PNNL-27557

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle Since 1965

Assessment of Existing Synchrophasor Networks

Final

April 2018

JD Taft

Prepared for the U.S. Department of Energy
under Contract DE-AC05-76RL01830

PACIFIC NORTHWEST NATIONAL LABORATORY
operated by
BATTELLE
for the
UNITED STATES DEPARTMENT OF ENERGY
under Contract DE-AC05-76RL01830

Printed in the United States of America

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information,
P.O. Box 62, Oak Ridge, TN 37831-0062;
ph: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Available to the public from the National Technical Information Service
5301 Shawnee Rd., Alexandria, VA 22312
ph: (800) 553-NTIS (6847)
email: orders@ntis.gov <http://www.ntis.gov/about/form.aspx>
Online ordering: http://www.ntis.gov


This document was printed on recycled paper.
(8/2010)

Assessment of Existing Synchrophasor Networks

Version 0.5

JD Taft[1]

April 2018

Prepared for
the U.S. Department of Energy
under Contract DE-AC05-76RL01830

Pacific Northwest National Laboratory
Richland, Washington 99352

[1] Chief Architect for Electric Grid Transformation, Pacific Northwest National Laboratory

Acronyms and Abbreviations

Apps	Applications
ARRA	American Recovery and Reinvestment Act
DFR	Digital Fault Recorder
DOE	Department of Energy
EMS	Energy Management System
GDOI	Group Domain of Interpretation
GETVPN	Group Encrypted Transport VPN
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
IRIG	Inter-Range Instrument Group
ISO	Independent System Operator
MPLS	Multiple Protocol Label Switching
MUX	Multiplexer
NASPI	North American Synchrophasor Initiative
PDC	Phasor Data Concentrator
PMU	Phasor Measurement Unit
PTP	Precision Time Protocol (also known as IEEE 1588)
QoS	Quality of Service
RAS	Remedial Action Scheme
RC	Reliability Coordinator
RFC	Request for Comments
RTO	Regional Transmission Operator
SCADA	Supervisory Control and Data Acquisition
SGIG	Smart Grid Investment Grant
SIPS	System Integrity Protection Scheme
SLA	Service-Level Agreement
SNMP	Simple Network Management Protocol
SONET	Synchronous Optical Network
TCP	Transmission Control Protocol
TO	Transmission Operator
UDP	User Datagram Protocol
VPN	Virtual Private Network
WAN	Wide Area Network

Contents

Figures

1.0 Background and Purpose

In 2007-09, DOE worked with NASPI to fund the development of an architecture framework for synchrophasor data networks. That "NASPInet" framework has been the prevailing guidance since that time, and various NASPInet elements were incorporated and tested in several of the SGIG synchrophasor projects. But the NASPInet framework is now outdated; synchrophasor technology has evolved; data volumes are increasing exponentially, and networking technology changes markedly every year. Many of the NASPInet design concepts remain useful, but implementations have fallen short; the NASPInet architecture must be sharpened and projected forward in terms of technology capability to enable the design and deployment of the next generation of synchrophasor data networks. The NASPInet 2.0 project is intended to update the NASPInet framework to address the new technology opportunities and needs for synchrophasor data networks. The main objectives of this project are to assess existing synchrophasor data networks, evaluate current and future networking technologies and synchrophasor network needs, perform a gap analysis, and recommend a path forward and specific technology and design recommendations for the next generation of synchrophasor data networks (henceforth to be called NASPInet 2.0).

The first step in this effort is to assess existing synchrophasor data networks (SGIG and other) to identify what works and what doesn't work and what falls short relative to current and likely future synchrophasor network requirements and current and forward-looking IT and networking practices. The primary elements of this task are:

- Assemble existing documentation, including ARRA update reports, the NASPI Reliability Coordinator data quality survey, the NASPInet network survey from March 2015 and various public presentations from organizations involved in PMU network deployment and operation
- Extract basic architectural prototype representations for a range of NASPInet implementations using existing documentation
- Specify evaluation criteria
- Analyze implementations using evaluation criteria and document operational issues, structural performance limitations, and security posture (both strengths and weaknesses) of existing North American synchrophasor communications networks
- Document interoperability strengths and weaknesses of existing North American synchrophasor communications networks (both at peer levels (e.g., TO <-> TO, ISO <-> ISO) and across hierarchies (e.g., TO <-> ISO) to the extent they can be determined

This document is the preliminary report on the above analysis. A companion document will build on this analysis to offer some forward-looking use cases and a framework for NASPInet 2.0, the next generation synchrophasor data networks.

Note that this document uses many acronyms; a glossary of these acronyms is provided at the end of the report.

2.0 Analysis of NASPInet Specifications

Conceived in 2007, NASPInet was developed as a guidance document and framework[1] for the design of synchrophasor data communications networks, which then existed only as a few small, research-grade deployments. The NASPInet framework was completed in 2009 after a round of industry review. NASPInet was recognized within the NASPI community as a set of design concepts rather than as a set of mandatory specifications for new synchrophasor networks. The NASPInet basic concepts informed the design of the various networks deployed through the North American SGIG projects, and three projects tested specific elements of the NASPInet specifications. All ARRA grants in this area were encouraged to include NASPInet by awarding higher technical merit for inclusion of NASPInet.
The NASPInet specifications used for this analysis are contained in several documents:

- Data Bus Technical Specifications for North American Synchrophasor Initiative Network
- Phasor Gateway Technical Specifications for North American Synchrophasor Initiative Network

These two documents are written in a parallel fashion, with much common material, as is to be expected given the relationship of the topic areas. The Data bus specification contains 155 detailed specifications in six major categories and 83 sub-categories. The Gateway specification contains 234 detailed specifications in eight major categories and 76 sub-categories. In addition, each document contains an attachment (Attachment II) with three additional large categories of requirements that amount to general best practices for hardware, software, and implementation/test/startup/training.

