Sensing and Measurement Architecture for Grid Modernization

Jeffrey Taft, PhD
Chief Architect for Electric Grid Transformation
Pacific Northwest National Laboratory

Paul De Martini
Visiting Scholar
California Institute of Technology

Pacific Northwest
NATIONAL LABORATORY

Proudly Operated by Battelle *Since 1965*

Caltech
Resnick Institute

February 2016

Sensing and Measurement Architecture for Grid Modernization

Jeffrey Taft   Paul De Martini

February 2016

Contents

Figures

1.0   Introduction

This paper addresses architecture for grid sensor networks, with primary emphasis on distribution grids. It describes a forward-looking view of sensor network architecture for advanced distribution grids, and discusses key regulatory, financial, and planning issues.

2.0   Sensing and Measurement Network Architecture Technical Issues

This section describes that standard approach to sensing and measurement for power grids, especially distribution systems. It then presents a different approach based on restructuring from the usual vertical organization to one that uses horizontal layering to break up silos and decouple applications.

## 2.1   Existing Approaches to Grid Sensing and Measurement

Grid sensors have generally been associated with specific systems or applications and have been deployed as adjuncts to those systems or applications. Consequently, they have not generally been treated as network structures with architecture and relevant standards. Sensor system architecture is a subset of grid architecture that cuts across electric infrastructure, ICT[1] networks, control and coordination structures, in addition to data management structures and starts with a consideration of requirements as driven by emerging trends and public policy, resulting in a set of desired grid qualities and necessary grid properties.[2]

Grid sensor architecture must consider the underlying physical system structure, the relationship to communications network structure, and the relationship or relationships to applications that make use of sensor data. It is helpful to view grid sensor and measurement systems abstractly in a layer format, as shown in Figure 1 below. As with other grid architecture work, these structures should be considered together, especially in the case where new communications networking is being developed along with the other structures, as would be the case in much distribution grid modernization. Existing legacy components and structures must be viewed as constraints as well as assets in the sensor architecture development and subsequent design processes.

---

[1] Information and Communication Technology

[2] J. Taft and A. Becker-Dippmann, Grid Architecture, available online at
http://energy.gov/sites/prod/files/2015/04/f22/QER%20Analysis%20-%20Grid%20Architecture_0.pdf
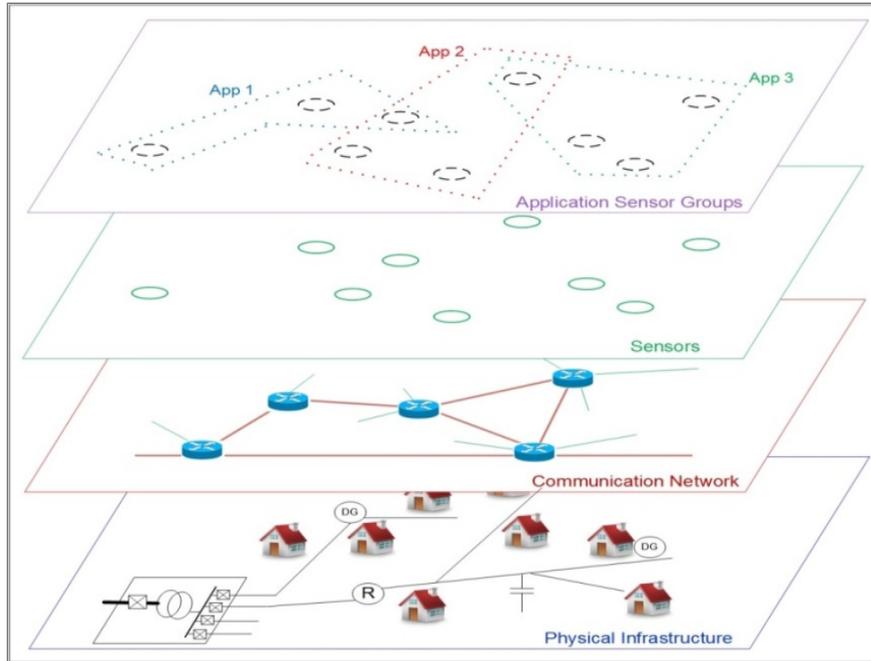
Figure 1. Sensor Architecture Abstraction Layer Model

