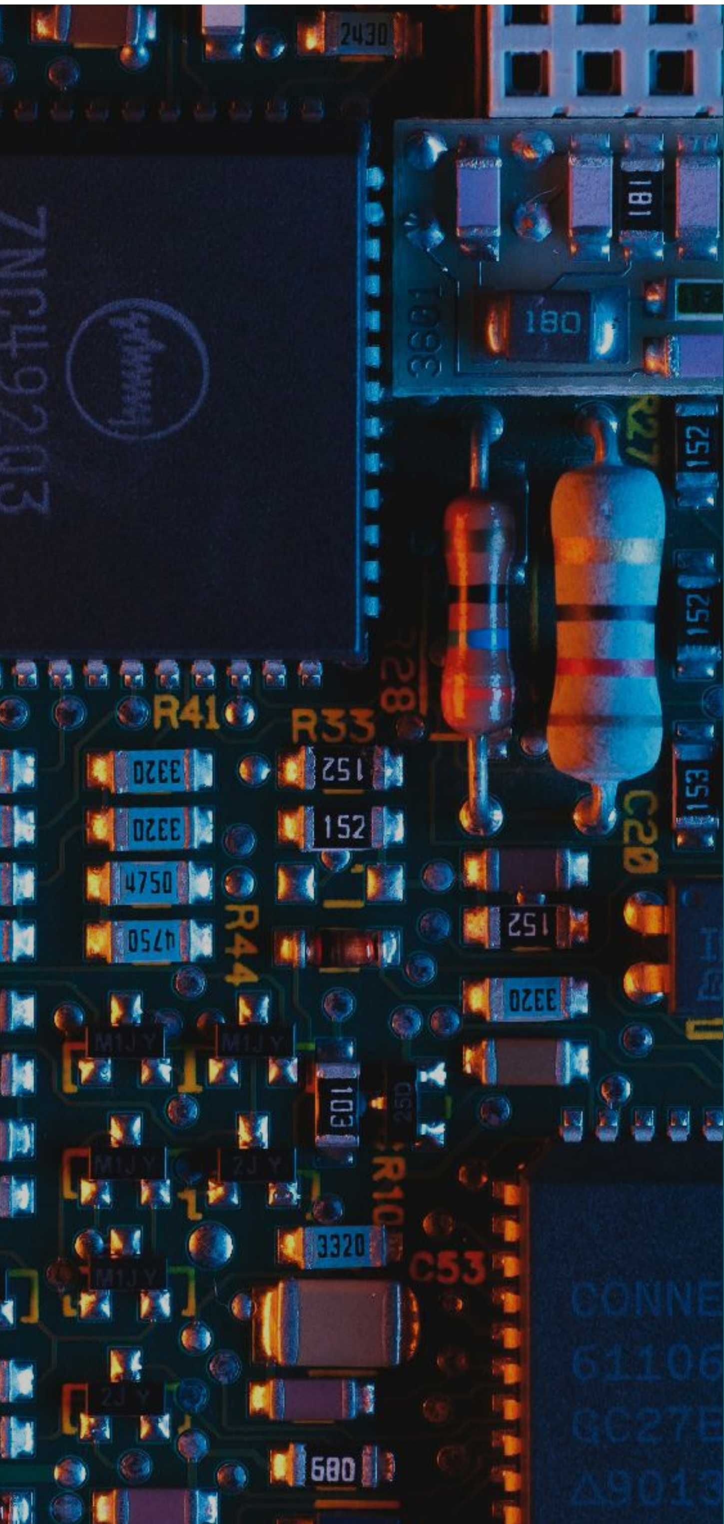

However, before VPAC can be widely used in substations, it must be tested to confirm its reliability, performance, and scalability under real-world conditions.

Unlike traditional protection systems, VPAC depends on third-party industrial servers, virtualization software, and shared computing resources, which introduce new performance considerations.

This white paper outlines a structured testing approach to evaluate:

- System resilience under component failures, such as fan malfunctions.
- Hardware limits to ensure protection applications run reliably.
- Near real-time performance to measure response times under different workloads.
- Virtualization techniques to reduce jitter and improve trip time consistency.

These tests provide data-driven insights to validate VPAC for mission-critical protection applications.

TEST ENVIRONMENT AND SETUP

The device under test is a Dell PowerEdge XR12 industrial server, configured as follows:

- **Processor:** Intel Xeon Gold 5318N (2.1 GHz base clock, 24 CPU cores).
- **Memory:** 256 GB RAM.
- **Cooling: Six cold-swappable fans**.
- **Network Interface Cards (NICs):**
  - 2x Intel E810-XXV Dual Port 25 GbE SFP28 (SR-IOV compatible).
  - 4x Broadcom BCM57504 NetXtreme-E 10G/25G adapters (SR-IOV compatible).
  - Redbox device for PRP-based Process Bus LAN connections.

Figure 1: The Industrial Server used in the tests

NETWORK CONFIGURATION

The test environment utilizes two network segments:

1. **PROCESS BUS LAN:** Handling Sampled Values (SV) and GOOSE messages for near real-time protection applications.
2. **MANAGEMENT LAN:** Used for configuring, monitoring, and managing virtualized applications.

ENVIRONMENTAL TESTING: RESILIENCE TO FAN FAILURES

Industrial substations expose VPAC servers to high temperatures and unpredictable conditions. Unlike traditional IEDs, VPAC servers rely on active cooling, making fan failures a key concern. This test assessed whether VPAC systems remain stable when cooling components fail.

Test Procedure

The Dell PowerEdge XR12 industrial server was placed inside a temperature chamber capable of controlled thermal stress testing to simulate real-world conditions:

- **Ambient temperature limits:** Tests were conducted at 45°C and 55°C the maximum rated ambient air temperature for the device.
- **Test duration:**
  - Initial 10 minutes recorded at room temperature.
  - Temperature gradually increased to 45/55°C over one hour.
  - Cooling period of 30 minutes after the chamber was opened.
- **Fan status monitoring:** Near real-time temperature and fan speed data was collected using the iDRAC condition monitoring interface.

GE VERNOVA

All the thermal parameters can be monitored using the iDRAC interface. In the iDRAC Web interface, Cooling overview, Temperature overview, Temperature overview, Fan Status, Temperature overview, Fan Status, Temperature probes status and the System Temperature Historical Data section can be seen.

