21, rue d'Artois, F-75008 PARIS
http : //www.cigre.org

CIGRE US National Committee
2022 Grid of the Future Symposium

Performance Testing of the Unified Grid Control Platform

P. MYRDA
EPRI
USA

H. FALK
OTB Consulting
USA

SUMMARY

EPRI recently completed the Unified Grid Control Platform – Utility Requirements Assessment project which developed functional requirements related to the aggregation of multivendor substation functions into a heterogeneous software environment based upon virtualization. The project identified a wide variety of virtualization methods that could be used to virtualize applications, such as automation, protection, SCADA, and asset management.

The project also identified a need to establish requirements for a virtualized protection, control, and monitoring system since these were expected to be the most challenging. Many of these virtualization methods are new to the electric power industry, and there is a need to test them to see how well they will perform the intended function.

The paper will present the proposed test plan to rigorously test the proposed hardware, software, and virtualization technologies and the associated architectures developed through the utility requirements assessment project.

This project will test the reliability, resiliency, and maintainability of select architectures agreed to by project participants. It will also develop a suitable test plan for future use by implementing utilities. Furthermore, the project will evaluate the test results for each architecture and determine the suitability for actual implementation in the field. The evaluation will encompass performance criteria of the architecture against specific performance requirements of the hardware, software, virtualization as well as environmental, physical and electrical requirements for the intended installations.

The UGCP intends to test approaches previously developed in the following business areas:

- Enabling remote upgrades to standard substation configurations.
- Leveraging the best OT and IT-based solutions to provide a more scalable and adaptable environment as part of the digital substation.
- Addressing NERC CIP compliance requirements where applicable.
- Advancing cybersecurity capabilities to mitigate evolving threats.
- Simplifying life cycle maintenance by leveraging proven
- IT-based tools to maintain the environment.

KEYWORDS

Virtualization, Architectures, Protection, Control, Monitoring, SCADA, Asset Management

pmyrda@epri.com

Background

The continued advancement of digital technologies into the electric power grid has created new challenges in maintaining reliable, secure operations. A recent project led by EPRI began the process to identify utility requirements for an advanced control platform that would seek to reduce the number and variety of digital platforms applied to the grid.

The Unified Grid Control Platform – Utility Requirements Assessment (CSP-URA) project developed functional requirements related to the aggregation of multivendor substation functions into a heterogeneous software environment based upon virtualization. The project identified a wide variety of virtualization methods that could be used to virtualize applications, such as automation, protection, SCADA, and asset management.

The project also identified a need to establish requirements for a virtualized protection, control, and monitoring system since these were expected to be the most challenging. Many of these virtualization methods are new to the electric power industry, and there is a need to test them to see how well they will perform the intended function.

The objective of this project is to rigorously test the proposed hardware, software, and virtualization technologies and the associated architectures developed through the CSP-URA. Going forward the CSP-URA will now be called the Unified Grid Control Platform (UGCP).

This project will test the reliability, resiliency, and maintainability of select architectures agreed to by project participants. It will also develop a suitable test plan for future use by implementing utilities. Furthermore, the project will evaluate the test results for each architecture and determine the suitability for actual implementation in the field. The evaluation will encompass performance criteria of the architecture against specific performance requirements of the hardware, software, virtualization as well as environmental, physical and electrical requirements for the intended installations.

The UGCP intends to test approaches previously developed in the following business areas:

- Enabling remote upgrades to standard substation configurations.
- Leveraging the best OT and IT-based solutions to provide a more scalable and adaptable environment as part of the digital substation.
- Addressing NERC CIP compliance requirements where applicable.
- Advancing cybersecurity capabilities to mitigate evolving threats.
- Simplifying life cycle maintenance by leveraging proven
- IT-based tools to maintain the environment.