In most grid systems, sensors are rigidly bound to specific systems and often form disjoint sets. This is illustrated in Figure 2 below, which is essentially an application/SCADA/communication network/sensor set stack. Note that the essential structures are vertical, leading naturally to silos. It is these silos which are the source of the essential limitations that the new architecture addresses.
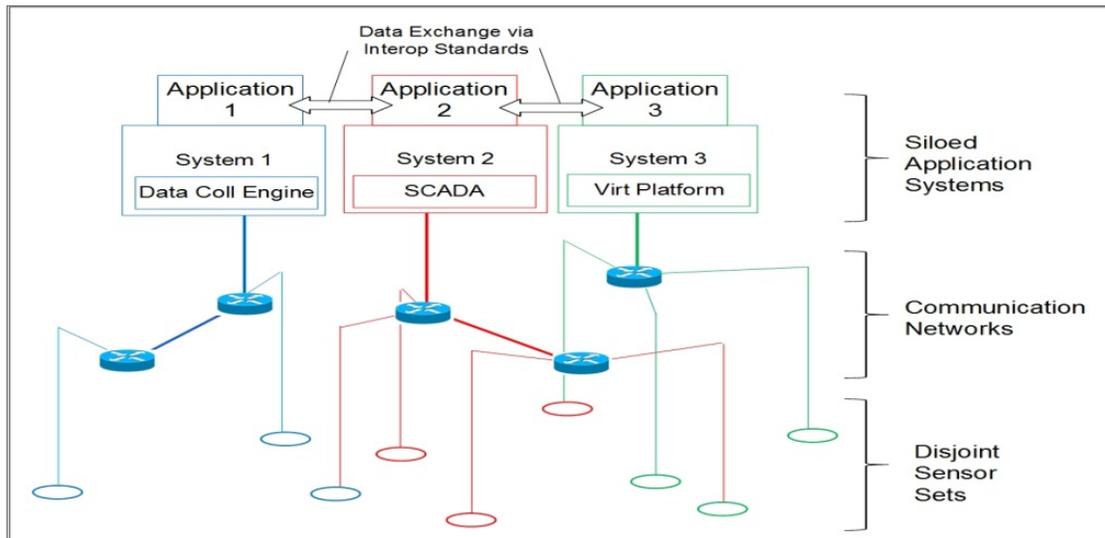

Figure 2. Traditional Grid Sensor System Structure

Communications networks for grid sensors are generally hub-and-spoke, or, in the case of AMI, local mesh to a hub-and spoke backhaul, which is still effectively hub-and-spoke. Such communication systems are often siloed along with the data collection head ends and application systems. The siloed

approach keeps data latency within a given silo system to a reasonable level, but adds significant latency to any other application that must request data from the system that initially acquired it. Each siloed system can perform data conditioning appropriate for the applications in that system, but this treatment may not be appropriate for other systems requesting the data.

In slow distribution systems, data communication latency is not much of an issue, except possibly for protection systems. As grid dynamics increase in speed, it becomes necessary to consider the impact of latency on the function and stability of closed loop controls. The same is true with respect to DER markets, should they exist, because of the close relationship to controls. In siloed system architectures, the need to pass sensor data from system to system via application level interfaces is a source of (often severe) latency.

Consider, for example, AMI data usage, as shown in Figure 3. Meter data is increasingly used by many applications that reside in various systems, both operational and in the back office. For the operational systems in particular, the access and latency issues become significant as distribution dynamics increase in speed. Such arrangements also constitute the imposition of a second low capability network on top of the basic communication network, which is known in network design to be a weak approach because the superimposed network dominates overall performance, chiefly by introducing unnecessary latency and complicating network management.
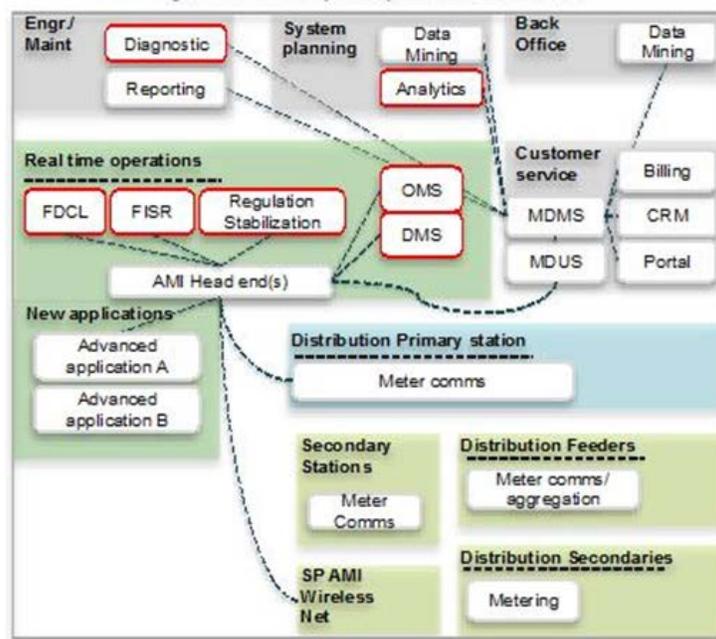
Figure 3. Multiple Uses of AMI Data

