

NASPInet 2.0 Architecture Guidance

Version 1.13 **DRAFT**
Oct 22, 2018

Contents

- 1.0 Background and Purpose of This Document 1
- 2.0 Scope 2
- 3.0 Purposes of PMU Networks 2
- 4.0 Overview of Applications 3
- 5.0 Emerging Uses for PMUs 3
- 6.0 Other Relevant Emerging Trends 4
- 7.0 Implications and Systemic Issues 4
 - 7.1 Systemic Issues 5
- 8.0 Key Definitions and Concepts 7
- 9.0 Foundational Principles 8
 - 9.1 Leading Network and System Practices 8
 - 9.2 Interoperability and International Open Standards 10
 - 9.3 Synchronization 11
 - 9.4 Function Allocation 11
 - 9.5 Some Implications of These Principles 12
- 10.0 Core Considerations 13
 - 10.1 Attributes of Capabilities and Functionality 13
 - 10.2 Cybersecurity 14
 - 10.3 Performance Characteristics 14
 - 10.4 Functional Flexibility 15
- 11.0 Synchrophasor Problem Domain Reference Model 16
 - 11.1 Regulatory Structure 16
 - 11.2 Industry Structure 17
 - 11.3 Market Structure 18
 - 11.4 Relevant Entities 19
 - 11.4.1 Planning/Operating Functions 19
 - 11.4.2 Reliability Service Functions 20
 - 11.4.3 Standards and Compliance Functions 20
 - 11.4.4 Others 20
 - 11.5 Bulk Energy Systems 20
 - 11.6 Distribution Systems 21

11.7	Key Constraints and Barriers	21
12.0	NASPInet 2.0 Architecture Principles, Objectives, Capabilities, and Functions	22
12.1	Architectural Principles for NASPInet 2.0.....	22
12.2	Objective of a NASPInet 2.0 Implementation.....	23
12.3	Capabilities of a NASPInet 2.0 Implementation	23
12.4	Function Classes for NASPInet 2.0 Implementations	24
13.0	NASPInet 2.0 Components and Structures.....	25
13.1	Key Component Classes for NASPInet 2.0 Architectures	25
13.2	Structures for NASPInet 2.0 Architectures	27
13.2.1	Observability Platform Structure.....	27
13.2.2	Platform Functional Mappings	32
13.2.3	Platform Mapping Summary Comments.....	36
13.2.4	Streaming Data Flows.....	37
13.2.5	Streaming Data Flow Summary Comments	40
13.2.6	Communication Network Structures	40
13.2.7	Wide Area Networks (WANs).....	41
13.2.8	Multiple Communication Network PMU Systems.....	41
13.2.9	Substations.....	43
13.2.10	Control/Data Centers	44
13.2.11	Network Services.....	45
13.3	Timing Distribution.....	48
13.4	Latency in NASPInet 2.0 Networks	50
13.4.1	Network Level Cyber Security	51
13.4.2	Registries	54
13.5	Registry Structural Issues and Alternatives.....	54
13.5.1	Failure Notification.....	56
14.0	Relevant Standards	58
14.1	IEEE C37.118.1 and C37.118.2-2011	58
14.2	IEEE Std C37.244™-2013	58
14.3	IEC 61850-90-5.....	58
14.4	IP protocol suite.....	58
14.5	IEEE 1588	59
14.6	IEEE C37.238.....	59
14.7	IEEE 1451	59
14.8	IEC 61968/61970	60
14.9	IEC 27002	60
14.10	IEC 62351	60
14.11	IEC 27040	60
14.12	NERC CIP x	60

14.13 NISTIR 7628	60
15.0 Summary General Guidance for NASPInet 2.0.....	61
16.0 Guidance on Newer/Emerging Technologies.....	61
16.1 Software Defined Networks	61
16.2 Cloud Services.....	63
16.3 Network Virtualization and Network Function Virtualization.....	64
16.4 Communication Protocols	65
17.0 Glossary.....	65
Appendix A – PMU Applications Lists	A.1
Appendix B – General Principles for Composing Architectures.....	B.1
Appendix C – NASPInet 2.0 Function Class Definitions.....	C.1
Appendix D – Standard Network Security Measures	D.1
Appendix E – Middleware and Sensor/Communications Layers	E.1

DRAFT

Figures

Figure 1. Vertical Decomposition of a System into Layers	9
Figure 2. Example IP Protocol Stack.....	9
Figure 3. Decomposition of a System into Platform and Applications.....	10
Figure 4. Mapping a Core/Edge Model to a Layer Model.....	10
Figure 5. Silo to Layer Conversion Process.....	12
Figure 6. Layer Re-Association Process	12
Figure 7. US Electric Sector Regulatory Structure Model.....	17
Figure 8. Example Industry Structure Diagram: Pacific Northwest	18
Figure 9. US Electricity Market Structure	19
Figure 10. Functional Description Stages	22
Figure 11. NASPInet 2.0 Platform Layer Model	27
Figure 12. NASPInet 2.0 Core/Edge Ring Model.....	28
Figure 13. NASPInet 2.0 Distributed Platform Model	30
Figure 14. NASPInet 2.0 Distributed Platform with Regional Communication Network.....	31
Figure 15. Data Acquisition Function Mapping	32
Figure 16. Data Transport Function Mapping	33
Figure 17. Data Synchronization Function Mapping.....	33
Figure 18. Data Integration Function Mapping	34
Figure 19. Data Curation and Q/A Function Mapping	35
Figure 19. Data Security Function Mapping.....	35
Figure 21. Device and Network Management Function Mapping.....	36
Figure 22. Single src case	37
Figure 23. Multi-src no storage case.....	38
Figure 24. Multi-src mandatory storage case.....	38
Figure 25. Multi-src mandatory aggregation optional storage case.....	39
Figure 26. Multi-src mandatory concentration case.....	39
Figure 27. Multi-src network aggregation case.....	40
Figure 28. Wide Area Network Multi-Aggregation Ring Structure	41
Figure 29. PMU Networking Across Multiple Communication Networks	42
Figure 30. Example PMU Data Flows Across Multiple Networks.....	43
Figure 31. Transmission Substation Communication Network	44
Figure 32. Control/Data Center PMU Network Structure	45
Figure 33. Distribution System Structure with microPMUs.....	46
Figure 34. Distribution System with microPMUs and Multi-Services Communication Layer	47
Figure 35. Structural Models for Timing Distribution.....	48
Figure 36. Timing Distribution Structure Illustration.....	49

Figure 37. Two Substation PMU Data Flow Example 50
Figure 38. Linear Communication Network String Example 51
Figure 39. Illustrative NASPInet 2.0 GDOI Security Data Flows, Single Communication Network 52
Figure 40. Illustrative NASPInet 2.0 Security Data Flows, Multiple Communication Networks 53
Figure 41. Federated PMU Registry Structure..... 55
Figure 42. Three PMU Failure Event Message Data Flow Options 57
Figure 43. Simplified Joint Grid/SDN Wide Area Control Structure..... 62
Figure 44. Redundant Cloud Communication Path Concept..... 64

DRAFT

Tables

Table 1. Present PMU Application Categories	3
Table 2. Selected Key Terms and Concepts	7
Table 3. NASPInet 2.0 Capabilities and Functions.....	24
Table 4. NASPInet 2.0 Component Classes	25
Table 5. Example Latency Bound Estimation	51

DRAFT

1.0 Background and Purpose of This Document

In 2007-09, The US Department of Energy 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.¹ Considerable practical experience was gained in the course of executing the SGIG PMU projects and it became clear that due to that experience and various emerging trends, including technological advances, that the NASPInet framework should be re-evaluated. In 2017, NASPI published an assessment of PMU deployment experience with a focus on NASPInet implementations.² This report may be used as a view of the present state of PMU networking in the US.

The assessment report concluded that a variety of direct and innovative approaches to NASPInet implementation had been created across the industry but that 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 important changes have occurred; implementation and operational experience has shown limitations in the earlier NASPInet guidance. NASPInet architecture can now be improved to support the design and deployment of the next generation of synchrophasor data networks and applications.

The purpose of this document is to apply the experience gained in the first generation of NASPInet implementations, along with good networking practices, understanding of relevant new technologies, and views of emerging applications for PMUs to develop a new architecture guidance for PMU networks. This NASPInet 2.0 document updates the original NASPInet framework to incorporate learnings from utility experience and to address the new technology opportunities and needs for synchrophasor data networks, including concerns such as cyber security.

The intended users of this document are:

- Utility engineers and IT specialists engaged in developing, deploying, and using PMU networks
- Developers and vendors of PMUs, PMU software, and associated support systems
- Developers and operators of communication networks that transport PMU data
- Developers and users of applications such as analytics, protection and control systems, and visualization tools for PMU data
- Standards development organizations
- Testing and compliance assessment organizations

In addition, researchers may find the document valuable in understanding how real PMU networks may be structured.

¹ Ref the ARRA grants for PMU deployment

² JD Taft, Assessment of Existing Synchrophasor Networks, PNNL-27557, April 2018, available online: <https://www.naspi.org/node/723>

2.0 Scope

This document provides a revised and updated guidance on the architecture of PMU networks. The scope of this document includes:

- Architectural principles for PMU networks
- PMU network design leading practices
- Function¹ allocations
- Performance characteristics and considerations
- PMU network structure
- PMU signals management
- Data management and support services
- Timing distribution
- Cyber security
- Communication and application protocols
- Standards for PMU networks
- New and emerging technologies

Note that this document is a guidance or framework; **it is not an architecture specification and it is not a design specification**. It is somewhat closer to being an architecture specification and so follows good architecture practice² but is not intended to be taken as a strict mandate. Rather it is intended to inform specific PMU network architectures and designs, recognizing that each region and utility will have constraints and objectives that may differ from others and so may require differences in architecture and design. However, those differences, important as they are, should fit into the general framework provided here. In that sense, this document is a meta-architecture.³

3.0 Purposes of PMU Networks

The utility industry has several key goals for the use of PMU networks. These are:

1. Providing observability of electric transmission system behavior, dynamics, and asset state
 - a. operator decision support
 - b. post-event analysis
 - c. planning
2. Enabling advanced grid capabilities

¹ Functions are actions that provide a specified result, usually as sub-elements of a larger, more complex system.

² See Appendix 2 for a brief list of general system architecture principles – these are applied in this document.

³ A set of high-level decisions that will strongly influence the structure of the system, but is not itself the structure of the system. The meta-architecture, through style, patterns of composition or interaction, principles, and philosophy, rules certain structural choices out, and guides selection decisions and trade-offs among others. See https://definedterm.com/meta_architecture .

- a. integration and management of grid resources that exhibit stochastic¹ behavior
 - b. improved and new forms of grid dynamic control and stabilization
3. Improving component and system integrity protection
 - a. flexible protection schemes
 - b. adaptive protection schemes
 4. Improving cyber-physical security

4.0 Overview of Applications

Table 1 lists a set of categories of existing PMU applications. Note that the use of synchronized measurement in power systems is an evolving area so that new applications are to be expected.

Table 1. Present PMU Application Categories

Application Category	Description
Grid State Monitoring	Use of synchrophasor data to support elements of grid state estimation, including phase angles, system frequency, and ROCOF
System Model Validation	Use of synchrophasor data to calculate or verify system model parameters
System Modal Monitoring	Use of synchrophasor data for online monitoring of system oscillatory and other phenomena
Post-Contingency Analysis	Use of historical synchrophasor data to analyze system events after the fact
Forensic Analysis	Use historical PMU data to identify and characterize faults and failures, including measurement and PMU operation issues.

Early applications have been focused on transmission systems. Because these provided monitoring and diagnostic capabilities, they were not considered to be part of critical infrastructure. As experience with PMU technology has grown, so has interest in applications involving real time closed loop automation.

5.0 Emerging Uses for PMUs

As of 2018, the primary uses for PMU data have been in non-real time applications and those with humans in the loop. Looking forward, uses will include more closed loop real time applications with no human in the loop. These uses will include wide area closed loop stabilization and oscillation damping, wide area closed loop backup protection, adaptive protection, low and variable inertia system control, and fault prediction.

In addition, the use of microPMUs and synchronized measurement at the distribution level is emerging. Proposed uses include asset characterization, power flow monitoring, voltage sag

¹ Containing a random element. Examples for the grid are wind and solar electricity generation and demand response.

detection/characterization, cyber security attack detection, and fast pre-fault and fault characterization.¹ Due to public safety concerns, some aspects of distribution level fault detection and characterization must be performed very quickly so that protective action can be taken immediately.

6.0 Other Relevant Emerging Trends

In addition to new PMU uses, there are a number of emerging network and data processing trends and technologies that can affect PMU network architecture and design. These include:

Software-Defined Networking: standard communication networks are somewhat closed and have routing logic built-in; SDN opens up the network by providing some observability and by making the control plane accessible to external applications (separating the control and data planes), thus providing opportunities to manage network operations more completely than just via the built-in adjustable router settings.

Network Virtualization: abstraction of network hardware and data flows so that virtual tunnels can be easily configured for logical data transport service on a per flow basis without making changes to underlying communication hardware. In effect, each tunnel appears to the user as a dedicated private network, but many such tunnels may run on the same underlying network hardware.

Network Functions Virtualization: abstraction of network services (such as firewalling) from underlying hardware and placing them on a tunnel so that they can easily be configured for individual data flows (as created via network virtualization).

Cloud Computing: use of shared sets of computing resources. These may be accessed via the internet and operated by third parties as a service.

New communication protocols: communication protocol development is a rich field, but for PMUs there are specific developments on several fronts:

- Extensions and modifications to IEEE C37.118
- Extensions and modifications to IEC 61860-90-5
- Development of a streaming telemetry transport protocol (STTP) motivated by PMU data transport issues

These emerging trends are considered in this document and should be factored into the development of any new or upgraded PMU network architecture.

7.0 Implications and Systemic Issues

For present day applications, two network performance criteria have been crucial to ensure delivery of accurate PMU measurements to applications:

- Low message packet loss

¹ JD Taft, Fault Intelligence: Distribution Grid Fault Detection and Classification, PNNL-27602, September 2017, available online: https://gridarchitecture.pnnl.gov/media/white-papers/FaultIntelligence_PNNL.pdf

- Precise synchronization

Going forward, a third criterion will become equally important:

- Low transport latency

Latency matters for two reasons: real time applications¹ have demanding timing constraints and closed loop controls can be made ineffective or even unstable by excessive latency in data or command delivery. The implication of this is that PMU networks must be designed and operated as strongly time-constrained systems. It is not enough to time-stamp data; it must actually be delivered in a strictly time-bounded manner according to the requirements of the application using the data for mission-critical purposes.

In addition, the use of PMUs at the electric distribution level will require the measurement of much smaller phase angle differences than has been the case with traditional transmission level applications. Applications such as feeder line segment or section impedance measurement and asset characterization will require highly precise and accurate phase angle determination, which in turn will depend on highly accurate and precise timing.

These requirements make the synchrophasor communication network an integral part of the grid sensing/measurement and protection/control system, not as a separate adjunct. This is especially true where PMU network performance, including performance of communication networks, data management systems, timing distribution, and applications are crucial to grid operation. Communication network performance is becoming a key component of grid performance, reliability and resilience. Implications of this growth in potential uses of PMU technology have led to the need for new guidelines in the architecture and design of PMU networks at all levels of electric power delivery.

The future use of PMU networks as part of real time protection and control for power grids means that such networks will be integral parts of grid operations and will turn these systems into critical cyber assets because they support real time grid operations.

7.1 Systemic Issues

The legacy situation described in the assessment report, combined with the emerging uses listed above leads to a set of systemic issues that should drive any new PMU network architecture. These are:

1. PMU networks must be capable of high performance in several areas simultaneously: low packet loss, low latency, and precise synchronization. Future PMU data flows may well increase in volume, so network capacity (bandwidth and throughput) requirements must also be treated with future uses in mind.
2. PMU data for a transmission network may be produced, transported, stored, and applied by multiple entities. Some organizations will carry out several or all of the steps, but each step may be carried out by a different organization. Performance of each entity affects overall system performance, as does the coordination among entities.²

¹ Real time applications are those in which the usefulness of a result is dependent not just on logical (algorithmic) correctness, but also upon the time at which the result is delivered, and when, in addition, timely delivery poses a significant design or implementation challenge.

² Coordination is the process or method by which multiple decentralized entities, systems, or devices are able to cooperate to solve a common problem.

3. While PMU data transport networks may not necessarily be owned and operated by electric utilities, PMU communication networks must be able to treat data transport as an integral system element in terms of performance, system management, and cyber security.
4. PMU signal bundling excessively limits flexibility in sharing and applying PMU data.
5. PMU networks will become part of the critical infrastructure for transmission systems and therefore will need strong cyber protection. This is especially an issue for devices in the field (PMUs in particular) and communication networks, both of which are viewed as being Achilles heels for the grid.
6. Timing distribution across PMU networks must support very high precision for phase angle measurement and for wide area multiple phasor comparison/calculation.
7. Timing distribution for PMU networks should not dependent on only one master clock source type. PMU networks need explicit recognition of the need for accurate, reliable, precise, cyber secure, and robust timing and therefore timing distribution.
8. PMU communication network are not physically homogeneous; nor are they composed of single Ethernet domains. Some portions of these networks (notably the substation links) may not support the capability and performance levels that core transport networks do. Network management is complicated by both of these factors.
9. Real time uses of PMU data are incompatible with multi-stage (or even single stage) data packet concentration and timestamp-based ordering processes from both latency and data loss standpoints.
10. PMUs do not necessarily incorporate advanced communication protocols, especially the legacy devices. PMU networks may contain mixed sets of legacy and new PMUs.
11. The same issue exists for protection and control devices; in addition they may not be equipped to receive and handle PMU data.
12. PMU networks are dynamic in structure, both physically and logically. New PMUs can be installed; old ones can fail or be taken off line. The same is true of communication links and applications (data sinks). The bindings of applications to PMU data sources can also be dynamic.
13. Constant streaming of PMU data results in large amounts of data in motion and large accumulations of data – all of which much be managed (routed, stored, curated, protected, and made accessible under controlled conditions). The original NASPInet guidance document presumed that PMU data streamed to control rooms and centralized PDCs, but future applications will require a more decentralized and edge-based structure, with applications that use PMU data potentially being located in substations and at various devices not located in control centers; hence NASPInet 2 must support peer-peer-like data flows.
14. PMU data may be used in multiple locations by multiple applications having differing latency requirements. Therefore the concept of discrete data transport services classes is not useful since any given PMU message may be subject to multiple simultaneous requirements.
15. Data management must not result in the deliberate destruction or unintended loss of data.

These systemic issues are significant drivers for future PMU network architectures.

8.0 Key Definitions and Concepts

A number of terms and concepts will be used extensively and immediately and so are defined in Table 2 here. A more complete glossary is provided at the end of this document.

Table 2. Selected Key Terms and Concepts

Term or Concept	Definition
Architecture	A high level depiction of a system, used to reason about the system's properties and behavior, and to specify key decisions about the shape (structure) of a system for the purpose of bounding and simplifying design decisions. Architectures are comprised of black box components, structures, and externally visible characteristics. Structures are the means by which components are related or connected.
Network	A group or system of interconnected people or things
communication network	An interconnection of communications components
electric network	An interconnection of electric components. This term may be applied to either transmission or distribution systems.
PMU network (Synchrophasor network)	An interconnection of PMUs, communication devices, and PMU data processing elements
Connectivity & topology	<p>Connectivity is the state of being connected, that is, the existence of at least one linkage between two endpoints or nodes. In graph theory it is quantified as the minimum number of links that have to be removed in order to separate two nodes. More generally it is the existence of some form of relationship, interaction or exchange between two or more components, elements, or nodes. For communication networks, network topology describes the specific connections among the communication devices and endpoints. Electrical connectivity is essentially the same thing for electric power flows and circuits.</p> <p>Electric power systems have dual connectivity: electrical and communication. This is an important factor in cyber security.</p>
Sample	A single discrete measurement of a physical quantity, typically quantized and represented digitally.
Signal	An impulse or fluctuating quantity, such as an electrical voltage or light intensity, whose variations represent coded information. The coded information may represent some type of behavior, or may indicate an event or other message. Signals may be analog (continuous in time and magnitude) or digital (discretely sampled and quantized and represented as number streams). An analog signal is any continuous signal for which the time varying feature (variable) of the signal is a representation of some other time varying quantity, i.e., <i>analogous</i> to another time varying phenomenon (e.g. power line voltage). A digital signal is the discretely sampled, quantized equivalent. The signal may represent something more complex than a simple physical fluctuation, such as a three phase current phasor. A sequence or stream of phasors may be considered a signal.
Data packet	A data packet is a unit of digital data made into a single package that can be transported via a digital communication network.
Message	A logical unit of information represented in digital form and being transported by a communication system. A message may fit into a single data packet, or may be

Term or Concept	Definition
	spread over more than one packet. A PMU message may be comprised of multiple synchrophasor signal samples.
Data stream	A coherent sequence of data packets (usually comprising a signal or signal set) being transported on a continual basis at a more or less uniform rate
Protocol (communication)	<p>A communication protocol is a system of rules that allow two or more entities of a communications system to exchange information via any kind of variation of a physical quantity. Generally, a suite of protocols is used to enable communication; such a suite may be represented as a stack.</p> <p>While some people want to consider certain protocols as communication protocols and others as application protocols, such a distinction is handled in a more sophisticated way using stacking via models like the Open Systems Interconnection Model, which uses seven layers (application, presentation, session, transport, network, data link, physical) to represent protocol stacks, with the top layer being where the Application level protocol is represented. In this manner, one may merely refer to the layer or layers at which a particular protocol operates. Another version of this approach is the IP model, which has four abstraction layers (application, transport internet link) plus the physical layer. Thus an application protocol is just one of a stack of communications protocols comprising a protocol suite.</p>

9.0 Foundational Principles

The following principles underlie the architectural guidance provided later in the document. In accordance with good architecture development practice, the principles are necessary to help ensure the conceptual integrity of the specification set.¹

9.1 Leading Network and System Practices

The management of PMU data is not fundamentally different from the management of other forms of streaming data (such as video data or financial transaction event messages) and there is nothing about the electric utility environment that makes it inherently exempt from network and data management methods developed for other industries and applications.

Layering is an architectural (structural) concept that has application in several ways and at multiple scales. Understanding the abstract concept of layering (as opposed to assuming that it only applies to some particular product or method) makes it possible to recognize how to use it for PMU networks. In general, layering is a vertical decomposition of system elements. One of the significant advantages of a layered architecture that has at least three layers (see Figure 1) is that a middle layer can isolate the layer above and the layer below from changes in the opposite layer. This provides resilience to changes in devices and technology, software, etc., and enhances future-proofing of investments for the same reason.

¹ See Appendix 2, item 2.


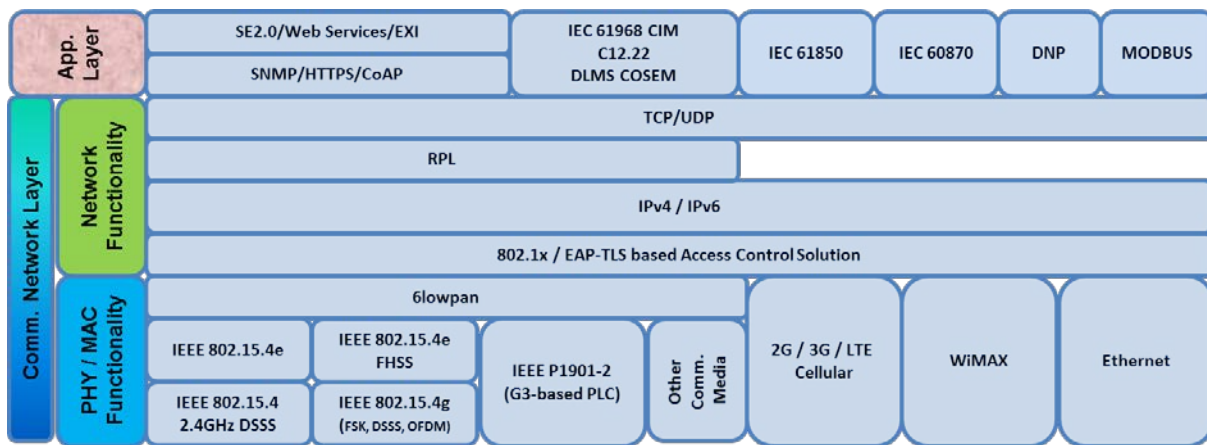


Figure 1. Vertical Decomposition of a System into Layers

The protocol suite model for the internet makes use of layering, although there it is usually referred to as a protocol stack – the stack itself being a layering. The OSI reference model is a seven layer model for communication between applications over a network.¹ The internet protocol suite uses a similar five-layer model.² Figure 2 shows an example of an actual IP protocol stack. Note the presence of the application layer in the stack.



Source: Cisco

Figure 2. Example IP Protocol Stack

Platform is another architectural (structural) concept whereby system elements are grouped into at least two layers: one consists of elements that supply a general set of capabilities and functions that are relatively static over time (the platform) and another that consists of functions that are used intermittently or change relatively often (the applications). A platform is a stable collection of components that provide fundamental or commonly-needed capabilities and services to a variable set of uses or applications through well-defined interoperable interfaces. A personal computer with an operating system is a platform; email, word processing, and spreadsheets are examples of applications that use this platform. Key properties of a platform include:

1. Separates foundation functions from end uses (“applications”) via layering
2. Provides a set of services and capabilities that are useful to many applications
3. Is stable over time, while the applications may change frequently
4. Provides decoupling of changes between applications and underlying infrastructure
5. May scale (adjust resources) to support variable demands from applications
6. Open: third parties can freely create applications that use the platform (needs open standard interfaces)

¹ Computer Network | OSI Model, <https://www.geeksforgeeks.org/layers-osi-model/>

² Internet Protocol Suite, https://en.wikipedia.org/wiki/Internet_protocol_suite

A platform usually has multiple functions and so a platform-based system may be thought of in terms of a vertical (layered) decomposition into platform and applications, followed by a horizontal (functions/component allocation) decomposition, as shown in Figure 3.




Figure 3. Decomposition of a System into Platform and Applications

This general concept can and should be applied not just to software or middleware, but to the underlying grid sensors and communications systems as well.¹

Core/edge structure is an architectural concept that many advanced communication networks use. The core is kept free of functions other than those related to data transport; other functions are connected at the edge, sometimes via an access layer. The internet and many IP-based networks are structured this way; this has proven to be capable of high performance and considerable extensibility and flexibility in supporting ever-evolving uses. In fact, the core/edge model is really a layer/platform model, where the core is the bottom layer, and the edge-connected device functions are the applications. Figure 4 illustrates the core/edge to layer mapping.

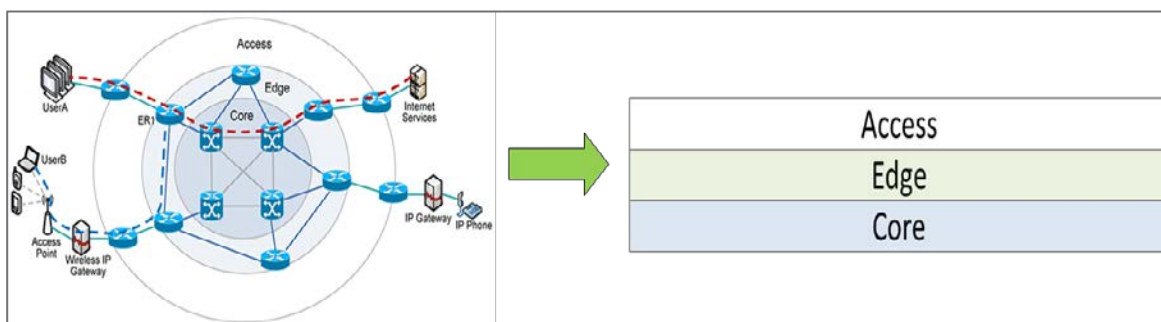


Figure 4. Mapping a Core/Edge Model to a Layer Model

9.2 Interoperability and International Open Standards

While the concepts of interoperability and open standards are widely accepted and understood, they are worth mentioning here as a point of reference for the guidance to be given later. The value of open standards for interface and interoperability includes:

- Reducing integration time and costs

¹ JD Taft and P De Martini, Sensing and Measurement Architecture for Grid Modernization, PNNL-25249, February 2016, available online:

<https://gridarchitecture.pnnl.gov/media/advanced/Sensor%20Networks%20for%20Electric%20Power%20Systems.pdf>

- Improved future-proofing and avoidance of stranded investments as technologies change
- Avoidance of vendor lock-in

Given the number of entities involved in the generation, transport, management, and use of synchrophasor data, interoperability at multiple levels is a crucial issue for NASPInet 2. From the discussion above, it should be clear that open interoperability standards are needed for the proper design of platforms as well (see platform property 6 above).