Benefits

The project is expected to provide the following benefits:

- Transform substations into intelligent hubs
- Integrate legacy equipment into a grid modernization architecture
- Consolidate physical devices resulting in reduced wiring and labor costs
- Provide cyber security services for grid modernization field and substation technologies
- Decrease maintenance and deployment costs with remote upgrade capabilities
- Improve reliability using virtual computing machines and associated tools to efficiently operate and manage the required computing environments

Project Approach

For each application or use case, a series of tests will be developed to assess the performance of the architecture and its elements. The complexity of protection, automation, and control functions within a UGCP system, utilizing data collected from intelligent electronic devices, will most likely be the most challenging. Features to be tested include:

- Selected protection, automation, and control functions in a substation including interconnecting circuits between devices with applicable communication protocols.
- Applying existing standards to protection, automation, and control applications for selected circuit elements.
- System architectures for typical substation configurations.
- Evaluation of expected resiliency, reliability, and maintainability of each architecture, along with their respective testing requirements.

UGCP demonstrations are proposed in three phases. This project is only addressing Phase I. Subsequent phases will be completed as separate projects, assuming each prior phase has a successful outcome. The overall objective is the development of performance metrics and test setups that can be utilized to test different hardware and software combinations. The tasks in each phase are outlined below.

Phase I: Develop baseline/foundational test set-ups and metrics for the following:

- Different architectures of container redundancy solutions and the impact of the container and I/O loading on recovery from failure.
- PTP/IEC 61850-9-3 inaccuracies introduced via various time synchronization solutions within a hypervisor environment.
- Impact of hypervisor and container loading on Sample Values subscription and utilization.
- Impact of hypervisor and container loading on GOOSE subscription, performance, and retransmission.
- Review of the performance and supervision capability of the hardware platform/hypervisor combination being utilized in this demo phase.

Tentative Phase II and III proposed testing scope assuming Phase I is successful.

Phase II: Apply the test set-ups from Phase I to assess:

- Test alternate hardware/hypervisor combinations.
- Impact on the performance of different guest Operating Systems.
- Different cyber security architectures.
- Performance of virtual cyber zone and security conduit solutions to develop and document methods to assess risk.
- Mappings of the cyber security conduits to NERC CIP requirements.
- Simple testing of a protection like application to measure performance versus hypervisor and container loading.

Phase III: Demonstrate application deployment in the field:

- Test a simple protection algorithm that subscribes to sampled values and protects through IEC 61850 GOOSE publication.
- Demonstrate different architectures for sharing information between containers.
- Other tasks to be determined through Phases I and II.

Focus of Phase 1 Testing

Phase 1 research is to determine specific baseline metrics regarding the impact of virtualization on several different aspects. The research will involve developing metrics for the following for specific research areas:

- General performance of virtualization versus non-virtualized assets

- Impact on performance and responsiveness as more containers are added to the virtual environment.
- Impact on performance and responsiveness as container resource and execution are increased.
- Ethernet I/O performance under various loading conditions
- Performance of resiliency and redundancy under different loading conditions

The main areas of research to develop these performance metrics are:

- GOOSE performance
- Sampled Value Performance
- PTP

In addition to developing these metrics, an additional research area related to the ability to monitor, supervise, and alert on failures or warnings will be investigated.

Performance Related Research

These are general steps to develop metrics that can be evaluated:

- Obtain base-line performance on a physical platform without virtualization.
- Obtain base-line performance on a single virtualized container without tuning of the virtual environment
- Measure different levels of performance based on additional containers and container loading.
- Work with vendors to tune the virtual environment.
- Measure different levels of performance based upon the "tuned" system.
- Vary I/O loading and run the tests within the hypervisor environment that offered the best overall performance previously (e.g., tuned or non-tuned).

In order to determine base-line performance there needs to be a test application developed that can receive and log pertinent information related to the metrics that are desired to be analyzed and an external stimulus to provide the signalling (e.g., messages) required. In special instances, there needs to be external monitoring of responses.

