Exploring a Tiered Architecture for NASPInet

Rakesh Bobba, Erich Heine, Himanshu Khurana, Tim Yardley

Information Trust Institute, University of Illinois

{rbobba,eheine,hkhurana,yardley}@illinois.edu

Abstract—One of the missions of the North American SynchroPhasor Initiative (NASPI) is to create a robust, widely available and secure synchronized data measurement infrastructure, called the NASPI network or NASPInet. Leveraging the Phasor Measurement Unit (PMU), a GPS clock synchronized measurement device capable of measuring the current and voltage phasors in the power grid, NASPI will improve reliability of the power grid with NASPInet providing data delivery. While a conceptual architecture of NASPInet and its functional requirements have been developed, in this work we address the challenge of designing a deployable architecture that can realize NASPInet on a continental scale. To do so, we explore a tiered architecture for NASPInet and analyze its impact on the Quality of Service (QoS), cyber security and network management services. Furthermore, we discuss the distributed computing opportunities afforded by our architecture.

Index Terms—PMU SynchroPhasors NASPI

I. INTRODUCTION

Wide-Area Situational Awareness (WASA) is recognized as a key enabling functionality for Smart Grids. It is one of the priority areas identified by Federal Energy Regulatory Commission (FERC) [1] and by the National Institute of Standards and Technology (NIST) Smart Grid Interoperability Standards Effort [2]. Time synchronized and precise grid measurements from Phasor Measurement Units (PMUs), called synchrophasor measurements, can provide a comprehensive view of the entire interconnection, when measurements from multiple utilities are combined, and improve wide-area situational awareness. Recognizing this potential, U.S. Department of Energy (DOE), North American Electric Reliability Corporation (NERC), and North American electric utilities, vendors, consultants, federal and private researchers and academics are collaborating on the North American SynchroPhasor Initiative (NASPI) [3], whose vision is to improve power system reliability through wide-area measurement, monitoring and control. NASPI is working to develop an "industrial grade", secure, standardized, distributed, and expandable data communications infrastructure, called the NASPI network or NASPInet, to support synchrophasor applications that depend on shared PMU data. A conceptual architecture of NASPInet and its functional requirements are captured in NASPInet specification documents [4], [5] commissioned by NASPI. However, a deployable architecture that can realize NASPInet on a continental scale is yet to be designed and developed.

The U.S. DOE recently announced selections for Smart Grid Investment Grant (SGIG) awards [6] which include projects that aim to install hundreds of additional PMUs (in total) across North America and develop associated communications and data sharing infrastructure. If the design of a realisticly deployable NASPInet architecture is available, then the implementations that come out of these projects have the potential to coalesce into the beginnings of a continental scale NASPInet. Furthermore, many PMU applications are being rolled out in North America and around the world as described in [7]. These developments provide the impetus for designing a deployable, secure, robust, expandable and continental scale distributed communications infrastructure that can realize the functionality of NASPInet.

However, the design of a deployable NASPInet architecture has to consider many challenges. First, there is the challenge of designing a distributed continental scale network. Second, it is challenging to meet Quality of Service (QoS) requirements of PMU applications some of which have very stringent latency requirements, for example 100*ms* for feedback control applications, over a continental scale network such as NASPInet. Third, it is challenging to provide cyber security for NASPInet and PMU data, *i.e.*, ensure availability of NASPInet, and availability, integrity and confidentiality of PMU data, even as NASPInet is expected to enable PMU data sharing among hundreds of entities. Specifically, NASPInet aims to enable data sharing based on unicast, multicast, and publish subscribe paradigms where the groups of recipients change dynamically in the latter two cases. This makes key and trust management all the more challenging. Finally, there is the challenge of managing a network with a large and diverse geographical footprint like NASPInet.

In this work, we explore the design of a realisticly deployable tiered NASPInet architecture that leverages the existing hierarchy formed by power grid operators, monitors and regulators. The use of tiered architectures for information exchange in large distributed systems is not uncommon and, in fact, is considered advantageous in other emerging critical infrastructures such as Electronic Health Records [8]. We discuss how this tiered architecture positively impacts the desired properties of Quality of Service (QoS), cyber security and simplified network management services specific to the goals and functional requirements for NASPInet. Furthermore, we discuss the distributed computing opportunities afforded by this architecture.