The siloed and closed system with back-end interoperability approach is a very expensive and difficult to maintain architecture and correspondingly increases system complexity and failure risks. The system integration costs for such systems deployed over the past ten years have run 3-5x of the operational software licensing costs. Ongoing maintenance of the evolving number of backend interfaces, as illustrated in Figure 3 also create significantly higher ongoing expense. These issues have been recognized for some time, but without a structural change to the sensor architecture, the result will be ever increasing cost and performance risk issues. Back-end information interoperability standards do not resolve this structural problem.

2.2   Data/Application Separation Architectures

A concept that has been proposed for use in grid system design is the separation of data from applications. In that model, all grid data is stored in a central multi-tiered data management system, instead of being stored in siloed sets by the applications systems. The software platform for this data management system runs on servers in the control center or data center. Each application accesses the common data store via open standard interfaces. The data platform approach is still vertically structured and the disadvantages of this approach are that it's essentially a centralized structure that does not accommodate distributed solutions well, and has high inherent latency due to indirect access to grid data, which must pass from data acquisition engines of head ends to data management platform to data stores before becoming available to applications. Such approaches have actually been developed[3] but most grid system vendors have tended to retain siloed architectures due to vendors using closed ("walled garden") product designs; the vendors have gradually added system level interfaces to support minimal interoperability requirements from utilities.

 In the sensor network architecture approach described below, the application silo constraint can be removed so that sensors may not only form application-specific groups with equally low latency to any application, but sensors may group *dynamically* and in *overlapping sets*. Such capability can greatly enhance overall grid adaptability by resolving the data access problem for applications in a manner that is more efficient and scalable than making use of interoperability standards to transfer sensor data from one application system to another. This architectural re-structuring transforms grid sensing and measurement from vertical sensor/data acquisition siloes to a sensor/communication network infrastructure layer that functions in a highly distributed fashion.

2.3   Advanced Distribution Sensor Network

In this section, we re-structure the sensor/network/data collection head end/application stack by partitioning horizontally to group the sensors and communications network into a single structure, and thereby separate the silos and decouple the applications from each other. This produces a *sensor network* for distribution grids that eliminates the need for exchange of sensor data among application systems and provides flexibility and scalability for both centralized and distributed systems. Multiple uses of field operational data are becoming increasingly necessary, and not just for AMI data as was depicted in Figure 3 above but for all kinds of grid and DER state data.[4]

The core concepts on which this new architectural approach is constructed are:

- Combination of streaming sensors and communications into a single structure (the sensor network)
- Dynamic sensor grouping and application binding
- Synchronized data sampling
- Use of network protocols and services as integral parts of the sensor network
- Network level cyber security

---

[3] Accenture staff, Accenture Launches Smart Grid Data Management Solution, March 2010, available online at https://newsroom.accenture.com/industries/energy/accenture-launches-smart-grid-data-management-solution-to-reduce-risks-and-costs-smart-grid-deployments.htm

[4] P. De Martini, L. von Prellwitz, and J. Taft, "Utility Data Management & Intelligence," available online at http://www.cisco.com/web/strategy/docs/energy/managing_utility_data_intelligence.pdf

It is possible to combine the sensors and communication network into a single structure that provides *grid data services* to applications in a highly flexible and scalable manner. Given sufficient network capacity, this structure can scale to large numbers of sensors, with incremental additions requiring minimal integration effort. Due to its inherently distributed nature, it can support multiple simultaneous centralized and decentralized models for application implementation, including Laminar Coordination, distributed intelligence, and multi-agent systems. While eliminating the silo effect, it allows each application vendor to retain control over the value that vendor creates when processing low level sensor data.