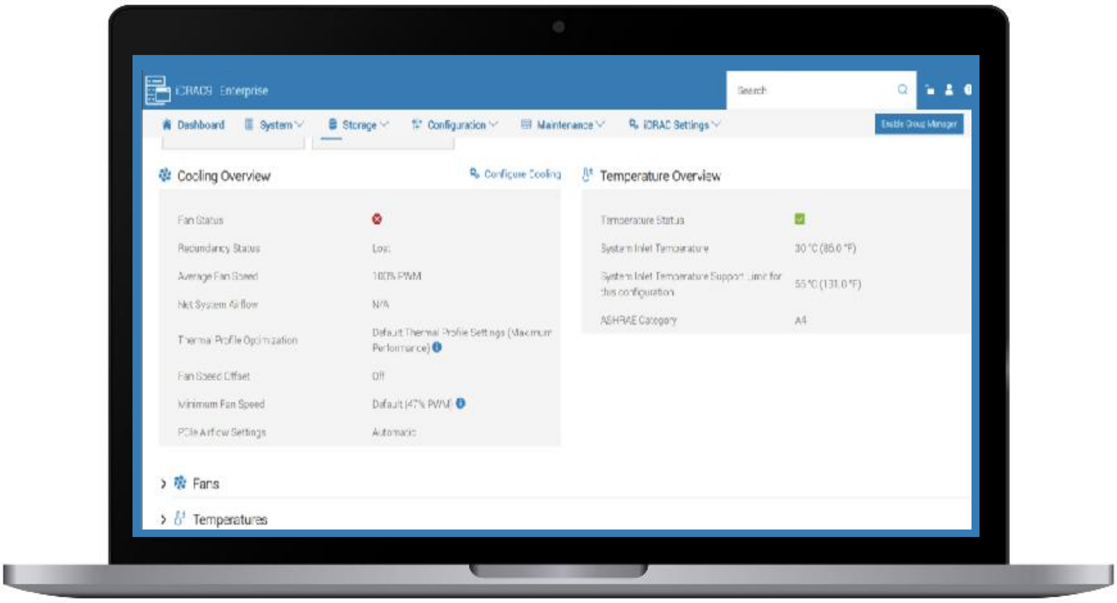

Figure 2: iDRAC condition monitoring interface

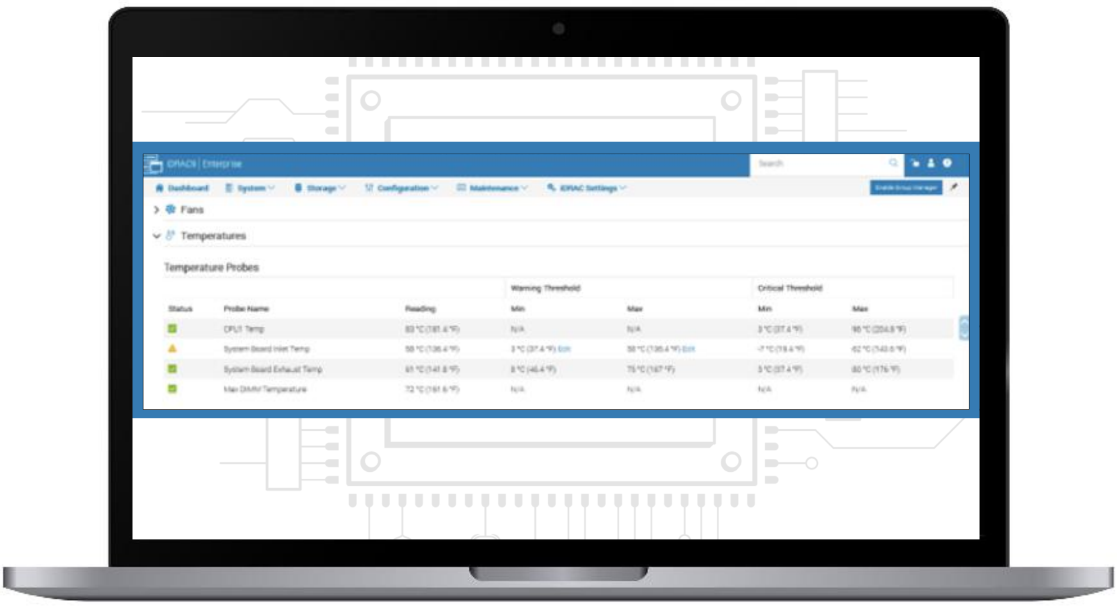

Figure 3: iDRAC temperature measurements including CPU and Memory as well as ambient exhaust temperatures

TEST CASES

TEST CASE	CONDITIONS	VENTILATION
TEST CASE 1	Normal operation (all fans working)	Reduced airflow
TEST CASE 2	One fan removed	Reduced airflow
TEST CASE 3	Two fans removed	Reduced airflow

Each test was performed under three system load conditions:

- **Idle State:** No active Virtual Machines (VMs).
- **Single VM Load:** One active VM running a protection application.
- **Full Load:** Ten VMs running protection applications with a simulated 90% CPU load.

FAN REMOVAL PROCESS

For Test Cases 2 and 3, fans were physically removed following standard electrostatic discharge (ESD) protection protocols.

1. SERVER POWERED DOWN AND DISCONNECTED.
2. RISER-1 MOUNTING SCREWS AND TOP PLASTIC HOUSING REMOVED TO ACCESS THE FANS.
3. FAN CONNECTOR DETACHED, AND FAN REMOVED.