9.3 Synchronization

Since the whole point of the use of synchrophasors is synchronized measurement, measures to assure synchronization, including robust timing distribution, are crucial. While it may seem more robust to develop asynchronous systems that are loosely coupled, and then to use some form of time stamping and concentration process to temporally align phasor signals, a key goal is to *sample synchronously* so as to capture a snapshot in which samples are temporally aligned, not merely grouped after the fact. Variations in latencies will cause PMU messages to arrive at usage points at different times, so processes will still be needed to assure that PMU signals are aligned, but for PMU-based calculations to be valid when the phasors are taken from widely separated locations in the grid, sampling synchronization is needed. Synchronization may be driven from a common master clock/timing source, or may be driven from a combination of master sources that operate in tandem or operate as primary and backup timing sources. A third option, firefly synchronization, is discussed in the timing section.

9.4 Function Allocation

Function decomposition is a part of the architecture task, but is only part of the job. A key aspect of system architecture is the allocation of functions across the architecture's components, something which is done in the context of the architecture's structures. The combination of structure and function allocation has a major impact on the location and nature of interfaces and the amount of integration necessary to implement a system. Well-chosen structures and allocations can greatly reduce the amount of necessary integration effort. In that regard, the determination of systems layers, platforms and applications (a preliminary and primary function allocation) is a crucial step in the specification of a PMU network architecture (or any system architecture) that is based on the relevant structures.

Four primary targets for function allocation are:

- Sensing and measurement
- Communication networks
- Data management middleware
- Applications

Simple approaches to function allocation tend to suffer from three problems: hyper-orthogonalization,¹ siloization,² and misallocation due to incomplete understanding of capabilities of the target components. Addressing these issues is exactly where the opportunity lies for the use of structuring and allocation to reduce and manage integration costs and risks.

¹ Hyper-orthogonalization is the overly rigid partitioning of functions for the sake of conceptual simplicity while failing to take into account practicalities of system design and implementation.

² Siloization is the stacking of functions into "vertical stovepipes" – this is generally done for the sake of a local optimization, such as to support a concept of product "completeness," regardless of the system implications.

Siloization (grouping functions and/or data into isolated organizational containers) may be addressed architecturally through silo-to-layer conversion as shown in Figure 5. Part of this process is layer re-association, which is also illustrated in Figure 6. Use of this and other system architecture processes provides the tools to develop advanced systems that exhibit strong adherence to good architecture principles, with a minimum of ad hoc structure and unnecessary clutter, while providing guidance to product developers whose products will support and be incorporated in the actual systems derived from these architectures.

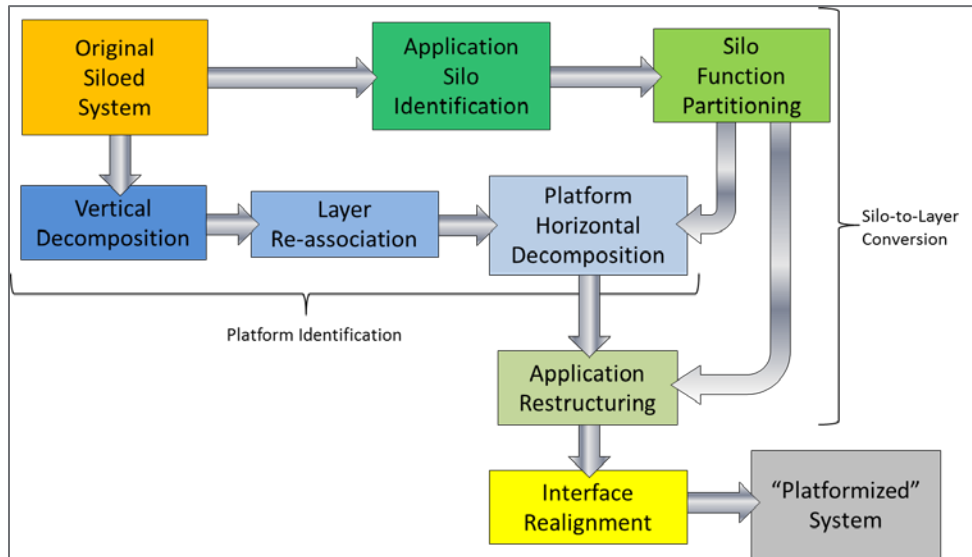


Figure 5. Silo to Layer Conversion Process




Figure 6. Layer Re-Association Process

Once the layer definitions are done it is possible to address the potential functional overlap between layers, such as communication networks and middleware. To do this means having a detailed understanding of the capabilities of each and viewing them in the context of the overall system architecture. Using the concept of layer re-association, it is possible to assign functionality in ways that clarify the value of each component in an objective manner.

9.5 Some Implications of These Principles

Given that communication networks are central to NASPINets, architects and system designers of PMU networks should give close consideration of the full capabilities of IP networks and the complete IP protocol suite. Similarly, close consideration should be given to middleware and associated protocols and standards. The potential functional overlap between these two in particular should be resolved in accordance with stated system architecture principles set by and for the organization developing the

architecture. Understanding the use of layer re-association to define platforms is an important element of this process.

The core-edge principle (a version of layering) supports a conclusion that has emerged from PMU network experience: phasor data concentration should be viewed as a function, not as a component, and the allocation of that function to actual system components is a significant architectural decision. In addition, over-the-top networks (networks imposed on an underlying core network) should be avoided, as applications and their associated functions should be located at the edge. Such considerations mean that PDC¹ stacking can and should be eliminated, as PDC stacking² imposes an unnecessary network on top of the core communication network.³ Finally, protocol conversion gateways should be minimized or eliminated by the use of well-structured interface definitions and open interoperability standards. Protocol conversion gateways solve interface problems in an ad hoc way but are also breaks in the network architecture that create problems in managing networks from operational, cyber-security, and upgrade perspectives.

10.0 Core Considerations

Among the drivers for this architectural guidance are a number of considerations derived from an understanding of the basic problem of synchrophasor data management, along with an understanding of issues derived from the practical experience of the ARRA program and analysis of emerging trends in the use of PMUs. In an actual architecture development, these would lead to a list of requirements that would become drivers for the architecture. Here we refer to them as core considerations.

10.1 Attributes of Capabilities and Functionality

A capability is the ability to execute a specific course of action in support of an objective. Functions are processes or actions needed to support capabilities. The attributes of capabilities and functions are early issues for NASPInet 2.0 architectures and lead to core considerations such as those listed below.

- Data transport – PMU networks must transport of PMU data to data stores and end use applications. A core issue here is the use of advanced but well-established communication network protocols and capabilities to provide transport capabilities in an environment with data streams from multiple data sources, each stream potentially having multiple destinations with widely varying latency requirements.
- Data persistence – persistence refers to the retention of data on any time scale. Therefore it includes not only traditional databases and historians, but also temporary buffers and, because of the potential for use in real time applications, even “data in motion.” PMU network architects and designers must

¹ Phasor Data Concentration – formerly Phasor Data Concentrator, this is now considered a functionality rather than a component.

² Note that the non-stacking principle applies to the PDC functionality in general. This includes temporal stacking, where PMU data has been concentrated and then stored, only to be re-concentrated with other PMU data later. Such re-concentration is also PDC stacking. This has implications for how PMU data is stored.

³ One of the factors that leads to the stacking of PDCs is subtle: they tend to be placed at organizational boundaries. Since PMU networks involve sharing of data flows across such boundaries, such use of PDCs arises due to fragmented system thinking. A core principle of system architecture that, absent a strong set of architectural principles and conceptual integrity, real architectures tend to mirror the structures of the organizations that produced them. This is not always a good outcome.

understand the multiple latency requirements that any given data stream may have, due to the potential for multiple simultaneous destinations for the data.

- Network management – this refers to both communication network management and PMU system management. In both cases, functions are needed to support network configuration, enrollment of devices, access control, performance monitoring, device fault management, software/firmware upgrades, security, data flow and link congestion management, and possibly accounting for billing purposes. A core consideration is whether and how PMU system management is integrated with communication network management.
- Application interface/integration – integration is the connection of elements into a functioning system, and the use of interfaces between components is a common aspect of integration. A core concern is the proper specification of interfaces (not just in terms of what flows across them, but also where they are located in the architecture). Architecture should inform interfaces, not the other way around. The bottom-up approach of specifying an interface between two boxes and then trying to create an architecture around it is a mistake.
- Data sharing – PMU data has wide applicability and the potential users of such data do not all reside inside the same entity that produces it. This is a direct consequence of utility industry structure (see the discussion industry structure in the Problem Domain Reference Model section below); it is also in part a consequence of industry regulation. A core concern is how to make appropriate data sharing available while observing constraints on data confidentiality and cyber security, and how to do so efficiently.
- Service classes – not all existing and potential uses of PMU data have the same requirements for PMU system performance. This is a consequence of the multiplicity of potential uses across multiple organizations that result in differing application time scales and data delivery needs. A core concern is how to provide the appropriate classes of service in PMU data acquisition, transport, and delivery without simply gold-plating everything. In addition, when telecommunications service providers are involved, then QoS is not entirely under the control of the utilities.

10.2 Cybersecurity

The ARRA PMU deployment projects paid some attention to cyber security issues, but on the whole, it was not made a priority. However, as PMUs are more fully integrated into grid operations, and especially as they become part of critical infrastructure functions such as protection, state estimation, and system stabilization, it will be necessary to incorporate comprehensive cybersecurity measures into NASPInets used in transmission systems. The degree to which this is done is a matter of determining appropriate levels of countermeasures to vulnerability and threat assessment. A substantial number of standard communication network-level security measures are well-known and are applied in at least some parts of the US electric utility industry, but threats continue to evolve and so it is important to structure PMU networks to be able to incorporate presently known defenses and be able to accept more as they are developed. Therefore the synchrophasor network architect should specify architecture structures with the concept of *structural securability* in mind, as well as sufficient structural flexibility to incorporate new approaches. Relevant architectural concepts include layering and platforms, modularity and module (de)coupling, and the use of graph theory to determine inherent securability characteristics of proposed structures.

10.3 Performance Characteristics

- Data Rates – existing PMUs mostly provide signal samples at either 30 or 60 samples per seconds in the US. Signals are encapsulated in data packet streams for transport across communication networks,

along with meta-data and routing and other communication network –related data. Emerging views about the use of PMUs indicates that much higher sampling rates may come into use, resulting in the need for higher data rates in the transport networks and in the PMU data processing systems. A core concern is that the NASPInet 2.0 architectures be capable of handling future data rates without the need to rip and replace systems based on the new architectures. In other words, it is desirable to use architectural methods to future-proof NASPInet 2 investments as regards data rate growth.

- Latency – applications can have differing latency requirements, depending on the functionality they provide. A core concern is that the PMU system be able to support all of the latency requirements, a potential challenge given the need to cross organizational boundaries, cover large physical distances, and likely traverse multiple heterogeneous communication networks.
- Throughput – this is different from raw data rates; in a communication network, there are potential issues of data overhead (some data must be transported along with the payload to manage its flow through the network and the network uses some of its raw bandwidth for network management purposes independent of specific payload data flows) and communication links may experience congestion, all of which can reduce (sometimes significantly) the average rate at which payload data is delivered. Networks can have latency curve “knees,” essentially breakpoints at which latency rises significantly, this causing variable performance. The communication networks must be designed (where possible) and managed to ensure proper data throughput.
- Availability – the PMU system be operational for an extremely high percentage of time. This will become more crucial as PMUs are used in real time grid operations functions.
- Reliability
 - Packet loss minimization – while packet loss is never desirable, it will become a much more severe issue for future PMU applications involving real time protection and control.
 - Data delivery assurance – in addition to simple packet loss, there are other reasons why PMU data may not be delivered to an application in a timely fashion. In addition, some applications may require means to assure that data is in fact delivered, despite the various ways that delivery can be disrupted or degraded.
- Failure notification – PMU networks are part of larger systems that exhibit Ultra Large-Scale (ULS)¹ complexity. One aspect of ULS system is that they experience “normal failures,” meaning that something is always not working and this is a part of normal operation, not an exception condition. A core concern is how to deal with these failures and an aspect of that is notification of failures to the appropriate system component or human that can deal with the failure, preferably “on the fly.”

10.4 Functional Flexibility

- Scalability – scalability has two forms: communication (data transport) scalability, which breaks down into communication network capacity and endpoint scalability (ability to increase the number of endpoints without major restructuring) and computational scalability, which is the ability to add computation resources as needed when computational problem complexity or data volumes increase. Since new applications for PMUs are contemplated, and faster data rates are also being considered (and there will be additional PMU deployments), the architecture must be able to support the various dimensions of scalability.
- Extensibility –as PMU uses evolve, PMU systems should accommodate extension of the functions or capabilities of these systems without massive rework. This is an aspect of future-proofing.

¹ Linda Northrup, et. al., Ultra Large Scale Systems, Carnegie Mellon University, June 2006, available online at www.sei.cmu.edu/library/assets/uls_book20062.pdf

- Resilience – for electric grids, resilience is a poorly defined concept in general, and typically only partially focuses on issues that would impact PMU systems. So the first concern is to have an appropriate way to understand what resilience means^{1,2} and then to consider architectural measures that act as countermeasures to vulnerabilities that would impact resilience. This concept has a close relationship to cybersecurity as well.

11.0 Synchrophasor Problem Domain Reference Model

One definition of a problem domain reference model is a high level depiction of the problem space – “a bird’s-eye view of a domain or problem space.”³ It essentially lays out the key elements that must be recognized in the development of an architecture. It provides a quick way to see and talk about a given problem space and helps in:

- Focusing problem definitions and requirements
- Framing stakeholder inputs
- Identifying systemic issues

Problem domain reference models are used as inputs to the architecture development process; architectures are mapped onto problem domain reference models. This section outlines the essential elements of the synchrophasor problem domain.

11.1 Regulatory Structure

The US electric utility industry is heavily regulated from a variety of sources and so regulation and regulatory structure are integral parts of the problem domain. Figure 7 shows the structure of US electric sector regulation. This diagram is not intended to show regulations; rather it depicts what entities are regulated, what entities do the regulating, and what the nature of the regulation is.

¹ JD Taft, Electric Grid Resilience and Reliability for Grid Architecture, PNNL-26623, March 2018, available online: https://gridarchitecture.pnnl.gov/media/advanced/Electric_Grid_Resilience_and_Reliability_v4.pdf

² S Widergren, et. al., Toward a Practical Theory of Grid Resilience, PNNL-27458, April 2018, available online: https://gridarchitecture.pnnl.gov/media/advanced/Theory_of_Grid_Resilience_final_GMLC.pdf

³ JD Meier’s Blog, Reference Models, Reference Architectures, and Reference Implementations, Feb 16, 2011, available online: <https://blogs.msdn.microsoft.com/jmeier/2011/02/16/reference-models-reference-architectures-and-reference-implementations/>.

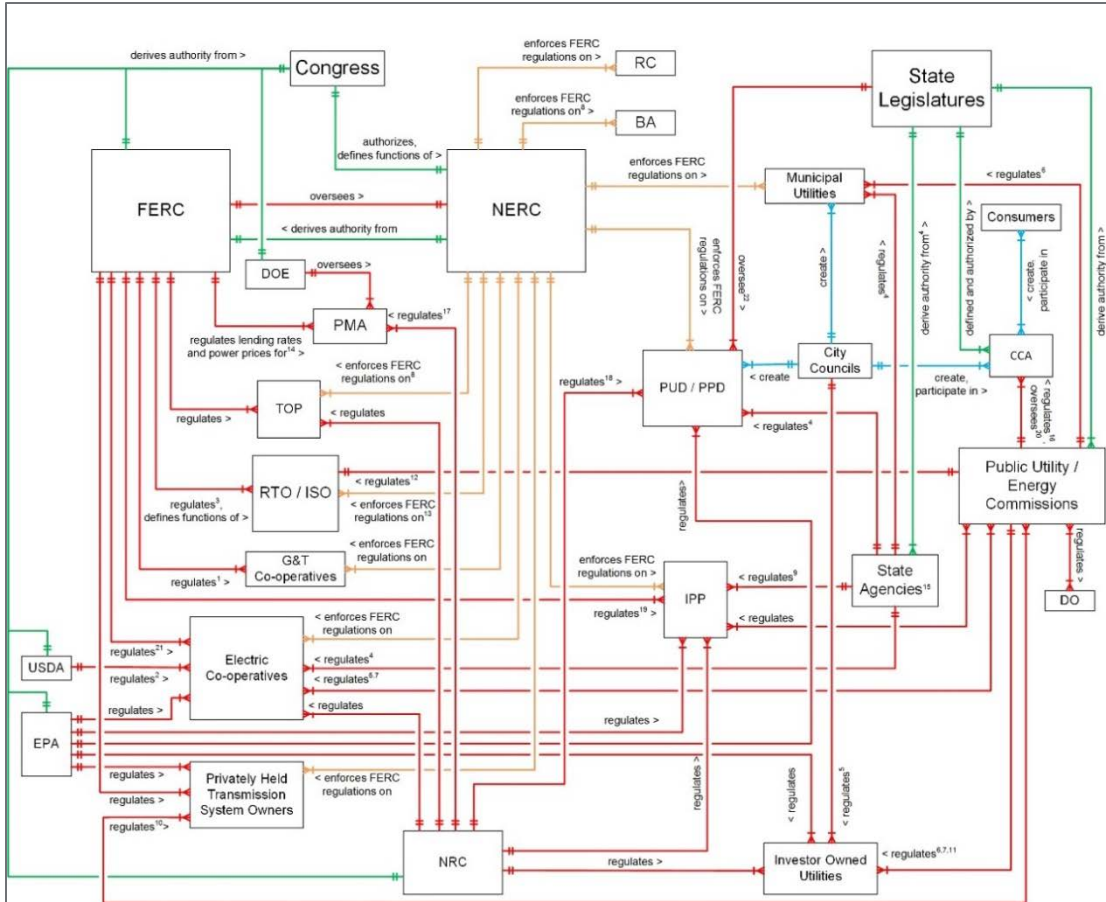


Figure 7. US Electric Sector Regulatory Structure Model

11.2 Industry Structure

Industry structure refers to the set of entities that participate in the electricity sector and how they are related. It varies widely by region in the US, as well as by utility type (IOU, co-op, municipal, power marketing authority, independent power producer, etc.). Figure 8 illustrates one example of utility industry structure in the form of an Entity-Relationship diagram. Boxes represent entity classes (not individual entities, unless a class has only one member) and lines represent the set of activities encapsulated in a relationship between the entities connected by the line. Note that this is a multi-layer diagram but here is shown as an example with all layers turned on. For practical work, NASPInet architects and designers should refer to the appropriate model(s) for the region in which they are working. Note: it is extremely important for the NASPInet architect to understand industry structure – this is the context in which NASPInets must exist.¹

¹ See Appendix 2, item 8.

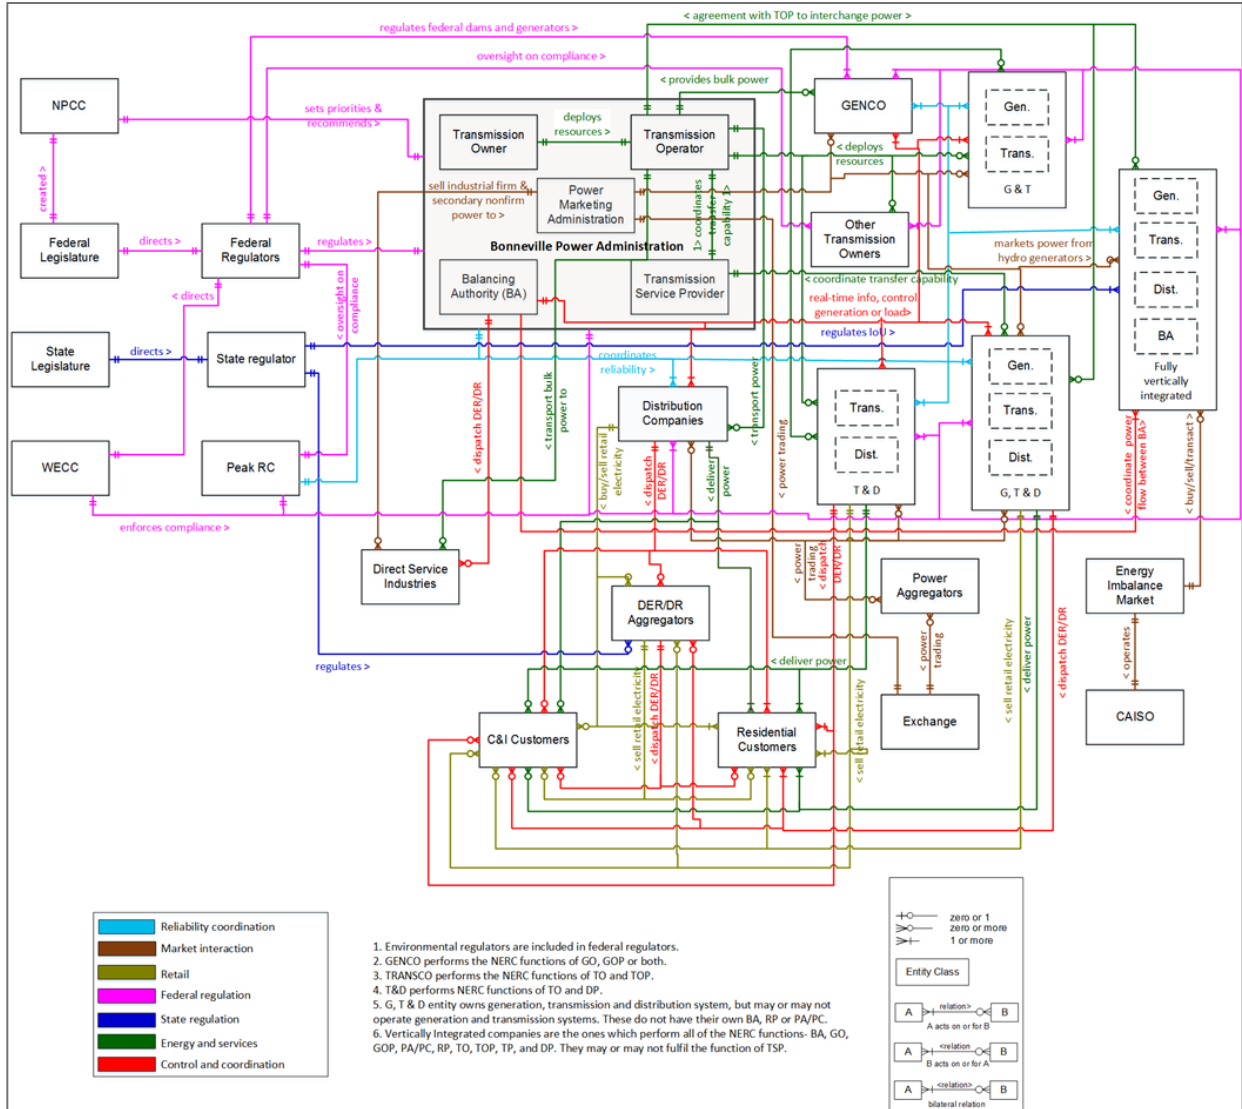


Figure 8. Example Industry Structure Diagram: Pacific Northwest

11.3 Market Structure

Market structure is technically a subset of industry structure but is sufficiently important and complex to warrant explicit models. Note that while organized wholesale markets exist in only part of the US, markets for electricity are ubiquitous. Figure 9 illustrates US electricity sector market structure.

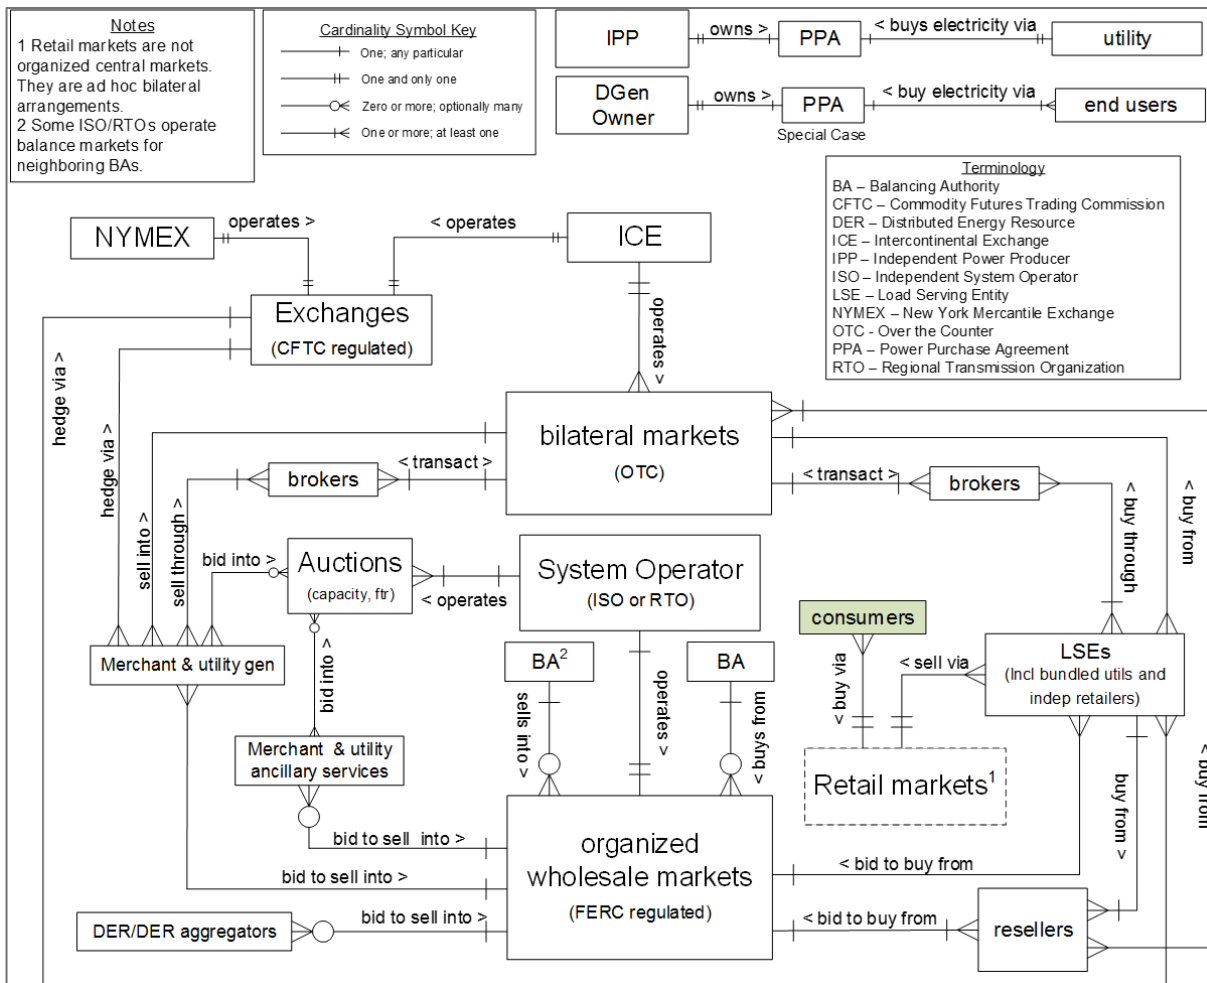


Figure 9. US Electricity Market Structure

11.4 Relevant Entities

A wide variety of entities are affected by the use of synchrophasor data and therefore in NASPInet 2.0. The following is a categorized list of such entities.