2.1 Comments on the Specifications

Given the level and number of specifications, it is unlikely that initial implementations would or could address every single requirement. Further, sufficient documentation on the existing implementations is not adequate to evaluate them against all of the detailed requirements.

Some specifications are vague, inconsistent, and in a few cases, improper. For example, the fault tolerance specification calls out the use of traffic management, but a specification of this type should list requirements, not determine how those requirements must be met – that is for the designers to determine. Some specifications are vague, such as calling for the system to be highly secure but with no definition of what that means. Some specifications are rather aspirational, but lead to implementation difficulties, such as “allow concurrent use of different naming conventions.” Some of the categories such as 9.8 Equipment Enclosures made the assumption that all the hardware needed to be field hardened which is not the case and actually irrelevant to the focus of the spec. This should have been left to the actual designers implementing within a specific environment.

Given the above, it was not practical for this analysis to try to evaluate the implementations against the detailed specifications. Instead, they were used as guidelines on what to look for, after being condensed into a more manageable set. It is not the goal of this document to grade individual PMU network implementations against this set of criteria, but rather to use them to aid in identifying systemic industry issues and practices.

Figure 1 below is extracted for the NASPInet Data bus specification document and indicates the scope of NASPInet (and synchrophasor networks in general). However, to fully understand the issues of this

[1] The NASPInet framework documents were developed as guidance documents, not as a binding, rigorous set of architecture requirements for North American synchrophasor networks.

analysis, it is necessary to look beyond the dashed line box in Figure 1 to understand the various ways that PMUs, concentrators, and other components are connected. The figure essentially outlines three ways in which TO PMUs and other components can connect to a synchrophasor network and two ways in which monitoring centers can. Understanding how this has actually been done will be useful for the network structure aspect of this analysis and for insights related to the development of an updated synchrophasor network specification.

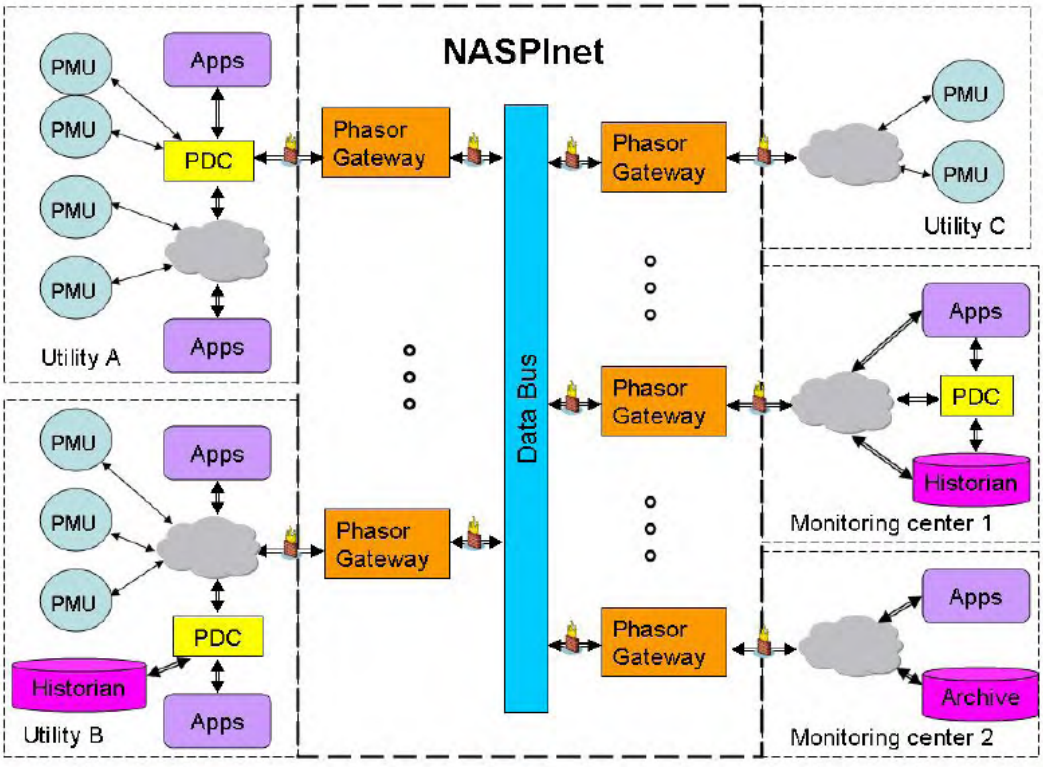

Figure 1. NASPInet Structure Model from the Specifications

Beyond the original specifications, this review adds the following criteria for evaluating the current synchrophasor network implementations:

- Network core and edge structure
- Use of advanced open standard protocols

The first addresses the issue of whether the network is inherently structured for flexibility (ability to accommodate changing functional requirements such as new applications) and scalability (the capability of a system, network, or process to handle a growing amount of work, or its potential to be enlarged in order to accommodate that growth) and the second addresses whether the network is positioned for future-proofing of investments (avoidance of stranded costs because the network could not support future requirements and therefore had to be replaced before the end of its design life) and best use of infrastructure capabilities without being limited by application component networks imposed on the underlying physical communication network, thus limiting network performance to the performance of the superimposed application network.