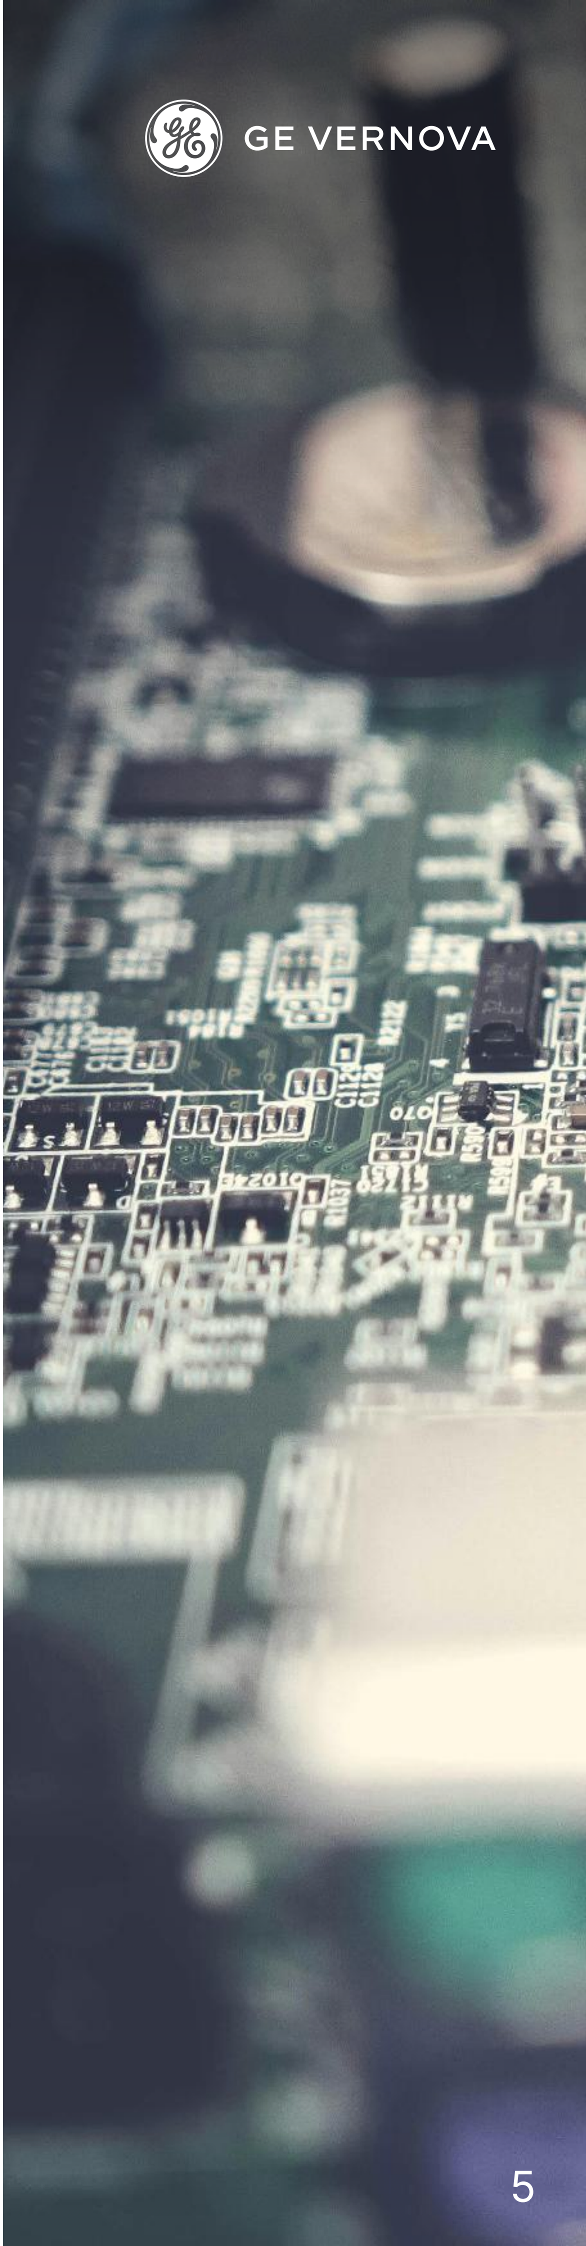

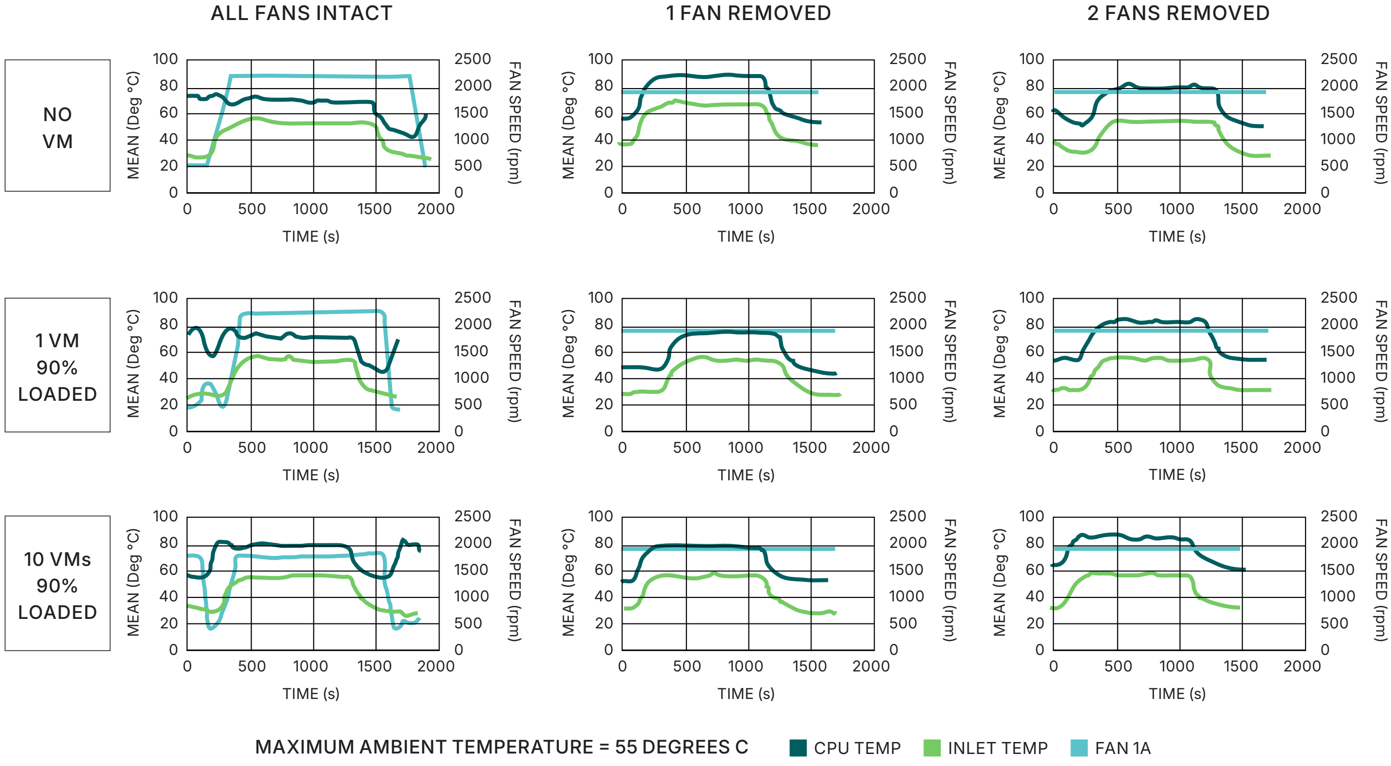

Figure 4: Temperature test at the stated maximum ambient temperature of 55 degrees

TEST RESULTS AND ANALYSIS

All Fans Operating (Baseline Performance)

- At 55°C ambient inlet temperature, all six fans immediately increased speed to 18,000 RPM (max speed).
- CPU temperature remained stable at approximately 70°C under low workloads (0 – 1 active VMs).
- Even under full system load (10 VMs, 90% CPU utilization across 20 cores), CPU temperature remained at 80°C, well below the maximum allowable 96°C.

One Fan Removed (Single Failure Scenario)

- Remaining fans compensated by running at full speed.
- CPU temperature remained below 80°C, even with 10 VMs running at 90% load.
- No significant performance impact observed.

Two Fans Removed (Worst-Case Failure Scenario)

- CPU temperature increased by 5 – 10°C across all scenarios.
- Even under maximum stress conditions (10 VMs at 90% CPU load, 55°C ambient temperature), CPU temperature remained below 86°C, well within safe operating limits.
- No unexpected shutdowns or performance degradation observed.

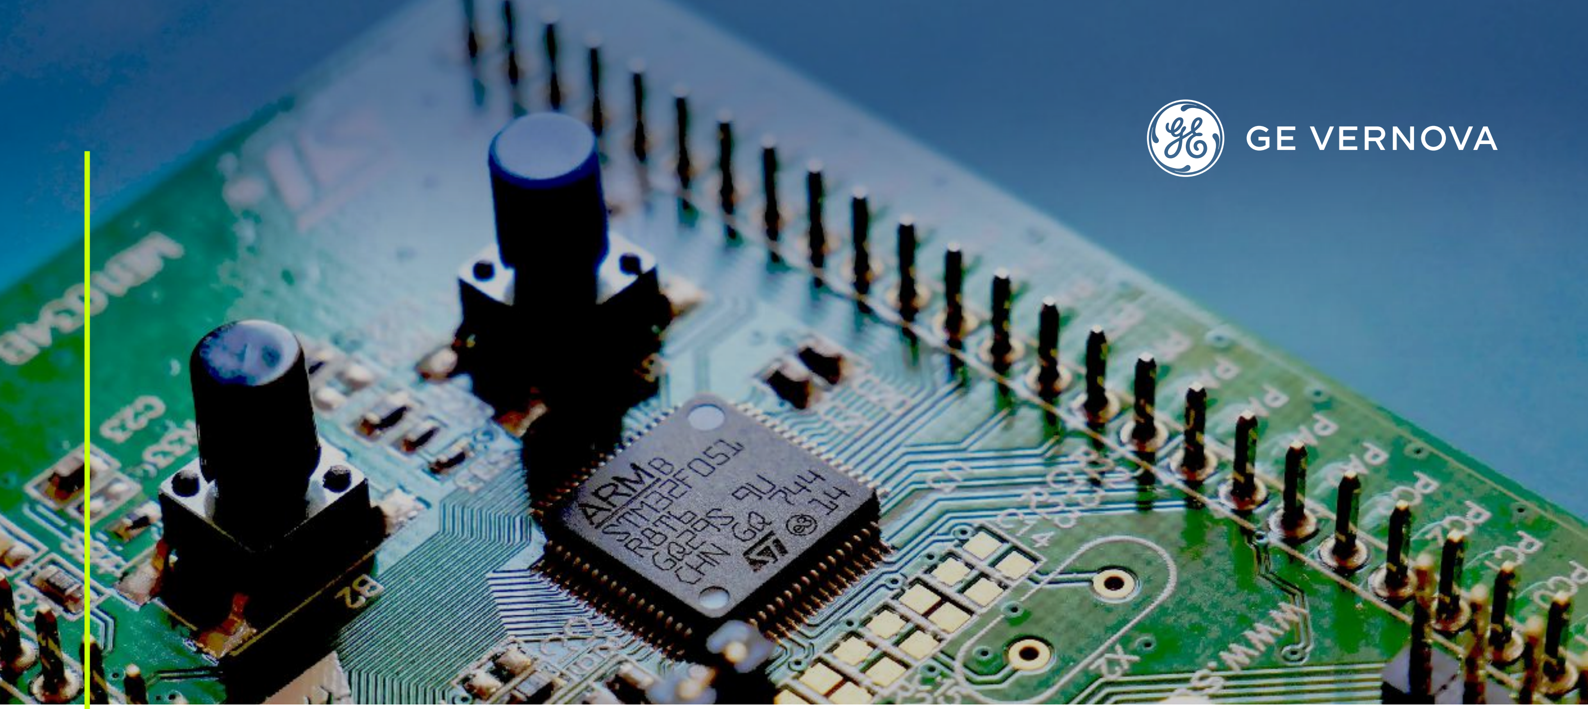

KEY TAKEAWAYS

- VPAC systems demonstrated high thermal resilience, maintaining stable operation even in extreme conditions.
- Redundant cooling design ensures stability—even with two fan failures, CPU temperature remained well below critical thresholds.
- 2N+2 level of redundancy confirmed for forced ventilation, making VPAC deployments feasible even in substations with limited cooling support.

Even under worst-case conditions (two fan failures, maximum system load, and extreme temperatures), the VPAC server maintained stable performance, proving its suitability for deployment in harsh substation environments.

NEAR REAL-TIME TESTING FOR PROTECTION PERFORMANCE

The performance of a Virtualized Protection System is highly dependent on software architecture. Unlike traditional hardware-based IEDs, low latency and deterministic response times are not inherently guaranteed in a VPAC system. Instead, they rely on the correct configuration of virtualization parameters to optimize protection trip times and minimize jitter.

This section evaluates key virtualization techniques and their impact on latency, determinism, and performance stability.

The following architectural choices were tested:

- Hyperthreading ON vs OFF
- SR-IOV vs vSwitch (Process Bus communication)
- Single Protection Function (Overcurrent) vs Multiple Protection Functions
- CPU Isolation vs CPU Isolation + IRQ Offloading

UNDERSTANDING VIRTUALIZATION TECHNIQUES

HYPERTHREADING (HT) ON VS OFF

- Hyperthreading (HT) allows a single physical CPU core to act as two logical cores, executing different workloads simultaneously.
- In many general-purpose computing applications, hyperthreading improves efficiency by increasing parallelism.
- In near real-time protection applications, hyperthreading may introduce unpredictable scheduling delays, leading to increased jitter in trip times.
- Testing was performed to assess whether disabling HT improves deterministic response times for VPAC applications.