2.4   Basic Structure of the Sensor Network

The sensor network architectural view treats sensors and the communication network as an integrated structure. Various services are inserted into this structure and, where possible, the structure employs advanced communication protocols to provide capabilities often either built into siloed applications or supplied via an abstraction layer software platform. Data can flow from streaming sensors (sensors that produce continual streams of data, much like PMUs[5] or video) to any authorized recipient application; in fact multiple devices or applications can receive such streams from the same sensor– applications merely need to be connected to the network at some point – in simple terms, "plug and play". In that sense, the sensor network can operate as a publish-and-subscribe data system. Such operation for sensors has been described and demonstrated in the context of PMU networks.[6]

For sensors that do not have streaming capability, data acquisition engines may be attached to the edge of the sensor network to perform more traditional polling and other modes of data collection. Hence both legacy SCADA and more distributed data collection can coexist on the same network. Similarly, distributed database data store nodes may be attached to the sensor network, or data may be accumulated into individual applications. Each application may associate sensors as needed, providing low-latency grid data access with great flexibility.

Various services can be integrated into the sensor network via attached servers or through integration into network management systems. These include standard network management and security functions as well as grid-specific capabilities such as sensor meta-data management, IEC 61850 CIM interface services, and grid topology/connectivity.

This structure can serve as a sensor data platform without the latency caused by passing through intermediate layers of software or transfer of data from one application system to another. It provides more flexibility than approaches that separate data from applications but then store the data in centralized data stores. New data sources (sensors) can be added by simple network attachment/admission, and new users of data can access data for which they are authorized in real time without intermediate virtualization layer and application or data collection system latencies. Compare this to the typical scenario where one system collects data from a set of dedicated grid sensors, then stores the data into an internal historian, and periodically copies the data over to a "shadow" historian, which then may be queried by another application system that wants the (by then very stale) data by way of an interoperability standard that may be several layers higher in abstraction than is needed by the system requesting the data.[7]

---

[5] Phasor Measurement Units

[6] 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

[7] Consider a voltage control application that accesses a smart meter directly through the network vs. one that sends a request to a meter data head end, waits for the head end to query the meter, and waits until the head end provides a value back to the voltage control application after the meter responds.

Legacy sensors and simple (non-smart) sensors and transducers can be used in an advanced sensor network by incorporating the concept of *micro-virtualization*. Distributed processing capability in the network provides local sensor abstraction for one or a few sensors attached at each of these points, with as many micro-virtualization platforms as needed in a given network. Processing capability for sensor micro-virtualization can be embedded in communication devices or can be attached to the network at or near the sensor site.

Sensor micro-virtualization can be employed in another manner as well. It enables the separation of smart sensors into two parts: a transducer with basic sampling, conversion, and communication capability, and a network node that performs processing to transform raw samples into more abstract quantities, such as phasors. In this manner multiple transducers can share a single computing resource on a localized basis so that Remote Terminal Units and PMUs can be virtualized across multiple sensing points.

2.5   Dynamic Sensor Grouping

Given a network of smart sensors, it is possible to have applications associate to sensor subsets in a general and flexible manner. Applications can merely subscribe to the data from the appropriate set of sensors and sensor sets do not have to be disjoint, but can all be treated independently and grouped virtually by applications as needed. The sensor network and the applications must have certain capabilities for this to work autonomously (it is always possible for a human network manager to specify associations but for real flexibility, the association process should be automatic). The set of capabilities that are needed include:[8]

- Discovery – applications must have ways to discover the sensors they need

- Advertisement – sensors must be able to advertise their presence on the network and to describe their capabilities and externally observable characteristics

- Binding and access control – sensors and applications must have means and protocols to agree on data stream subscription, including function, location, and security criteria

- Precedence resolution – for sensors that require control inputs to set parameters, a mechanism is needed to resolve which application has precedence when sensors can be shared; how precedence is established, when it expires, etc.

- Security – there must be means to manage data security across multiple overlapping groups and applications, as well as means to determine when sensors are lying or are malfunctioning

The sensor network physical and logical structures are illustrated in Figure 4 below.