3.0 Process

The full set of NASPInet specifications was aggregated into the following set of higher level criteria by following the essential outlines of the two specification documents plus the additional network structure criteria described above:

- Network Structure and Communications – basic architecture and operational modes
- System Integration – mechanisms for connecting devices and systems to perform complete functions
- Sizing, Performance, and Availability - design capacity for handling expected data volumes; ability of the network to carry out data transport well in terms of network available bandwidth, latency, jitter, and packet loss; percentage of up time and reachability of input and output ports in a network
- Security – ability of the network to protect data integrity, privacy, and confidentiality; ability to control access to the network, ability to maintain device, network, and application integrity; ability to resist intrusion and to detect and mitigate intrusions when they happen
- General/Sys Admin/Ops and Functional Specification – Generic capabilities and best practices in the design, deployment, and operation of networks

Based on available documentation, several generic network structures were identified and the existing North American synchrophasor networks were grouped according to these structures. These groups were then analyzed according to the above criteria. Additional analysis was performed by inspecting the available information on current synchrophasor networks to identify gaps and potential weaknesses on a systemic basis (not on a per case basis). This analysis will form part of the basis for developing the NASPInet 2.0 specification in a later phase of this project.

4.0 Source Materials

Materials used in this analysis include:

- NASPI Working Group SGIG Update presentations
- NASPI Work Group presentations
- NASPI Work Group Success Story presentations
- NASPI Reliability Coordinator Data Quality Survey (March 2016)
- NASPI 2014 Survey of Synchrophasor System Networks – Results and Findings (July 2015)
- Various presentations from utilities

In all, data involving 47 utilities and related organizations comprising system operators, reliability coordinators, transmission owners, has been used for this analysis. In some cases, utilities have provided information on their implementations; in others, utilities were included in descriptions of implementations involving multiple organizations. Some information came from surveys that did not identify individual utilities, whereas in others they were identified. This report does not identify individual utilities or synchrophasor networks, since this analysis is more focused on overall trends and systemic issues, not specific network deployments.

5.0 Analysis of Current Synchrophasor Networks

5.1 Network Structure and Communications

A small but representative set of network structures was derived from analysis of 21 structure or architecture diagrams from the above-listed literature. Rather than try to analyze all 21 individually here, and in order to keep confidentiality, this study reviewed the 21 instances to identify common elements and approaches. While individual small variations exist as would be expected, there are primarily five network forms that have been used to date. The analysis presented here is based on the five prototype forms that span the set of solutions examined for this study. These five are shown in the figures below.

Figure 2 illustrates a regional PMU network with two variations. In the top diagram, each TO connects a PDC via WAN to the regional entity Super PDC. From there data is made available to the regional entity's applications and systems. In the bottom diagram, rather than connect the Super PDC to other regional entities, a separate set of regional PDCs is placed between the Super PDC and the PDCs belonging to the other regional entities. In the case of Variation B, PMU data must pass through a chain of as many as six PDCs, which can impose high latency costs upon the data flow before it reaches the receiving applications. PMU data must pass through the entire chain of PDCs before reach the application, thus accumulating latency that increases with the number of PDCs transited. Analysis of experience at one utility indicated that this also led to some data loss due to timeouts. Such accumulated latency is not crucial for offline analytics, but data loss is, and for future real time operations, latency even without data loss will become problematic, so a structure that locks in such latency should not be viewed as future-proofed.