VIRTUAL MACHINE (VM) PINNING TO SPECIFIC CPU CORES

- VM Pinning ensures that each Virtual Machine is permanently allocated specific CPU cores, preventing other workloads from using those resources.
- This isolation prevents CPU resource contention, making performance more predictable and repeatable.
- VPAC systems were tested with CPU Pinning enabled vs. dynamic CPU allocation to measure its impact on trip times.

SR-IOV VS VIRTUAL SWITCH (VSWITCH)

- Single Root I/O Virtualization (SR-IOV) allows direct access to physical network hardware, bypassing the hypervisor's virtual networking layer.
- Virtual Switch (vSwitch) processes network traffic in software, introducing additional delays and increasing communication jitter.
- The hypothesis: SR-IOV should provide lower latencies and better near real-time determinism compared to vSwitch.

CPU ISOLATION AND IRQ OFFLOADING

- CPU Isolation dedicates specific CPU cores exclusively to near real-time protection tasks, preventing background OS processes from interfering.
- IRQ Offloading assigns interrupt handling tasks (such as network packet processing) to separate cores, ensuring that critical protection functions are never interrupted.
- The hypothesis: Combining CPU Isolation with IRQ Offloading should improve trip time stability by reducing latency spikes.

TEST SETUP

Testing was conducted using a Doble F6150sv test kit to simulate fault scenarios and measure VPAC system response times.

The test system included:

— Dell PowerEdge XR12 Server (Intel Xeon Gold 5318N, 24 cores, 256 GB RAM).

— ESXi 8 Hypervisor, running multiple VMs with Photon Real-Time OS v5.0.

— Process Bus Network: 1 Gbps Layer 2 Ethernet switch.

— PTP Grandmaster Clock for time synchronization.

— Doble F6150sv Test Kit for generating IEC 61850-9-2 LE Sampled Values.

— Linux Packet Capture Machine for monitoring GOOSE messaging latency.

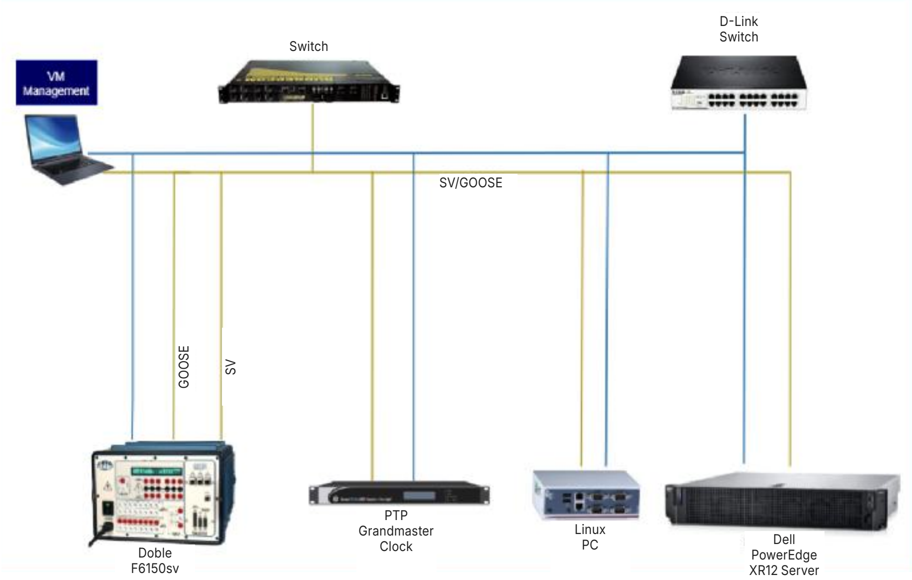

Figure 5: Test set up for near Real-Time Performance tests

VIRTUAL MACHINE CONFIGURATION

PARAMETER	CONFIGURATION
OPERATING SYSTEM	Photon RTOS v5.0, Kernel 6.1.10-6.ph5-rt
MEMORY ALLOCATION	2 GB per VM Reception (PTP Sync Enabled)
CPU CONFIGURATION	VM Pinning, HT On/Off
DISK STORAGE	25 GB per VM
NETWORK 1	VM Management Interface
NETWORK 2	Sampled Value

The hypervisor was tuned for near real-time performance by:

- Disabling CPU C-States to prevent CPUs from entering low-power modes.
- Enabling Direct I/O Virtualization (VT-d) for efficient interrupt handling.
- Setting CPU Power Management to “Performance Mode.”

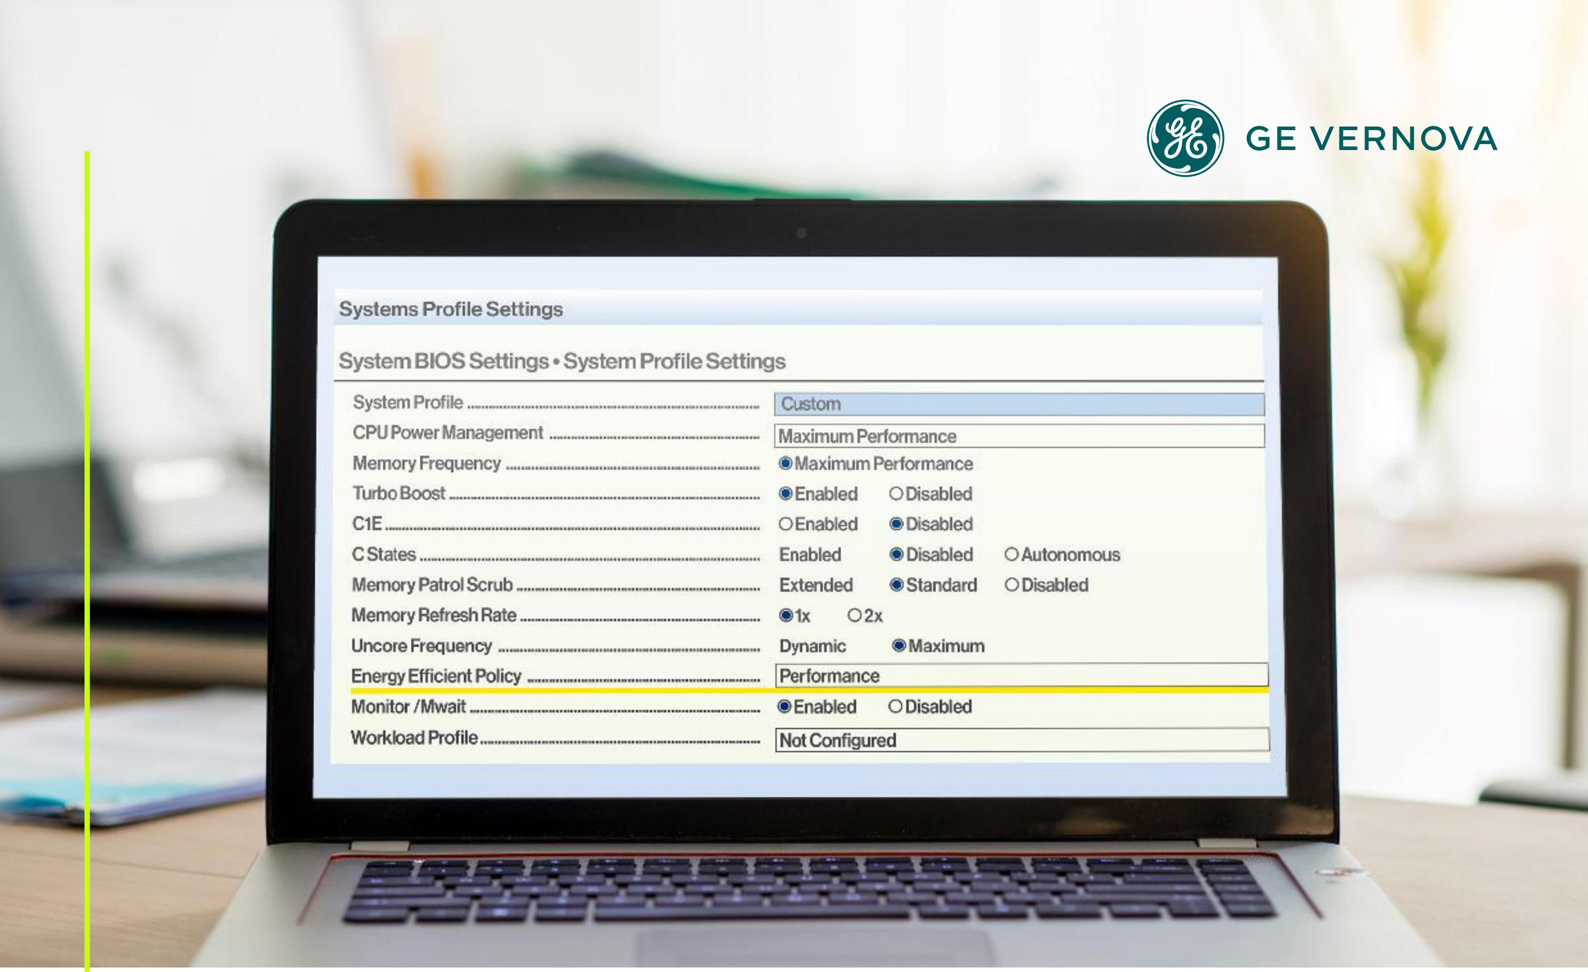

Figure 6: BIOS settings applied for near Real-Time Performance tests

CYCLIC TEST: MEASURING CPU RESPONSE TIME STABILITY

A Cyclic Test was conducted to validate the near real-time performance of the virtualized environment. This test measured the delay between a thread's scheduled execution time and its actual execution. Since VPAC systems must operate with consistent timing, it was crucial to verify whether the system could maintain deterministic response times under heavy computational loads.