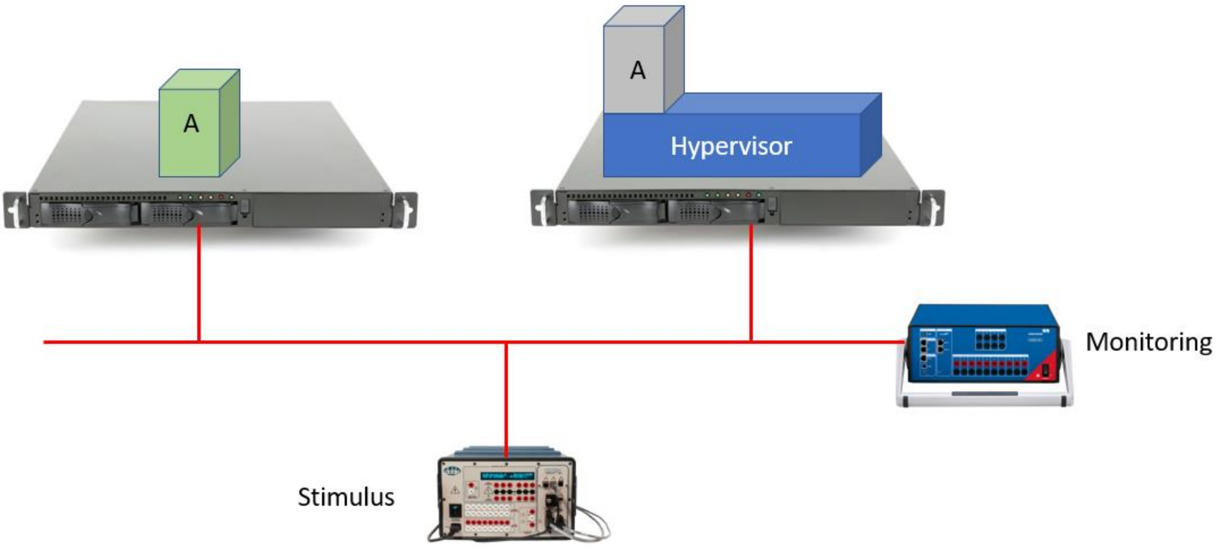

General Set-up for Baseline metric testing

The figure above shows the conceptual test setup with a physical platform being used to baseline relative performance and then a container utilized to execute the same set of tests. To have the ability to compare similar results, the number of cores in the physical system is suggested to be no more than 4 cores and 8G of memory. The container would then be configured with the same number of cores and memory. Depending upon the area being tested, there may be one or more baseline tests.

Once the baseline test is completed, the next step would be to do a loading test for an additional baseline.

The next step would be to add additional load to the baseline configuration. Loading would consist of three separate areas: network interface traffic, CPU, and memory. Other areas may need to be added based upon results. A program will need to be developed that can increase CPU load and memory utilization.

The next general test would be to add one or more containers and see what the performance of the original virtual machine (VM) will be as more containers are added to the system at various load levels. The basic philosophy of this test is to develop a curve so that the impact of context switching can be determined.

GOOSE Testing

The areas that need to be tested:

- GOOSE reception performance
- Impact of different containers subscribing to different GOOSE messages through the "same" adapter.
- GOOSE reception performance using different Ethernet adapters with one adapter per container.
- GOOSE Publication Performance
- Impact of different containers publishing different GOOSE messages through the same adapter.
- Single Container GOOSE Publication Time Allowed to Live variance
- Maximum ping/pong performance ingress-to-egress
- Maximum ping/pong performance egress-to-ingress

It is anticipated that a Hybrid Signal Analyzer for Power Utility Automation Systems will be needed as well as a source to simulate multiple GOOSE messages. A ping-pong responder may also need to be developed if one can't be located.

GOOSE Reception Performance

The purpose of this test is to determine the maximum number of GOOSE messages that can be received under various loads in a single container.

Various numbers of GOOSE messages need to be subscribed to by the application in the container. It is suggested to test subscriptions for: 1,10,100, and 500 different GOOSE messages. The 100 and 500 level will probably require the use of multiple GOOSE simulation test sets.