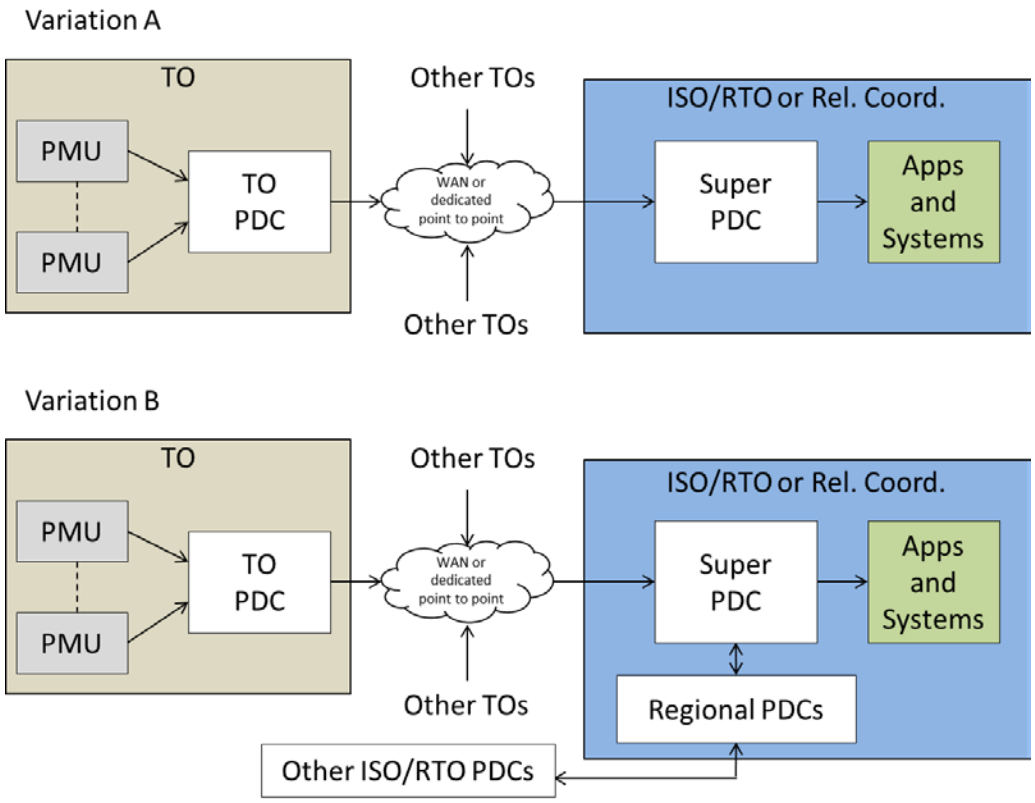

Figure 2. Regional PMU Network Model 1

Figure 3 illustrates a second approach to regional PMU networking. In this model, the main and backup operations centers are supported via separate and redundant WAN connections linking redundant sets of PMUs. This network arrangement is a typical arrangement for reliable communications for large scale data centers, even when there is only one data center site, in which case both WANs connect to the single site. This redundant structure can provide very high network security, availability and reliability needs required to support future mission-critical real-time synchrophasor applications for grid operation. It does however have the same latency issues as the structure of Figure 2.[1]

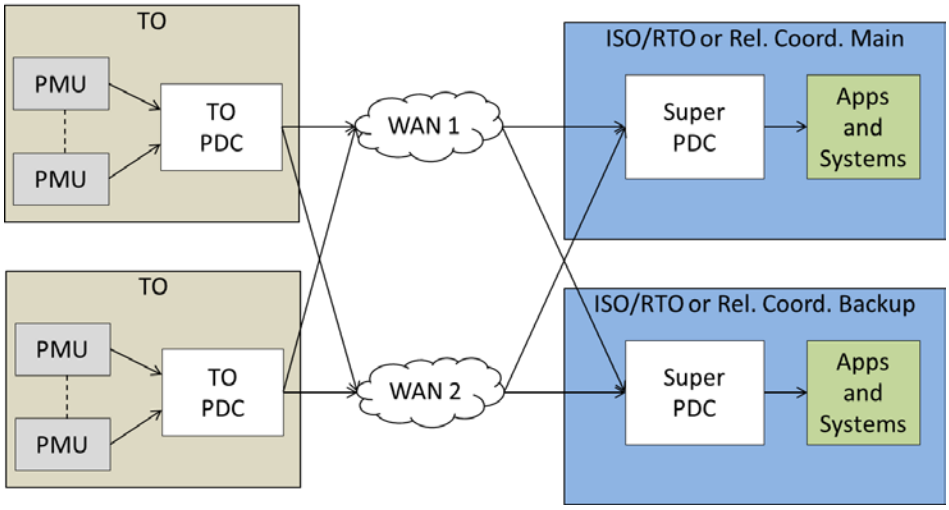

Figure 3. Regional PMU Network Model 2

Of course, Regional Model 2 may well include connections to other regional entities, as indicated in Regional Model 1 Variation B.

Figure 4 illustrates two variations on a model for PMU networking at a TO. In the upper diagram, substation PMU data is transmitted directly to a PDC housed at the TO data center.

[1] Note that some of the PDCs may have been installed before the IEEE C37.244-2013 guide for PDC requirements was released. However, this guide does not recommend against PDC stacking; in fact the implication of Figure 2 in the IEEE document is that an arbitrary number of PDCs may be stacked.

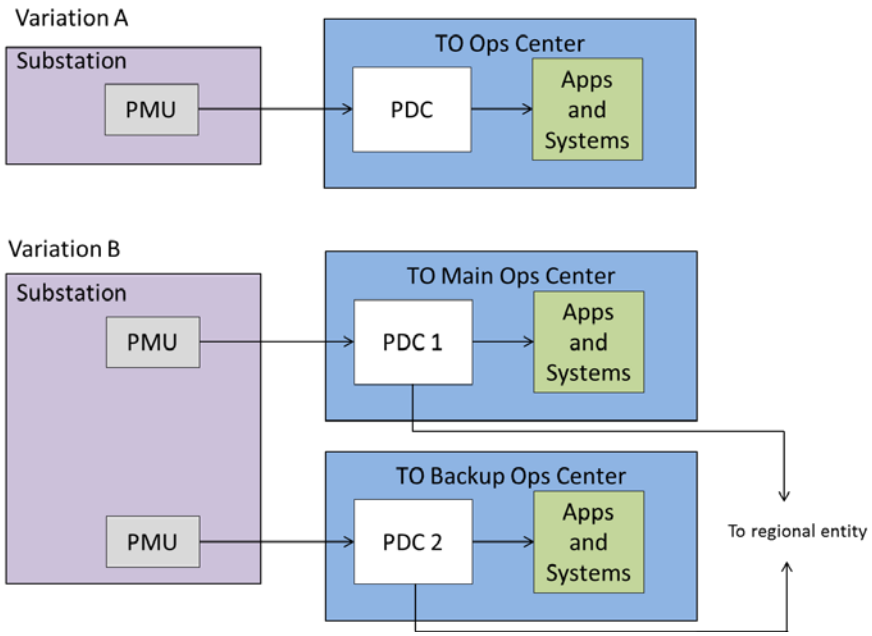

Figure 4. TO PMU Networking Model 1

In the lower diagram, redundant PMUs in the substation transmit data to separate main and backup Operations Centers. Each can send data on to a regional entity such as a system operator. Neither version incorporates a PDC at the substation. The models of Figure 4 essentially use the Utility C/Monitoring Center 1 structure from Figure 1, except that no gateways are employed and all PMU data must flow through the Monitoring Center PDC.