The Stress-NG utility applied artificial loads to push CPU usage to 90%, 95%, and 100%. The test ran for 10 minutes, during which the system recorded timing deviations. Even at full CPU load, the worst-case jitter was only 47 microseconds, confirming that the VPAC system handled near real-time scheduling with minimal variation.

This result demonstrated that VPAC applications can operate reliably under high processing demands without significant timing delays, which is critical for ensuring fast and predictable protection response times.

KEY FINDINGS

— The VPAC system maintained stable near real-time scheduling, even under full CPU load.

— The worst-case jitter was only 47 microseconds, proving high accuracy in scheduling.

— High CPU utilization did not degrade near real-time performance, confirming system reliability.

PARAMETER	90 % Simulated CPU Load	95 % Simulated CPU Load	100 % Simulated CPU Load
Total number of latency measurements	6000000	6000000	6000000
Min Latencies (in µs)	1	1	1
Avg Latencies (in µs)	1	1	1
Max Latencies (in µs)	38	32	47
Histogram Overflows	0	0	0
Histogram Overflow at cycle number	NIL	NIL	NIL

Figure 7: Cyclci Test latencies in microseconds

ROUND-TRIP TESTING: EVALUATING END-TO-END PROTECTION RESPONSE TIME

The Round-Trip Test measured the total response time from fault detection to GOOSE trip command issuance. This test assessed how various virtualization settings impacted latency and jitter in VPAC systems.

The test used a Doble F6150sv Test Kit to inject IEC 61850-9-2 LE Sampled Value (SV) streams into the VPAC system over a Process Bus network. Upon detecting a fault, the VPAC Protection Relay (VPR) sent a GOOSE trip message, which was received by a Linux-based test system, simulating a circuit breaker.

The total round-trip time was recorded, from fault occurrence to trip reception.

Results showed that Hyperthreading increased latency and jitter, making it less suitable for power system protection.

When disabled, the roundtrip time improved from 6.4 ms to 5.6 ms, and jitter was significantly reduced. SR-IOV performed better than vSwitch, reducing the trip time from 6.2 ms to 5.3 ms. The best performance was achieved with CPU Isolation and IRQ Offloading, lowering the trip time to 5.1 ms and nearly eliminating jitter.

These results confirm that disabling Hyperthreading, using SR-IOV, and applying CPU Isolation with IRQ Offloading optimize VPAC response times for near real-time protection.

KEY FINDINGS

- Disabling Hyperthreading improved trip times from 6.4 ms to 5.6 ms.
- SR-IOV outperformed vSwitch, reducing trip time to 5.3 ms.
- CPU Isolation + IRQ Offloading provided the most stable performance, lowering trip time to 5.1 ms and minimizing jitter.

The performance with hyperthreading off is consistently better than hyperthreading on. VSwitch has better average but higher maximum times and more jitter than with SR-IOV.

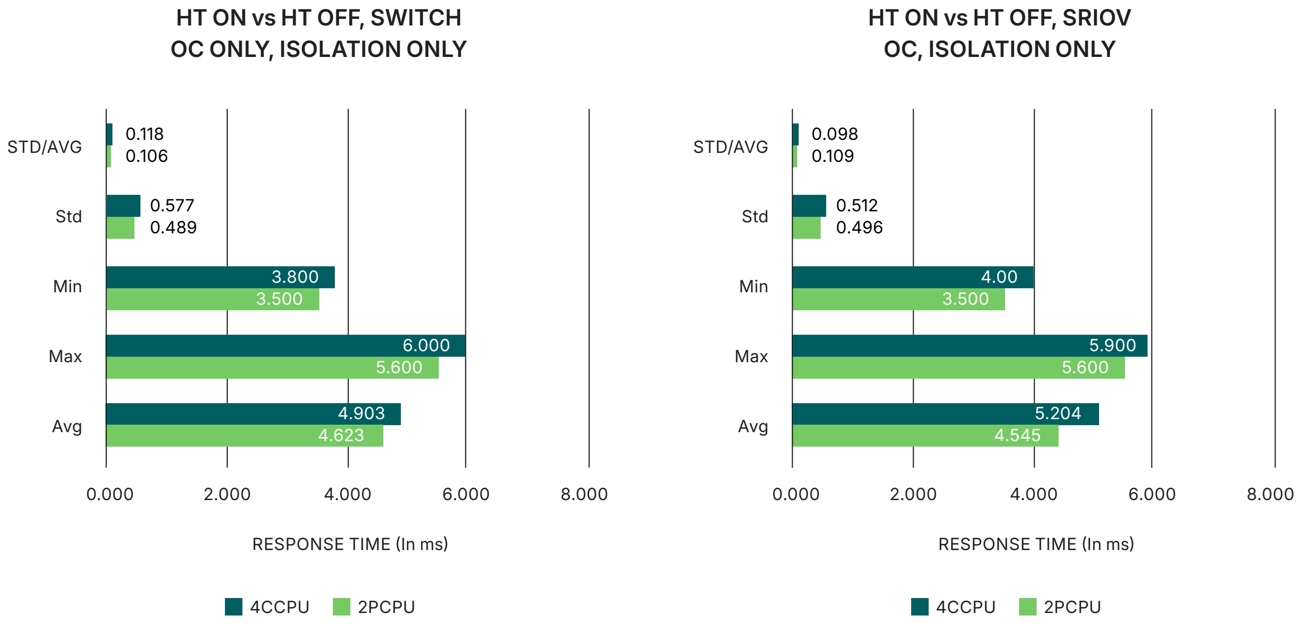

Figure 8: Round Trip test results with Hyperthreading on and off and with Virtual Switch and SR-IOV