---

[8] Based on a discussion with Rick Geiger of Cisco Systems in the context of Internet of Things (IoT).
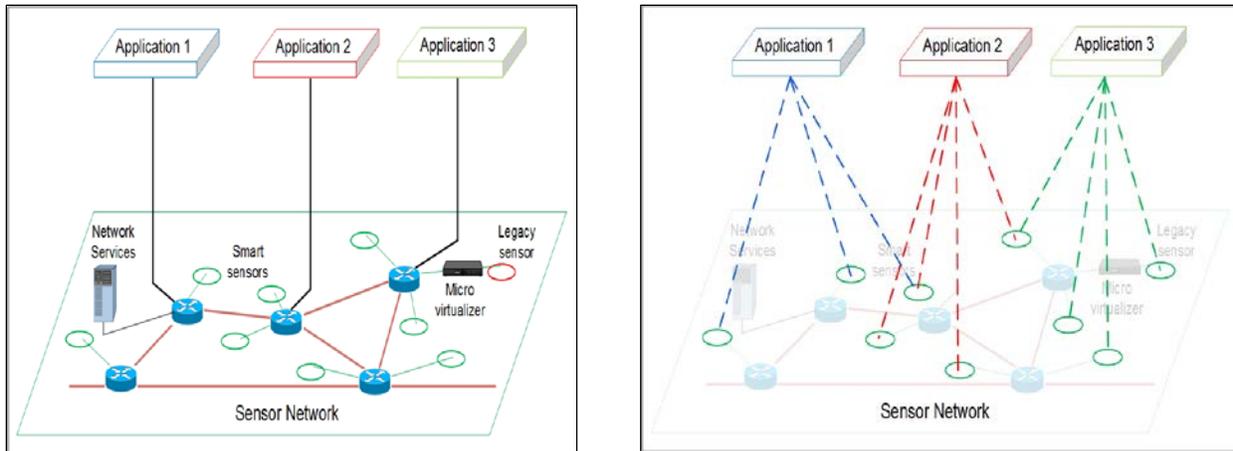
Figure 4. Physical and Logical Sensor Network Structure

The diagram on the left of Figure 4 shows a simplified physical structure in which a mix of smart and legacy sensors is integrated with a communication network and network services. Legacy sensors are micro-virtualized and applications are connected to the network at convenient locations, based on where computing resources are located. This may be in a substation, at a utility pole-mounted processor, in a grid control device, or in a communications device that supports application software.

The diagram on the right of Figure 4 shows how the sensor network appears to various applications. Each application subscribes to data streams from the sensors it needs, thereby obtaining grid data at the lowest possible latency short of direct hard wiring. Micro-virtualization and network services are seamless, so that the applications do not have to be concerned with data acquisition details. In effect, the sensor virtualization platform concept has been moved to the communication network. The original sensor/network/data collection head end/application stack structure shown in Figure 2 has now been re-partitioned into an application set and the services/communication net/sensor plane structure of Figure 4.

2.6 Sensor Network Protocols and Services

The sensor network can use existing communication network protocols along with additional services supplied via attached processing to form a complete sensing and measurement platform. Among the key protocols and services are:

- Direct access to sensors via the IPv6 protocol suite, including broadcast modes

- MPLS and PIM/SSM[9] for handling streaming data and providing a publish and subscribe mechanism for sensor data[10]

- Standard network management functions with extensions for sensor monitoring

- IEC CIM interface services – Generic Data Access, High Speed Data Access, Time Series Data Access, Generic Eventing and Subscription

- Sensor registry service or service advertisement

---

[9] Multi-Protocol Label Switching, Protocol Independent Multicast, and Source Specific Multicast

[10] P. Myrda, et. al., Recommended Approach to a NASPInet Architecture, 2012 45th Hawaii International Conference on System Sciences, January 2012, pp 2072-2081. Available online at https://www.computer.org/csdl/proceedings/hicss/2012/4525/00/4525c072.pdf

- IEEE 1451.4 TEDS (Transducer Electronics Data Sheets) service[11]

- Electrical network connectivity service

- Timing distribution

- Software Defined Networking (SDN) interface services

The sensor registry service facilitates discovery of sensors. Alternately, service advertisement (by each sensor) can be used. The registry approach is easier to manage but is a centralized capability, whereas service advertisement is highly distributed and scalable but is not available from legacy sensors.