Figure 5 shows a common arrangement where each substation has a PDC. Data from the substation PDC is transmitted to the Operations Center Super PDC. From there it is made available to applications and optionally an EMS as needed.

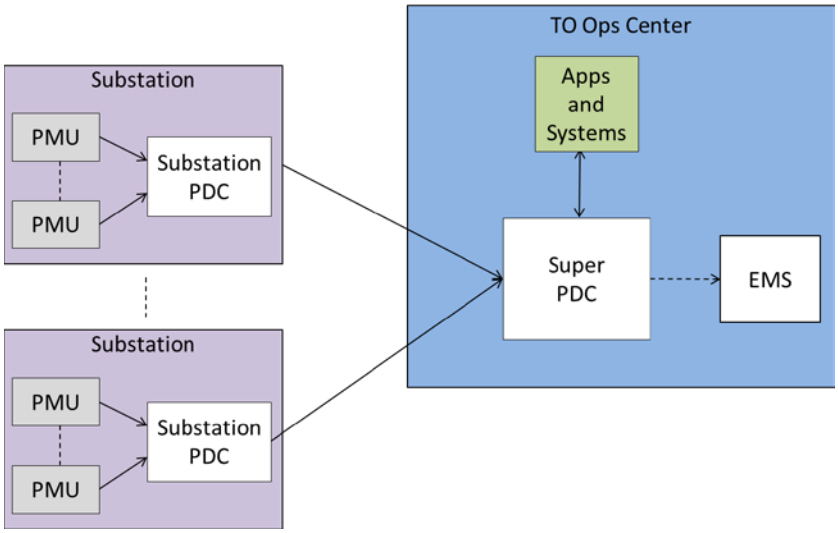

Figure 5. TO PMU Networking Model 2

The model of Figure 5 uses the Utility A/Monitoring Center 1 structure from Figure 1, except that no applications exist in the substation, no gateways are employed, and all PMU data must pass through the Monitoring Center PDC. Figure 5 differs from Figure 4 at the substation level in an important way: in Figure 4 there are no PDCs in the substations, whereas in Figure 5 the substations have PDCs.

Finally, in Figure 6, substation PMU data are multiplexed and transmitted to the Operations Center without any substation PDCs, where they are de-multiplexed. A set of PDCs in the Operations Center performs the concentrator functions that are handled in the substations in the model of Figure 5.

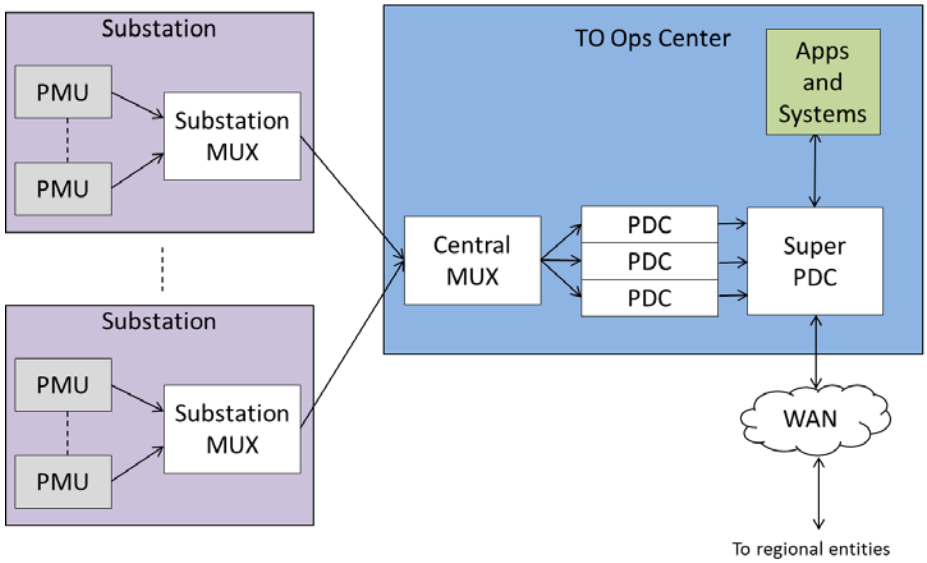

Figure 6. TO PMU Networking Model 5

The output streams from the set of PDCs are sent to a Super PDC, also located at the Operations Center. This arrangement has, in effect, virtualized the substation PDCs to the Operations Center. The model of Figure 6 uses the Utility A/Monitoring Center 1 structure from Figure 1, except that no applications exist in the substation (and the substation PDCs have been moved to the Monitoring Center), no gateways are employed, and all PMU data must pass through the Monitoring Center PDC.