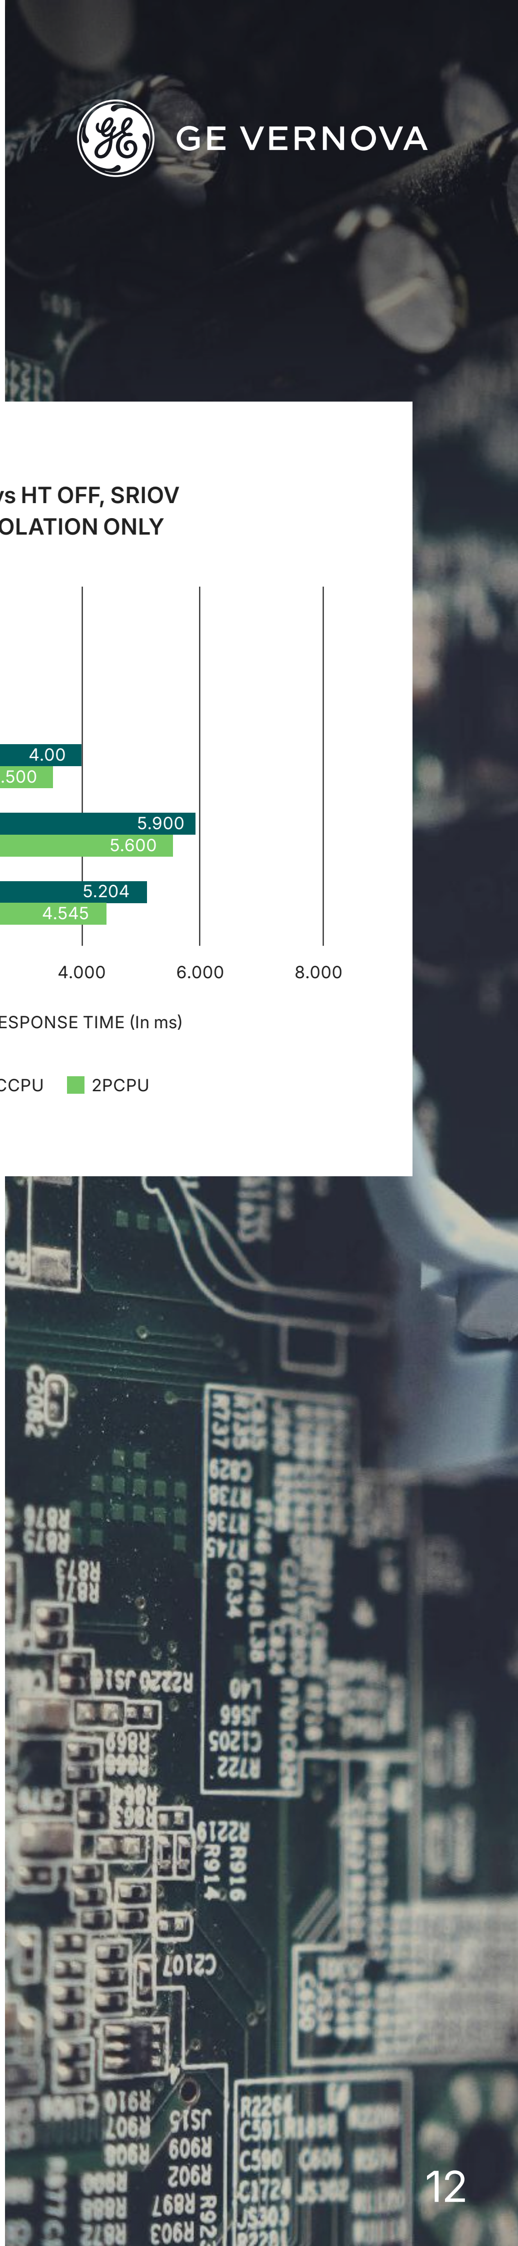

The next chart illustrates the impact of hyperthreading when combined with CPU isolation and IRQ offloading. It compares performance using 4 logical CPUs (Hyperthreading ON) vs. 2 physical CPUs (Hyperthreading OFF) to assess how hyperthreading affects protection response times.

IRQ offloading reduces some maximum latencies. SR-IOV still has less jitter and shorter maximum times than VSwitch

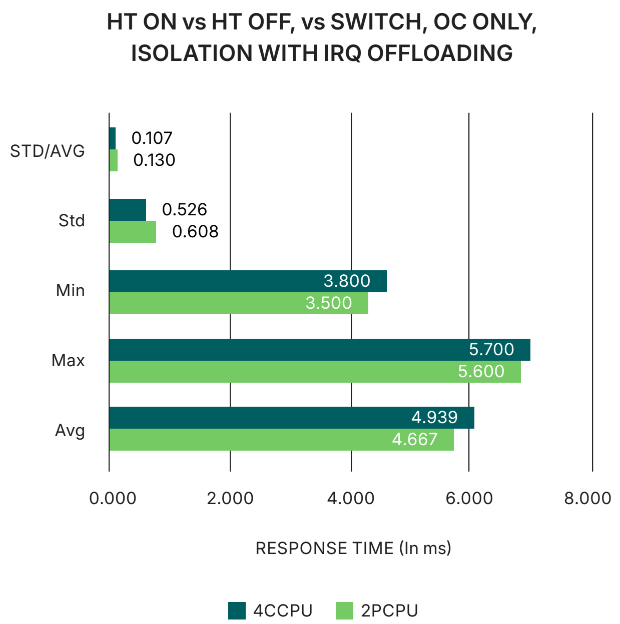

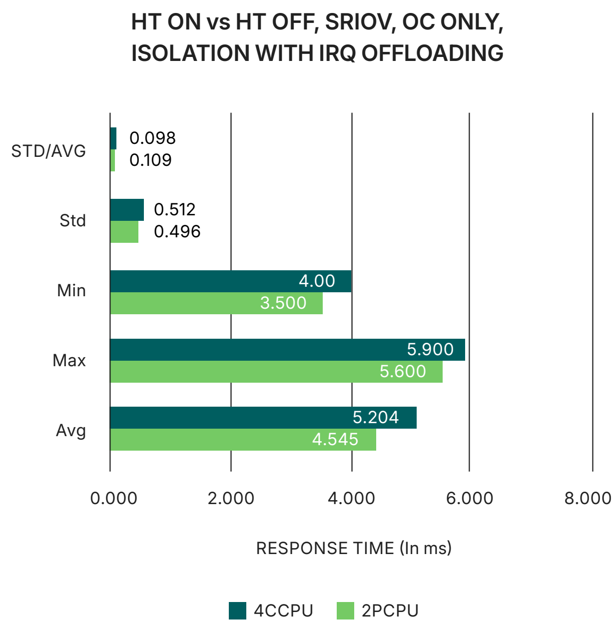

Figure 9: Round Trip test results with IRQ offloading

This chart examines the effect of running multiple protection functions within a single Virtual Machine. It compares performance with 2 and 4 physical CPUs, showing a slight increase in maximum trip times as additional protection functions (Distance and Differential Protection) were activated.

However, the Instantaneous Phase Overcurrent protection trip times remained consistent, indicating that the additional functions had minimal impact on response time.

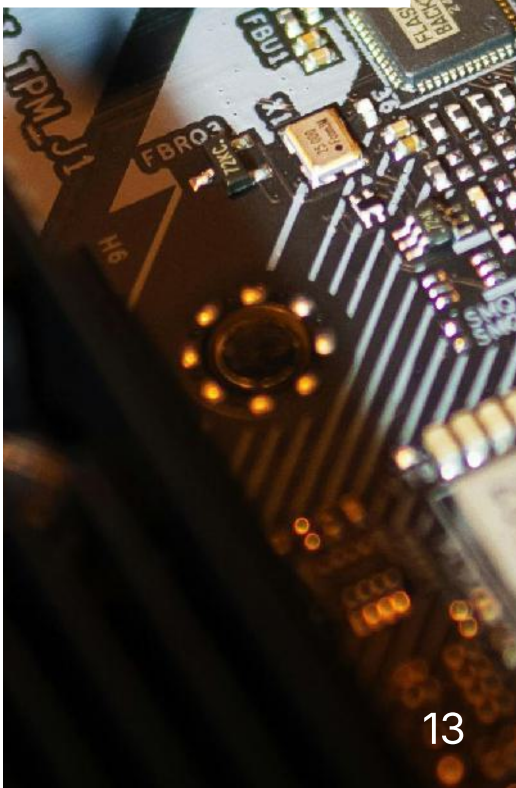

Assigning the VM statically to four physical CPU cores instead of only two produces a marginal benefit in the reduction of Max times.

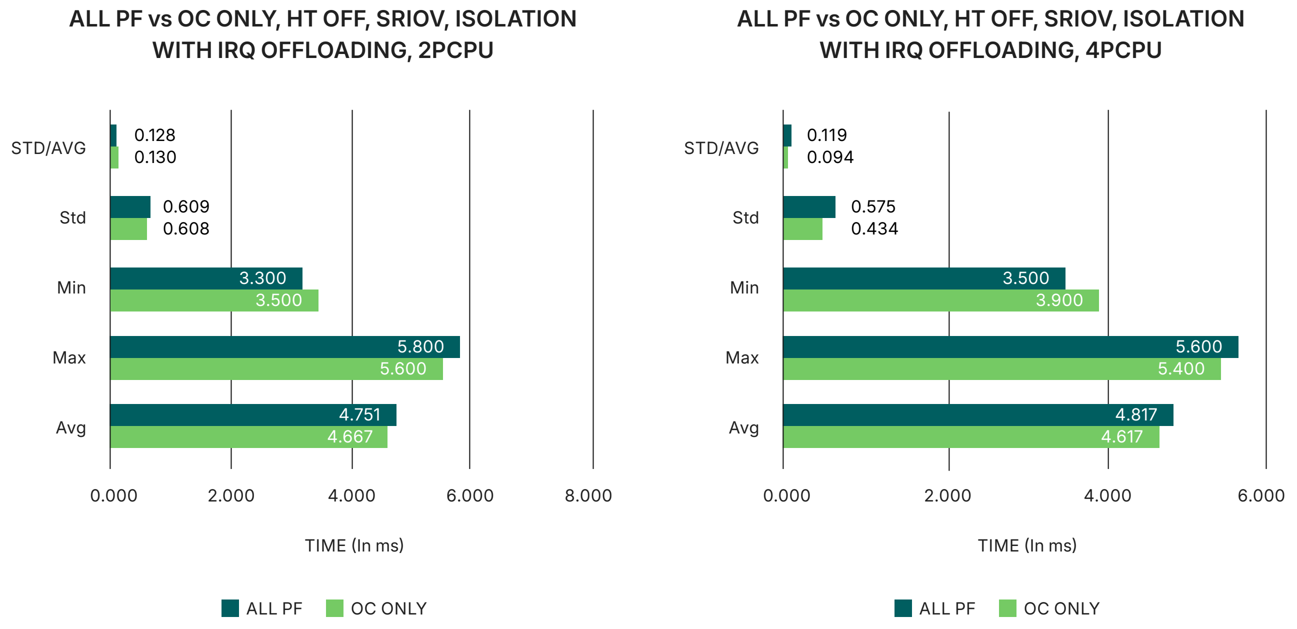

Figure 10: Round Trip test results with adding more cores

These tests used basic protection applications to effectively measure latency and jitter introduced by virtualization. Despite varying configurations, all tests resulted in total round-trip times of 6 ms or less, proving that VPAC can achieve reliable near real-time performance.