11.4.1 Planning/Operating Functions

- Distribution Provider
- Distribution Owner
- Distribution Operator
- Generator Owner
- Resource Planner
- Transmission Planner
- Transmission Operator
- Generator Operator

- Transmission Owner

11.4.2 Reliability Service Functions

- Balancing Authority
- Interchange Authority
- Planning Coordinator (also called – Planning Authority)
- Transmission Service Provider
- Reliability Coordinator

11.4.3 Standards and Compliance Functions

- Standards Developer
- Compliance Enforcement Authority [Compliance Monitoring and Enforcement]

11.4.4 Others

- Product, application, and system developer/suppliers
- Telecommunication service suppliers
- Third party data management/analytics service suppliers
- Researchers

Ultimately, of course, the retail consumers of electricity have a stake in NASPInets, but indirectly through their desires for reliable and affordable electricity service. They are not directly concerned with the architecture of NASPInets.

PMUs are not just for bulk energy systems; the use of PMUs at the distribution system is emerging, so it is important to recognize differences in the problem domains as well as in the requirements for PMUs for these two sub-domains.

11.5 Bulk Energy Systems

Transmission grids are generally comprised of interconnected transmission systems, wherein the systems are strongly meshed. Most are three phase AC systems that are operated in a balanced mode, although some high voltage DC links exist. Transmission systems may include transmission substations and switching stations, equipment for flow control such as switching, phase shifters, and variable frequency transformers; and equipment for regulation such as FACTS¹ (static VAR compensators, unified power flow controllers, etc.). The US transmission systems are organized into three major interconnections, with various inter-tie points connecting them. There is a small interconnection in Alaska which has no inter-tie to other interconnections. The larger interconnections have interchanges with Canada (Quebec Interconnection) and Mexico (CFE).

¹ Flexible AC Transmission System

By 2017, there were over 2,500 PMUs deployed on transmission systems in North America, many due the ARRA projects.¹ Many utilities developed NASPInet implementations, based to some extent on the original NASPInet guidance document developed in 2007-2009. The recently developed assessment report on the work done under the ARRA projects encapsulates much of the valuable experience developed in that work.²

11.6 Distribution Systems

Electric distribution systems differ from transmission systems in a number of important ways: while primary distribution is usually three phase, the circuits are generally unbalanced due to the fact that most loads are single phase. Many primary distribution circuits are simple radials, but some are partially meshed in urban and suburban environments. In certain dense urban areas, the distribution secondaries are dense meshes (fed at multiple points through network transformers and network protectors) or are spot networks (fed from multiple points into what is essentially a bus). For partially meshed distribution, the topology of the network can change on both short and long time scales and may not be well documented or modeled. Many distribution system do not have distribution SCADA and some do not have complete substation SCADA. While many have wireless meters and meter communication networks, most of the older communication networks are unsuitable for handling high speed telemetry. This is relevant because unlike for transmission systems, the use of PMUs at the distribution level would likely require PMUs to be located not just in distribution substations but also at locations on feeders outside of the substations. Limitations in communication network may also lead to the placement of PMU data analytics and applications in locations other than operations or data centers.

In the 20th Century, there was not much need for high speed measurement distribution systems because except for protection, nothing happened on time scales faster than about five minutes. With the rise of Distributed Energy Resources and internet-connected remotely controllable loads, distribution grid dynamics have been shifting to much shorter time scales. In addition, focus on asset management and fault intelligence for reliability purposes has introduced new measurement requirements for distribution systems. The question as regards PMUs is whether the need is for full synchrophasor measurement, or just for finely grained amplitude only or phase angle difference only measurements. We can distinguish among the three; there are arguments for each view. Given how the angle measurements are made though, it would seem logical to implement full phasor measurement so that phasors are available when phase angles needed.

Distribution-level applications of synchronized measurement may require PMU-like capabilities but existing PMUs may not be suitable for use here, so “micro-PMUs” have been suggested for this purpose. In this guidance the general term PMU will be used for both transmission and distribution devices.

11.7 Key Constraints and Barriers

A general barrier to the use of PMU data is the need to share data across organizational boundaries. While the technology to do so is available, there are issues of data confidentiality and cyber security, as well as more traditional organizational boundary problems. Certain regulatory issues regarding unequal access to information can come into play with regards to electricity markets and have led to issues with the structure of message formats for PMU signals.

¹ A Silverstein, Synchrophasors and the Grid, NARUC Summer Meeting 2017, (NASPI-2017-P-005), available online: https://www.naspi.org/sites/default/files/reference_documents/naspi_naruc_silverstein_20170714.pdf

² JD Taft, Assessment of Existing Synchrophasor Networks, PNNL-27557, April 2018, available online: https://gridarchitecture.pnnl.gov/media/white-papers/Synchrophasor_net_assessment_final.pdf

A barrier to use of synchrophasors in distribution systems is that while transmission PMUs can be placed in transmission substations, distribution PMUs must be placed not just in substations but at points along distribution feeders. This has implications for device packaging and cost, as well as installation and maintenance costs. Since many distribution circuit control devices have sensing and measurement (i.e. voltage and current waveform sampling and processing) capability built in, as well as communications interfaces, it may be helpful to add synchrophasor functionality to them in order to avoid the cost of separate device installations on the feeder circuits.

A second barrier to the use of PMUs at the distribution level is that phase shifts over the length of a distribution feeder, or worse, feeder segment are *much* smaller than the phase angle differences found in transmission systems. This means that phase angle measurement must be much more precise than those in transmission PMUs.

12.0 NASPInet 2.0 Architecture Principles, Objectives, Capabilities, and Functions

In general, the development of the functional definition of a system at the architecture stage progresses from objectives to increasingly more granular and voluminous detail, as shown in Figure 10.

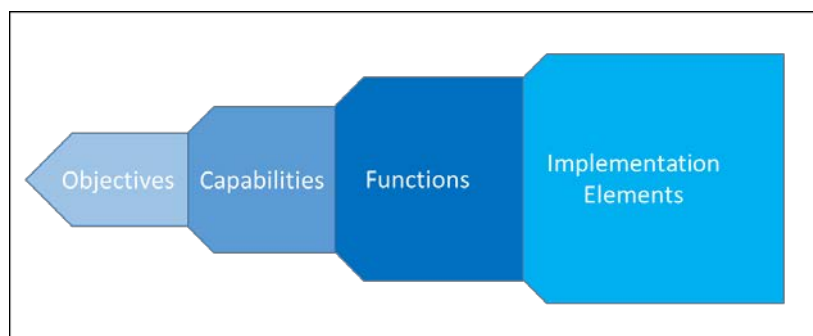


Figure 10. Functional Description Stages

In this guidance, we will describe aspects of the first three stages, but not the last, as that transitions from architecture into design. First however, we will list a set of key architectural principles needed to help establish the conceptual integrity of any NASPInet 2.0 architecture.

12.1 Architectural Principles for NASPInet 2.0

The following is a short list of guiding principles specific to NASPInet 2.0:

1. Use leading communication network practice. Do not presume that PMU networks are somehow so different from other information networks so that these practices and principles are not applicable.
2. Do not unnecessarily constrain NASPInet 2.0 architectures with structures that originate in organizational boundaries or functional/application overlays.
3. Use known resilient architectural structures (layering, platforms, core-edge structure, etc.) to provide functional flexibility and future-proofing. Employ redundancy/critical component backup to eliminate single points of failure.

4. Use modularity principles, namely module strength and module (de)coupling, to limit propagation of cascading failures.
5. Minimize dependencies to limit system brittleness (anti-resilience)¹ and performance compromise (avoid cascading components, for example).
6. Address cyber security systematically and comprehensively.
7. Enable graceful degradation² and fault tolerance³ in the event of component or system failures.
8. Use international open standards to maximize interoperability, simplify system integration, and avoid vendor lock-in.

12.2 Objective of a NASPInet 2.0 Implementation

The primary objective of a NASPInet 2.0 implementation is to provide the synchronized measurement-based *observability platform for electric power systems*. Since present power systems are primarily AC-based, synchronized phasor measurements or elements of such measurements (such as relative phase angles) may be used to meet this objective, in support of applications including protection and control, system modeling, fault management, asset management (life cycle and utilization), and cyber-physical security.

In accordance with the architectural definition for platforms described earlier in the guidance, this means that a NASPInet 2.0 implementation does not include the *processing* of PMU data. That is the role of PMU applications, and that role cannot be fully known because PMU applications are still undergoing a good deal of evolution and validation. In accordance with the core-edge architecture principle, applications should *exist at the edge of the NASPInet, not internal to it*.

That said, a NASPInet 2.0 design should recognize presently understood and anticipated applications in terms of data usage, latency requirements, system location, etc., in order to select structures meet presently understood requirements while future-proofing investment in the NASPInet 2.0 platform. This guidance will inform the structure for the platform, including how applications connect to it, but will not specify the applications or means to process PMU data in applications. That is for the application architects and developers to specify.

12.3 Capabilities of a NASPInet 2.0 Implementation

A NASPInet 2.0 implementation must have the following set of capabilities to meet the objective described above.

1. Data acquisition – the sensing and measurement of physical quantities (voltage and current), coupled with the determination of local computation of relevant mathematical representations and derivations (synchrophasors, system frequency, ROCOF, harmonics, etc.)

¹ A brittle system is a system characterized by a sudden and steep decline in performance as the system state changes. A system is brittle if it is unable to effectively absorb stress-induced shock.

² Graceful degradation is the ability of a system or network to maintain limited functionality even when a portion of it has been rendered inoperative. In graceful degradation, the operating efficiency or speed declines gradually as an increasing number of components fail. It is somewhat analogous to ductile failure in materials.

³ Fault tolerance is the property that enables a system to continue operating properly in the event of the failure of (or one or more faults within) some of its components. A highly fault-tolerant system may continue at the same level of performance even though one or more components have failed.

2. Data transport – moving data from the sensor location to any of many possible destinations where the data will be stored or consumed; also moving data from storage to consumption point; secondarily transporting meta-data, device and system management data, commands, configuration data and commands, performance monitoring data and cyber security information as needed
3. Data synchronization – ability to synchronize data acquired at differing locations so that valid phasor arithmetic and other calculations may be carried out and so that coherent views of system state may be assembled
4. Data integration – persistence of data over various time scales as determined by application requirements; means to make data sets and streams available to applications in a timely fashion and via mechanisms (interfaces) that are standardized and open to any authorized device, system, or user
5. Data curation/quality assurance – means to organize and manage the life cycle of data and to determine the quality of the data (correctness of representation of the physical or system variables it is intended to represent or fitness for use in the intended applications), prevention of accidental or careless loss
6. Data security – means to protect data whether in motion, in storage, or in use, from compromise via cyber-attack; also means to similarly protect devices, systems, and applications
7. Device and network management – means to administer/manage devices and the communication system; the ISO Telecommunications model for this is referred to as FCAPS (Fault, Configuration, Administration, Performance, Security) which are the management categories defined in the ISO model (for billing organizations administration may be replaced by accounting)

12.4 Function Classes for NASPInet 2.0 Implementations

Functions are the set of processes or actions that support capabilities needed by the system to meet the objectives set for it. The capabilities listed above decompose into a set of function as outlined in Table 3. See Appendix 3 for definitions of the function classes.

Table 3. NASPInet 2.0 Capabilities and Functions

Capability	Functions
Data acquisition	Sensing (transduction), sampling and discretization, compensation, units conversion, signal processing computations (synchrophasors, frequency, ROCOF, harmonics, etc.), data formatting/packaging (for transport), timestamping
Data transport	Routing and switching, data flow management, device/system interface to communication networks, protocol execution
Data synchronization	Timing source interface, timing distribution, timing synchronism, time stamping
Data integration	Buffering, aggregation, time alignment, conversion, databasing, management of data in motion (streaming), data presentation to applications; signal registry functions; device meta-data management; data sharing across organizations
Data curation/quality assurance	Data source quality analysis; signal validation, editing & estimation; signal correction, archiving, data life cycle management
Data security	Data integrity, data confidentiality and privacy, device/system integrity, access control; data accessibility assurance
Device and network management	FCAPS, incl. PMU configuration, control, software/firmware update, remote diagnostics; network management, network policy management

13.0 NASPInet 2.0 Components and Structures

System architectures consist of three types of elements: black box components, structures, and externally visible characteristics. Functions are allocated to the component and structures as needed. This section outlines the component classes (types) and structures and shows mappings of capabilities and functions to the structures.

13.1 Key Component Classes for NASPInet 2.0 Architectures

Components are the “boxes” in a block diagram. At the architectural level, these are black boxes, meaning we are not concerned with the internal implementation – that is for the developers; architecture is agnostic about technologies, products, and tools. Consequently, we categorize components into component classes that are characterized by externally visible characteristics, such as functions and performance. In an actual architectural specification the component classes would each have a definition document, but here we only provide short descriptions of some selected key component classes.

Table 4. NASPInet 2.0 Component Classes

Component Class	Description
Sensor	Device for creating an electronic representation of a physical variable.
transducer	Device that converts a physical variable variation into an electrical signal that varies as the physical variable does.
line sensor	Device that takes one or more transducer outputs and converts them to a form (usually digital) that can be recorded or transmitted. May locally computed derived quantities from the basic signal or signals.
PMU	Phasor Measurement Unit – a sensor that converts specific transducer outputs into digital phasor form using a common time source to create synchrophasors that can be recorded or transmitted.
Communication network	A connected set of elements that can transport signals or digital data from one or more sources to one or more destinations.
router	Mechanism that determines paths through a network from source to destination.
switch	Mechanism that relays data packets from an incoming channel to an outgoing channel so as to follow a path network determined by a router.
hub	Device that provides a common connection point in a network. Incoming packets are copied to all other ports.
network appliance	A specialized device attached to a network that is optimized to provide some specific function or service to the network or network users.
(edge) nodes	Devices and systems connected to the network so as to use the data transport services, including phones, tablets, servers, computers, control devices, and PMUs. See the discussion on core/edge structure above.
physical media	Physical element that carries the lowest level electrical or electromagnetic signals from one device to another (e.g. coax cable, optical fiber, radio channel, etc.)
network management system	Computer system that identifies, configures, monitors, updates and troubleshoots network devices, and displays status and performance data.

Component Class	Description
network policy management system	Central repository for rules and services where rules define the criteria for resource access and usage.
Data persistence device	Mechanism for storing data and providing means for retrieval
buffer	Mechanism for storing small amount of data, typically for short periods of time
database	A data store that includes mechanisms for data insertion, structure and organization, and retrieval
stream (data in motion)	Very short time scale storage of data as it moves through a communication network or processing system; persistence is incidental to the main purpose but can be treated as storage for stream-type processing (“data in the pipe”).
Cyber security system	A collection of related or interacting elements intended to implement cyber security measures in digital communications networks.
firewall	A network security device that monitors and controls incoming and outgoing network traffic based on predetermined security rules.
SIEM	Security Information and Event Management – system that combines security information management with security event management functions to provide real-time analysis of security alerts generated by applications and network hardware.
security appliance	A specialized device capable of hosting or implementing network security functions.
key manager	A system for generating, distributing and managing cryptographic keys for devices and applications. It may cover the secure generation of keys over the secure exchange of keys, secure key handling and storage on the client, backend functionality for key generation, distribution, and replacement and the client functionality for injecting keys, storing and managing keys on devices.
identity services engine	A system that enables the creation and enforcement of security and access policies for endpoint devices connected to a network’s routers and switches
certificate authority	An entity or system that issues digital security certificates.
encryption	Encoding of data into unrecognizable forms to protect it from access or use by unauthorized parties.
link level	Encryption between pairs of network devices.
end-to-end	Encryption between endpoint devices without relying upon network devices.
Applications server	Computer hardware and software for the development and execution of applications.
Visualization system	Application software and display devices used to help humans understand complex data and analytics.
Middleware	Originally, software that acted as a connector/translator between two computer systems (see for example, CORBA). More recently and appropriate to NASPInet 2, software that provides a set of services as a layer between applications and a foundational layer which might be a computer operating system or some physical infrastructure such as a sensor network.
Gateway	A device that functions as an entrance from one network to another. A communication network gateway is a network node, often a router. In the old NASPInet Data Bus model, Phasor Gateways were used as the means to connect any particular organizations PMU-related devices and systems to the NASPInet “Data Bus” but that model is deprecated in this Guidance.

13.2 Structures for NASPInet 2.0 Architectures

This section outlines guidance for various structural elements of a full NASPInet 2.0 architecture.

13.2.1 Observability Platform Structure

The original view of NASPInet was that it was structured like a data bus. The data bus view included multiple PDCs that could and did end up being stacked and did not take much advantage of the capabilities of modern communications networks. In this Guide, NASPInet 2.0 is viewed as *platform* rather than a data bus. The platform view lends itself well to the use of layering, modularity, and other architectural constructs that are recognized for scalability, functional flexibility, and the ability to secure futureproofing of technology investments. The concepts of layering and platforms have been discussed above; Figure 11 shows a basic layered platform for NASPInet 2.0.

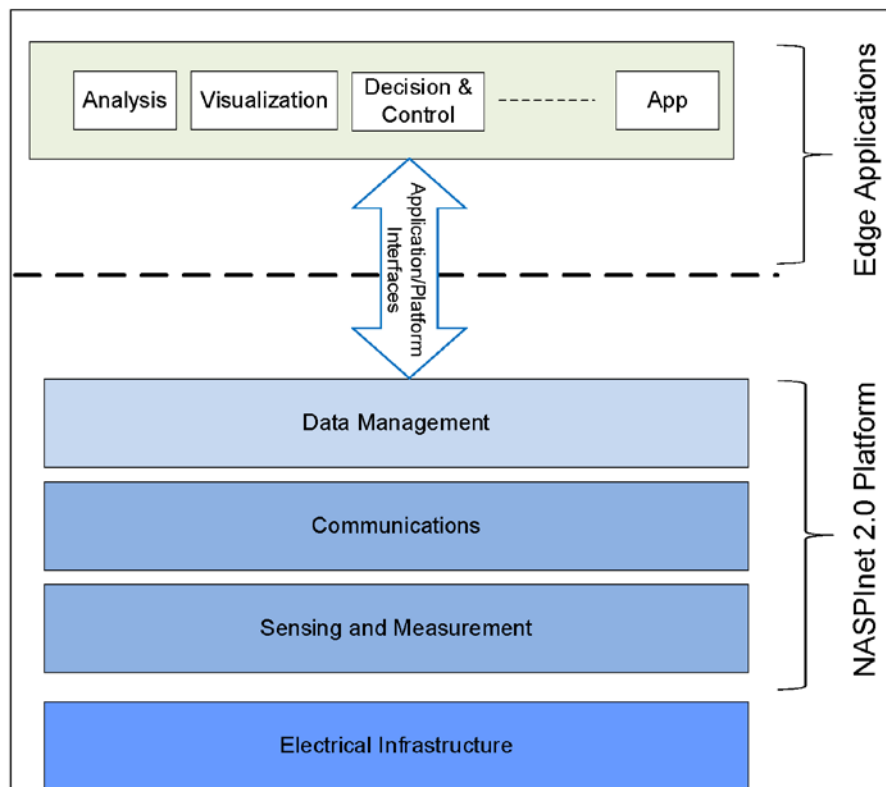


Figure 11. NASPInet 2.0 Platform Layer Model

This model contemplates three primary platform layers: sensing and measurement, communications, and data. The function groups of the seven capability classes map onto the platform layers in a manner detailed below.

The three primary NASPInet 2.0 platform layers have been stacked above the core infrastructure and the applications layer (not part of the NASPInet 2.0 platform) is depicted as well but will not be detailed. Another way to depict this structure is to use the concentric ring model shown in Figure 12. This model more clearly illustrates the core/edge aspect of the model, but as pointed out above, this is equivalent to the platform/layer model.

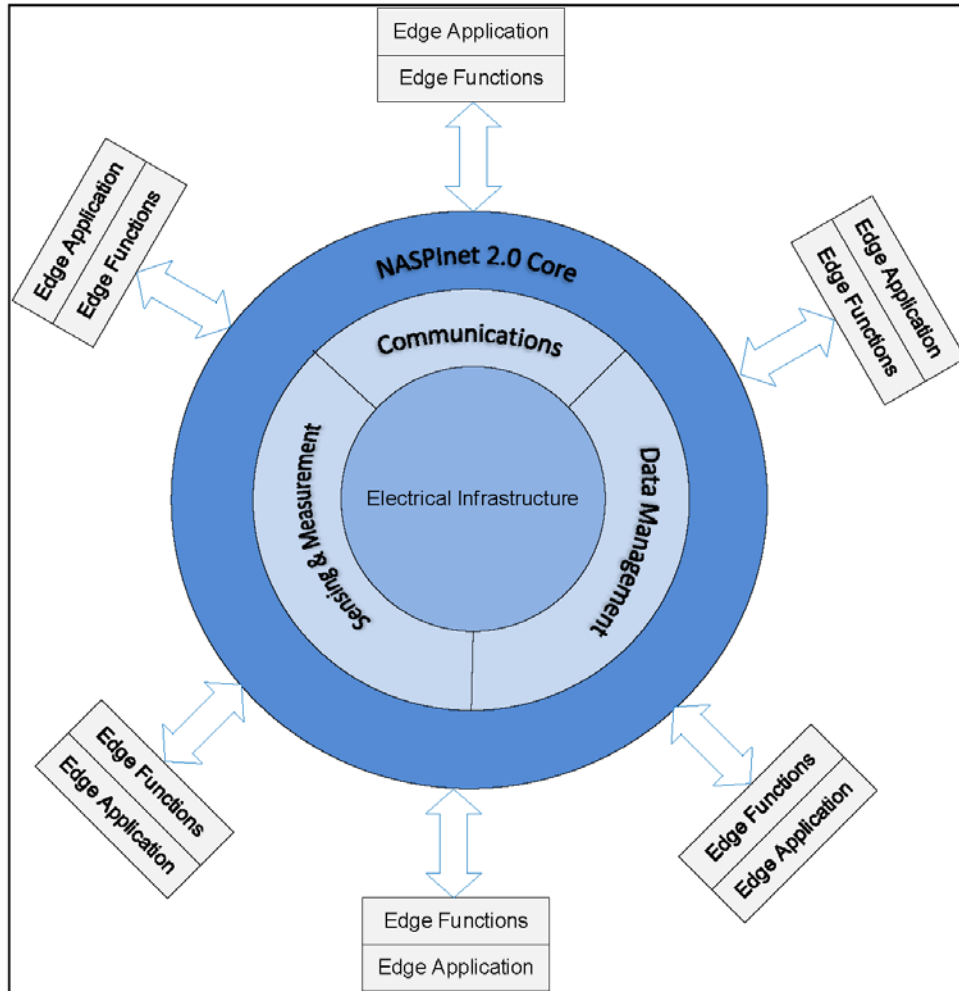


Figure 12. NASPInet 2.0 Core/Edge Ring Model

Since the use of transmission level PMUs is essentially a community enterprise (due to the interconnected nature of electric power systems), it is actually helpful to view a NASPInet 2.0 as a *distributed platform*. Figure 13 shows this concept. Each entity that participates in the PMU system owns and operate a portion of the distributed platform. Hence the structure of the platform benefits from community use of open interoperability standards as this makes it easier for an entity to join with others in the sharing of PMU data.

That does not preclude having well-defined interfaces between organizations and in fact this is to be desired for cyber security reasons. However, the boundaries and security measures can be applied at the communication network level and do not need to (nor should they) involve placing NASPInet 2.0 functional blocks such as PMU data concentration at organizational boundaries. Access control, data integrity, data security and confidentiality, and device and system integrity can all be implemented without the need to compromise the core/edge concept of NASPInet 2.0 functionality and structure. The communications layer may be several communications networks that have interconnections as shown above. In this case, each entity will manage its own network and have an agreement to interconnect with neighbors as needed. For the regional entities such as system operators and reliability coordinators, the problem becomes one of needing essentially a hub-and-spoke arrangement to each of the transmission operators in its region. In this model, both peer-to-peer and hub-and-spoke point-to-point links may

proliferate, which is not very efficient from either interconnection/integration, data flow management, or investment perspectives. The transport of data from one entity to another in this approach could be done control center-to-control center, but this introduces system latencies that can easily outweigh any network latencies. The establishment of all the necessary physical point-to-point links among the various entities is tedious and likely to lead to any given entity (especially the system operators and reliability coordinators) having to support multiple interconnection approaches.

Figure 14 shows another alternative. This alternative is a common regional communication network operated by a telecommunications service provider. This may be a publicly available network or may be a privately operated network with limited access. In either case, each entity will have its own interface to the common communication network and will still have its own internal communication network, but management of data flows on a regional basis may be done by the communication network operator. The tradeoff is that each utility entity will have less control over communication network options and operations and must contract for data transport service with the same telecommunications service provider.

This model introduces another entity into the architecture. That has impacts on the allocation of responsibilities for such functions as cyber security, data transport performance management, and data transport reliability/availability. One of the tasks to be carried out in the development of an architecture is the allocation of roles and responsibilities to entities, systems, and devices. Clearly these two structural models will differ in those allocations.

The platform concept is inherently modular in nature and implementations should maintain this modularity. Doing allows different entities to implement the portions they need while retaining compatibility with other entities. It also allows some functions to be allocated to a designated entity operating on behalf of all of the entities involved. This entity could be any of the organizations on a given NASPInet 2.0 system, since the allocated functions may be implemented as network-based services. Examples of such candidates for allocation to services are signal registry functions, data transport QoS monitoring, and some cyber security measures.