All of these models make use of PDCs (either physical or functional) but the literature has very little in the way of depictions of actual gateways as shown in Figure 1. In some cases substation gateways were referenced, but in the context of supplying PDC functionality. Some PDCs were described as providing data storage for PMU data for some period of time, with historians providing longer term storage. None of the examined models illustrated multi-tier storage beyond what was just mentioned.[2]

The use of PDCs in these deployments mostly align with obvious system and organizational boundaries; although in a few cases extra PDCs are used as a means to partition data flows. In most cases PDCs were treated as single entities, with one exception depicting a High Availability local network of PDCs and a few cases (see Figure 3 and Figure 4) PDC redundancy was provided due to the use of main and backup operations centers. Most survey respondents indicated that PDC redundancy was a prime resilience measure, as was the use of redundant communication networks. More than one approach to communications redundancy appeared in the implementations, with the dual independent WAN version representing one of the stronger approaches; this is a design that is widely used for data centers. Large

[2] One of the reviewers of this document commented that a model of this type was presented in a DNMTT session in March 2017.

scale data centers are useful models for some aspects pf synchrophasor network design because the face and have dealt with the same issues of connectivity, security, scalability, and operational practices that face PMU network operators and users.

Overall, PDCs have been deployed in obvious ways, meaning that the PDCs act as physical system and organizational demarks, meaning they define essential physical and logical boundaries that align with organizational boundaries. While this is a clear way to lay out a system, it violates the core and edge structure that makes the internet so flexible[3] and leads to situations where PMU data may have to flow through a number of PDCs serially (anywhere from three to as many as seven PDCs are reported). While the present round of applications is not overly sensitive to latency (except for the data loss issue, which is of course significant where it may have occurred), this structure can present severe latency issues (far beyond network latency) and will be problematic for wide area closed loop controls and other latency-critical applications. Some survey respondents indicated PDC-related latencies greater than 10 seconds. Fundamentally, the PDC structure places a second low performance network on top of the base high performance communication network. The low performance network will always dominate performance.

Some PMU networks (typically TO PMU-to-control room networks) carry other data besides PMU data (especially DFR and SCADA data with video showing significantly as well) and most implementations make use of multiple WAN transport technologies. TCP and UDP are used about equally for transport, while TCP was the majority choice for PMU configuration and control. Early in the PMU network development work, some utilities had packet loss problems with UDP and felt that it could not be used but investigation showed that some router buffers were not configured correctly and that correcting this eliminated the UDP packet loss problem. MPLS and SONET were cited most often as transport mechanisms in the NASPI 2014 Survey, but the two thirds of survey respondents indicated that they deployed no QoS mechanisms, and only 40% of those using third party WAN providers had SLAs in place and many indicated they did not monitor latency or jitter performance per application. Some of this is due to the use of third party networks (telco, private third party, even internet) but many of the TOs provide their own networks and so should have been able to provide QoS measures and monitoring. Some survey respondents indicated that their communication links were solid and that jitter and packet loss were not a significant problem. Most PMU networks use GPS timing and most TOs use IRIG-B for clock distribution. None indicated more sophisticated approaches to clock distribution (IEEE 1588/C37.238,[4] boundary clocks, transparent clocks, GPS timing integrity monitoring with automatic switchover to Cesium atomic clock backup in case of GPS failure, etc.). Network management is generally limited to SNMP and some use of centralized network management applications, but this appears to be less than half of the implementations.

While it is understandable that initial PMU network implementations would focus more on basic functions, it will be very important going forward to address QoS and network performance monitoring and management thoroughly, and for those using communication service providers, to put SLAs in place. This is because the use of PMU networks for (future) real time decision support, and protection and

[3] While the internet is often thought of in terms of its layered protocol model, there is in fact another architectural principle that is largely responsible for the internet's excellent scalable performance characteristics: the core of the internet is kept devoid of anything except high performance data transport elements. All application elements are kept to the edge of the internet, thus avoiding the issue of having a weak application element causing wide impact on network performance for other users. In the case of synchrophasor networks, this means keeping devices like PDCs and gateways strictly at the edge of the network, and not forming a de facto "overlay" network by cascading PDCs and/or gateways.

[4] IEEE Std 1588-2008 (Revision of IEEE Std 1588-2002) - IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, IEEE, 2008, available online: https://standards.ieee.org/findstds/standard/1588-2008.html. Note that this standard is being updated in 2018.

control will entail stringent network performance requirements appropriate for these mission-critical functions.

5.2 System Integration

System integration is the connection of various components and subsystems so that the resulting overall system can deliver a specified set of capabilities and optimized value. The concept applies at any level where some form of interconnection of interaction must occur, so specifically at the electrical connection level, at the communication network level, and at the level of data exchange between devices and applications or just between applications. Integration is typically an expensive activity for utilities, but the costs are mitigated somewhat through the use of interoperability standards and through the use of various tools, including enterprise service buses and other middleware. At the communication network level, integration includes establishing physical connectivity, and configuring protocols, addressing, buffers, etc.

System integration approaches have varied by type of organization. TOs have relied less upon middleware but system operators and reliability coordinators have made use of this technology in a manner consistent with their general approach to operating environment design. Since system operators and reliability coordinators do not own extensive physical infrastructure, they have tended to outsource networking, whereas TOs have tended to operate their own networks, so that it is reasonable to expect different approaches to system integration from those two groups. [5]

5.3 Network and Data Interoperability

Most implementations have used IEEE C37.118 as the standard for PMU data exchange. As of the time of the surveys, no utilities were using IEC 61850-90-5, but it was still in formation when the designs were being done. IEC 61850-90-5 provides for use of GDOI[6] as a security measure, something that is lacking from the NASPInet specifications and from most of the existing NASPInet designs (it was referenced in one). All of the designs relied upon GPS timing and some specified IRIG-B timing distribution, but none indicated any means of distributing timing across organization boundaries, instead relying on sample data time stamping and PDC alignment.[7] This and the extensive use of stacked PDCs are issues to be re-examined going forward for future applications that include wide-area closed loop protection and control and SIPS/RAS; future implementations will need controlled access flow of PMU data across organizational boundaries without PDCs at every boundary.