In Section II we provide some background on the control hierarchy of the power grid, NASPI and NASPInet. In Section III we describe our proposed architecture. In Sections IV, V and VI we discuss how the proposed architecture impacts the implementation of Quality of Service (QoS), cyber security and network management services respectively in NASPInet. In Section VII we discuss the distributed computing opportunities provided by our architecture and conclude in Section VIII.

II. Background

A. Power Grid

The North American electric power grid is a highly interconnected system hailed as one of the greatest engineering feats of the 20th century. However, increasing demand for electricity and an aging infrastructure are putting increasing pressure on the reliability and safety of the grid as witnessed in recent blackouts [9], [10]. Furthermore, deregulation of the power industry has moved it away from vertically integrated centralized operations to coordinated decentralized operations. Reliability Coordinators (RCs) such as Independent System Operators (ISOs) or Regional Transmission Operators (RTOs) are tasked by Federal Energy Regulation Commission (FERC) and North American Electric Reliability Council (NERC) with overseeing reliable operation of the grid and providing reliability coordination and oversight over a wide area. Balancing Authorities (BAs) are tasked with balancing load, generation and scheduled interchange in real-time in a given Balancing Authority Area (BAA). A BAA is a geographic area where a single entity balances generation and loads in real-time to maintain reliable operation. BAA's are the primary operational entities that are subject to NERC regulatory standards for reliability. Every generator, transmission facility, and end-use customer is in a BAA.

B. NASPI

In order to improve the reliability of the power grid while meeting the increased power demand, the industry is moving towards wide-area measurement, monitoring and control. NASPI was formed with this in mind and it's mission is to create a robust, widely available and secure synchronized data measurement infrastructure with associated monitoring and analysis tools for better planning and reliable operation of the power grid. NASPI envisions deployment of hundreds of thousands of Phasor Measurement Units (PMUs) across the grid that send data at 30 to 120 samples/second to hundreds of applications in approximately 140 BAAs across the country. The GPS clock synchronized PMUs are sensors that can read current and voltage phasors at a substation bus on the transmission power network. These PMUs give direct access to the state of the grid at any given instant in contrast to having to estimate the state as is done today. Phasor Data Concentrators (PDCs) at substations or control centers receive and time align the data from multiple PMUs before providing them to historical archives or applications.

NASPI applications envisioned to utilize PMU data have varying requirements classified into four classes based on their data requirements as shown in Table I. Typically feedback control applications like transient stability control fall into Class A, open loop control applications like state estimation fall into Class B, situational awareness applications like visualization and monitoring fall into Class C and post event analysis applications like disturbance analysis fall into Class D.

TABLE I
PMU APPLICATION CLASSES

	Class A	Class B	Class C	Class D
Low Latency	Critical	Fairly Critical	Somewhat Critical	Not Critical
Reliability / Availability	Critical	Somewhat Critical	Not Critical	Fairly Critical
Data Accuracy	Critical	Somewhat Critical	Not Critical	Critical
Time Alignment	Critical	Critical	Somewhat Critical	Not Critical
Message Rate	Critical	Somewhat Critical	Somewhat Critical	Critical
Sample Application	Small Signal Stability	State Estimation	Visualization and Monitoring	Disturbance Analysis

C. NASPInet

Sharing PMU data widely, *i.e.*, with other utilities, provides wide area situational awareness which is identified as a priority both by FERC and NIST. NASPI is working to design a wide area network infrastructure, dubbed NASPInet, to enable wide area sharing of PMU data. Figure 1 shows a high-level conceptual architecture envisioned for NASPInet [4], [5] . NASPInet will be composed of Phasor Gateways (PGWs) and a Data Bus (DB). The DB includes a Wide Area Network (WAN) and associated services to provide basic connectivity, QoS management, performance monitoring, and cyber security and policy enforcement over data exchanged through NASPInet. PGW is the sole access point of entities like utilities and monitoring centers (*i.e.*RCs) to the DB. The PGW will manage the connected devices on the entity's side, manage QoS, administer cyber security and access rights, perform necessary data conversions and interface the entity's own network with the DB. NASPInet is intended to facilitate the secure exchange of both real-time streaming data and historical data. PGWs are expected to support both one-to-one unicast data sharing and one-to-many publisher-subscriber based data sharing in an efficient manner.