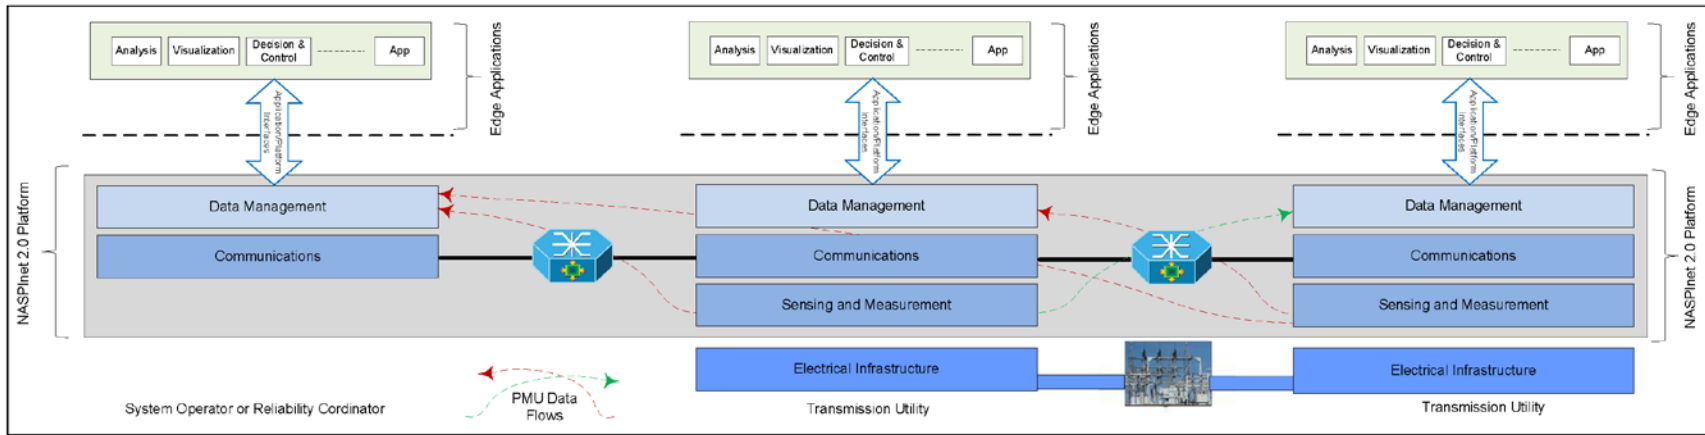


Figure 13. NASPInet 2.0 Distributed Platform Model

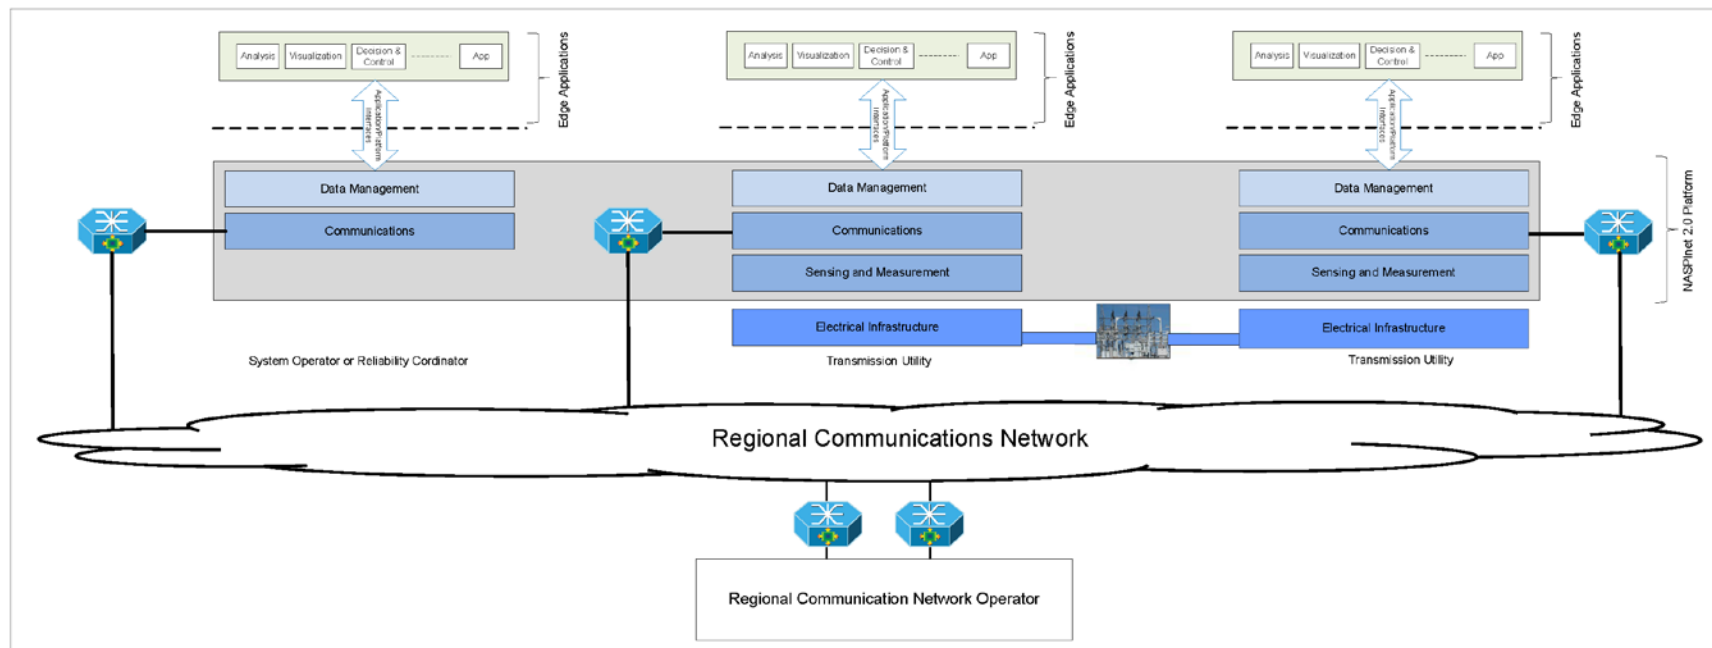


Figure 14. NASPInet 2.0 Distributed Platform with Regional Communication Network

13.2.2 Platform Functional Mappings

The function sets associated with each capability category can and should be mapped to the various layers of the NASPInet 2.0 platform. These mappings indicate how functions can be allocated to the various layers without constraining the implementations of the functions. Hence no reference is made to specific systems or products (that would be system design, not architecture). Note that some functions may have mappings to more than one layer – this is because more than one layer may have a proper role in some aspect of the function.

Figure 15 shows the Data Acquisition function mapping.

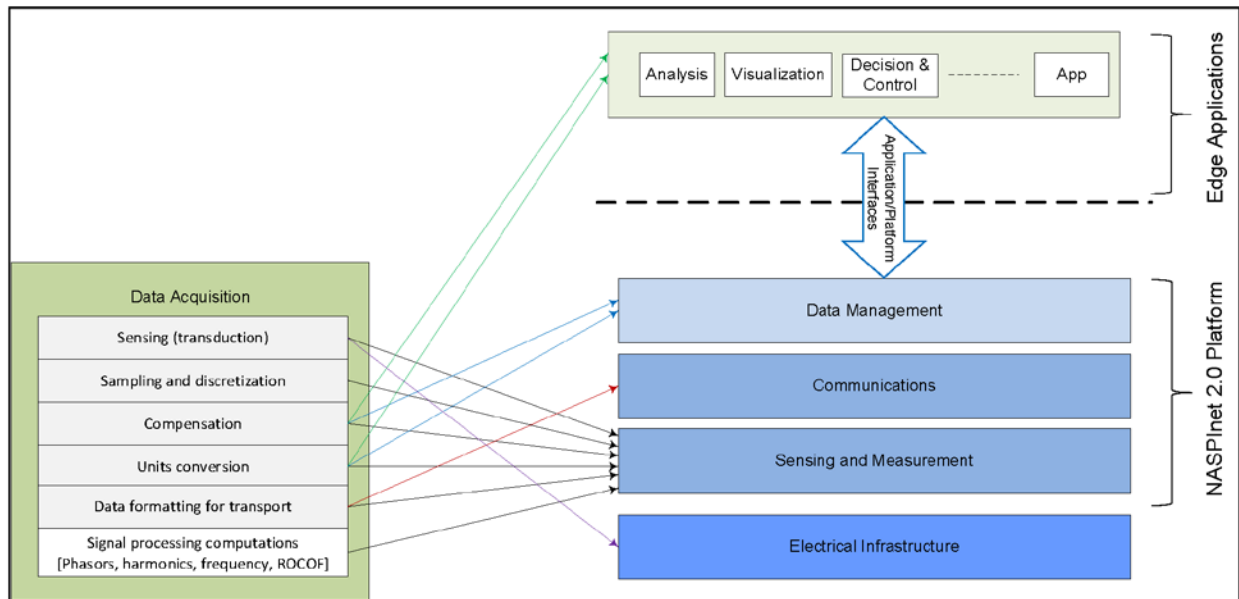


Figure 15. Data Acquisition Function Mapping

Compensation for factors affecting sensor accuracy and units conversion may be carried out in the data management layer, or it may be carried out at the application level. The latter is likely in the case of real time closed loop control, so that excess latency is not introduced by having data go through the data management layer before reaching the protection or control device or application. Note that some aspects of data formatting may actually be accomplished in the communication network.

Figure 16 shows the Data Transport function mapping.

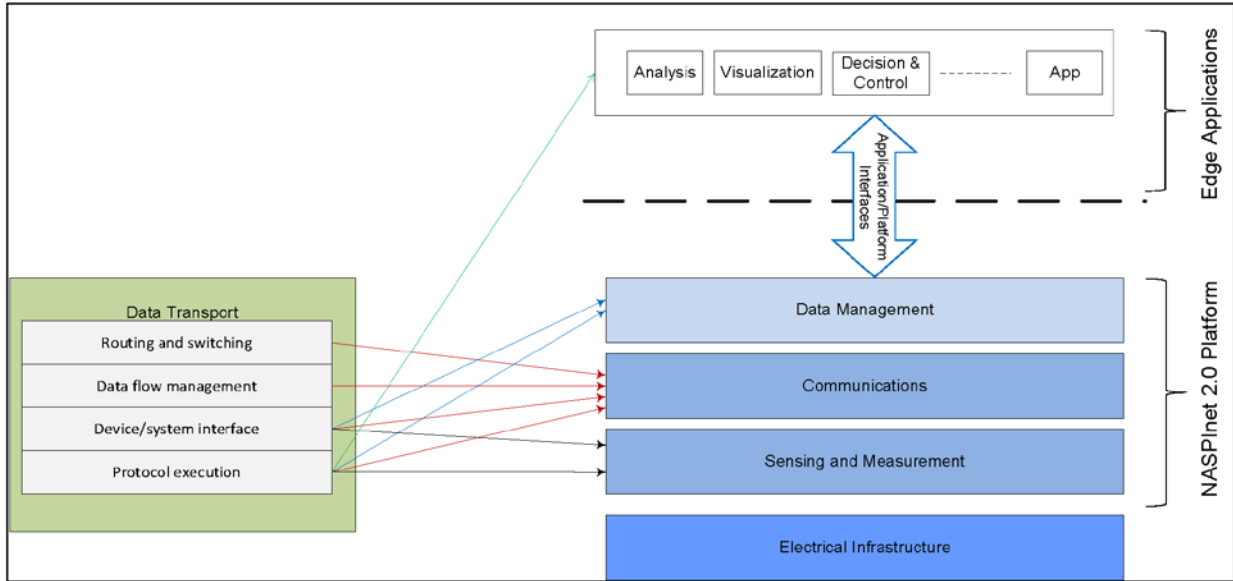


Figure 16. Data Transport Function Mapping

Data transport is clearly a Communications layer capability, but the endpoints (data sources, data storage, and data-consuming application devices/systems must also participate by supporting the appropriate communication protocols. Routing and data flow management are shown as Communications layer functions, but this does not preclude network management approaches such as Software Defined Networking (discussed later in this Guidance).

Figure 17 shows the Data Synchronization function mapping.

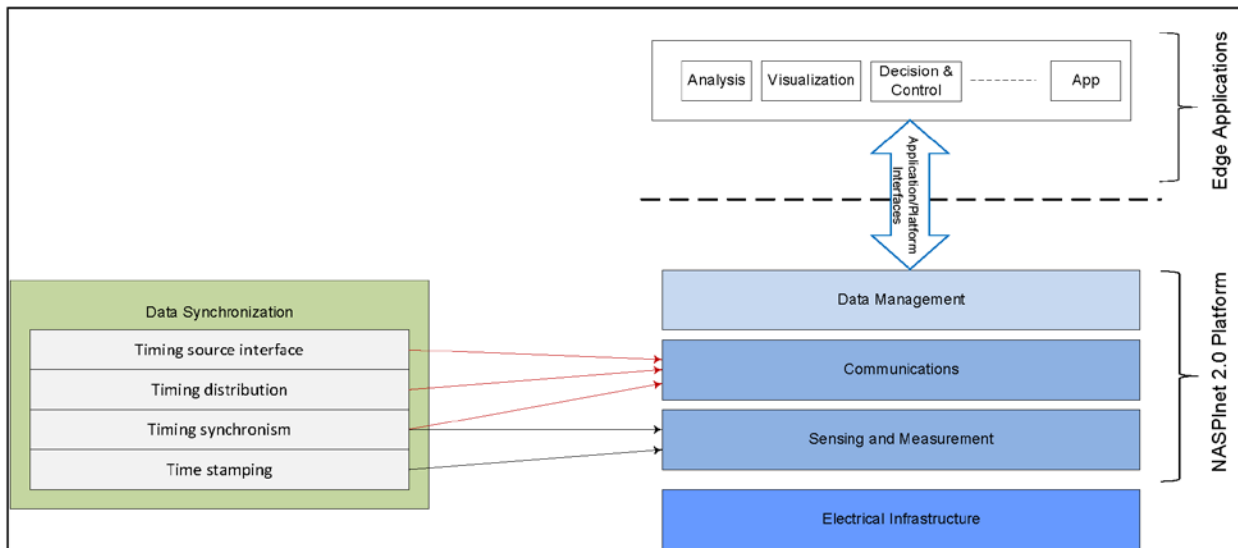


Figure 17. Data Synchronization Function Mapping

Synchronization is allocated to the Communications and Sensing layers only because alignment of signal messages (data concentration) is located in the Data Integration capability mapping below. This is done to satisfy the core/edge architectural principle and structure. Timing synchronism refers to the means by

which time (or time stamp value) is agreed upon by the various element of a NASPInet (see the section on Timing Distribution below).

Figure 18 shows the Data Integration function mapping.

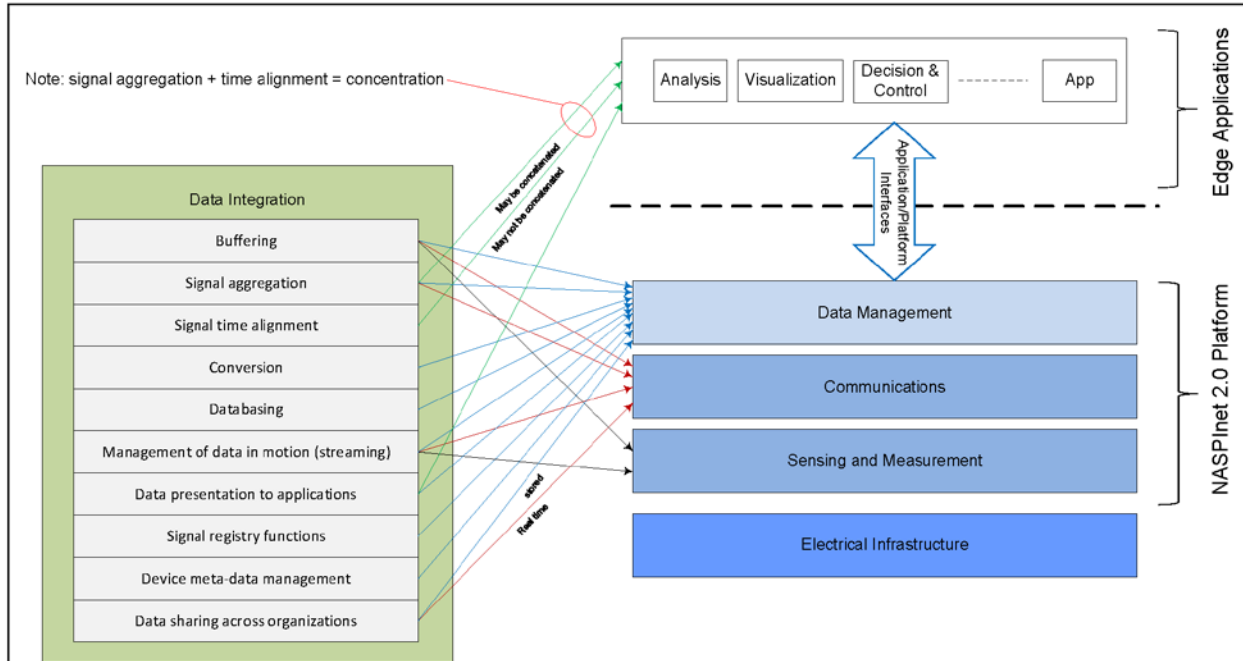


Figure 18. Data Integration Function Mapping

The bulk of functionality for the Data Integration capability resides in the Data Management layer, but note that modern communication networks are capable of supporting many data management functions and so such allocations should be considered in the development of any specific NASPInet architecture and design. This can reduce the need for Data Management layer software and also minimize latency for real time applications. Use of the Communications layer is necessary when the data persistence method involves data in motion (streaming). Also, communications networks can provide data aggregation, data access, and publish-and-subscribe capabilities for streaming data, including streams with the possibility of multiple destinations.

Note that signal alignment, by which we mean time alignment of PMU signal sample messages, is included in this capability and is mapped to the applications layer. This is consistent with the principle that synchrophasor data concentration is a function, not a device or system, and should not occur more than once per application; in other words should be confined to the edge close to the application using the data so as to avoid PDC stacking.

Figure 19 shows the Data Curation and Q/A function mapping.

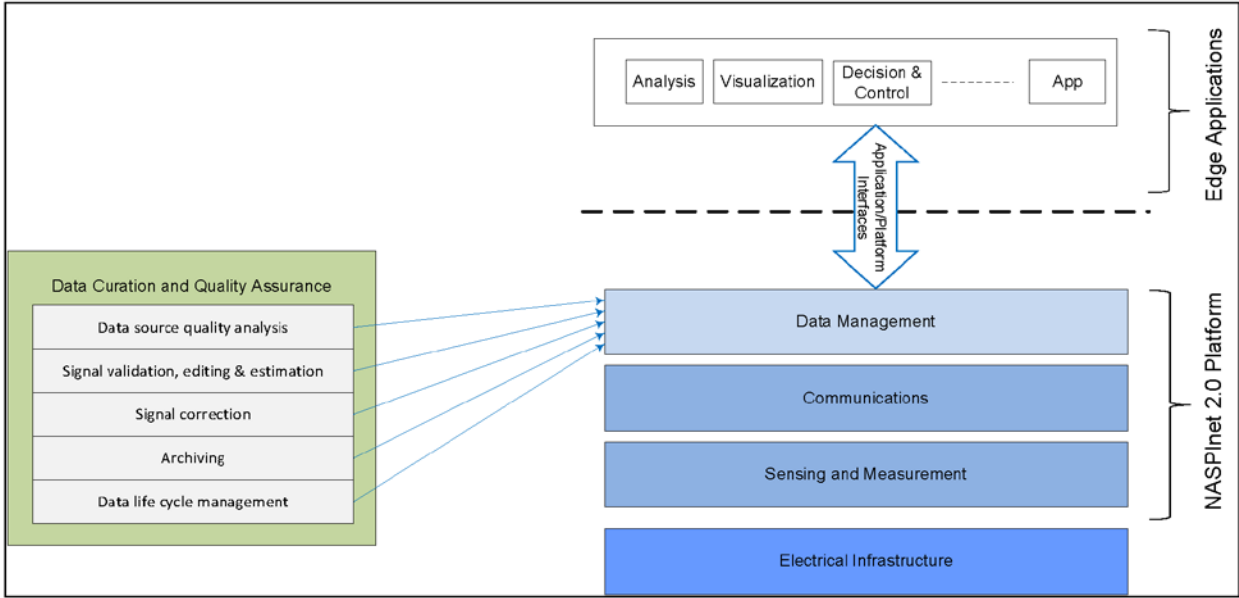


Figure 19. Data Curation and Q/A Function Mapping

Signal correction operates on a slower time scale than compensation and units conversion. Compensation and units conversion is needed on all time scales and so is included in the Data Acquisition capability, whereas correction is an ex post facto adjustment of values based on after-the-fact determination of a systematic and quantifiable error in the measurement.

Figure 20 shows the Data Security function mapping.

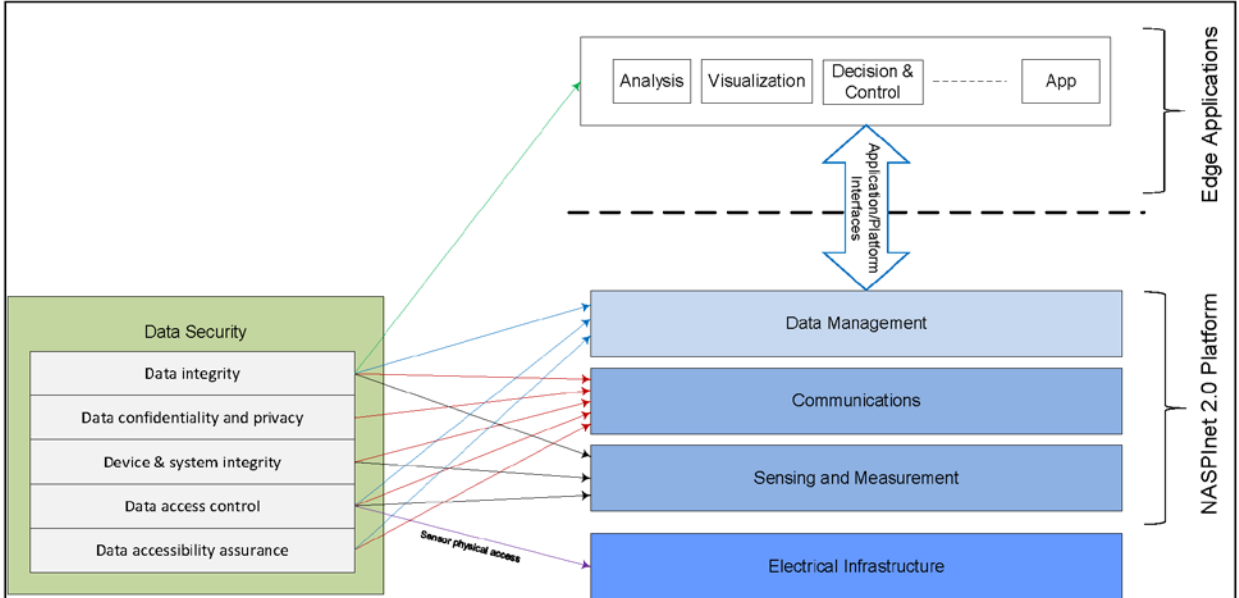


Figure 20. Data Security Function Mapping

Data security functions are generally allocated to communications networks and data management systems and devices. However, physical access is an aspect of cyber security, which is why there is an allocation to the Electrical Infrastructure layer (for six wall security at substations, for example). Also, the

application layer may have a role in maintaining data integrity since it may be necessary to use end-to-end encryption in the case where third party communication networks are involved in PMU data transport as per the NERC CIP guidelines.

Figure 21 shows the Device and Network Management function mapping.

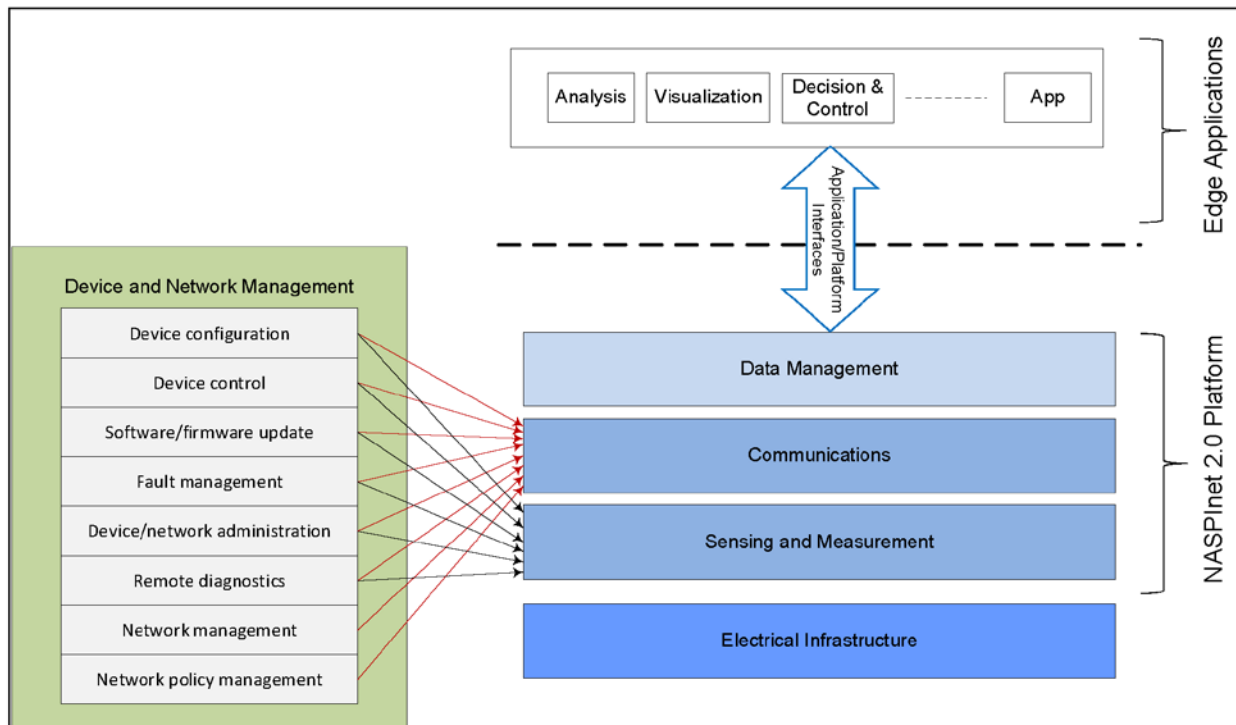


Figure 21. Device and Network Management Function Mapping

In general, FCAPS functions map to the Communications and Sensing layers. Traditionally, FCAPS has been applied to communication network devices, but here extends to the sensing and measurement devices because a) they are also network devices and b) the device management aspects relating to function as a grid sensor are very similar to standard FCAPS functions and so can be implemented separately or as part of a comprehensive approach to network management (i.e. integrated with communications FCAPS implementations).

13.2.3 Platform Mapping Summary Comments

The original approach to NASPInets was based on the concept of a data bus. That model has limitations that become evident when considering the set of systemic issues, core considerations, ARRA experience, and foundational principles listed above.

The platform structure concept described in this Guidance is widely used in information systems but has broader applicability as shown above. The platform is not limited to information systems but can and should include sensing and measurement, data management, communications, and physical infrastructure as layers. Among its advantages are that the platform model cleanly separates applications from general support capabilities and functions and provides a structure onto which those capabilities and their constituent functions can be mapped. Using this model and the appropriate mappings, development of a specific NASPInet 2.0 architecture follows directly.

13.2.4 Streaming Data Flows

Because the platform model allocates applications to the “edge” or application layer, we can group use cases according to fairly broad categories based on data flows and then consider data flow models instead of attempting to catalog every use case. The use case categories used in this Guidance are:

- I. Non-real time analysis, planning, and modeling(no time constraint)
- II. Near real time analytics, visualization, and operational decision support (seconds to minutes)
- III. Real time close loop protection and control (future: sub-second to sub-cycle time constraint)

In the following diagrams, “src” refers to the measurement data source (e.g. PMU but may be other synchronized sensors). Note that time stamping is for the time at which the measurement was *made*, not the time at which the data transport occurred. Therefore time stamping should be determined by the src device, based on when the raw waveform samples were captured. “Persist” refers to any form of data storage.

Time alignment of multiple signal streams occurs only once for each application. When aggregation and time alignment are combined, this constitutes phasor data concentration in the sense of the old PDCs, but here it is a function, not a device or system. This means that concentration may be implemented flexibly by the application, as a service, or otherwise, as long as the non-stacking constraint is observed. In general, the functions of aggregation and time alignment are separated in this Guidance, providing more flexibility in implementation and in avoiding PDC stacking.

The set of diagrams below are simplified data flow models where we are considering only the streaming PMU signal data, beginning at the sensors (src) and terminating at applications that consume the data. Clearly there are other data flows for management of the communication and sensing devices, security, etc. These are considered elsewhere in this Guidance. The diagrams illustrate several different potential approaches to data flow structure. They are not all equally good but are provided to illuminate the structural issues associated with PMU data flows. Note that the term “Persist” in the diagrams refers to any of the modes of data storage described in the Core Considerations section above except data in motion, which is treated in the “no storage” models.

Model 1: Single src; direct use and optional storage for delayed use; Use Case Categories I, II, III




Figure 22. Single src case

The single src model is straightforward but not very useful since the point of synchronization is to be able to combine data taken from multiple points on a geographically dispersed dynamic system. It is provided here as a baseline to illustrate streaming of PMU data to multiple destinations that will use the data on different time scales. An implication of this model and similar but more complex ones that follow is that the data typically will have multiple uses and therefore multiple destinations. This can include multiple direct applications (only one is shown but it could be many) as well as storage for later use. This concept has implications for communications network structure.

Model 2: Multiple src; aggregation and time alignment combined (data concentration); no storage; Use Case Categories II, III




Figure 23. Multi-src no storage case

Persistence may exist in the form of data in motion, but no explicit databases or other buffers are used. Note that aggregation could be accomplished in the data management layer or by the consuming application. This model is useful for real time applications but in practice the PMU data will almost certainly also be stored for other uses as well. This is one reason why it is important to consider how to structure communications to support streaming to multiple destination simultaneously. Consideration of the mix of latency requirements that may apply to any given data stream leads to an architectural systemic issue (and therefore eventually a design requirement) to handle the multi-destination data streaming problem with as little latency as possible, meaning at the communication network level.