GOOSE Reception Performance (different containers)

The purpose of this test is to determine the maximum number of GOOSE messages that can be received under various loads to a single container as well as the inefficiencies introduced by the hypervisor.

Various numbers of GOOSE messages need to be subscribed to by the application in the container. It is suggested to test subscription for: 1,10,100, and 500 different GOOSE messages. Different GOOSE messages need to be subscribed to by different containers. All GOOSE messages should be received through the same physical Ethernet adapter.

Various numbers of GOOSE messages need to be subscribed to by the application in the container. It is suggested to test subscription for: 1,10,100, and 500 different GOOSE messages. The 100 and 500 level will probably require the use of multiple simulation test sets of the GOOSE Blaster.

It is suggested that the number of containers used for subscription be: 2,3,4, and 8 if resources allow. Please note that the results of GOOSE Reception Performance provides the results for a single container and therefore does not need to be retested as part of this test campaign.

This test will probably require the GOOSE Blaster.

<u>GOOSE Reception Performance (different containers and Ethernet adapters)</u>

The purpose of this test is to measure the impact/improvement of GOOSE reception segregation by container and adapter pair.

This is the same test as GOOSE Reception Performance for different containers except the application in the container(s) are bound to a different Ethernet Adapter for reception.

<u>GOOSE Publication Performance</u>

The purpose of this test is to determine the maximum number of GOOSE messages that can be published under various loads from a single container.

Various number of GOOSE messages need to be published. It is suggested that 1,10,100, and 500 different GOOSE messages are published with a TAL/repetition set so that there is no retransmission. Using the Hybrid Signal Analyzer for Power Utility Automation Systems, or equivalent determine and record the variance.

<u>GOOSE Publication Performance (different containers)</u>

The purpose of this test is to determine the maximum number of GOOSE messages that can be published under various loads to a single container as well as the inefficiencies introduced by the hypervisor.

Various numbers of GOOSE messages need to be published to by the application in the container. It is suggested to test publications for: 1,10,100, and 500 different GOOSE messages. Different GOOSE messages need to be published to by different containers. All GOOSE messages should be published through the same physical Ethernet adapter.

Various numbers of GOOSE messages need to be published to by the application in the container. It is suggested to test publication for: 1,10,100, and 500 different GOOSE messages.

It is suggested that the number of containers used for subscription be: 2,3,4, and 8 if resources allow. Please note that the results of GOOSE Publication Performance provides the results for a single container and therefore does not need to be retested as part of this test campaign.

<u>Single Container GOOSE Publication TAL variance</u>

The purpose of this test is to determine the variances incurred due to loading that may inhibit meeting a low TAL value (important for protection applications). Using a container-based application that is publishing a variable number of GOOSE data streams, adjusting the TAL, and having an external monitor to determine the distribution of packets that meet TAL, exceed TAL, and the amount of deviation.

It is suggested that the test be run at 1,10, and 20 publication groups with TALs of 1,2, and 5 msec.

<u>Maximum Ping/Pong Performance Ingress-to-Egress</u>

The purpose of this test is to determine the impact of loading on the responsiveness of an application executing in a container.

The concept is to have an external (e.g., non-container) based application publish different numbers of GOOSE data streams and measure the distribution of time of reception of the “pong” (e.g., response from the container).

It is suggested that the number of groups be 1,10.20,50, and 100.

Maximum Ping/Pong Performance Egress-to-Ingress

The purpose of this test is to determine the impact of loading on the responsiveness of an application executing in a container.

The concept is to have a container-based application publish different numbers of GOOSE data streams and measure the distribution of time of reception of the “pong” (e.g., response from the container).

It is suggested that the number of groups be 1,10.20,50, and 100.

SV Testing

The areas that need to be tested:

- SV reception performance
- Impact of traffic through a single Ethernet adapter
- Impact of traffic through multiple Ethernet Adapters

For this test, a subscribing application needs to be created that can measure the network latency variance of reception and calculate the distribution of the variance (e.g., jitter).