When optimized with CPU isolation, IRQ offloading, SR-IOV, CPU pinning (4 physical cores per VM), and Hyperthreading OFF, the trip time was reduced to 5.6 ms, even with all protection functions activated.

While pinning VMs to four physical cores instead of two provided only a small improvement in maximum trip times, IRQ offloading significantly improved deterministic performance. In contrast, Hyperthreading introduced jitter and occasional longer trip times, reinforcing the recommendation to disable it for near real-time protection applications.

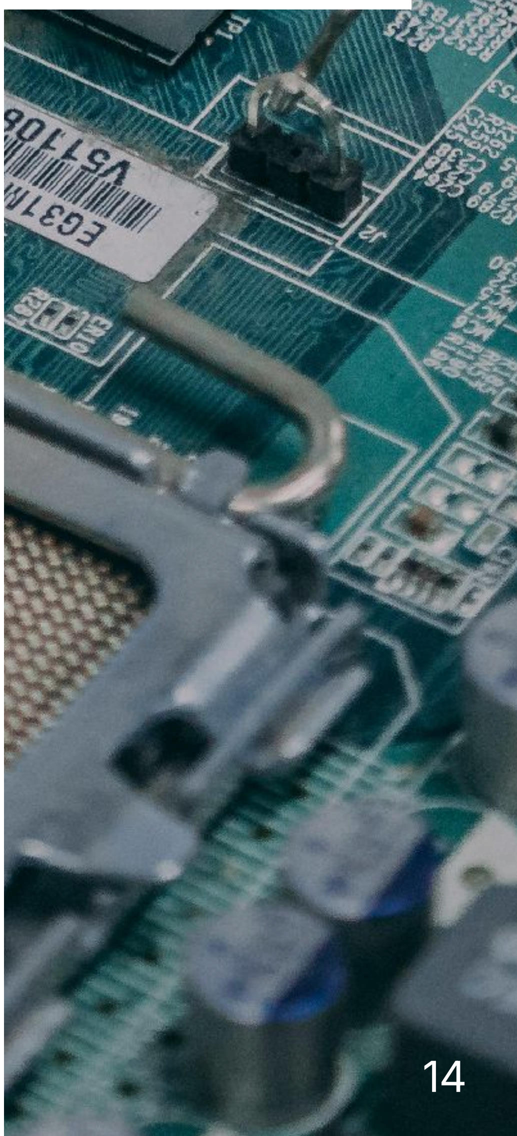

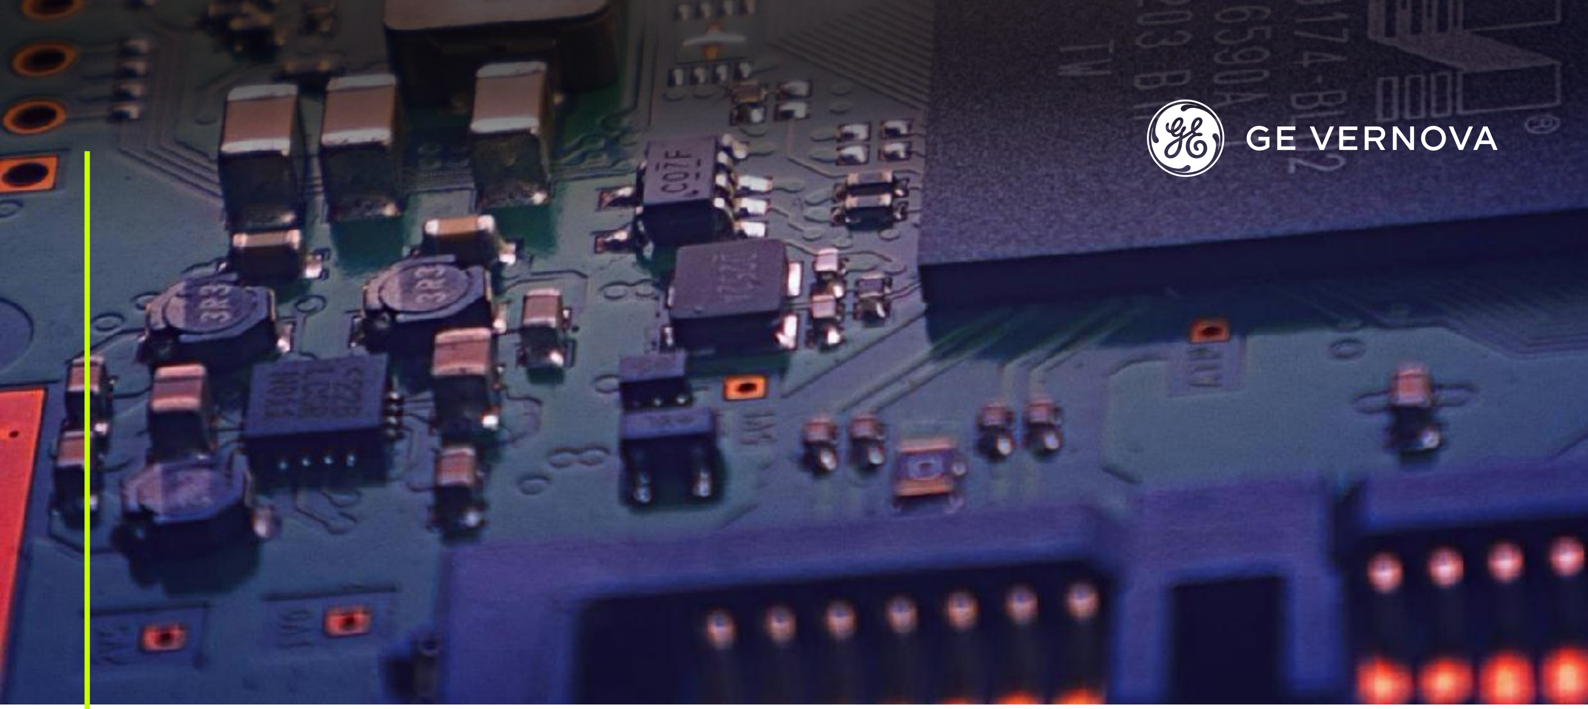

BOUNDARY CONDITION TESTING: DETERMINING VPAC SYSTEM CAPACITY

VPAC systems can run multiple virtualized protection applications on a single server, so it is critical to define hardware boundary conditions. This ensures system sizing meets operational requirements.

This testing aimed to identify the maximum number of Sampled Value (SV) streams that could be processed by a single Virtual Machine (VM) and the total system capacity across multiple VMs.

This test assessed the maximum processing capacity of the VPAC system by evaluating:

1. How many Sampled Value (SV) streams a single VM can process reliably.
2. How many Sampled Value (SV) streams a single NIC can take without packet drops and latency spikes.
3. The total system capacity across multiple VMs.

To eliminate local area network bandwidth as a limiting factor, a 1 Gbps link was used for testing. The Doble F6150sv injection test set generated a maximum of three simultaneous SV streams.

Additional SV streams were injected using TCPreplay, an open-source tool that replays captured network traffic.

A Python-based application was developed to create PCAP files containing specific SV stream parameters, including SVID, frequency, voltage, current phasors, VLAN ID, and capture duration, ensuring accurate simulation of large-scale substation environments.

Initial tests evaluated the maximum SV streams a single VM could handle. As the number of streams increased, roundtrip testing measured mean trip time, maximum delay, and jitter.

The results showed that a single Virtual Machine could reliably subscribe to 25 – 30 SV streams before performance deteriorated significantly.

Beyond this threshold, latency increased, and jitter became unstable, indicating that processing limitations had been reached.