Model 3: Multiple src; mandatory data storage followed by data concentration during usage; Use Case Categories I, II

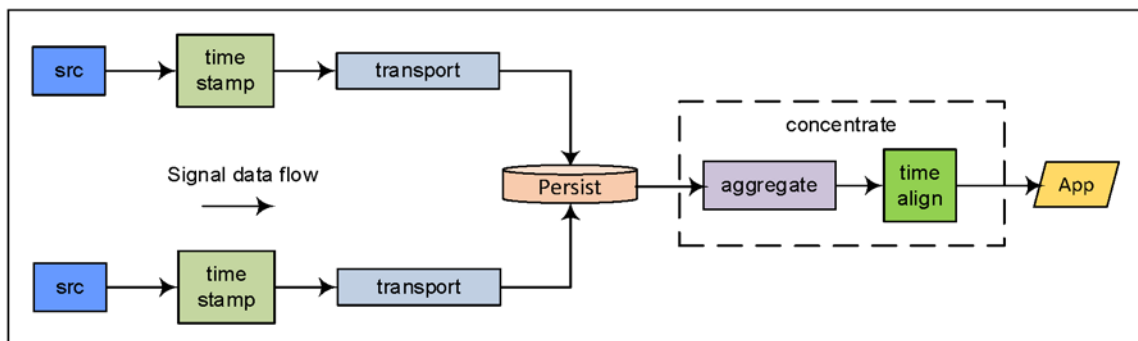


Figure 24. Multi-src mandatory storage case

Model 3 always employs data storage and all applications get their data from storage, not from raw signal streams. For any particular application the concentration function will immediately precede the application processing. This model presumes persistence is in the form of a database, so Use Case Category III is not listed for this model due to the latency that would be incurred by passing through storage. Some types of storage technologies may permit a “pass-through” capability, but this can be described in the next model. This model is intended to address the case where all data goes into storage and must be retrieved from it for use by any application.

Model 4: Multiple src; data storage after aggregation but before alignment; Use Case Categories I, II, III

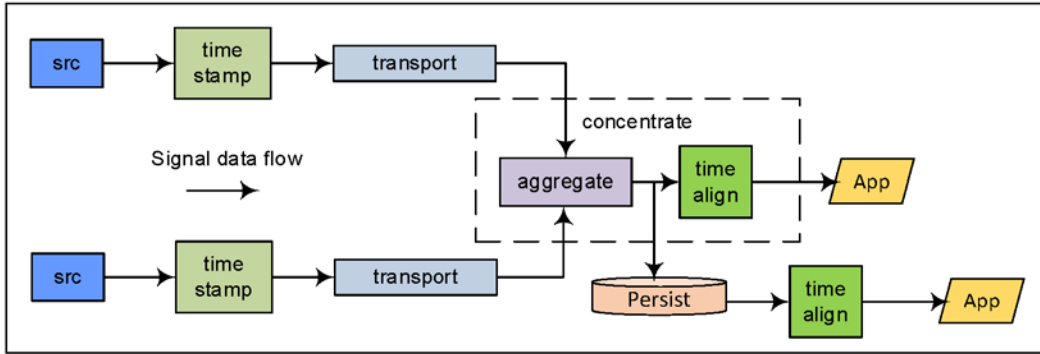


Figure 25. Multi-src mandatory aggregation optional storage case

Model 4 can store data only after aggregation by sending to multiple destinations, one of which is storage. The stored data is aggregated but not time-aligned. A disadvantage of this approach is that data may need to be aggregated differently eventual for non-real time applications, meaning that where multiple sets of aggregated signals exist, new applications may require different sub or super sets of the aggregated data, potentially making it necessary to disaggregate and then re-aggregate if this approach is used. If the streams are aggregated for transport purposes but the signals are stored individually, this problem is eliminated and after the fact aggregation can be done as needed per application.

Model 5: Multiple src; optional data storage after aggregation and alignment; Use Case Categories I, II, III

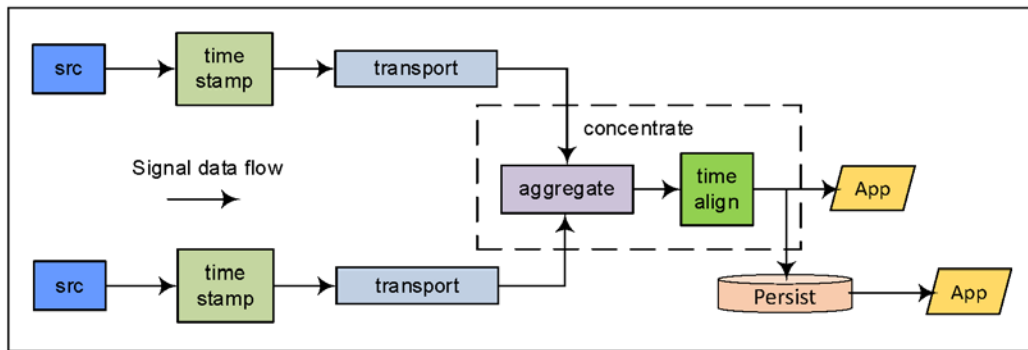


Figure 26. Multi-src mandatory concentration case

Model 5 stores data only after aggregation and time-alignment (concentration). A disadvantage of this approach is that data may need to be synchronized with other data not aggregated at the time of storage for non-real time applications. This approach may be efficient but can pose the problem described above: new applications may require different sub or super sets of the aggregated data, potentially making it necessary to disaggregate and then re-aggregate to perform the new time-alignment. While the post processing is certainly possible, the work to re-process will have to be done for each new Category I application. At the present state of development of PMU applications and uses, it seems difficult to envision how much effort that could eventually entail.

**Model 6: Multiple src; optional data storage after aggregation *in the network* but before alignment;
Use Cases I, II, III**




Figure 27. Multi-src network aggregation case

Model 6 aggregates data *in the network* and stores data optionally only after aggregation. Aggregation possibilities are determined by communication network structure. In practice, most communication networks have aggregation points for their own reasons – these may or may not be useful for NASPInet 2.0 purposes.

13.2.5 Streaming Data Flow Summary Comments

The choice of streaming data flow models is an important early architectural decision that has consequent impacts on much of the rest of the NASPInet architecture. Some of the foregoing models (notably Cases 3 and 5) are likely not suitable, but the others may be useful separately or in combination. The NASPInet architect should develop a streaming data set flow model or models early on (driven by systemic issues, some of which are the Use Case Categories) and then test them in the context of the rest of the architecture (when it has been developed) against representative use cases as a means of validation. This approach is much more manageable than trying to assemble a catalog of uses cases to drive the architecture because the set of use cases is evolving and so cannot be completely catalogued.

The architectural specification may use a single comprehensive flow model or an ensemble of models, depending on how the architect wishes to proceed. In either case, these models set important structural bounds on the entire NASPInet. The physical communication systems must support these logical data flows and many others, so data flow models are important inputs to the communications layer architecture. In general, models 2, 4, and 6 are preferable to the others for the reason stated above.

13.2.6 Communication Network Structures

Communications networks are generally treated on two levels: logical data flows and physical networks. The physical networks must be capable of supporting the entire ensemble of logical data flows, as well as having various desirable characteristics such as scalability. They may also provide support for other platform capabilities, including functions such as data flow management and cyber security.

While it is common to discuss utility communications networks as if there are three tiers (WAN, FAN, NAN), in fact it is useful to recognize a more granular model involving Wide Area Networks (WANs), substation internal networks, control or data center internal networks, and distribution level networks.

13.2.7 Wide Area Networks (WANs)

In most cases, transmission utilities and other organization such as system operators and reliability coordinators will use existing communication networks. When a network must be created however, the structure shown in Figure 28 is a useful prototype. It is based on multiple levels of ring aggregation, but also provides for linear aggregation (such as with dedicated microwave links). This structure is highly scalable and provides redundant pathing in order to be resilient against link or device failures.

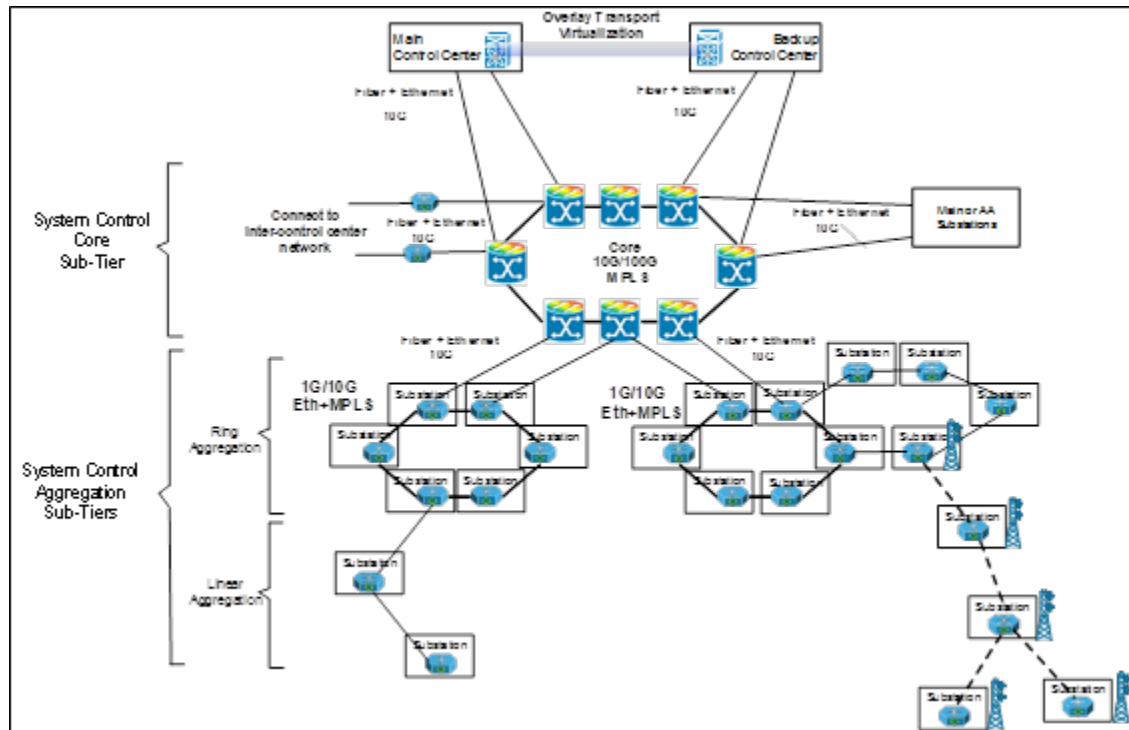


Figure 28. Wide Area Network Multi-Aggregation Ring Structure

This structure may be used in part or in whole, as it essentially encompasses many sub-structures that have been used for wide area networking. Such a network is often depicted as a cloud.

13.2.8 Multiple Communication Network PMU Systems

When PMU data must be shared across utility organizations, each may have its own WAN, as depicted in Figure 29. In such a case, the two WANs may be connected via border routers that provide the point of interconnection and any related network services.

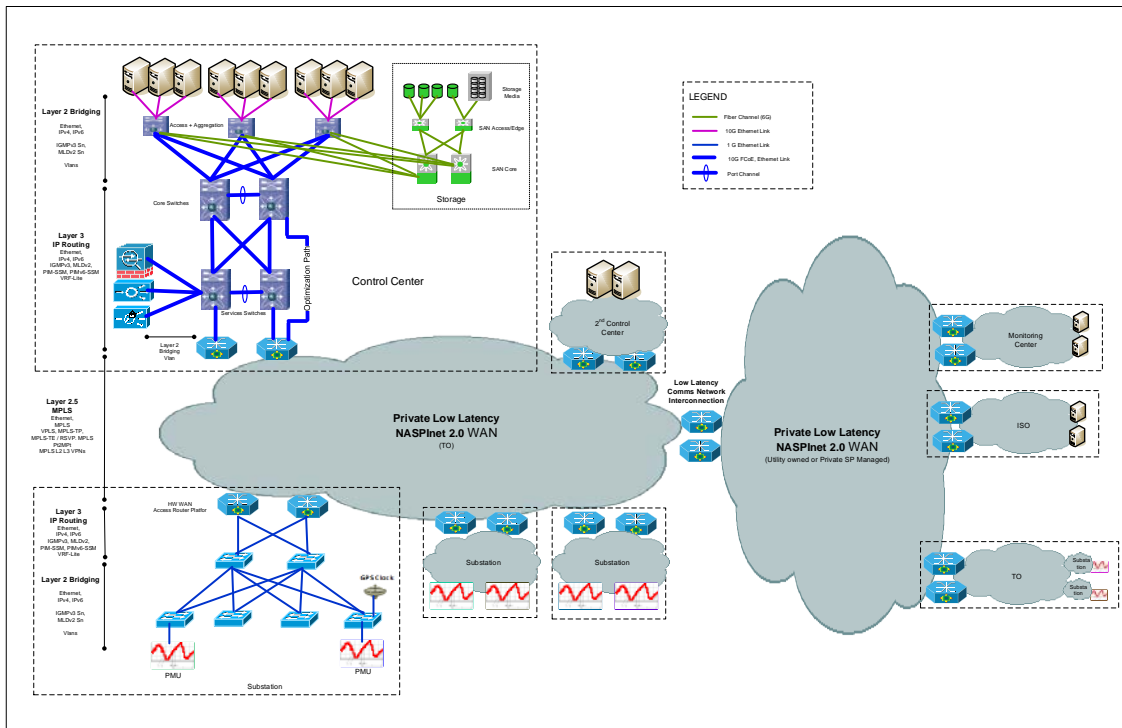


Figure 29. PMU Networking Across Multiple Communication Networks

Figure 30 shows an example of PMU data flow from one utility NASPInet 2.0 WAN to another that connects multiple entities (such as a reliability coordinator, an ISO, and another transmission utility). Note that the same PMU signal stream may have multiple destinations inside and external to the source utility's WAN. This points to the usefulness and efficiency of communication protocols that can manage multi-destination flows, as opposed to the use of over-the-top networking of concentrators and gateways and point-to-point logical connections.

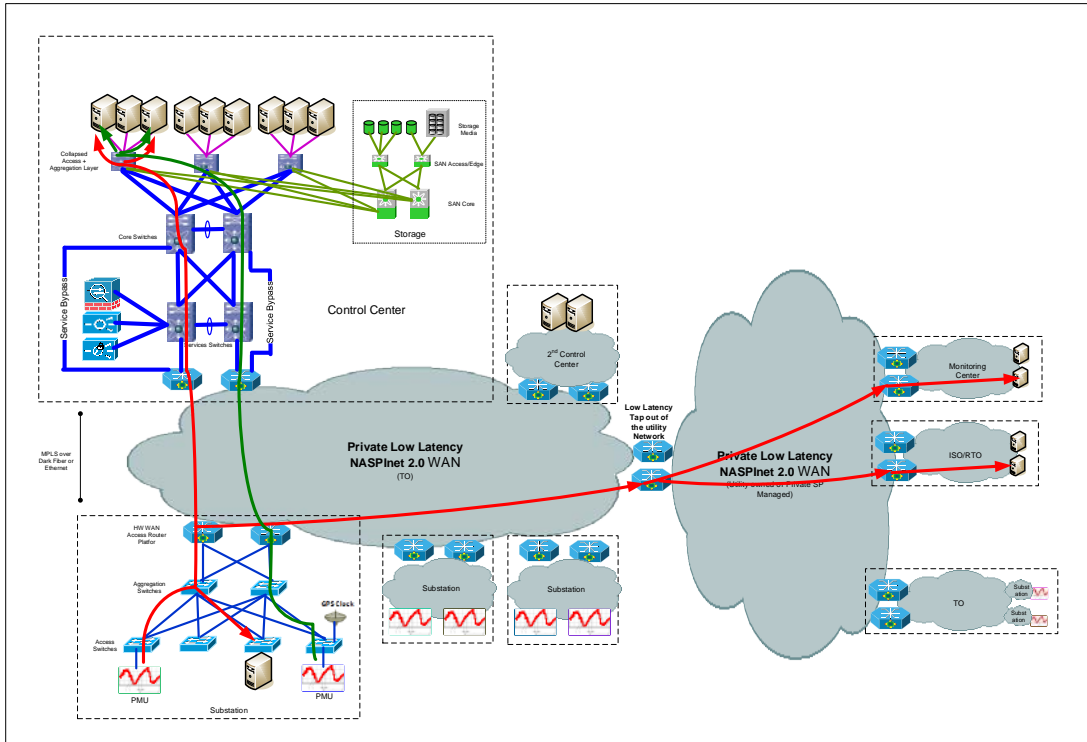


Figure 30. Example PMU Data Flows Across Multiple Networks

13.2.9 Substations

Since transmission-level PMUs reside inside transmission substations and PMU data may be used inside substations, NASPInet 2.0 must consider intra-substation networking. Once again, many utilities will use existing substation networks, but here there may be more flexibility in designing or restructuring communications. Figure 31 shows a simplified model for intra-substation communications networking. In this case, the network has a high degree of redundancy and makes use of tree structure rather than rings for high performance (low latency), fast fault recovery, and high resilience. This arrangement has internal redundancy and uses redundant connections to the WAN.

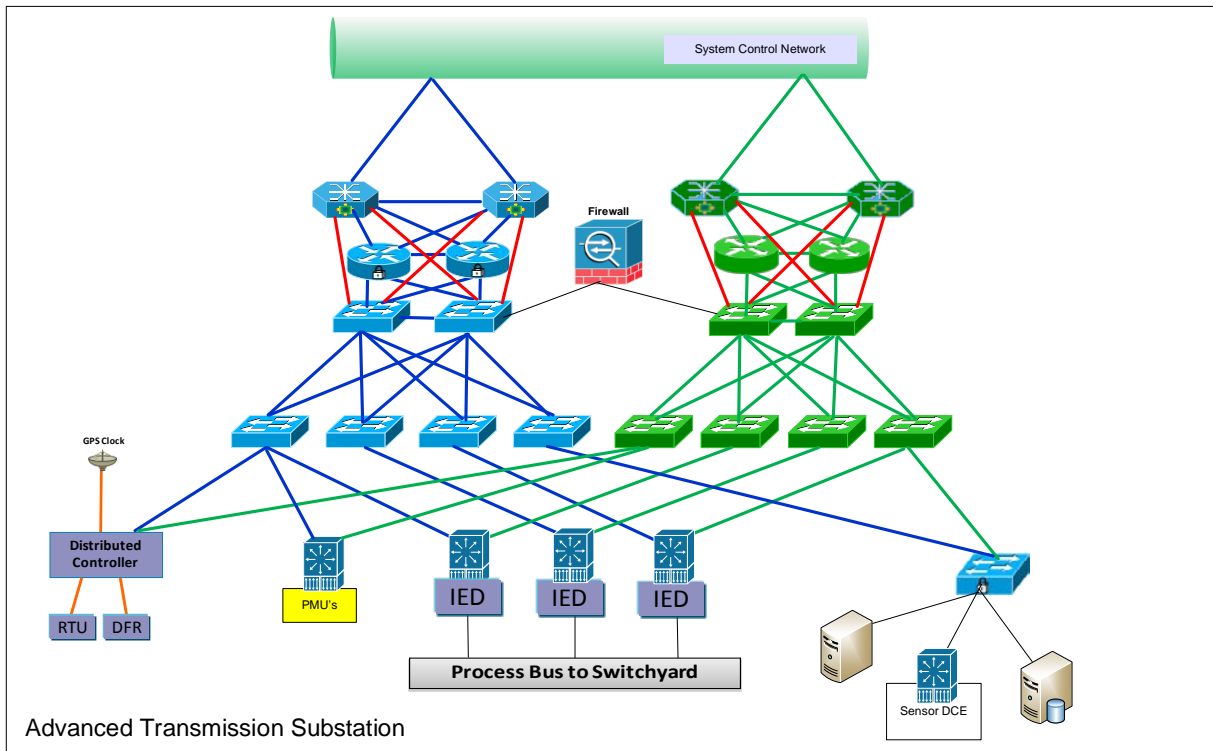


Figure 31. Transmission Substation Communication Network

13.2.10 Control/Data Centers

Just as with substations, it is necessary to consider intra-control center (or data center) communication network structure. Once again, most organization will already have networks but may have flexibility and/or need to restructure. Figure 32 shows a simplified structure that provides for efficient and redundant connection of PMU data flows into the center, with provision for real time flow to applications, as well as data flow to/from a storage sub-network.

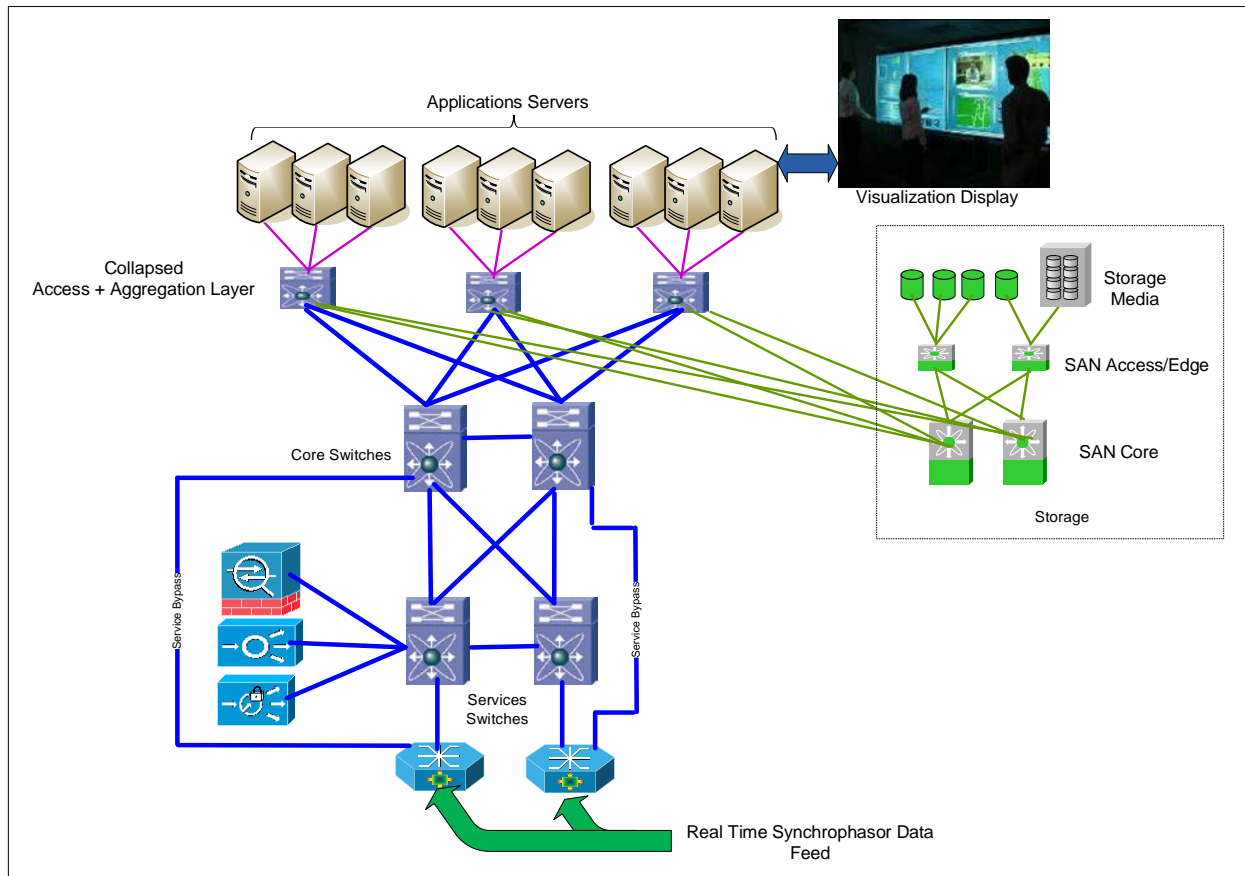


Figure 32. Control/Data Center PMU Network Structure

13.2.11 Network Services

A variety of communication network management and related services must be provide and hosted, and a key architectural issue is the determination of where and who does this hosting. For this purpose it is necessary to have a proper model of industry structure as described in the section on Problem Domain Reference Models above and to have an allocation of roles and responsibilities related to synchronized data management in a NASPInet 2.0 architecture. This allocation is an architecture level task and deliverable item. The question is complex when some functionality such as data transport is placed partially or completely in the hands of a non-utility third party such as a regional telecommunications service provider. The architecture must provide clarity of roles for telecommunications network management and network services hosting functions.

Distribution Systems

Various use cases and potential applications for PMUs on electric distribution systems have been identified.¹ In some approaches, either amplitude-only or phase angle-only measurements are deemed sufficient as opposed to full phasors, but in any event the measurements must be synchronized and for many distribution-level applications, the expected requirement for phase angle precision and accuracy is far more stringent than for most transmission applications.

¹ NASPI DNMTT, Synchrophasor Monitoring for Distribution Systems: Technical Foundations and Applications, January 2018, available online: <https://www.naspi.org/node/688>

Distribution systems pose some special concerns related to PMU (or phase angle sensing unit) locations, data transport, and location of application functions. Unlike transmission systems, distribution applications may require that sensing be located at various places on the distribution feeders, not just in substations. Furthermore, where the concept of distributed intelligence is being pursued, it may be the case that synchronized sensor data processing make take place in other than control or data centers – specifically in substations or at grid devices located outside of substations. Transport of measurement data in such cases can be complicated by the structure of distribution system communications. Referring to Figure 33, it may be the case that the distribution utility has multiple disjoint communication systems, some of which may be convenient for providing sensor connectivity but may not have the necessary performance capabilities or even the right connectivity to deliver sensor data to necessary destinations. At some utilities, communications to substations is separate from communication to distribution feeder level devices and the systems at the control center(s) may be separated for cyber security purposes. Some electric utilities do not allow substation-to-substation communications for cyber security reasons.

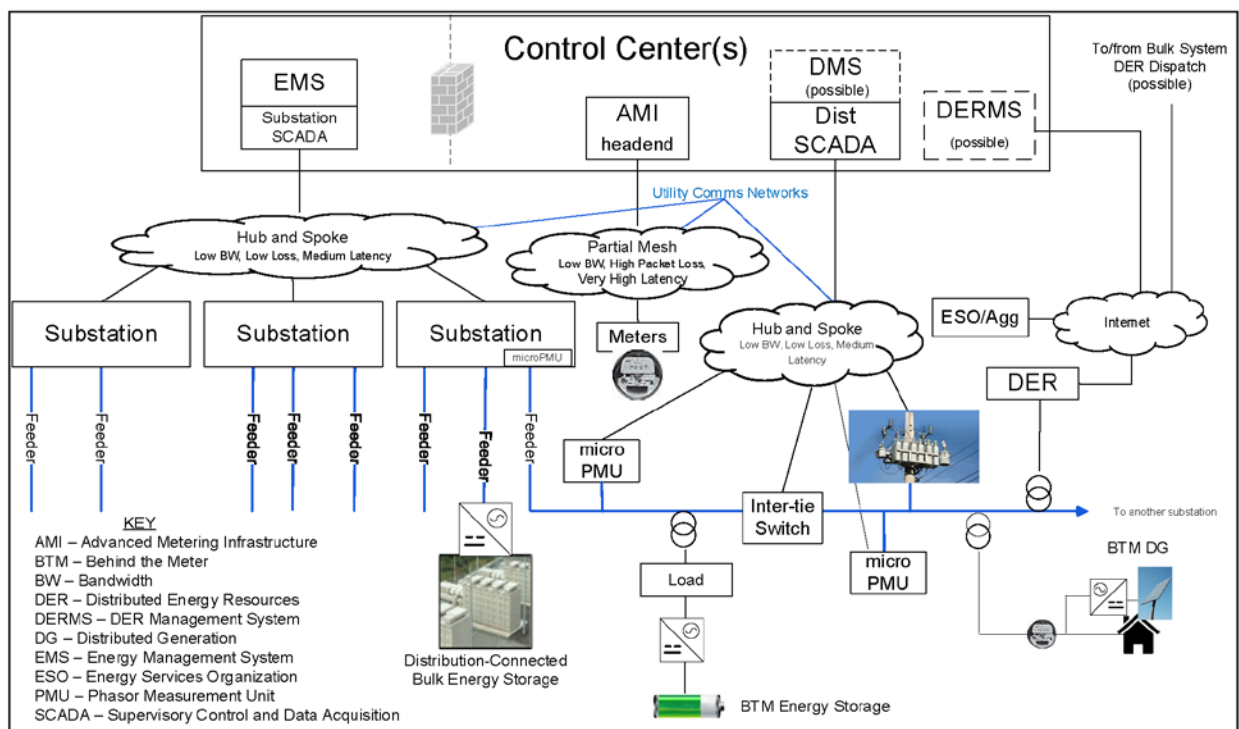


Figure 33. Distribution System Structure with microPMUs

Finally, if the distribution utility is using a telecommunication service provide to transport sensor data, it may be the case that the actual network path traversed by the data is unknown to the utility and is variable due to the manner in which the service provider operates its communication network.¹ This means that latency may be unpredictable, variable, and can be large if the service provider gateways are remote from the utility sensors and control center or other data destination.