SV Reception Testing

The purpose of this test is to determine the maximum number of SV messages that can be received under various loads to a single container.

Various numbers of SV messages need to be subscribed to by the application in the container. It is suggested to test subscription for: 1,10,15, 17, and 20 different GOOSE messages.

SV Reception Performance (different containers)

The purpose of this test is to determine the maximum number of SV messages that can be received under various loads to a single container as well as the inefficiencies introduced by the hypervisor.

Various numbers of SV messages need to be subscribed to by the application in the container. It is suggested to test subscription for: 1,10,15, and 17 different SV messages. Different SV messages need to be subscribed to by different containers. All SV messages should be received through the same physical Ethernet adapter.

It is suggested that the number of containers used for subscription be: 2,3,4, and 8 if resources allow. Please note that the results of GOOSE Reception Performance provides the results for a single container and therefore does not need to be retested as part of this test campaign.

SV Reception (different Containers, different Ethernet Adapter)

The purpose of this test is to measure the impact/improvement of SV reception segregation by container and adapter pair.

This is the same test as SV Reception Performance for different containers except the application in the container(s) are bound to a different Ethernet Adapter for reception.

Precision Time Protocol (PTP) Inaccuracy Testing

The general purpose of testing is to determine the best architecture to be used to distribute time. Key metrics to be investigated are inaccuracies and jitter. This test setup will require a GPS simulator and a PTP grandmaster.

There are two deployment architectures that need to be tested:

- PTP Client within each container (would require accurate time client/server synchronizing system
- PTP Client supplied within Hypervisor

An application needs to be designed so that the inaccuracies of time synchronization within various containers can be recorded.

This application would be utilized to test:

- Hold-over time drift
- Discrepancy of time between various containers.
- Accuracy of OS PTP Clients

Critical to holdover and discrepancy testing is the ability to interrupt the PTP being emitted by the grandmaster. However, the network must not have any boundary clocks on the network to ensure that PTP is totally disrupted. There also needs to be a non-PTP time source on the network. The following diagram shows a possible network setup using both NTP and PTP being supplied by a GPS simulator.