Interface between TOs is essentially PDC-to-PDC for PMU streams. Not all implementations show this however, so it is not clear if that is implemented in all cases, but several designs included this capability. Interfaces between TOs and system operators or reliability coordinators are included widely, of course; these are also via PDCs for the PMU data streams. Likewise, system operators and reliability coordinators provided similar capabilities for exchange among themselves.[8]

---

[5] Integration to EMS had been minimal as of the time of the report, so not much detail was available on this topic.

[6] GDOI (Group Domain of Interpretation, IETF RFC 6407): a cryptographic protocol for group key management called out in IEC 61850-90-5 for use in PMU network security implementations.

[7] Cross-organizational time distribution would require organizational agreements. For situations where country boundaries must be crossed, this would also require agreement on whether the time basis is UTC or TAI (International Atomic Time). UTC and TAI differ by 37 seconds.

[8] The NASPInet concept of Signals was intended to facilitate flexible data sharing. In addition, a registry was intended to facilitate role-based access control for signal subscription.

While some organizations have implemented sophisticated tools for retrieval of PMU network event data within their organizations, little information is presently available on data exchange mechanisms, especially as relates to historical data, or how one organization arranges to obtain data flows from another (meaning the technical methods, not the data sharing agreements). The NASPInet 2.0 specification should provide more guidance in these areas.

5.4 Sizing, Performance, and Availability

The survey results indicate that PMU networks are performing well in terms of sample rates (30 and 60 samples per second are routine, and some sites report even high rates). Some survey respondents indicated that most of the communication link problems occurred between substations and TO control centers, whereas TO-system operator links were evaluated as performing well. Most substations did not have redundant communication paths. While jitter and packet loss are not widely monitored, these do not seem to be serious issues in any systemic way (local specific issues are always possible). The one area of concern for the future in terms of performance is latency, and this is largely a systemic (architectural) issue. While gateways do not seem to have been used as much, PDCs and PDC stacking are quite common (although a few survey respondents indicated no PDCs until the PMU data reached the control room). The latencies inherent in such structures will be problematic for future real time operational applications. This is not an issue that can be addressed though faster routers or bigger application servers-it is inherent in the nature of the PDC function coupled with the stacked PDC structure. The structural issue has come about by violating the core-and-edge principle of network structure.

5.5 Security

The available view of security measures in existing North American synchrophasor networks is sketchy. While various measures are in place, including link level encryption, VPNs, access control lists, port disablement and sticky MAC, firewalls, and intrusion protection systems, it appears that there is no comprehensive or consistent approach to network level security for PMU networks.[9] Firewalls are often depicted in system diagrams and one presentation listed link level (router-to-router) encryption, but use of device authentication, network level signature analysis, network segmentation, or supply chain security management was not disclosed. Some implementations described data integrity methods, but these were aimed at detecting faulty data, not tampering or intrusion. One presentation indicated use of the GetVPN commercial implementation of GDOI. The list of measures above is far from complete and any given organization appears to be using only one or a few of them.

The apparent weakness of cyber security approached in PMU networks may be due to lack of information from the surveys, but also may be taken in light of the fact that many organizations *do not presently classify PMU networks as critical cyber asset*s. This classification issue must change going forward. For forward-looking designs, the potential interplay with NERC CIP guidelines should be considered.
In practice, network level security should be viewed as a multi-layer, multi-measure framework based on four pillars:

- Access control
- Data integrity, privacy, and confidentiality
- Intrusion resistance, detection, and mitigation

[9] It should be noted that there is a gap between the original NASPInet security guidance and an" ideal" security posture, so the comments here should be taken in the light of considering how to improve the next round of guidance.

- Device and platform integrity

Within these pillars, there are a large number of available tools, devices, and technologies that are commonly available to be built into security frameworks and solutions (see Appendix A for a list of typical measures). In addition, there are many processes that must be applied as well. There is very little indication that any of the organizations took a comprehensive view of PMU data security. The survey data had no information from which to evaluate either physical security or people/process issues.

For the NASPInet 2.0 specification, considerable attention should be paid to laying out comprehensive guidance for PMU network cyber security, using something like the four pillars above. It will be important for the industry to adopt a rigorous approach to cyber security for synchrophasor networks.

5.6 General/Sys Admin/Ops and Functional Specs

While the NASPInet specifications have extensive sections on these topics and a very long and detailed list of specific functions, not much information was available on how these were implemented in the PMU projects.

6.0 Summary of Results

Current North American synchrophasor networks have dealt with a great many complexities and constraints in order to get systems in operation to provide proof of concept and operational experience with PMUs. Not every specification of the NASPInet documents could be implemented; some were ambiguous, some were ambitious, and some were even obsolete well before the projects could be completed.

The NASPInet specifications provided very little in the way of architectural guidance and this is ultimately the biggest obstacle to the use of the NASPInet specifications. This is not to blame the creators of the specifications and certainly not the project implementers; creation of these new networks involved a great many unknowns and legacy constraints and all while the technology was in flux. Given all this, the creation of the specifications and the development of the current synchrophasor networks have been remarkable.