Figure 34 illustrates a different communications structure for electric distribution. In this approach, distribution sensing and communications are treated as an infrastructure layer, and the communication

¹ JD Taft, The Impact of 5G Telecommunications Technology on US Grid Modernization 2017-2025, October 2017, PNNL 27068, available online: https://gridarchitecture.pnnl.gov/media/advanced/Communications_final_GMLC.pdf. See especially Figure 2 LTE Communications.

network is a multi-services network.¹ If this communication layer is a private network owned by the electric utility, then considerable flexibility exists to arrange and manage PMU data flows because the communication network can be used as a publish-and-subscribe mechanism that automatically handles packet duplication optimally (from a network traffic standpoint) when the data stream has multiple destinations. If the communication network is owned and operated by a third party such as a telecommunications service provider, then special considerations may come into play. For example, if the communication network is wireless (such as an LTE network) then a capability known as “hairpinning” may be needed at the towers to direct data flows locally back to substations or control centers without flowing to remote packet data gateways first, as a means of limiting and controlling transport latency and jitter (latency variation).

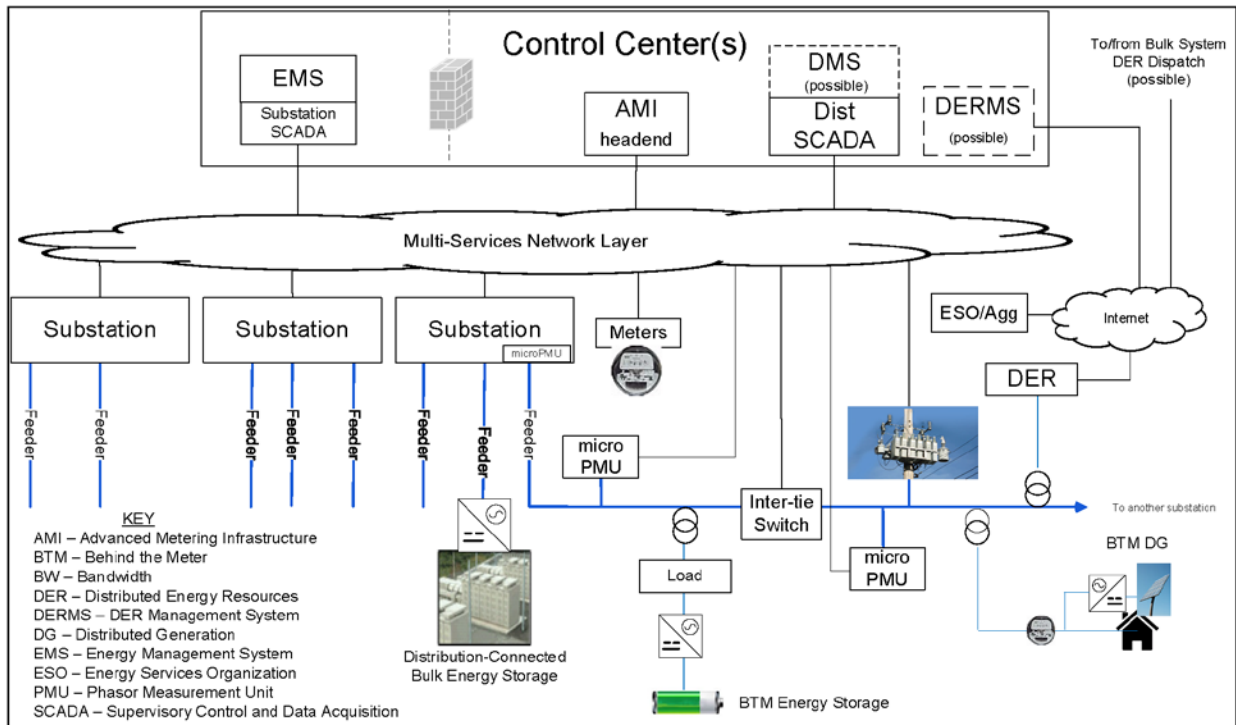


Figure 34. Distribution System with microPMUs and Multi-Services Communication Layer

While utilities may be constrained to use existing communication networks, the present (2018) trend toward distribution system modernization generally involves new investment in communications, and so the change to a sensor/communications layer structure may be practical as part of a larger effort than just the deployment of synchronized sensing. Note that the distribution level communication structure may actually be two-tier, where the upper tier is a high performance multi-services network and the lower tier is composed of some number of purpose-built networks that connect to the upper tier.

¹ JD Taft and P De Martini, Sensing and Measurement Architecture for Grid Modernization, PNNL-25249, February 2016, available online: <https://gridarchitecture.pnnl.gov/media/advanced/Sensor%20Networks%20for%20Electric%20Power%20Systems.pdf>

13.3 Timing Distribution

Synchrophasors require timing distribution or another mean of establishing synchronism (agreement among the elements of the present time or time stamp value). The typical solution for this has been to use GPS timing signals but this is recognized as being weak from a resilience standpoint. While timing may be derived from multiple satellites, the GPS system still represents a potential single point of failure, especially in the event of RF interference, space weather problems or cyber-attack.¹ Consequently, additional sources and methods of synchronization are desirable in line with the NASPInet 2.0 architecture principles.² In Figure 35, the middle illustration shows a structure in which GPS timing is combined with two other sources (cesium clock and Loran) whose timing signals are distributed through the NASPInet communication network.

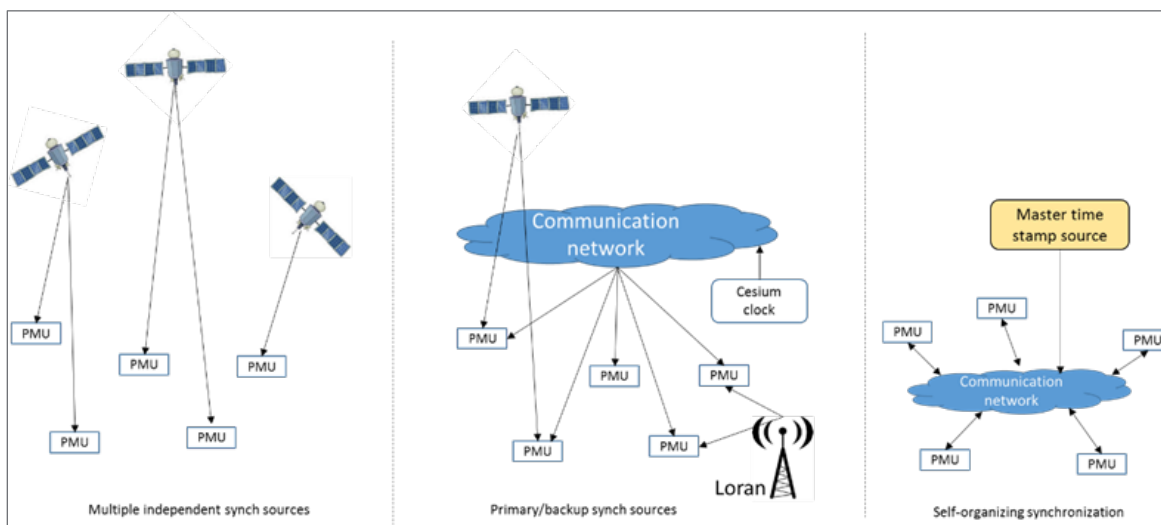


Figure 35. Structural Models for Timing Distribution

Alternatively, a sensor network such as a PMU network may be self-synchronizing³ such that sampling temporal alignment is achieved internally in the network and then external sources are only needed to inject information for time/date stamping into the network as a whole, rather than at each sampling point. Such an arrangement requires that the PMUs be able to communicate with each other, including across organization boundaries.

The distribution of timing signals via communication network requires much more than just the use of SNTP.⁴ Given the multi-organization nature of the Transmission level PMU platform, a timing distribution architecture also needs:

- Time distribution at every hop in the network
- Transparent Clocks (TCs) at every node as per C37.238

¹ M S Almas, et. al., Vulnerability of Phasor-Based WAMPAC Applications to Time Synchronizing Spoofing, available on IEEE Xplore. See also

https://www.naspi.org/sites/.../04_Vanfretti_VulnerabilityWAMPAC_20180425.pdf

² NASPI Time Synchronization task Force, Time Synchronization in the Electric Power System, NASPI-2017-TR-001, March 30, 2017, available online: <https://www.naspi.org/node/608>

³ G. Brandner, et. al., "Firefly Synchronization with Phase Rate Equalization and Its Experimental Analysis in Wireless Systems," Elsevier Computer Networks 97, June 24, 2016, pp. 74-87.

⁴ Simple Network Time Protocol. See also PTP (Precision Time Protocol – IEEE 1588).

- This clock measures delay over the node and inserts the information into the correction field in order to compensate for the variations in the propagation delay
- TCs allow also addition of slave clocks into the node, necessary in order to bring accurate timing to this node
- Ability to monitor the integrity of the GPS reference, and switch to an alternate source if compromised
- Redundant primary source of time such as military grade multi-standard Global Navigation Satellite System (GPS, GLONASS, Galileo) at key points
- Failover and recovery logic for use of timing sources
- Boundary clocks (BCs) where the networks of different organizations connect
 - BC on the sending side would have an indication what time is sent to the receiving organization
 - BC on the receiving site terminates the time and has an indication of what it receives
 - Receiving BC distributes time further in own organization
- Atomic clock as a backup source for GPS
 - With 1E-11 accuracy the commercial cesium clock can be within 1 millisecond (1E-3) accuracy for 1E8 seconds or about 1,157 days (3.17 years)
 - If 1 us is required between different systems then it can be kept by the atomic clock for 27.77 hours assuming initial time offset of 0

Figure 36 illustrates a multi-organization network architecture for timing distribution. This is not intended to be a specification or a design; it is an illustration of a timing distribution network structure.

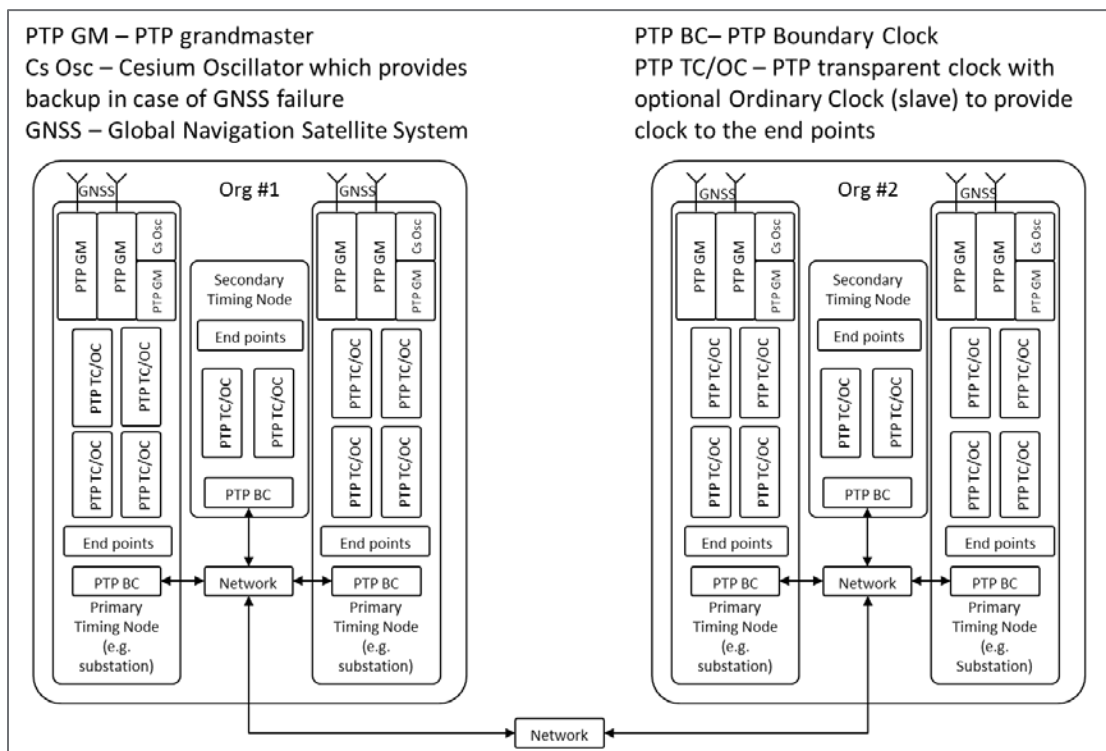


Figure 36. Timing Distribution Structure Illustration

Note: To be useful as a platform for data acquisition and control, the network must deliver precise time to applications. Common time in distributed systems is required for not just to the end devices but also throughout the distributed structure.

Note: communications networks have *very* precise internal timing signals, but a link between network timing and computing timing is currently missing. Many CPUs still rely on the SNTP time delivered via software. SNTP time can be replaced by more accurate PTP time but resolution still won't be as good as having a direct hardware timing interface in the CPU.¹

13.4 Latency in NASPInet 2.0 Networks

Latency requirements are determined by the applications that make use of the data. Evaluating latency in complex communication networks is difficult if the network is not accessible for testing, but it is possible to put a lower bound on latency with a modicum of information about the communication network. Consider the network of Figure 37. While the details of routing and switching are only determined at run time, a standard way to assess minimum latency is to assume momentarily that a path has been determined, such as shown in green. The path may be taken as static and from it the number of and types of communication elements can be tabulated. If physical media are known and propagation path lengths can be estimated, then analysis can proceed as shown below.

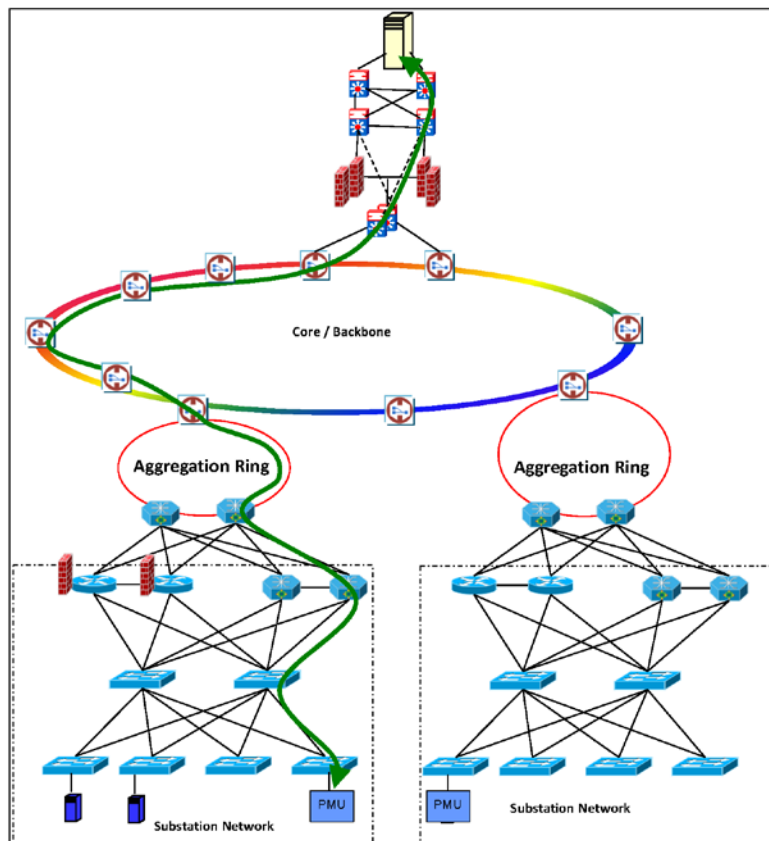


Figure 37. Two Substation PMU Data Flow Example

¹ At one time NIST had a project underway to create a standard for getting highly precise timing from communication networks into application CPUs.

Given the information and assumptions described above, it is possible to extract a linear (in the geometrical sense) model, such as the one shown in Figure 38.




Figure 38. Linear Communication Network String Example

Using such a model and with the assumption of no congestion, one may calculate latency as shown in Table 5 for a case with end to end optical fiber distance of 200 kilometers and 50 μ sec of communication device switching latency.

Table 5. Example Latency Bound Estimation

	64 byte pkt (μ sec)	256 byte pkt (μ sec)	512 byte pkt (μ sec)	1500 byte pkt (μ sec)
Propagation Latency	990	990	990	990
Switching Latency	1100	1100	1100	1100
Transmit Latency	2.0992	8.3968	16.7936	49.2
Queuing Delay	4.1984	16.7936	33.5872	98.4
Total end to end Latency	2096.2976	2115.1904	2140.3808	2237.6

In this example, a 1500 byte packet can be transported over 200 kilometers of optical fiber via a communication network path involving 22 communication devices in under 2.3 milliseconds. Note that this is a lower bound and that many factors can increase the actual latency, such as congestion and the related phenomenon of knee effect due to increasing the number of active endpoints, and changing routing caused by communication network management algorithms and network supervision (especially in the telecommunications service provider case). Similar analysis can be made for any communications network as long as the appropriate parameters are available, based on networking technology, physical media, etc. Such an analysis can help determine if a particular application will be able to function over selected distances and networks, or alternatively can help determine the necessary location of an application to avoid communication network-based delay problems. This kind of analysis becomes most important when PMUs are used in real time closed loop control and protection applications.

13.4.1 Network Level Cyber Security

Cyber security at the communication network level can involve a very larger number of potential techniques, tools, and components. Use of these measures creates the need for data flows beyond the actual PMU data, such as for device registration, identity and certificate/key management, and event

notification. Figure 39 illustrates a data flow model for a multi-cast enabled network using GDOI¹, such as specified in IEC 61859-90-5 in a single communication network.

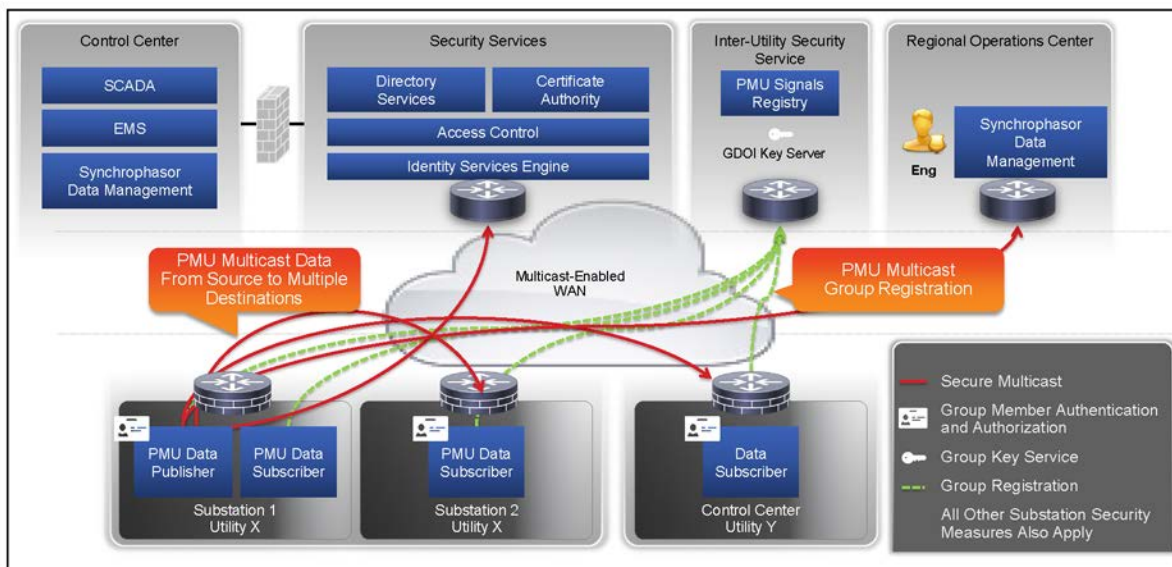


Figure 39. Illustrative NASPInet 2.0 GDOI Security Data Flows, Single Communication Network

In multi-domain networks, as is the usual case when more than one organization is involved in sharing PMU data, security data flows are necessarily more complex, as illustrated in Figure 40. In such a case group key management must be done across network and organizational boundaries, a capability that can be difficult to implement in an ad hoc manner. The need to manage group keys while enabling each PMU owner to ensure cyber security of its assets is one of the factors that leads to the layered, distributed platform model described above. That model makes it easier to employ capabilities of modern IP telecommunication networks to manage security in complex environments efficiently. Note that in this data flow model, three functional groups are involved beyond the sources and uses of PMU data. These are the security services group, the PMU data management function, and the inter-utility services group that includes the PMU signal registry functions. In the figure above, PMU data management is shown as being allocated to a regional operations center separate from the Transmission Control (or Operations) Center, but in fact these functions may be allocated to and housed within one entity. They can however, be split apart as shown if needed. The allocation of these roles and functions is an architectural decision that has clear design implications and so must be a part of a NASPInet 2.0 architecture.

¹ Group Domain of Interpretation, IETF Standard RFC 6407, and Internet Security Association and Key Management Protocol, IETF RFC 2408 for group key management.

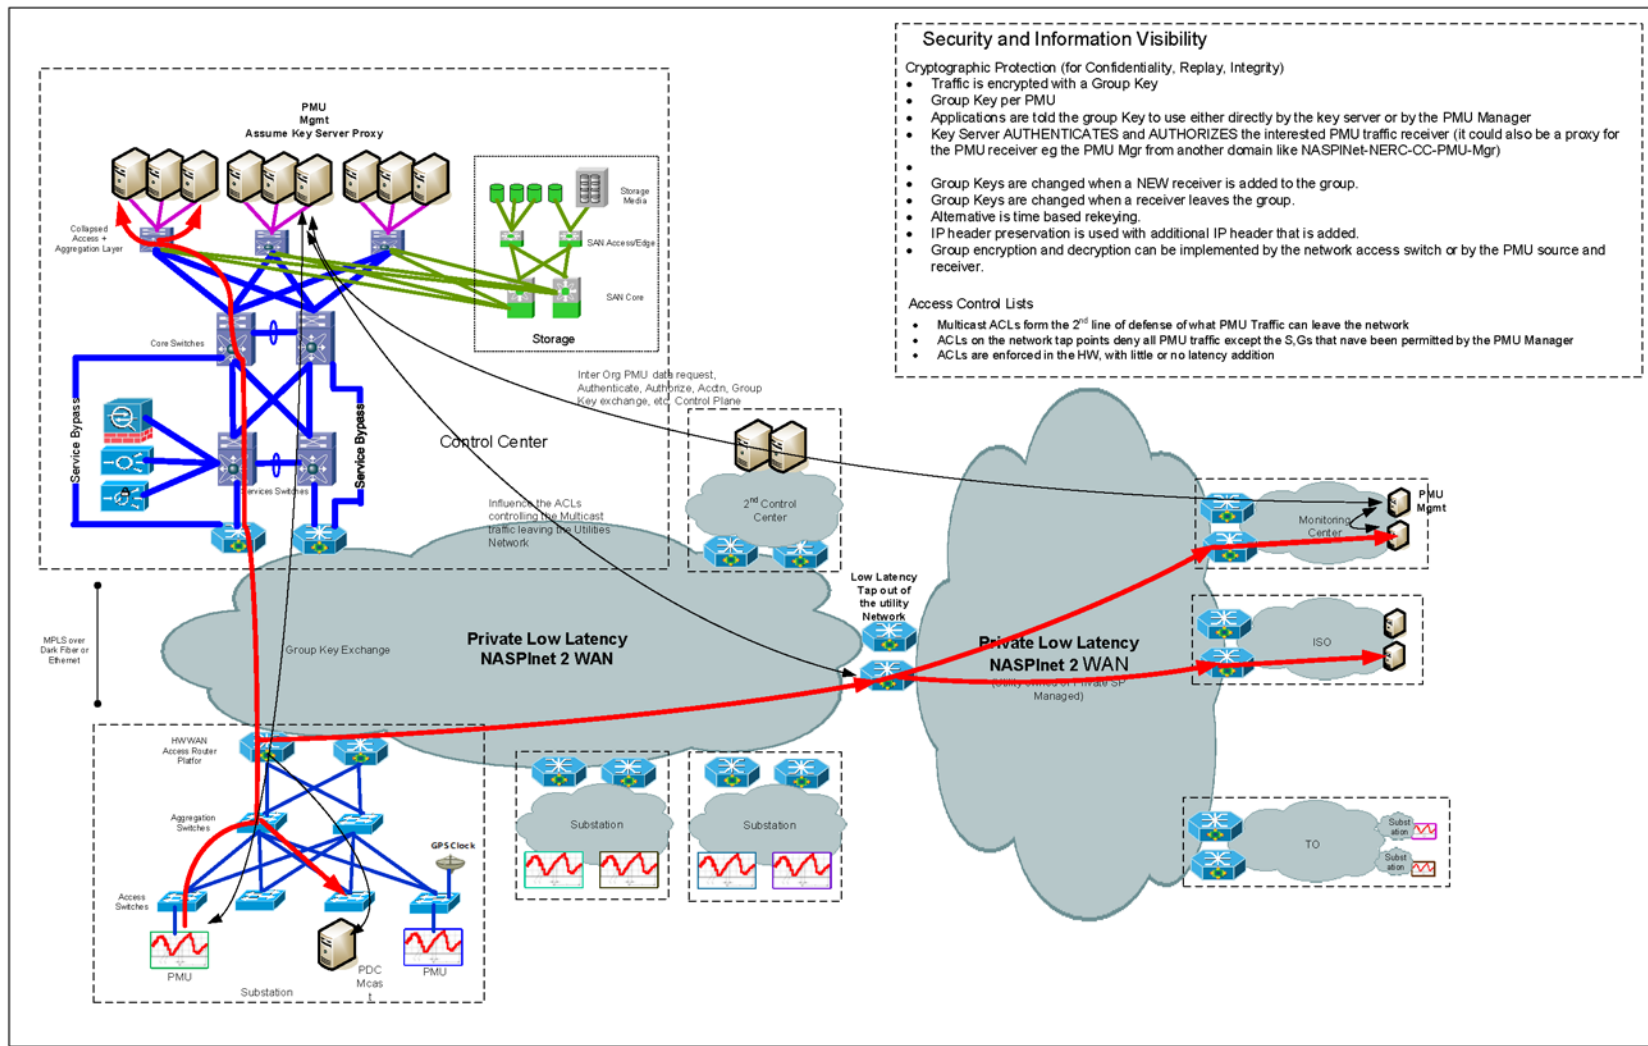


Figure 40. Illustrative NASPINet 2.0 Security Data Flows, Multiple Communication Networks

13.4.2 Registries

The sharing of PMU data by multiple entities, especially by different organizations, requires a means for the user organization to identify, locate, and communicate with the data source, and means to understand the data from such a source. This addresses a classic interoperability challenge, framed in the context of a distributed multi-owner/multi-user infrastructure. The essential functions of such a registry are:

- Register devices and signals by the “publisher” of PMU data
- Provide a unique identification of any registered signal
- Enable discovery of registered devices and/or signals by authorized parties
- Enable bilateral subscription processes that are initiated by the subscriber to PMU data
- Secure access to the registry

As illustrated in Figure 46 above, the function of a PMU registry is closely related to the cyber security functions of data security and data access, so much so that the registry functions and the data security functions should be considered a functional group and should also be grouped with management of bilateral data access agreements.