In order to separate the grid systems into a sensing network and other structures, sensors must be generally accessible. In conventional systems, sensors typically "belong" to a specific vendor closed system and that system manages the sensor meta-data such as calibration curves. Here the IEEE 1451 standard for smart sensors can be used to help decouple sensors from applications. Smart sensors can incorporate the necessary self-characterizing information directly, but legacy sensors cannot. This is where the network service for TEDS comes into play: it provides the necessary sensor meta-data in a store accessible as a network service to any authorized application.

The electrical connectivity service is needed to provide context for sensor data and control actions. This implies continual re-discovery of electrical connectivity, since connectivity in most distribution systems changes on both short and long time scales. Hence this service has two parts: re-discovery, and application access support.

2.7   Network Level Cyber Security

Security for the sensor network must be an interlay in the network, as opposed to an overlay. In other words, security must be an integral part of the network, not something that is added on after the fact of network deployment. The set of capabilities and services consists of four categories:

- Data integrity, confidentiality, and privacy – includes encryption (inter-nodal or end-to-end), key and certificate management, IPSEC, etc.

- Device and platform integrity – methods to ensure devices and systems have not been comprised at the hardware or code levels, including methods for ensuring supply chain integrity, tamper resistance/detection, signed firmware images, posture assessment, secure software life cycles, etc.

- Access control – identification, authentication, network access control, nodal access (sensor binding to applications), subscription control

- Intrusion detection and mitigation – signature and behavioral analysis, traffic analysis, node exclusion, non-repudiation, network segmentation, VLAN, etc.

Modern IPv6 networks provide multiple layers of capabilities in each of these categories.

---

[11] https://standards.ieee.org/develop/regauth/tut/1451d4.pdf

2.8   Synchronized Data Sampling

To support advanced grid applications involving fast dynamics, synchronized sampling is needed so that sample skew can be minimized or preferably eliminated. Three elements are needed to accomplish this:

1. High precision time distribution

2. Synchronized data sample acquisition

3. Time stamping

Note that time stamping alone is not sufficient. The samples must be *acquired* at as nearly the same time as possible, not just annotated with the time of acquisition. In order to accomplish this, each sampling device, whether it is a smart sensor or a data acquisition engine, must have the same sense of clock time to a high degree of precision, which can be accomplished via network time protocols. The distribution of timing in networks is somewhat complex but is well understood.[12] The complexity arises from the need to operate across network domains, leading to multiple classes of network clocks and protocols for timing distribution.

3.0   Sensing and Measurement Architecture Regulatory and Business Considerations

The architectural approach proposed in this paper raises several regulatory and business considerations related to utility investment planning and regulatory review and authorizations. Specifically, operational communications infrastructure should be treated as a core investment in a modern grid and not subject to a separate cost benefit analysis and rate cases.

As has been discussed for over ten years[13], telecommunications is an integral aspect of a modern grid that supports the sensing, protection and control systems needed for safety and reliable operation of the electric system in the 21^{st} century. It has become clear over this period that the evolution of telecommunications technologies and service provider offerings can enable a unified WAN/FAN communications architecture and infrastructure that supports multiple uses. This is analogous to how utilities use enterprise telecommunications to support voice, data and video. A unified, multi-use operational field network comprised of wide area and field area networks is the best practice in the industry today.

As such, consistent with utility enterprise information and telecommunication investments and physical grid infrastructure, operational communication network investments should be included in general rate cases as core investments in a modern distribution grid that evolves from electro-mechanical systems to integrated, intelligent digital systems. This means that AMI business cases should not include costs for telecommunications WAN/FAN components and costs. There are several reasons for not taking this approach. Focusing operational communications planning on AMI functions tends to not include full capability required for modern grid leading to suboptimal technology selections and investments. Also, to include operational communications, as described in this paper, in an AMI business case disproportionately burdens a smart meter case with costs for a common system that is shared across multiple applications. Finally, AMI case approval processes tend to be controversial, lengthy, and could

---

[12] P. Gaspar, Cisco, Frequency and Time Synchronization In Packet Based Networks, available online at: https://www.cisco.com/web/YU/expo2010/pdfs/Prenos_frekvencije_i_faza_preko_paketskih_mreza.pdf

[13] Modern Grid Initiative, "A Systems View of the Modern Grid", National Energy Technology Laboratory, 2007.

delay deployment of utility operational communications networks needed for T&D safety and reliability, as well as DER integration.