But it was recognized as early as 2011 that the NASPInet specifications and architecture shown in Figure 1 were becoming outdated. It is clear that Figure 1 represents a conceptual view of the NASPInet data bus and gateway concept but this in itself is not enough to fully support the extensive specifications and aid in preventing significant gaps that the utilities were left to address. Some of the gaps in the specification include:

- Lack of clear guidance about network architecture, with more focus on administrative functions and data access use cases than on network structure, but structure sets the essential bounds on system performance and so must come first
- Strong network design principles and best practices are not well reflected in the specifications as guidance[1]
- Vague and incomplete specification of security architecture left the issue to individual projects to resolve
- Limited recognition of the role of network management in PMU network performance and security
- Limited guidance on network level security

In addition, much experience has been gained through the PMU implementation projects; also network technology has continued to evolve during this period of time. The technology evolution has resulted improvements in network devices, network security, network management tools, and the emergence of Software Defined Networking from concept to the beginnings of a practical tool; whereas the experience gained has shed new light on NASPInet requirements and practical constraints.

[1] Note that the original NASPInet document intentionally did not include network design principles, on the presumption that this would encroach on each utility's flexibility in choosing an implementation.

7.0 Final Comments

The original NASPInet Specifications were necessary and useful to help jump start the large scale deployment of PMUs. They were published in May of 2009, and so are now more than eight years old. Eight years is not considered long in the electric utility world, but it is a lifetime (or two) in the communication networking world.

A new NASPInet 2.0 specification is needed. The new specification should address the architectural gaps discussed above and without attempting to prescribe design, should reflect recent field experience while accommodating many new use cases and new networking technologies and practices. A re-examination of the gateway concept seems in order, as does the practice of cascading PDCs. The NASPInet 2.0 specification should pay more attention to:

- network logical level structure
- data management architecture including storage tiers and sharing of tiered stored data;
- comprehensive cyber security
- distributed registry structure
- performance monitoring/performance management measures

Timing distribution should be addressed in a more comprehensive manner, as should the use of advanced but off-the-shelf open standard networking protocols.

A new specification will greatly aid the inevitable second wave of PMU network development and deployment, which will need to support real time analysis for operator decision support and closed loop protection and control functions; these will require superb PMU network performance, availability, reliability, and security. As with all decentralized infrastructures, the communication networks are absolutely crucial to the success of the control systems and so having a set of NASPInet 2.0 specifications backed by strong architectural views and real industry experience is a crucial risk management and success factor for grid modernization.

8.0 Resource References

NASPI Synchrophasor Starter Kit: https://www.naspi.org/sites/default/files/reference_documents/4.pdf

NASPI 2014 Survey of Synchrophasor System Networks Results and Findings: https://www.naspi.org/sites/default/files/reference_documents/8.pdf

Synchrophasor Technology Glossary: https://www.naspi.org/sites/default/files/reference_documents/58.pdf

Synchrophasor Applications in Transmission Systems: https://www.smartgrid.gov/recovery_act/program_impacts/applications_synchrophasor_technology.html

PMU Networking with IP Multicast: https://www.cisco.com/c/en/us/products/collateral/routers/2000-series-connected-grid-routers/whitepaper_c11-697665.html

Building an architecture based on IP-Multicast for large phasor measurement unit (PMU) networks: http://ieeexplore.ieee.org/document/6497794/

Appendix A

Typical Network Security Measures

Appendix A

Typical Network Security Measures

The following is a list of typical network security measures, technologies, and methods. It is not exhaustive and does not address people/process issues.

- Crypto: link layer, group, and application layer; GDOI, as it has been incorporated into IEC 61850-90-5 specifically for PMU network encryption
- RBAC (RADIUS and TACACS; AAA; NAC)
- Mutual authentication; EAP and media independent identity protocols
- Posture assessment
- X.509, secure key generation and management, scalable key management (DMVPN, GETVPN for example)
- SIEM, firewalls
- IPS, including SCADA IPS signatures
- Containment: Virtualization and Segmentation (VRF – virtual routing and forwarding, MPLS VPN and VLAN); data separation
- Tamper resistant device design, digitally signed firmware images, firmware/patch authentication and integrity verification
- Digitally signed commands
- Rate limiting for DOS attacks
- Wire speed behavioral security enforcement
- Packet tamper detection, replay resistance
- SUDI 802.1AR (secure device identity)
- Access control: VLANs, ports
- Storm detection and traffic flow control: traffic policing and port blocking
- ARP inspection; DHCP snooping
- Honey pots/honey nets/sinkholes
- Unicast reverse path forwarding (IP address spoofing prevention)
- Hierarchical QoS
- Security policy managers
- MAC layer monitoring
- Control plane protection (coarse packet classification, VRF-aware control plane policing)
- Secure code development and code hardening (against buffer overflow, self-modification; remove unnecessary protocols)

- Six wall physical security for devices and systems; access detection and mitigation (i.e. port shutdown)
- Manufacturing supply chain security management
- Data quality as tamper detection
- Anti-counterfeit measures

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by **Battelle** *Since 1965*

U.S. DEPARTMENT OF
ENERGY

902 Battelle Boulevard
P.O. Box 999
Richland, WA 99352
1-888-375-PNNL (7665)

www.pnnl.gov