While the older view of the PMU registry capability was that it should be a system-wide registry of Phasor Gateways, PMUs, and IEDs, as well as service components, processes, and other entities required by Phasor Gateways and/or other Data Bus components, the change in paradigm from Data Bus (and implicitly, a single centrally managed communication network – something that was never the actual case) with centralized management to distributed grid observability platform allows for a more distributed structural approach for registry capability. Phasor Gateways, PDCs, and centralized management of security and QoS on the communication network are not appropriate to a modern NASPInet.

In the context of NASPInet 2.0, registry structure must be aligned with the essential distributed nature of the NASPInet 2.0 observability platform, while respecting organizational boundaries and roles.

13.5 Registry Structural Issues and Alternatives

Three approaches to registry structure and operation for NASPInet 2.0 are:

- Centralized
- Decentralized
- Federated

In a centralized registry model, some entity would have the role and responsibility to host, manage, and secure a registry database and provide the necessary functionality to make it usable. PMU owners must register devices and provide the necessary associated data. This is essentially the old NASPInet model and is included here to account for legacy implementations. Open questions for this model are:

- What entity maintains and operates the registry?
- How is the data integrity of the registry assured?
- How is compliance with any specification for data formats and access processes managed?

In a decentralized model, each organization that owns PMUs would create and maintain its own registry and manage access to it and the PMU data streams. A key question here is how do subscribers find the information they need on PMUs without having to access many separate registries? Does registry structure have to be common (standardized) and if so, by whom? How is shared cyber security managed? In a federated model, each PMU owner would still create and maintain its own registry and manage access to it and the PMU data streams, just as in the decentralized model, but then the collection of registries would be federated into a mega-registry. This poses two questions:

- What entity owns and operates the meta-registry?
- What happens if the meta-registry fails?

Figure 41 illustrates the essential structure of a federated registry system.

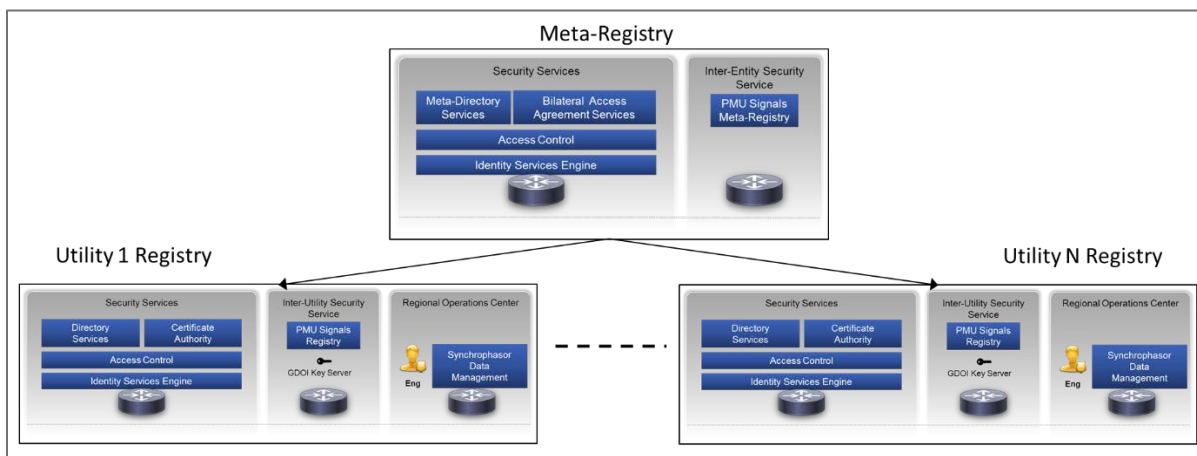


Figure 41. Federated PMU Registry Structure

The federated model inherently encompasses both the centralized and decentralized models as subsets, especially if the meta-registry operates by virtualizing the individual registries. In the full federated model, the meta-registry could be operated by a regional entity such as a reliability coordinator. By allocating the role of *broker of bilateral data sharing agreements* to this same entity, the functions of access by subscribers to individual registries can be managed efficiently. Each PMU owner would be responsible for the content of its registry and ultimately control access to it and to its own PMU data. If the federation mechanism fails, the individual registries can still function, which provides a degree of resilience. Each utility would still be able to use its registry for internal purposes and could open up access to other authorized utilities if needed on temporary basis while the meta-registry is off line.

The foregoing discussion indicates that the registry, in whatever form, is a major point of interface for NASPInets. Consequently, it should be the focus of a significant effort on definition and standardization since there is very little in the way of standards for this now. It is beyond the scope of this Guidance to specify such definitions or standards, but some of the areas needing attention are:

- Content and format of registry databases including:
 - Device identification, location, configuration (see IEEE 1541.4 in this regard)
 - Signal descriptions, including type, reporting rate, data format (see IEEE C37.118, IEC 61850-90-5), signal processing that has been applied, signal origin if signal has been processed
 - Device location information

- Relevant system model information
- Device ownership
- Protocols for registration, update and de-registration of devices, signals, etc.
- Protocols for subscription and de-subscription (see PIM/SSM for some support)
- Protocols for establishing and registering bilateral data sharing agreements

Development of the indicated definitions and standards should be done via a recognized Standards Development Organization such as IEEE or IEC.

13.5.1 Failure Notification

While PMU networks are not in themselves Ultra-Large Scale (ULS) systems, they are attached to and will eventually be integrated with the grid, which is ULS. It is appropriate to consider ULS properties in connection with NASPInets, especially the property of normal failures. In ULS theory, failures in highly complex systems are considered part of normal operations as opposed to being exception conditions. This implies that NASPInets should be structured to deal with failures as routine events rather than special cases. In addition, future applications of PMUs are likely to include real time closed loop protection and control. Consequently, while notification of failures may be done in the typical FCAPS manner, it may also be useful or even necessary to distribute failure event messages more widely to devices and systems that would be affected by the failure. From a data flow standpoint, the NASPInet 2 architecture should determine whether failure event messages will be:

- relayed through a network manager
- be broadcast generally
- publish for subscription by “interested” devices and systems

Figure 42 illustrates these three event data flow models.

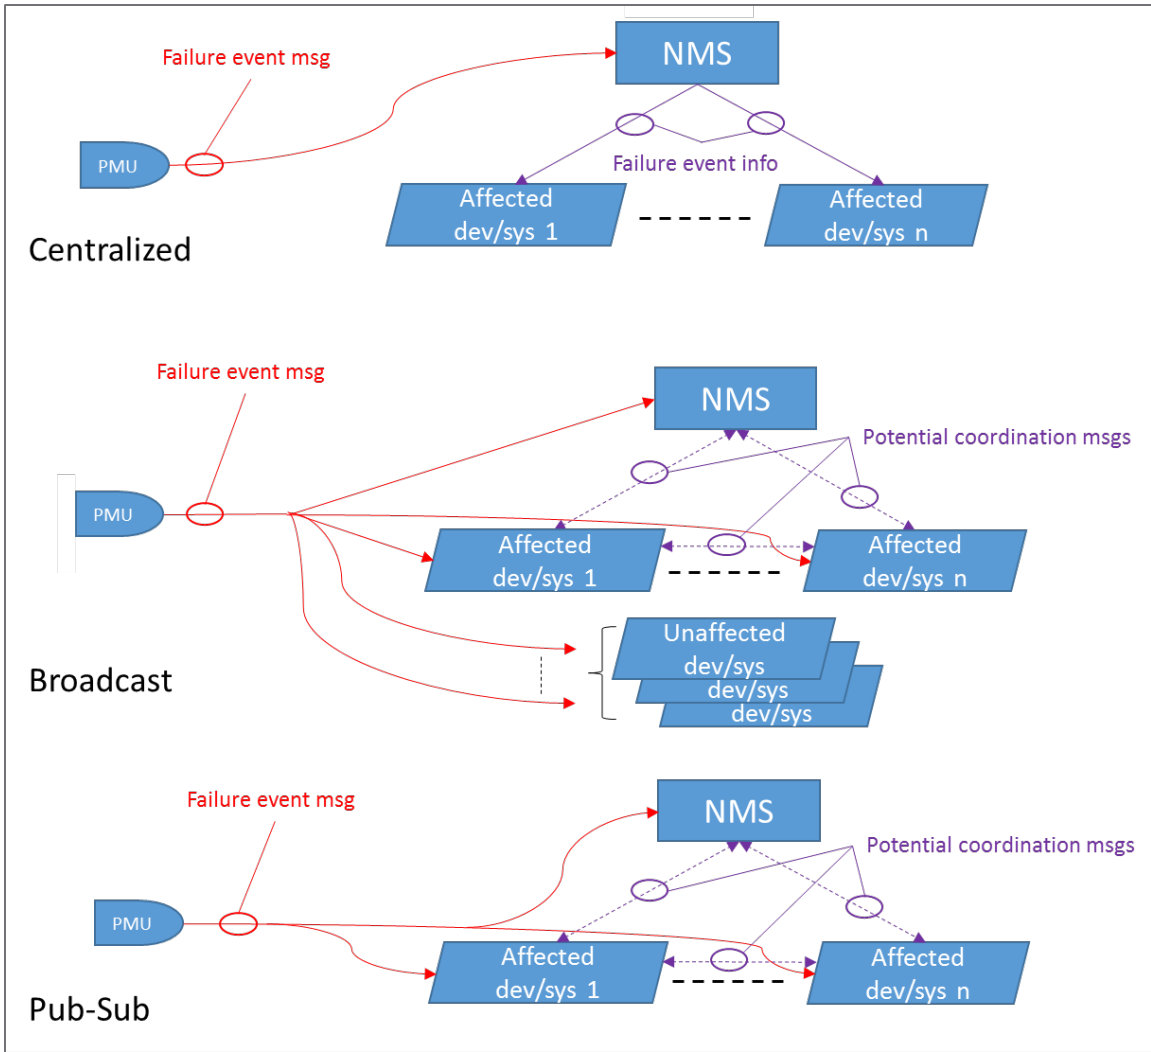


Figure 42. Three PMU Failure Event Message Data Flow Options

The first option introduces the most latency but can provide convenient means to implement centralized logic that determines how to manage and distribute failure event information. The second option relies upon each receiver to decide how to handle the event, including those that must determine they are not interested, and also leads to the most network traffic. The third allows for low latency in distribution the event information while limiting network loading. In the second and third cases, the logic for handling failure events is decentralized and provisions may be needed to allow for collaboration and coordination among responding elements to avoid conflicting responses. The second and third cases are harder to design but are potentially faster (lower latency) and more resilient than the centralized approach.

The “pub-sub” (third) model requires that interested devices and systems subscribe in advance of the failure event message being generated. A modification of this would include posting the failure event messages to a blackboard so that newly interested devices and systems can find this information when this discover they need it. Blackboard systems are shared information repositories having knowledge sources, knowledge users, the blackboard itself (knowledge repository) and a control monitor that manages use of the blackboard. They have been used in multi-agent artificial intelligence applications.

14.0 Relevant Standards

The following standards are or can be relevant to NASPInet 2.0.

14.1 IEEE C37.118.1 and C37.118.2-2011

This is the original standard for transport of PMU data and is still in wide use. There are two parts:

- C37.118.1 covers synchrophasor measurement requirements including measurement definitions, dynamic performance, and error quantification.
- C37.118.2-2011 covers synchrophasor data transfer requirements, including representations, message contexts, and message formats, timing, and error detection and correction. This part contains an annex (F) that briefly discusses communication over IP networks but is too brief to convey much guidance. Also, the discussion on synchrophasor networks strongly suggests PDC stacking, which clearly should be avoided.

This standard have gone through various revisions since 1995 and is widely used in one version or another in the US. C37.118 is explicitly a streaming data protocol.

14.2 IEEE Std C37.244™-2013

This is the IEEE Guide for Phasor Data Concentrator Requirements for Power System Protection, Control, and Monitoring. This standard defines a PDC as “a node in a communication network where synchrophasor data from a number of PMUs or PDCs is processed and fed out as a single stream to the higher level PDCs and/or applications.” The standard depicts PDC stacking and the use of PDCs as an over-the-top-network. Given the guidance on network principles above, this standard should be updated or deprecated.

14.3 IEC 61850-90-5

This standard extended the extensive IEC 61860 suite specifically for routable protocol support for synchrophasor exchange. Much of the content has to do with alignment to the IEC 61850-7-2 (Abstract Communications Interface) but it also introduces GDOI for group key management.

14.4 IP protocol suite

The Internet Protocol Suite is a large set of standards developed through the Internet Engineering Task Force (IETF) for protocols used in digital communications. The standards are described in documents referred to as RFCs (Request for Comments) and implementations (protocol stacks) are widely available and are incorporated into standard routers and switches used to build IP networks. IP networks use many protocols so it is common to consider an IP stack for any particular network. The IP model defines a layered (or stacked) approach to communication function that is somewhat related to the core/edge principle described earlier.

While the number of IP protocols is large and continues to grow, a few of special interest to NASPInet 2.0 systems are:

TCP – Transmission Control Protocol. A primary element of IP suite that provides connection-based reliable, ordered, and error-checked delivery of a stream of bytes between applications running on hosts communicating via an IP network, with the ability to request re-transmission for lost packets and delivery failure notification back to a source.

UDP – User Datagram Protocol. A connectionless method for sending messages without prior communications to set up connections. It is often used where re-transmission is not a good option and for situations where establishing connections is problematic from a scaling perspective.

MPLS – Multi-Protocol Label Switching. The name is a bit misleading because the purpose of this was changed from protocol switching to the use of packet labels to determine forward, thus eliminating the need to examine packet contents. MPLS is viewed as operating between Layers 2 (data link) and 3 (network) of the OSI model and so is sometimes referred to as a Layer 2.5 protocol. It is used widely in the core parts of large service provider communication networks. MPLS can carry many types of traffic besides IP packets.

SSM – Source Specific Multicast. A protocol in which packets are delivered only to a receiver that has specifically requested them from the source. Multiple receivers may request from the same source because they essentially subscribe to the source and therefore can enable the communications network to act as a publish-and-subscribe mechanism at the data transport level. Where multiple destination devices request data for the same source, the network will determine the best location for packet duplication based on network topology, thus minimizing demands on the network. SSM generally makes use of UDP and some enhancement products exist to deal with packet loss where necessary.³⁹

14.5 IEEE 1588

IEEE 1588 defines the Precision Time Protocol for distributing clock information in communications networks. The standard is based on a hierarchical or master-slave architecture involving multiple clock types. Various profiles exist for PTP applications, including a Power Profile for use in electric power systems. See the discussion on clock and timing distribution in the Structures section above.

14.6 IEEE C37.238

C37.238 (IEEE C37.238-2011 - IEEE Standard Profile for Use of IEEE 1588 Precision Time Protocol in Power System Applications) defines the profile for the use of Precision Time Protocol (PTP) of IEEE Std 1588-2008 in power system protection, control, automation, and data communication applications utilizing an Ethernet communications architecture.

14.7 IEEE 1451

This standard for smart transducers has not been widely recognized as being useful for PMU networks but has elements of value. In particular, IEEE 1451.4 (which was developed mostly for process control and similar transducers and sensors) defines the concept of and a standard for Transducer Electronic Data Sheets. The original intention was that information on device manufacturer, model number, version number, serial number, physical units of the physical variable being sensed, as well as calibration information (sensitivity) could be embedded in the transducer itself, to allow for automated

³⁹ UDP packet loss can be due to improper buffer configuration and so should not automatically be viewed as an inherent shortcoming of UDP.

interoperability. Since not all sensors have embedded TEDS, the same format could be used in a data repository hosted on behalf of the devices. PMU TEDS files could be included in device or signal registries.

14.8 IEC 61968/61970

The IEC Common Information Model (CIM) is a standard for vocabulary and basic ontology for aspects of the electric power industry ranging from enterprises down to individual wires and their characteristics. CIM is maintained in the form of UML models and includes definition of an XML representation for network model exchanges. Efforts have been initiated for harmonization of CIM and SCL (Substation Configuration Language – another parallel effort of IEC for describing substation systems) and for harmonization of CIM and MultiSpeak, which has been viewed as having a better approach to representation and exchange of power state information than CIM.

14.9 IEC 27002

This is a general standard addressing aspect of security broadly, with 14 general topics areas, some of which are people/process oriented (Humans Resource Security), some of which are technical (cryptography), and some of which are business oriented (legal and contractual compliance requirements). This is a general framework, with other IES standards focusing on narrow aspects (see below).

14.10 IEC 62351

This standard addresses authentication of data transfer through digital signatures for authenticated access, prevention of eavesdropping, prevention of playback and spoofing, and intrusion detection. Its relevance here is its connection to IEC 61850 (see above).

14.11 IEC 27040

IEC guidance for storage security that is generally in alignment with ISO/IEC 27002. It provides an overview of key storage and storage security concepts, describes the controls that support storage security technical architectures, and provides guidelines for the design and implementation of storage security.

14.12 NERC CIP x

Presently NERC CIP consists of a large ran evolving standards and requirements for protection of critical infrastructure. Synchrophasor networks have not been considered critical infrastructure up to this point (2018) but future applications are very likely to make PMU networks critical and so architects and designers should plan to meet these requirements.

14.13 NISTIR 7628

This is a guide, not a standard, but is the most comprehensive compendium or principles and practices for power grid/electric utility cyber security.

15.0 Summary General Guidance for NASPInet 2.0

The following are general principles to keep in mind when developing a NASPInet 2.0 architecture.

- Don't confuse organization structure and boundaries with NASPInet 2 structure.
- A NASPInet 2.0 system is a grid observability platform, not a data bus. The platform is intended to provide the synchronized measurement information need to support grid operations, planning, modeling, and analysis. There is no "Data Bus" and there is no need for Phasor Gateways. To the extent that the organizational boundaries must be respected this should be done at the communication network level, not by layering an over-the-top network of devices or systems on the core communication network.
- Map functions to the right place in the platform.
 - As a corollary, consider what functions should map to the communications systems (which are platform layers).
- Keep layer definitions and functions clean (this is an aspect of modularity principles, namely module strength and de-coupling).
 - Networking is integral to NASPInet 2.0 but there should be no "over-the-top" (of the communications systems) networking.
 - As a corollary, this means no PDC stacking –there should only be one time alignment function per application.
- Cyber security should be integrated into the NASPInet 2.0 architecture. Treat cyber security in terms of vulnerability and securability of a NASPInet, not as some overlay structure or separate architecture.
- When developing the NASPInet 2.0 architecture, proceed in an orderly and organized fashion.
 - Keep design separate from architecture and develop the architecture **first**.
 - Objectives -> Capabilities -> Functions -> Architecture -> Design -> Implementation
- Avoid "wiring" a particular product, technology or tool into the architecture. If one is wired in, then it must be treated as a constraint when developing the architecture.

16.0 Guidance on Newer/Emerging Technologies

In the decade since the original NASPInet Guidance was written, a number of developments have arisen that may be useful for NASPInet 2 networks. Some are well-established in other computing environments whereas others are emerging but are not yet fully established.

16.1 Software Defined Networks

In many communication networks, the control (routing) and forwarding (switching) functions are distributed throughout the network. These functions are generally decoupled in modern networks but are built in to the network devices in a manner that provides only limited accessibility to the network operator. Software Defined Network (SDN) technology evolved as a means to make the network control functions accessible to external applications. Much of the motivation for SDN development originally

came from changing data flow patterns *inside data centers*. In the case of PMU networks, the issues are focused on *Wide Area Networks* and related to network congestion, packet delay variation, packet loss, and service outages. These issues will become increasingly important as PMUs are used for near real time decision support and for wide area closed loop protection control applications. Figure 43 illustrates a simple view of the combination of a joint grid and communication network control application, a universal power flow controller, and two PMUs in a configuration that makes use of SDN (actually SD-WAN) to measure and manage communication network performance for the purpose of maintaining stability of the closed loop control.

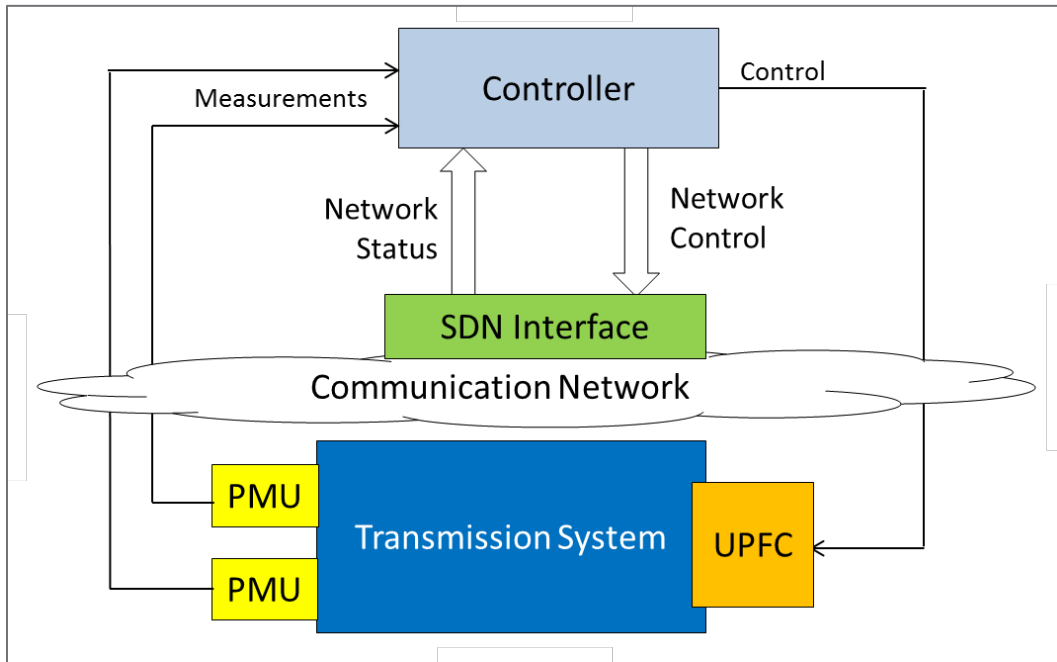


Figure 43. Simplified Joint Grid/SDN Wide Area Control Structure

SDN and SD-WAN are relatively new and still evolving; existing networks do not support the necessary interfaces and functionality. For PMU networks, if the communication network is operated by a third party provider, then it may not be possible to apply SD-WAN since the communication network is being managed by the network operator for multiple purposes and may only allow requests from a utility as opposed to direct control. In such a case where PMUs are being shared and PMU data may flow across multiple networks, there arises a coordination issue in terms of which entity has the role and responsibility for network performance management via SDN. If the communication network is shared (as will often be the case) there must be a means to resolve conflicting objectives for network management. Unless SD-WAN is being applied to a single private network, it may not be possible to resolve the conflicting objectives.

SDN raises additional operational issues by introducing a centralized network control that that reduces resilience (single point of failure) and increases cyber vulnerability. Also, for existing networks, the entire network must be reconfigured to support SDN. Finally, existing SDN does not scale well with the number of network devices being managed.

16.2 Cloud Services

The term Cloud Services covers a wide array of computing and IT capabilities delivered to end users over the internet from server facilities hosted by cloud computing providers. Offerings include Software as a Service (SaaS), Platform as a Service (PaaS), and (computing) Infrastructure as a Service (IaaS). Use of cloud computing is well established in many industries and provides a variety of standardized benefits, including:

- Proportional and on-demand use of computing capacity, storage, and software applications with monthly billing instead of licenses
- Access to continually modernized computing resource and up-to-date software versions
- Access to scalable computing resources without capital expense
- Cloud vendor cyber security for the elements they host

Use of cloud services requires that the utility have an internet connection to the cloud service and that utility data be present in the cloud servers for some amount of time. If PMU systems and data are to be considered part of critical infrastructure for NERC CIP purposes, then the responsibility for cyber security partly rests with the Cloud Service provider, who may not agree to meet NERC CIP requirements.

Some additional issues:

- Pay as you go pricing may make costs somewhat unpredictable.
- Configuration options are limited compared to owned infrastructure
- Data privacy and confidentiality terms of service may not be consonant with utility needs
- There are open questions on the ownership of data stored on cloud servers
- While cloud providers claim high cyber security (and are highly motivated to use excellent physical and cyber security), significant breaches have occurred. Due to the nature of such services (storing data for thousands of companies) cloud providers are highly targeted by hackers
- Data losses (leakages) have occurred and cloud services can exit the business, leaving stored data in limbo.
- The cloud provider servers may not be located in the service territory of the utility; and may not even be located in the United States. Further, the cloud service operator may move data and computing from data center to data center in order to manage electricity expenditures by doing the computing in locations where electricity rate are lowest, even on a daily basis.⁴⁰ This is known as a “follow-the moon” strategy and a consequence is that the user cannot know where the data is physically located. It may be possible to address data location contractually.

From an architectural standpoint, if an electric utility, system operator, or reliability coordinator plans to use a cloud service for PMU data, then communication should be via a redundant secure network arrangement, as Figure 44 illustrates.

⁴⁰ Stacy Higginbotham, Google Gets Shifty With Its Data Center Operations, GIGAOM, June 16, 2009, available online: <https://gigaom.com/2009/07/16/google-gets-shifty-with-its-data-center-operations/>

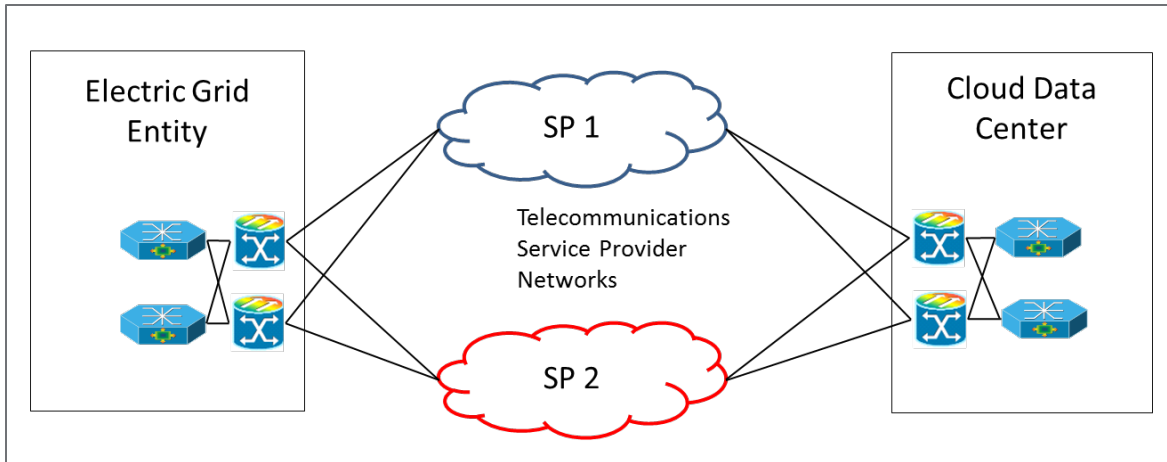


Figure 44. Redundant Cloud Communication Path Concept

16.3 Network Virtualization and Network Function Virtualization

Network virtualization is the combination and abstraction of physical network elements into a software-based logical network entity that can be administered so as to provide flexible networking to multiple processes in a manner that appears to each user as if it has its own network, similar to the way a hypervisor provides multiple virtual processors to software applications or processes by sharing the underlying physical hardware.

Network virtualization (NV) refers to two slightly different forms:

- External network virtualization is the combination of many networks or network parts into a virtual (abstract logical) unit.
- Internal network virtualization is the provision of network-like functionality to software containers on a single network server.

Network virtualization is the virtualization of *network hardware* and can improve utilization of physical networks and avoid the need to build multiple physical networks. It can also be used to provide simulation environments for testing purposes. NV can be used to support Network as a Service (NaaS) and so may be used by network infrastructure owner/operators to supply virtual network services to owners of PMUs or users of PMU data.