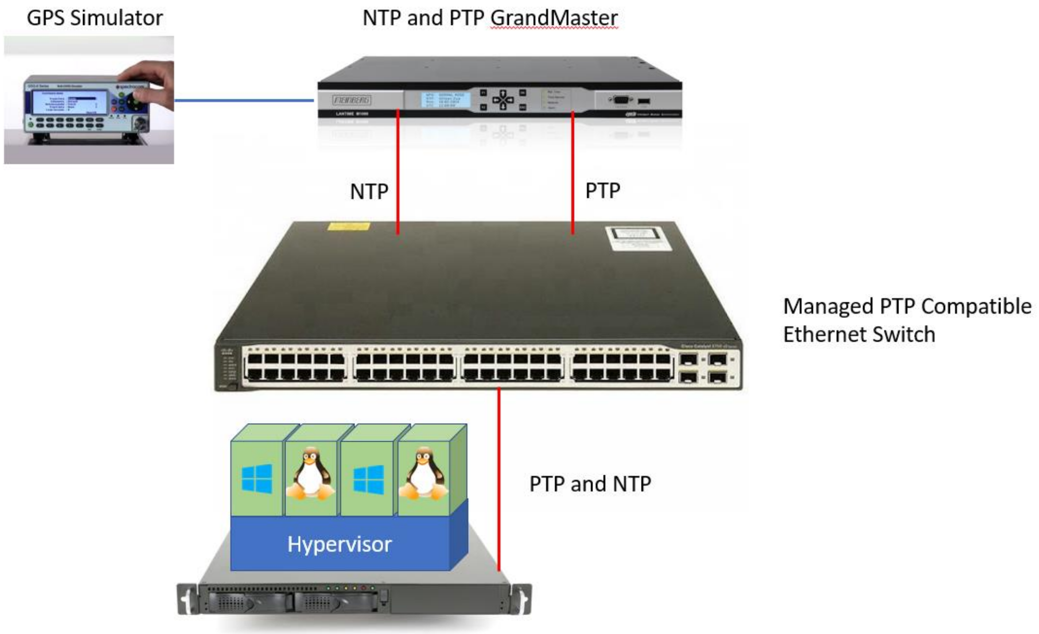

A program needs to be developed utilizing a NTP/SNTP client that DOES NOT synchronize the time in the container. This application needs to optimize the NTP offset through averaging multiple round trips of NTP exchange and thus the NTP server is required to not have a limit on the number of requests within a time period.

The NTP client would be part of a container application that utilizes the NTP time versus container time to provide information relevant to the testing (e.g., such as the delta-time and the relative time from which application started executing). The delta-time vs relative time will allow a graph to be created showing inaccuracies and drift. Since NTP can only be precise to 500 µsecs, there may be inherent errors that need to be filtered out.

Note: Hypervisor support of PTP needs to be investigated. If the Hypervisor does not support PTP, one aspect of PTP testing will not be possible.

PTP Security and Resiliency

There are two potentially interesting aspects to PTP that could be tested:

- Detection of GPS signal tampering and potential recover and enunciation.

  This would probably require the creation of something similar to a common hypervisor service to take in multiple clock sources and determine if clock/GPS tampering is occurring. An algorithm could be developed to detect this tampering and to adjust the clock inaccuracy accordingly and to enunciate the detection of the issue.

- Testing of the new IEEE1588/IEC 61850-9-3 construct for PTP security.

Redundancy Testing

The purpose of this test campaign is to determine the overall performance or redundancy mechanisms should a container fail. The various redundancy deployments that have been identified are:

- Active-Active (does not need testing)
- Active-Standby (with active input processing by the Standby)
- Active-Standby where container needs to be loaded from a SAN when the Active fails.

The concept would be to use the GOOSE ping/pong application in the various redundant platforms so that external monitoring of the time-delta between active-to-standby execution transfer can be measured externally.

For the deployment architecture of Active-Standby (with active input processing by Standby) an application extension would need to be provided to allow the standby to detect when the primary application/container is terminated. Research is needed with the Hypervisor supplier to determine if the hypervisor management system can provide this signal. This research must analyze how the signal is generated and the impact of the hypervisor management system not being present.

Comparison Against Physical Devices

Protection is one of the most time critical applications. During discussions with other potential stakeholders, it was proposed that a migration strategy towards virtualization could be done using a primary/secondary system construct where one system is based upon virtualization and the other based upon physical devices (e.g., as is done today).

Dedicated Hardware Pros: It is expected that the physical system would have better performance characteristics based upon dedicated hardware that typically have Field Programmable Gate Arrays (FPGA)s to process Sampled Values and Digital Signal Processing (DSP) chips that are dedicated to produce certain mathematical quantities.

Dedicated Hardware Cons: Physical devices have pre-allocated resources and are designed to process a predefined number of signals (e.g., Sampled Values and GOOSE) and have limited logic capabilities. Physical devices also have limited expansion capability when compared against virtual environments.

It is proposed to baseline test a selected set of physical protection devices related to:

- The performance of Sampled Values as the number of subscriptions is increased until the maximum number is reached.
- The performance of GOOSE processing as the number of GOOSE subscriptions is increased until the maximum.
- GOOSE latency (e.g., publications) as the device load is increased through GOOSE subscription and Sampled Value Subscriptions.

Performance of Software Equivalents to Physical Device Specific Purpose Integrated Circuits

Physical devices typically have FPGAs to handle Sampled Values and DSP chips to produce mathematical results based upon Sampled Values. To provide similar dedicated capabilities, cores or CPUs would need to be dedicated to performing specific functions. The easiest function to test is that of the DSP chip.