III. Architecture

While NASPInet specification documents [4], [5] capture the functional requirements for both PGWs and the DB, a realistic architecture for realizing the DB is yet to be developed. We observe that the inherent structure of the power grid provides two key benefits for the design of such an architecture. The first benefit is that the areas of concern for a given utility are generally geographically localized rather than globally present. Utilities primarily care about their own electrical network and the buses that are connected to their network but not necessarily owned or controlled by them. This concept of locality generally implies that PMU data from nearby utilities is likely to be used more often than data from ones farther away. While there are examples of special protection systems (SPSs) that require real-time data from sensors thousands of miles away, there are many more applications that use regional data. Furthermore, data from distant PMUs may be used in less stringent real-time applications such as post event analysis and planning applications. The second benefit is that utilities already share data with their reliability coordinators (RCs) on a regular basis and have already established trust relations. Reliability coordinators are
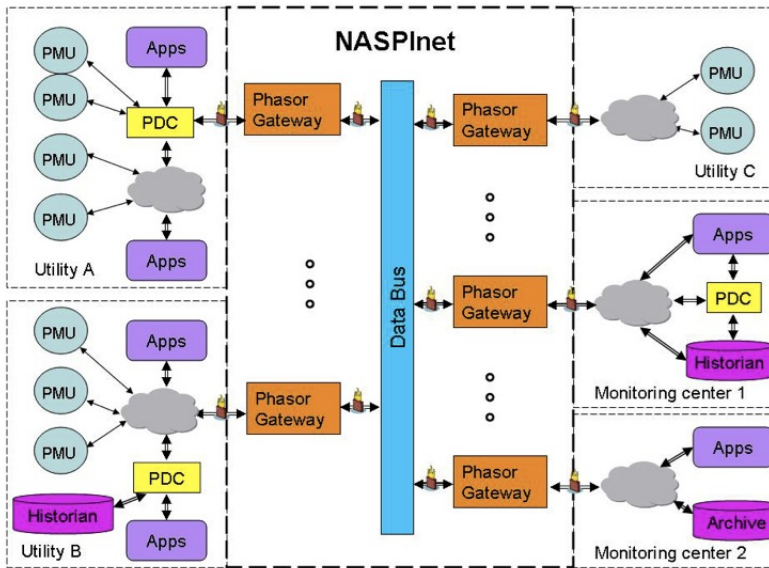
Fig. 1.   NASPInet Conceptual Architecture [4], [5]

also increasingly sharing control responsibilities with utilities and coordinating actions for a given geographical area when needed. Thus, these trust relations are deeply connected to regulatory issues and motivate further leveraging of these relations rather than building entirely new ones. Keeping these two aspects in mind we explore a hierarchical or tiered approach to designing NASPInet architecture.

A. The Proposed Architecture

Figure 2 shows the proposed NASPInet architecture. Keeping to the vision provided by the conceptual architecture shown in Figure 1, the proposed architecture is composed of PGWs that are connected by a DB. In the proposed architecture the DB consists of Hubs and a managed real-time secure wide area network, called the core network, that connects the Hubs. PGWs will connect to their local Hub using dedicated, managed, real-time secure high speed links. We refer to this as the Hub network. We envision that RCs, or delegated/outsourced entities acting on behalf of RCs, will act as the Hub for the BAs and other grid entities under their jurisdiction. The choice of RCs acting as Hubs leverages the existing trust relations between BAs that usually host PGWs and the RCs.

This hierarchical structure leverages a simple hub-and-spoke topology in the Hub networks which simplifies providing QoS guarantees for real-time data sharing between PGWs in the Hub network. However, the hub-and-spoke topology also means that the Hub is a single point of failure. Even though this failure impact will be localized to the Hub network it is not desirable. To mitigate this we propose replicating the Hubs and provisioning backup links that connect PGWs to the Hub for improved resiliency. For the case of a PGW having to share real-time data to be used for Class A applications with another PGW in a different Hub network, we propose an additional direct dedicated link to be installed

between the PGWs. Furthermore, the proposed hierarchical architecture provides natural points of aggregation (*e.g.*for publish-subscribe paradigm), potential benefits of coordinated control, and opportunity to leverage the locality to reduce costs (*e.g.*, bandwidth, latency, provisioning) and increase efficiency. We will discuss the specific advantages provided by this hierarchical architecture for implementing QoS, cyber security, and network management in the following sections. We propose that Hubs in this architecture be sophisticated nodes with a lot of computation and storage capabilities and not just simple routers or gateways. This affords us a lot of distributed computing opportunities which will be discussed in Section VII.