Network function virtualization (NFV) uses software to abstract classes of network node functions, such as load balancers, session border controllers, firewalls, and intrusion protection, which then run on standard servers and can even be cloud-based (see discussion on cloud services above). In other words, it is the virtualization of various network *services*, as opposed to virtualization of network hardware. In the sensor/communication layer model describe above for the PMU-based observability platform, NFV can be used to provide network services for a NASPInet, as well as additional data services.

The uses of NV and NFV are primarily design and implementation issues, not architectural ones.

16.4 Communication Protocols

The primary communication protocols used with PMU networks are described in the section in relevant standards above. Experience with PMUs has shown that it can be necessary to segment or partition PMU data according to intended destination and usage, something that the original protocols did not provide. As a consequence there is an effort to develop a new Streaming Telemetry Transport Protocol (STTP). The definition of the protocol can be found in the STTP definition document.⁴¹

An annex to this Guidance, entitled “A Comparison of Phasor Communications Protocols” provides a comparative analysis of IEE C37.118, IEC 61850-90-5 and the proposed STTP. NASPInet architects and designers with interest in this topic should consult the three standards documents and the comparative analysis document.

17.0 Glossary

Term	Definition
Application (app)	Generically, a use to which something is put. In computing, where the short term app is often used, a software program that supplies or implements a narrow end user function (as opposed to a platform function). In Grid Architecture, an application may involve more than software, but the same idea applies.
Architecture	A high level depiction of a system, used to reason about the system’s properties and behavior, and to specify key decisions about the shape (structure) of a system for the purpose of bounding and simplifying design decisions.
ARRA	American Recovery and Reinvestment Act – 2009 legislation that provides funds which were instrumental in purring the deployment of PMUs in US electric transmission systems.
Cybersecurity	The state of being protected against the criminal or unauthorized use of electronic data, or the measures taken to achieve this.
Data packet	A data packet is a unit of digital data made into a single package that can be transported via a digital communication network.
Data stream	A coherent sequence of data packets (usually comprising a signal or signal set) being transported on a continual basis at a more or less uniform rate
Galileo	European Union satellite navigation system.
GLONASS	Globalnaya Navigazionnaya Sputnikovaya Sistema, or Global Navigation Satellite System. This is the Russian satellite navigation system.
GPS	Global Positioning System – the US satellite navigation system.
Interoperability	The ability of two or more digital systems or components to exchange information and to use the information that has been exchanged.
ISO	International Organization for Standards based in Geneva, Switzerland.
ISO/RTO	Independent System Operator/Regional Transmission Operator – roles in the operation of bulk energy system in the US as defined by FERC
IOU	Investor Owned Utility
Latency	In telecommunications, delay in the transfer of data packets.
LSE	Load Serving Entity – any of a class of organizations responsible for acquisition of electricity on behalf of end-users and wholesale customers.

⁴¹ Grid Protection Alliance, STTP Streaming Telemetry Transport Protocol, available online: <https://github.com/sttp>

Term	Definition
Message	A logical unit of information represented in digital form and being transported by a communication system.
Network	A group or system of interconnected people or things
Observability	Intuitively grid observability is temporal, geospatial, and topological awareness of all grid variables and assets. A more formal definition is the ability for any combination of dynamic system state and inputs to determine the system state in a finite time using only measurement of system outputs. Observability is the technical basis for “visibility.”
Phasor	Phase angle vector – a mathematical object consisting of magnitude and phase of a sine wave of specified frequency.
Phasor Measurement Unit (PMU)	A device that samples AC power line voltage of current waveforms and calculates the phasors for each power line phase being sampled.
Platform	A stable collection of components that provide fundamental or commonly-needed capabilities and services to a variable set of uses or applications through well-defined interoperable interfaces.
Protocol (communication)	A communication protocol is a system of rules that allow two or more entities of a communications system to exchange information via any kind of variation of a physical quantity. Generally, a suite of protocols is used to enable communication; such a suite may be represented as a stack.
Protocol stack	A group of protocols that run concurrently and that are employed for the implementation of network protocol suite.
Pub/sub	Publication/subscription – this refers to a data flow management mechanism in which data is “published” by data sources such as streaming sensors, and “subscribed to” by any authorized device or system that wants to use the data stream. Mechanisms that provide such capability include message brokers, enterprise service buses, various kinds of middleware, and communication networks.
Registry	Generally, an authoritative list of one kind of information. For NASPINets, a database and associated functionality that describes either PMUs or PMU signals, along with sufficient information to allow authorized users to access data from the registered elements.
ROCOF	Rate of Change of Frequency – a useful power grid system variable, when the frequency in question is the system frequency.
Signal	An impulse or fluctuating quantity, such as an electrical voltage or light intensity, whose variations represent coded information.
Synchrophasor	A phasor measured using a reference time base for synchronization so that multiple phasors measured at different locations on a power grid can be combined or compared mathematically.

Appendix A – PMU Applications Lists

Appendix 1 contains lists of PMU applications consolidated from several sources.^{42,43,44,45}

A.1 PMU Applications (Transmission)

Number	PMU Application (Transmission)
1	Situational Awareness Dashboard
2	Real Time Compliance Monitoring with Reliability Standards (Angle of Separation, Voltage & Angle Profiles, MW, MVAR flows, Load-Resource Imbalance)
3	Frequency Stability/Islanding
4	Real Time Equipment and Performance Monitoring and Trending
5	Anomaly Characterization and Alarming (Real time alarming on hard limits & "out of normal" conditions, suggest preventive action)
6	Sensing for State Estimation
7	Small-Signal Stability Monitoring
8	Voltage Stability Monitoring/Assessment
9	Thermal Monitoring (Overload)
10	Real Time Performance Monitoring and Trending
11	Sensing for Anomaly Characterization and Alarming (Real time alarming on hard limits & "out of normal" conditions, suggest preventive action)
12	Baseline Normal Phase Angle Trends
13	Disturbance Analysis
14	Frequency Response Analysis
15	Model Validation
16	Automatic Alarming of Remedial Action Schemes
17	Out of step protection
18	Short-term stability control (e.g. transient stability)
19	Long-term stability control (e.g. wide area frequency, voltage stability)
20	FACTS feedback control, smart switchable networks
21	Multi-terminal protection
22	Negative sequence pilot protection
23	Fault impedance distance relaying
24	Directional relaying
25	Mode meter and mode shape meter
26	System frequency and ROCOF monitoring
27	Protection system operation verification
28	Dynamic line rating

⁴² D. Novosel, et. al., IEEE Power and Energy, Jan/Feb 2008, pp. 54-55.

⁴³ A von Meier, et. al., Micro-Synchrophasors for Distribution Systems, available online: https://www.researchgate.net/publication/269306862_Micro-synchrophasors_for_distribution_systems

⁴⁴ D. Novosel, Tutorial on PMU Technology and Applications, International Conference on Synchrophasor Measurement Applications, Brasil, 2006.

⁴⁵ NASPI Distribution Task Team, Synchrophasor Monitoring for Distribution Systems: Technical Foundations and Applications, NASPI-2018-TR-001, January 2018, available online at <https://www.naspi.org/>

29	Black start and system restoration management
30	GMC/GIC monitoring
31	Cyber-attack detection

A.2 Distribution PMU Applications

Number	PMU Application (Distribution)
1	Transient condition analysis (pre-fault)
2	Fault detection/classification/location, including open phase fault detection
3	High impedance fault detection and arcing detection
4	Asset utilization monitoring and equipment health diagnostics
5	Distribution grid state determination and power flow monitoring, including reverse flow and phase unbalance
6	Renewable generation monitoring
7	Distribution asset characterization and utilization monitoring and optimization
8	Droop control for inverter-interfaced DG (esp. on rural grids)
9	Microgrid protection via communication-assisted relays
10	Unintentional islanding detection
11	Fault induced delayed voltage recovery discovery
12	Circuit topology detection and phase identification
13	Load and generator models; line segment impedance measurement
14	Disaggregation of DG and load
15	Voltage sag detection/analysis
16	Cyber-attack detection

Appendix B – General Principles for Composing Architectures

The following is a list of general principles describing good architecture practice.

1. A good architecture is one that meets the needs of the stakeholders (especially the users) to their satisfaction, does not violate established principles of system architecture, and takes into account the relevant qualities and properties as the stakeholders require.
2. Good architectures have conceptual integrity: the architecture is based on a set of principles that guide architectural decisions in a traceable manner, should be clean of unnecessary complexities or exceptions; should produce enforceable constraints; similar problems should be solved in similar ways, and the architecture should have an overall consistency of approach.
3. Conceptual integrity is best achieved by a small cohesive team of like-minded architects. Architecture should be the product of a single architect or small team with an identified leader.
4. Essential functionality drives complexity, not architectural “elegance.
5. Architectural structures should have formal bases where possible to minimize ad hoc configurations with unknown properties.
6. Architecture should not depend on or be constrained by a particular commercial product, tool, or business model.
7. The architect must be cognizant of the global system when optimizing subsystems.
8. Stakeholders should be involved in the process as much as possible, giving frequent and honest feedback on all aspects of the system architecture.
9. Each component should be responsible for only a specific feature or functionality, or aggregation of cohesive functionality.
10. Components should be coupled only through explicit structure, avoiding hidden coupling where possible.
11. Architecture defines interfaces, not vice versa.

Appendix C – NASPInet 2.0 Function Class Definitions

C.1 Capability Class Function Definitions

Function	Capability Class	Definition
Sensing (transduction)	Data Acquisition	Conversion of a physical variable such as voltage or temperature into an analog signal (typically a small voltage) that varies in proportion to the physical variable
Sampling and discretization	Data Acquisition	Conversion of a continuous time and continuous magnitude analog signal into a digital signal (sequence of integers) by capturing the analog signal magnitude at regular and discrete time intervals and converting to discrete quantized values represented as integer numbers
Compensation	Data Acquisition	Correction of sensor signals for factors such as temperature that can cause a transducer output to vary for reasons other than variation in the physical parameter being measured
Units conversion	Data Acquisition	Conversion of the raw digital integer values from the sensor analog-to-digital converter into scaled values that represent the physical variable in engineering units, such as Volts or Amps
Data formatting for transport	Data Acquisition	Packaging of digital sensor data into packets suitable for use by telecommunication protocols
Signal processing computations	Data Acquisition	Computations carried out on the digital signals to extract representations or determine values of variables not directly sensed, such as frequency or real and reactive power (computed from voltage and current signals)
Routing and Switching	Data Transport	Processes used in telecommunication devices to determine paths through a network from source to destination (routing) and to relay data packets from an incoming channel to an outgoing channel (switching)
Data flow management	Data Transport	Determination of data flow paths through a network – commonly known as routing.
Device/system interface	Data Transport	The physical and logical means to connect a device (such as a PMU) of a system (such as a server set) to a communication network.
Protocol execution	Data Transport	The operational use of telecommunications protocols by a device or system.
Timing source interface	Data Sync	The means to connect a communications network to a source of timing information such as a GPS receiver or a Cesium clock.
Timing distribution	Data Sync	The means to provide timing information from a source to the devices that need it, including edge devices such as PMUs but also including device internal to the communications network such as routers.

Function	Capability Class	Definition
Timing synchronism	Data Sync	The means by which time (or time stamp value) is agreed upon by the various elements of a distributed system.
Time stamping	Data Sync	The process of attaching a data item to a set of data to indicate temporal order among a set of events.
Buffering	Data integration	Providing means to hold variable amounts of data temporarily so as to accommodate variations in channel flow capacity or data processing resource availability; means to cushion the shock of fluctuations in data flow
Signal aggregation	Data integration	The process and means to accumulate signals or data from multiple sources or flows into a single data set or flow. ⁴⁶
Conversion	Data integration	Change in representation of a data set, such as from one standard from to another.
Databasing	Data integration	Means and methods for storage and retrieval of data in databases.
Streaming	Data integration	Moving data from one device or system to another in a regular sequence. Data streams from PMUs by sending a continual sequence of messages into the communications network.
Data presentation	Data integration	Means and processes to display data such as PMU data informs useful to humans.
Signal registry functions	Data integration	Tools and processes for tabulating the existence and characteristics of PMU signals, and for controlling access to those signals.
Device meta-data management	Data integration	Collection and use of information that describes a data set such as PMU data. Can include such items as source, data size, description of the content, how it was created, how it is formatted, etc.
Cross organization data sharing	Data integration	Means to make data available between separate organizations. Can include sharing data descriptions (see registry functions) as well as protocols for moving data between organizations and permitting access.
Data source quality analysis	Data Curation/QA	Processes to check on the reliability of data by performing analyses on data sets for temporal consistency, variance and other potential indicators of data source issues.
Validation/editing/estimation	Data Curation/QA	Processes for dealing with basic errors in data acquisition that include simple checks on data plausibility (range checks for example), correcting simple data errors, and filling in gaps with values calculated from time-adjacent or otherwise related data.

⁴⁶ Aggregation two has common definitions in regard to data. One is the *consolidation* of data sets; the other is the *summarization* of data sets. In the former case, the result is larger in volume than any of the inputs, whereas in the latter it is smaller than the tabulation of the input data. We are using the former case definition.

Function	Capability Class	Definition
Signal correction	Data Curation/QA	Adjustment of a signal to compensate for some distortion, such as modification of a phasor to offset for system frequency variation.
Archiving	Data Curation/QA	Storage of data for long term preservation.
Data life cycle management	Data Curation /QA	Method and processes for managing the flow of data throughout its entire life cycle.
Data integrity	Data security	Means such as encryption that prevent the corruption of data.
Data confidentiality & privacy	Data security	Means to determine what data can be shared with third parties (privacy) and mean to allow authorized access to sensitive data (confidentiality).
Device/system integrity	Data security	Means to ensure that devices and system have not been compromised by attack, invasion, or failures.
Data access control	Data security	Communication network level means to limit access to the network itself or any of the data traversing it.
Data accessibility assurance	Data security	Means to assure that data is available in a timely manner to the various users at the proper levels of identity or role-based accessibility.
Device configuration	Device/Network Mgmt	Means and processes for providing settings and parameters to devices that determine their operation.
Device control	Device/Network Mgmt	Management of devices (usually remotely) to direct or supervise their operation.
Software/firmware update	Device/Network Mgmt	Means and processes for sending new software or firmware, or incremental changes (patches) to software and firmware in a device or system.
Fault management	Device/Network Mgmt	Tools and methods to detect and handle failures in network and other digital devices.
Device/network administration	Device/Network Mgmt	The processes and means for managing network, including monitoring and adjusting performance, dealing with faults, and provisioning (configuring) networks.
Remote diagnostics	Device/Network Mgmt	Tools and methods for determining the nature of problems with a device or system in a distant location.
Network management	Device/Network Mgmt	Supervision and direction of a communication network to deal with issues such as congestion and maintaining quality of service.
Network policy management	Device/Network Mgmt	Use of a set of technology-independent policies (rules) that automatically govern the behavior of a communications network through interpreted logic that enhances the underlying hard coded behavior of the network devices.

Appendix D – Standard Network Security Measures

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
- Device and platform integrity

Using these as basic principles, an extensive variety of technical measures may be applied at the communication network level. These are all available in product form and so can be applied at the design stage. An understanding of these is helpful to developing the architectural specifications that facilitate or require such measures.

The following is a list of standard network cyber security measures that any communication network operator should include as the minimum to be considered as part of a comprehensive cyber security program. It does not address people and process issues, but those are vital as well. Two processes in particular to consider are manufacturing supply chain security management and secure code development and code hardening (against buffer overflow, self-modification; remove unnecessary protocols).

The technical measures list:

- Crypto: link layer, group, and application layer
 - GDOI as it has been incorporated into IEC 61850-90-5 specifically for PMU network encryption
- Role-based Access Control (RADIUS and TACACS; AAA; NAC)
- Mutual authentication; media-independent identity authentication protocols
- X.509, secure key generation and management, scalable key management (DMVPN, GETVPN for example)
- SIEM, firewalls
- Intrusion Prevention System (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)
- Six wall physical security for devices and systems; access detection and mitigation (i.e. port shutdown)
- Anti-counterfeit measures
- Air gapping (physical network isolation, data diodes)
- Structure for securability

Some related security elements that are not network components or functions include:

- Secure code development and code hardening (against buffer overflow, self-modification; remove unnecessary protocols)
- Manufacturing supply chain security management
- Data quality as tamper detection
- Anti-counterfeit measures
- Security posture assessment

In addition, there are a variety of people and process elements that must be included to secure NASPINets but are beyond the scope of this Guidance.

Appendix E – Middleware and Sensor/Communications Layers

Two approaches to creating structure in the architecture of a NASPInet 2 platform are the use of middleware and the use of sensor/communication infrastructure layers. While these two approaches can seem equivalent, there is a significant underlying structural difference that should be clearly understood.

Middleware got its start as a means to connect two monolithic software systems running on separate computers. Since these systems had not been developed with the means to share information or processing, special software was developed fit in between them, as Figure E.1 shows.




Figure E.1. Early Middleware Structure

Later the concept was generalized to connect and facilitate cooperation among multiple systems. Figure E.2 shows the essential structure of such an arrangement.




Figure E.2. Multi-System Middleware Structure

The middleware may serve multiple systems but the essential structural characteristic of data flow through the middleware remains. The middleware may supply additional service besides data transfer.

For the enterprise IT environment in the context of web services and SOA, the evolution of middleware resulted in the development of enterprise service buses, as illustrated in Figure E.3.

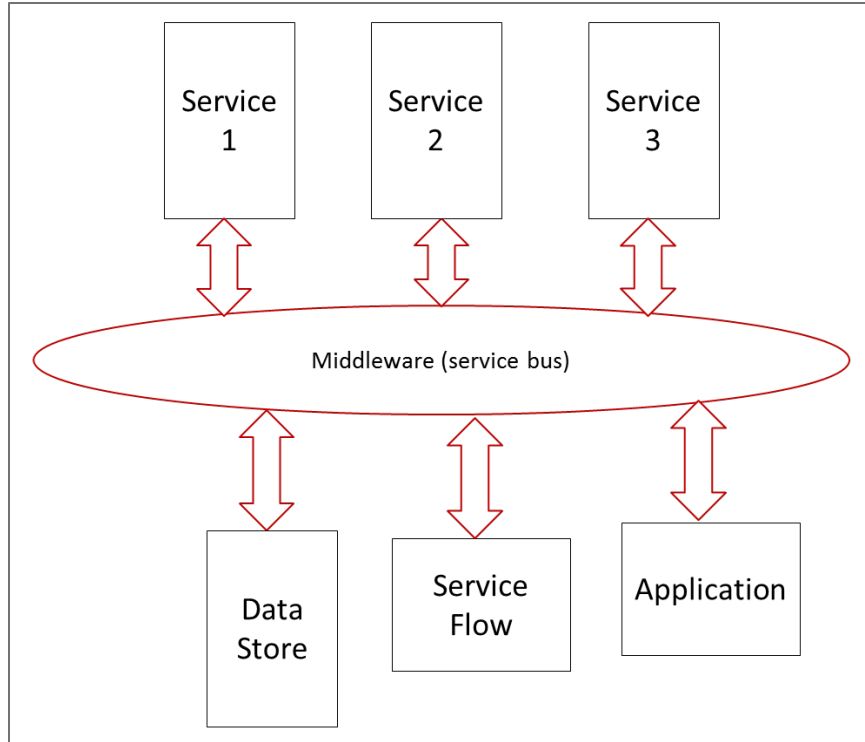


Figure E.3. Enterprise Service Bus System Structure

The nodes in Figure E.2 may be individual services, databases, or applications running on some collection of servers. Note that the middleware is not so much a bus as it is a bus station, with all data flows passing through the service bus (middleware) server or server cluster.

For sensor systems and for emerging IoT systems, middleware has been recast as a platform that provides an abstraction layer capability that separates the sensors and communications from the applications that use it and hides the details of the physical system from the applications. While this is consistent with the layering principle, the underlying structure of actual implementations contains the same structural constraint as the previous version of middleware, as Figure E.4 shows.




Figure E.4. Middleware Platform Logical Structure

Middleware in this context provides a logical data management layer and is often depicted in a manner that makes it appear to be a distributed structure. In the NASPInet case it can isolate applications from the details of the physical and communication elements of the PMU network. It may supply various services, including pub/sub capability and other data management functions. To accomplish this, the new middleware software is inserted in between the physical system and the applications. The middleware must be hosted on one or more servers and must itself be managed in addition to the other elements of the PMU network. PMU data must flow to the middleware and be redistributed from there. Drilling down into the platform model a bit more reveals the structure illustrated in Figure E.5.




Figure E.5. Platform Middleware Structure Detail

The platform middleware resides on a server or server cluster in a data or operations center, whereas the data sources and control elements are distributed throughout the electrical infrastructure in the power grid case. This middleware model derives from the basic IoT scenario, which presumes the field devices are independent and decoupled. This model can work for the older PMU use cases but does not match the forward-looking use cases related to closed loop protection and control, as well as potential use of distributed analytics and grid intelligence. Figure E.6 illustrates the essential mismatch between the IoT model and the likely future grid model.

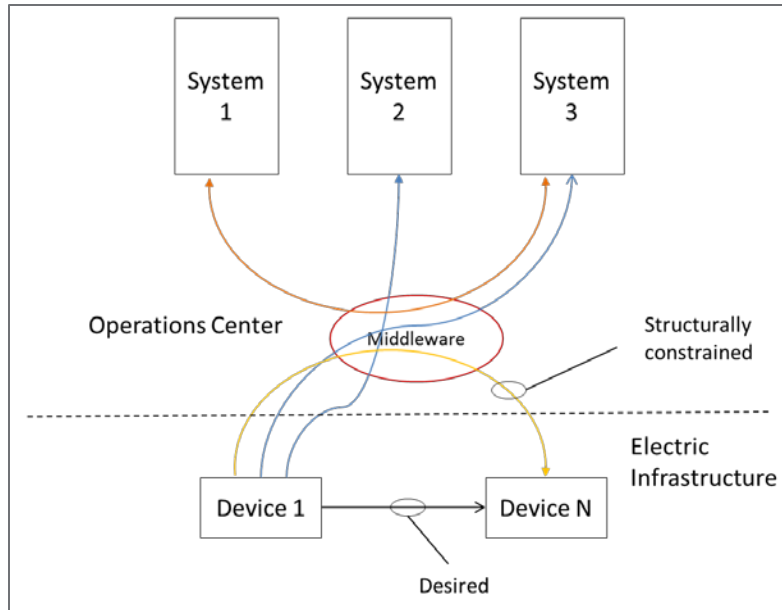


Figure E.6. PMU Model vs. IoT Model

An essential difference between the IoT and future PMU models is the need for peer-to-peer data flow. Forcing that flow through a middleware system located at an operations center results in the middleware becoming a bottleneck, a source of excess latency, and a new cyber vulnerability.

The problem here is that while an abstraction layer is helpful from a software standpoint, the underlying structure imposed by the implementation is not appropriate for a system that involves a highly decentralized infrastructure, such as the electric grid. But another approach is possible. In the grid infrastructure, the communication network *will be present in any event and the PMU data will transit that network*. It is possible to use the built-in capabilities of the communication network to provide essential services associated with middleware without the introducing the essential bottleneck structural constraint.

Instead of introducing new software components (the middleware), the sensor/communication infrastructure approach takes advantage of modern networking technology to implement the layer directly. In this approach, the sensor/communication layer is highly distributed and has multiple connection points for devices and applications. Many data flow paths are possible because there is no need to move all the data to any one point first and so low latency data flow paths for real time uses are available. Figure E.7 shows the sensor/communication model and the middleware model.

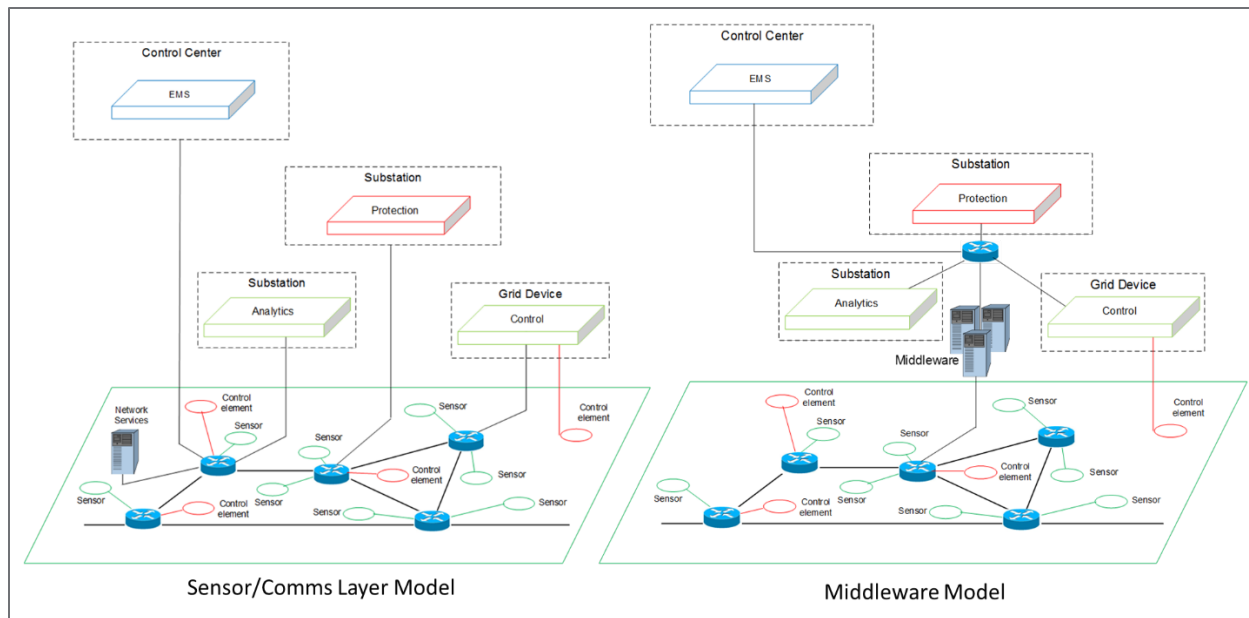


Figure E.7. Sensor/Communications and Middleware Platform Layer Models

Comparing the two structures, it is clear that while the middleware acts as an abstraction layer from a logical function perspective, from the physical perspective it is not so much a layer as it is a structural data path constraint. The middleware approach provides a mechanism for implementing a flexible set of data management services and physical system abstraction, but the structure introduces a bottleneck. It also introduces coupling vulnerability among the applications that use the middleware (if the middleware fails, so does the set of applications using it) and adds latency to the data flows due to the requirement that the data pass through the middleware system.⁴⁷

By pushing certain data management functions “downward” in a sense to *the already distributed infrastructure of the communications network*, it is possible to relieve the structural constraint that could be induced by the use of middleware. The sensor/communications layer can supply data flow management including pub/sub operation, and can supply various data functions with network-attached services.⁴⁸ As long as the communication network has multiple access points and suitable path redundancy, the sensor/communication layer model is a true distributed model at both the logical and physical levels and so does not contain an essential bottleneck or single point of failure. Data flows with low latency requirements are not constrained by the need to pass through a single point and no structural application coupling is introduced. This means that the sensor/communications layer approach can have a greater degree of structural resilience than the middleware approach.⁴⁹

Finally, while the middleware might be able to provide FCAPS capability for the PMUs, the communication Network Management System can do this for both the communication network and the PMUs.

⁴⁷ Note that the middleware server may actually be a server cluster, which helps with fault tolerance.

⁴⁸ For a description of this approach as demonstrated to NASPI, see Cisco staff, PMU Networking with IP Multicast, available online at http://www.cisco.com/c/en/us/products/collateral/routers/2000-series-connected-grid-routers/whitepaper_c11-697665.html.

⁴⁹ S Widergren, et. al., Toward a Practical Theory of Grid Resilience A Grid Architectural Approach, PNNL 27458, April 2018, available online: https://gridarchitecture.pnnl.gov/media/advanced/Theory_of_Grid_Resilience_final_GMLC.pdf, see especially Appendix C.

Middleware may make sense for use inside a control or data center if the conduit effect does not pose a problem, but in an environment where PMU data usage is highly distributed and may be part of low latency closed loop functions, the structural issues discussed above should be considered.