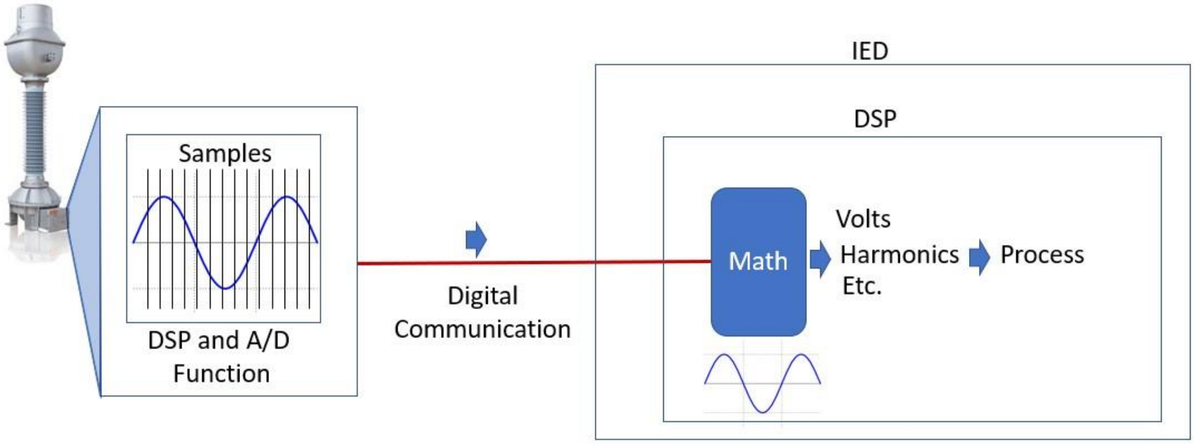

Figure provided courtesy of OTB Consulting Services.

The figure shows that even though Sampled Values is communicated digitally, it is the DSP in the physical device that produces the calculated values required for:

- Phasors
- RMS Values
- Harmonics
- Other protection related values that need to be calculated for protection.

It is proposed to develop a software equivalent of the DSP and to benchmark its performance versus various Sampled Value loading.

There are several ANSI "C" mathematical packages that appear to be appropriate for this experiment:

- https://liquidsdr.org/blog/raspberry-pi3-benchmarks/
- https://iowegian.com/scopefir/

There are others that could be researched.

Therefore, this research would be:

- Select a set of quantities that need to be produced through use of the software DSP.
- Design an experiment so that the performance of quantities produced can be measured including latency under various Sampled Value loading.
- Select a software library to be utilized to produce the selected quantities.
- Implement the DSP function to consume Sampled Values and to produce the quantities including timestamps and Quality values.
- Implement what is required to produce experiment measurements/metrics.
- Execute the experiments.

Development Activities

- Develop a program that can be configured to make sure that execution is consuming a percentage of CPU and memory.
- GOOSE Application Development:
    - Develop a GOOSE subscriber that can record number of packets received in and record the number of packets, out of sequence packets, and Time Allowed to Live (TAL) timeouts detected per a specified time interval.

- Develop a GOOSE publisher so that the impact of context switching by the hypervisor can be observed on TAL and GOOSE transmission jitter.
- Develop ping/pong subscriber/publisher and publisher/subscriber

- SV Application Development:
  - Develop a SV subscriber that can record number of packets received in and record the number of packets, out of sequence packets, and TAL timeouts detected per a specified time interval.
- Redundancy
  - An application may need to be developed to aid in the timing of failover.
- PTP – NTP client to measure inaccuracies and drift

Conclusion

This paper summarizes the test plan of the proposed hardware, software and virtualization technologies and the associated architectures developed through the Unified Grid Control Platform (UGCP). The paper identified a plan for testing the reliability, resiliency, and maintainability of select architectures agreed to by project participants. It also proposes developing a suitable test plan for future use by implementing utilities. Furthermore, the project proposes and approach to evaluate the test results for each architecture and determine the suitability for actual implementation in the field. The evaluation encompasses performance criteria of the architecture against specific performance requirements of the hardware, software, virtualization as well as environmental, physical and electrical requirements for the intended installations.