IV. QUALITY OF SERVICE

Supporting class-based data delivery over a wide area network like NASPInet requires the network to provide Quality of Service (QoS). That is, it should be able to provide different priorities to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow. For a given data flow this may mean ensuring a required bit rate, delay, jitter, maximum packet dropping probability and/or bit error rate, and timely delivery. For a given network this may mean supporting dedicated bandwidth, resource provisioning and allocation, avoiding and managing network congestion, shaping network traffic and managing priorities in an end-to-end manner. Networks that provide these capabilities are collectively referred to as QoS-managed networks, or simply managed networks. QoS managed wide area networks are no longer a rarity today. In fact, even Internet systems can provide QoS to a limited extent. However, providing QoS for a continental scale network such as NASPInet with class or application-specific delivery assurances described above is a difficult challenge that faces hurdles in cost-effective design and deployment, allocation and provisioning, and priority management. In order to achieve QoS for applications over
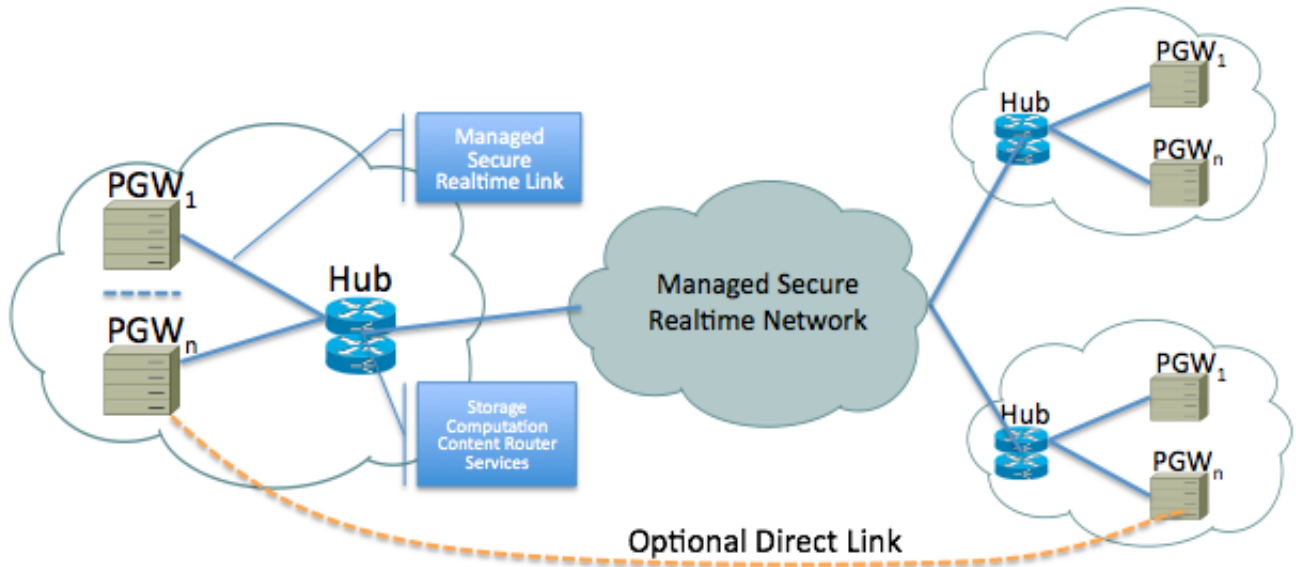
Fig. 2. Proposed Tiered NASPInet Architecture

NASPInet, *each* network element and segment needs to provide certain QoS management functions. These functions will vary depending on the element characteristics and capabilities as well as the segment they are part of.