Given these considerations, we make the following recommendations for utility planning and regulatory assessment and authorizations.

3.1 Utility Recommendations

3.1.1 Operational Communication Architecture

- Utilities should develop operational communications architecture to support organizational and policy objectives over a 10+ year horizon

- Identify utility and policy objectives and related attributes that drive functional requirements to support substation and distribution automation, grid sensors, protection schemes, distributed device control, smart metering, and integration and control of DER.

- Perform a review of the current WAN/FAN operational network infrastructure and planned investments in relation to expected connectivity, bandwidth, and latency required, based on industry references, to achieve objectives.

- Analyze the extent to which third party data, information and/or communications infrastructure could be integrated into utilities' information and communication architecture.

3.1.2 WAN/FAN Communication Roadmap

Deployment of an operational field communications system roadmap based on the architecture developed above. This involves:

- Identifying relevant WAN/FAN communications technologies that can address the functional and industry reference requirements.

- Evaluate utility grid modernization and other relevant T&D investment plans (near and long term) to identify required timing and functionality for a WAN/LAN communications deployment in its service area/s.

- Evaluate DER and time-varying rate diffusion scenarios for patterns that would inform the timing and functionality for a WAN/LAN communications deployment.

- Based on the above, provide overall incremental conceptual cost estimate and timing of investment options that provides a vision for the deployment of a WAN/FAN operational communications network to support objectives.

3.2 Regulatory Recommendations

- Operational field communications deployments generally span more than one rate case period and continuity of investment authorization is essential to achieve full functionality for all desired applications.

- Timing and specific investments will evolve due to technological innovation, and with changes in the use of the distribution system to support customer adoption of DER and use of DER for wholesale markets and distribution needs. Utility flexibility in technology deployment post authorization is important.

- Consider a utility's grid architecture[14] as a foundational requirement for investment authorization; the sensing and measurement layer in the overall grid architecture should address the elements described in this report and not just interoperability and cybersecurity.

- As such, any prudency reviews should be based on policy and business outcomes as well as conformance to the utility's architecture, not specific technology or hardware, to allow a utility to evolve with innovation and customer needs.

4.0   Summary and Conclusions

This sensing and measurement architecture has significant implications for electric utilities and regulators in terms of planning, financing, and road mapping for grid modernization, as well as deployment strategies and realized system performance. This is because it offers flexibility, low latency, and scalability at significantly lower cost and operational risk.   This is achieved by significantly changing the view about where interoperability, and related standards, should apply by decoupling sensors from application systems and coupling them to communication networks instead. In summary, this architectural approach:

- Allows for the integration of new sensors while accommodating the integration of existing sensors within a unified sensor network structure.

- Allows for publish and subscribe data access flexibility to address the timing and latency requirements unique to each of the growing number of operational applications employed in a modern grid.

- Allows for a simpler design and implementation devoid of "band-aid" investments to achieve back-end integration and interoperability at a significantly lower cost.

- Creates a sensor network with *structural future-proofing* (the use of architectural structure inherently to safeguard infrastructure investment) that substantially reduces the risk of potential for asset stranding due to accelerating information and communication technology improvement.

Finally, a 21st century grid should consider the sensing network and related communications infrastructure as core investments architected, designed, and implemented with the requirement to support multiple existing and future operational applications over its asset life. This also means that sensing and communications networks should be treated as core investments in rate cases – not as a part of a single application like AMI. Including operational communication investments within the business case of a single application, like AMI or distribution automation, typically leads to a suboptimal point solution and ultimately higher costs. A thoughtful grid architecture that considers the proposed sensing and measurement approach in this paper will serve customers and utilities well into the future.

---

[14] Architecture identifies the entities and components of a system and the structures (interconnections and relationships) among components; the accompanying descriptions document the responsibilities of each entity or component in a technology-neutral manner. It also defines the processes that perform functions and the information or data flows that are shared among these processes. Architecture allows for multiple possible implementations and so does not specify particular equipment or software – that is the role of design, which yields exactly one specific implementation. Architecture does include any definitions of business services, applications, and other relevant information for needed for development purposes.