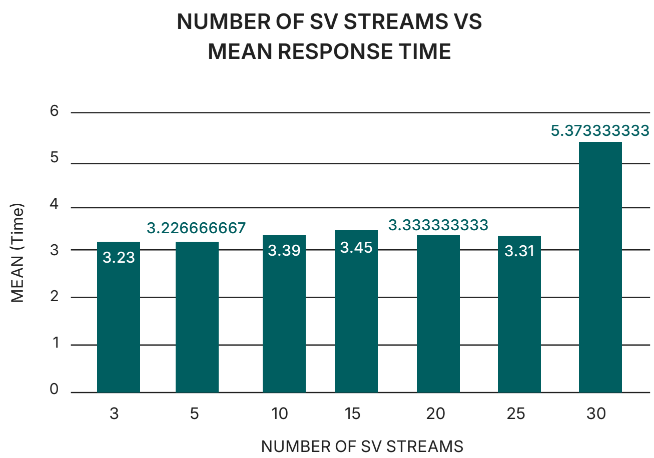

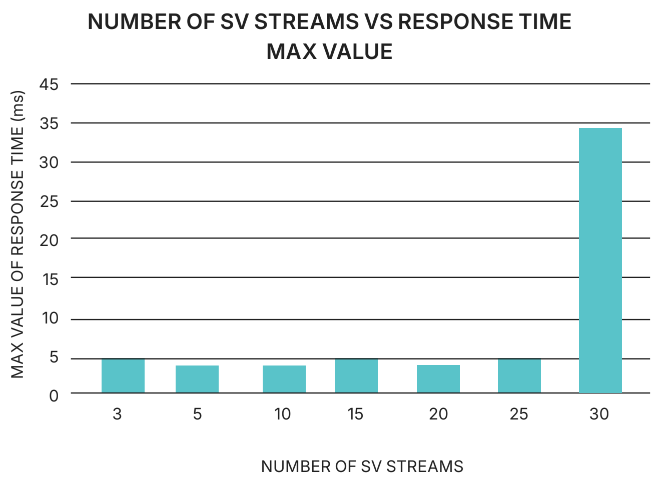

Figure 11: Virtual Machine Boundary Condition Results showing Latency and Jitter spike when subscribing to more than 25 streams

Packet Drops become increasingly problematic from 50 Sampled Value Streams onwards.

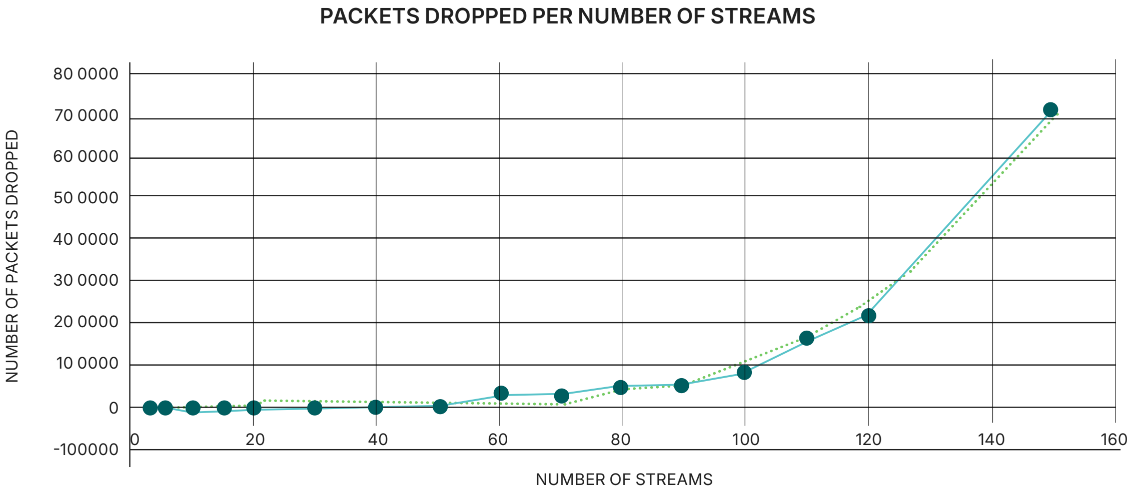

Figure 12: Boundary Condition testing of Network Interface Card

To assess system-wide scalability, additional tests examined the impact of multiple active VMs on overall performance. The test was repeated with two, four, and six VMs, each subscribing to 30 SV streams. Even at the maximum configuration of ten active VMs processing 30 streams each, no packet loss occurred, and the mean packet-to-packet latency remained stable.

PARAMETER	2 VM 30 STREAMS	6 VM 30 STREAMS	10 VM 30 STREAMS
Mean (in µs)	208.331	208.331	208.331
STD (in µs)	6.02558	6.10403	8.10993
STD/Mean	0.028923108	0.029299672	0.0389281
Number of packets dropped	0	0	0
Total packets	27301170	27303060	27306600
Packet Loss rate	0	0	0

While CPU and memory utilization remained within expected limits, further testing focused on Network Interface Card (NIC) capacity. A Packet Capture Utility running on a VM with PCI Passthrough monitored packet drop rates as the system approached its limits. Results showed that packet loss increased exponentially beyond 50 SV streams, confirming that the 1 Gbps NIC was the primary bottleneck rather than CPU or memory constraints.

> **Average packet to packet latency also increases in a similar fashion to the packet drops.**

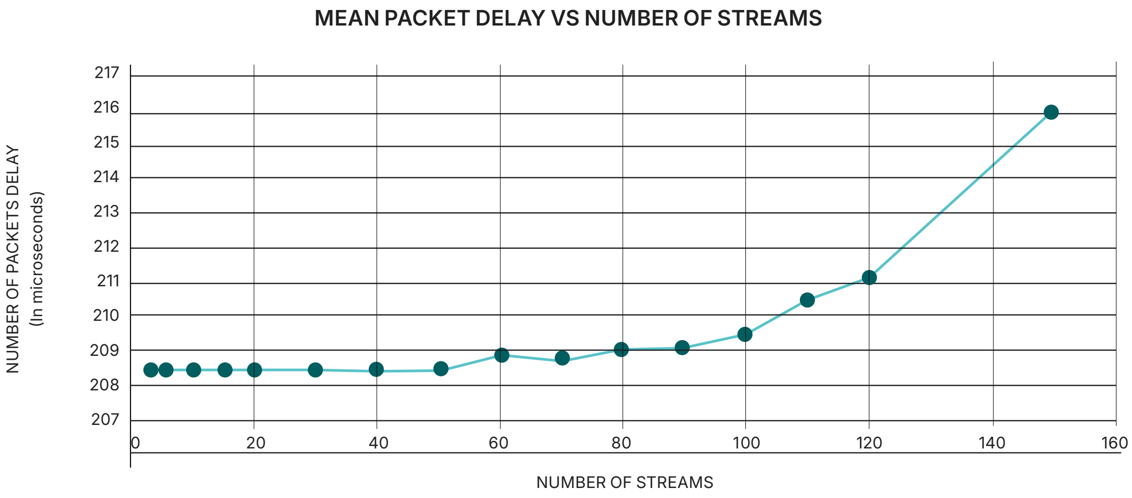

Figure 13: Packet interval delay

These findings indicate that NIC bandwidth is the key limiting factor in VPAC scalability. A server with two 1 Gbps NICs could support up to 100 Sampled Value streams, distributed across four to ten VMs, before packet loss and performance degradation became a concern. This suggests that higher-bandwidth NICs (e.g., 10 Gbps) or additional network interfaces would be necessary for further scalability while maintaining near real-time protection performance.

KEY FINDINGS

- A single Virtual Machine can process up to 25 – 30 SV streams before latency degradation occurs.
- NIC bandwidth, not CPU capacity, is the primary scalability constraint.
- Beyond 50 SV streams per 1 Gbps NIC, packet loss increases exponentially.
- Adding more Virtual Machines (up to 10 VMs) did not degrade performance, proving system scalability.
- Upgrading to 10 Gbps NICs or adding additional NICs is necessary for further scalability.

CONCLUSION

Testing confirmed that VPAC systems provide stable, deterministic performance when properly configured. Environmental stress tests showed that industrial servers remained operational despite fan failures and maximum CPU loads, proving their resilience in substation environments.

Performance tests demonstrated that disabling Hyperthreading improved trip times, SR-IOV reduced network latency, and CPU Isolation with IRQ Offloading minimized jitter under heavy workloads. Boundary condition testing identified NIC bandwidth as a key constraint, with packet loss occurring beyond 50 Sampled Value streams on a 1 Gbps NIC, highlighting the need for higher-bandwidth interfaces in future deployments.

These results validate VPAC's ability to support near real- time protection functions in a virtualized environment, with further testing needed on long-term stability, scalability, and interoperability.

Learn more

To learn more about VPAC implementation, testing methodologies, and next-generation protection solutions, visit GE Vernova or contact our experts for further insights.

GE VERNOVA