A. Benefits of the Tiered Architecture approach

If NASPInet is viewed as a single network providing data exchange capabilities between n nodes (i.e., PGWs) then ensuring adequate QoS between any two nodes quickly scales to be an $O(n^2)$ problem. In practice, however, such networks are designed using a hub-and-spoke model with end nodes being connected using point-to-point links to managed wide-area backbones. Such an approach creates natural boundaries for addressing each of major hurdles, namely, design and deployment, provisioning and allocation, and priority management. The next set of questions for a network like NASPInet then emerge: how big should such a managed wide-area backbone be? Should it be a single network or multiple connected ones? If there are multiple ones then how can end-to-end properties be ensured?

In this work we address some these concerns by proposing a tiered architecture that offers a natural hierarchy for simplifying these challenges. Specifically, we argue that by structuring the network into two main tiers, one based on current regulation-driven data-paths between utility PGWs and Hubs and the other based on static multi-path connections between Hubs, we offer a vision that simplifies QoS support.

Network links between utilities and RCs are carrying increasing amounts of data associated with a range of applications. To handle this traffic most utilities are investing in high-speed, high-bandwidths links. This links can also be leveraged for sending PMU data thereby simplifying design and deployment. Careful provisioning and allocation of resources is still needed to ensure that these links can support data needs for the foreseeable future. However, since these are point-to-point links upgrading them in the future becomes easier to handle.

Priority management needs to leverage globally designed services and face an interoperability challenge, however, careful provisioning of these links can simplify the process for priority management.

There exist networks like NERCnet that already connect RCs and other power grid entities. Such networks can be leveraged to establish a new (or upgraded) core multipoint Hub network for PMU data sharing. This higher tier network will involve significant effort but because of fewer and relative static end points we argue that it will be easier to provide adequte QoS. Relevant QoS issues have been explored in the literature, in general [11], and, in particular, for Smart Grids [12], [13].

V. Cyber Security

Applications utilizing PMU data shared over NASPInet range from feedback control applications such as out-of-step protection to situational awareness applications such as real-time compliance monitoring. It is crucial to secure NASPInet and the data that traverses it to ensure the availability and integrity of the applications relying on the data which in turn affect the reliability of the power grid. Specifically, NASPInet should ensure the availability, integrity and confidentiality of PMU data. Ensuring these security properties over a wide-area distributed network such as NASPInet involves implementing many security functions and mechanisms some of which are discussed below.

Intrusion Tolerance, Detection, Recovery and Response or Infrastructure Security: In order to ensure availability of PMU data, NASPInet should provide a reliable data sharing infrastructure that is highly available, *i.e.*, always up and running. While a fault tolerant design that avoids single points of failure will help with this goal, such a design alone will not be sufficient. NASPInet, being part of a critical infrastructure, will likely be a target for many cyber attacks with adversaries ranging from script kiddies to nation states. Thus,

the NASPInet architecture should be resilient against cyber attacks and intrusions. That is, there should be mechanisms in place to, 1) protect NASPInet from cyber attacks, 2) monitor for and detect cyber attacks and intrusions, and 3) recover from and respond to cyber attacks and intrusions in a timely fashion.

Network Admission Control: As the name indicates, Network Admission Control (NAC) refers to controlling admission into the network. In the case of NASPInet, only authorized PGWs should be allowed to connect to the DB and communicate with other PGWs. Without such admission control, rogue entities might access the DB and create unwanted traffic that will consume network resources and make it difficult for NASPInet to meet the QoS requirements of legitimate data flows or worse it could cause denial of service to legitimate data. NAC is a protective mechanism and one tool to ensure infrastructure security.

End-to-end Integrity and Confidentiality Protection for PMU Data: Inadvertent or malicious modification of PMU data traversing NASPInet must be detected. Without such protection, operators or applications can be led into making catastrophic decisions based on false data. Furthermore, the confidentiality of PMU data should also be protected from malicious eavesdroppers as PMU data can reveal sensitive information about the state of the grid which can be leveraged by malicious entities in disrupting grid operation. PMU data should be integrity and confidentiality protected all the way from the source PGW that is sending the data to the destination PGW that is receiving the data, *i.e.*, end-to-end. Furthermore, mechanisms used to protect integrity and confidentiality of PMU data should not adversely affect the availability of PMU data and should be designed to be as lightweight as possible to satisfy the necessary requirements. Apart from unicast data sharing, *i.e.*, one sender to one receiver, NASPInet aims to support multicast and publisher/subscriber based data sharing with dynamic recipients, *i.e.*, one sender to multiple receivers that change dynamically. Integrity and confidentiality protection should be provided for such cases as well.

Logging and Auditing: Logging and auditing refers to the process of securely logging all activity and events on NASPInet and then auditing those logs to identify anomalies and suspicious activities or events. Activity and event logging in NASPInet will provide accountability and act as a deterrent to miscreants. Logging and auditing forms a good starting basis for intrusion detection.

All the security functions and mechanisms discussed above depend on two core security services, namely trust management, and key establishment and management. For example, to implement network admission control there should be a trusted entity or a trusted collection of entities that define the admission policies and enforce them. Even for simple admission policies based on Access Control Lists (ACLs), entities requesting access need to be authenticated. This supposes that the requesting entity has a digital credential, *e.g.*, identity certificate issued by a trusted entity, or a cryptographic key shared with the network administrator. Thus network admission control depends on pre-established trust and keys. Similarly, end-to-end integrity and confidentiality protection

depend on pre-established cryptographic keys between the source PGW and destination PGW. These keys are presumably established after the two PGWs get to know each other, *i.e.*, after trust establishment.

A. Benefits of a Tiered Architecture

In this section we discuss how the implementation of the security services, functions and mechanisms benefit from the tiered NASPInet architecture proposed in Section III.

Without a tiered architecture NASPInet will presumably be a managed real-time network, administered by a single entity, connecting all the PGWs. For such a network, a reasonable design would be to delegate the management of network infrastructure security, *i.e.*, monitoring, logging and auditing, intrusion detection, recovery and response, to the network provider itself. To implement the other security functions such as key and trust management discussed above in such a network there are two potential options. The first option is to leverage a trusted entity to provide the security services. That is, the trusted entity will, 1) distribute credentials to PGWs for network admission control, and 2) mediate trust and key establishment among PGWs for end-to-end integrity and confidentiality protection of PMU data. It has been suggested [4], [5] that NERC could play this role. However a disadvantage of this option is that it exposes a single point of failure, namely the central trusted entity that provides key and trust management. Compromise of this trusted entity or any one of the services provided by this trusted entity could significantly degrade the security and availability of NASPInet. Furthermore, given that NASPInet is envisioned to support hundreds of PGWs, key and trust management for such a potentially dynamic population could be a significant undertaking for a single entity. An alternative option is to implement these security services and mechanisms in a distributed manner based on peer-to-peer paradigm where all the PGWs connected to NASPInet collectively assume the role of the trusted entity. This certainly removes the single point of failure and improves resiliency. However, achieving consensus and maintaining consistency across hundreds of PGWs is not easy and could be very complex to manage.

The tiered architecture is a middle ground that provides a manageable alternative while eliminating the single point of failure. In the tiered architecture the Hubs are responsible for providing security services for the Hub network. That is, Hubs enforce admission control on Hub networks, and provide key and trust management services. Trust and key establishment between PGWs under two different Hubs is enabled with the cooperation of both the Hubs. Trust and key establishment between Hubs could be mediated by a central trusted entity like NERC or could be established in a peer-to-peer manner. Since the group of Hubs would be small is size and relatively stable, *i.e.*, not many changes in membership, both options become more feasible.

VI. NETWORK MANAGEMENT

Network management is critical to the success of NASPInet and as such is an important point to properly address. The

International Organization for Standardization (ISO) defines the five major functional areas of the Network Management Model as performance management, accounting management, configuration management, fault management, and security management.

Performance management is generally defined as monitoring, assessing, and adjusting the available bandwidth and network resource usage in order make a network run efficiently. Accounting management monitors and assesses the use of those networking resources for the purpose of billing or accountability. Configuration management is used to track and identify hardware and software configuration that affects the operation of the network. Fault management is the process by which a system detects, logs, and alerts management to problems that are affecting the network. Lastly, security management primarily deals with addressing network access control, intrusion prevention and detection, recovery from and response to intrusions, and policy definitions.

It is typically assumed that a single operator controlled network is centrally managed and should provide a homogeneous monitoring and management environment. The reality is that even with a single provider, when the network is continental scale, there are often multiple systems operating in a heterogeneous environment. Generally this is a result of multiple acquisitions, legacy systems, or even regional preferences throughout an organization. Thus our proposed tiered architecture with a core network and separate Hub networks does not add any more complexity than a single managed network. To the contrary, the proposed architecture pays specific attention to the geographic diversity and attempts to leverage the inherent tiered hierarchy of the grid to simplify network management. Communications and configuration constraints are localized to the Hub network. This will allow the operators of the Hub network to operate with a more homogeneous configuration and monitoring solution, while sharing network monitoring data with the rest of the network in an interoperable manner for increased wide area situational awareness.

This tiered hierarchy also structures the escalation path and points of contact when there is a network outage, loss of visibility, or other network management issue. The hub management entity provides this point of contact and then can coordinate as necessary with other entities to resolve any issues that may arise, or resolve them locally without complicating the process to the end nodes.

The key components of network management include monitoring and appropriate management to obtain certain levels of performance, accounting, configuration management, fault management, and security management. With the proposed architecture, each of these areas benefits from the leveraged locality. Performance benefits due to the geographic proximity of the edge devices to the hub. Configuration management and accounting are simplified due to the homogeneous network environment and single regionalized provider managing and providing resolution to issues. With the distributed nature of the architecture, it allows for better fault tolerance inherently in design and the same or higher level of fault management available in a centralized approach with a reduced number of nodes to monitor. Security management also benefits from localized key management, trust relationships, potentially less complicated firewall rule-sets, and reduced scale of the regionalized approach versus the global connection paradigm by minimizing the need for multi-party (or potentially multi-vendor) interactions.

VII. Distributed Resources

A. Network Technologies

As shown in Figure-2, the proposed tiered architecture consists of several interconnected Hubs, with multiple PGWs connected to each Hub. Any implementation of this architecture will most likely place Hubs in or near data centers, due to modern computing and networking realities. This is particularly true in light of the data storage arrays required to archive all of the PMU signals. Such a system begins to represent other large-scale systems in computing, notably those of large Internet companies, such as Google or Facebook. These systems are known as web-scale architectures.

Web-scale is a term which describes systems that have very large data sets, large numbers of client requests, and/or very large distributed networks. Such web scale systems are often characterized by "horizontal" solutions rather than "vertical" ones. They often take advantage of fast LAN technologies and clusters of commodity hardware to provide near linear scaling. They have fostered research and architecture redesign to address scaling problems and to create a more flexible system. Several distributed solutions have evolved or been adapted to provide services on these systems. Further, distributed computing models such as MPI and Map-Reduce are often leveraged to allow extensive computation as well as particular optimizations of the applications running on this flexible architecture.

Below we describe these distributed service and computing notions, and how they can benefit NASPInet.

B. Distributed Services

Many mechanisms exist by which services can be distributed [14] including election mechanisms, distributed hash table based protocols, and agent based services. Each of these techniques have benefits and drawbacks, therefore the mechanisms used for a service may vary depending on the specific requirements of that service.

A distribution of NASPI services, with each Hub hosting a node in the service network, provides many advantages important to a critical network. The advantages provided by distributing these services come in three primary areas: administration, reliability and availability. Administratively, the advantages are primarily in leveraging existing regulatory and reporting relationships and entities. Each Hub entity would maintain services for utilities already within its domain. The Hub could maintain a node for each service within its Hub network or aggregate those services where appropriate, and the nodes would communicate with each other as peers through the Hub. This allows each Hub entity to maintain autonomy, while still cooperating and maintaining an interoperable network.

The latter two advantages, reliability and availability, are similar, but differ in subtle and important ways. Reliability describes the resiliency of a service, such as during node failure. If a single node goes down, the impact is non-existent or minimal due to the core design of distributed services. With an event such as a server failure or software bug, there would be little or no impact on the overall operation, at most affecting the service consumers local to the failure. This is in contrast to a centralized service where a failure would impact all service consumers.

Availability, is the ability of a service to be consumed in cases of a low level failure, such as network link failure which partitions the network. In these cases, a consumer may be able to reach only a subset of the nodes providing a service. Because the service is distributed amongst the nodes, the consumers maintain the ability to function (within the limits of remaining network connectivity). Even more notable in such a split, both parts of the split operate with a given service letting them continue to function. Compare this to a network split or failure with a centralized service, where the service would be lost for any consumers not on the same side of the split as the centralized service.

Further benefits come in utilizing the natural aggregation and distribution points presented by the Hubs. Using the Hubs to consolidate signal subscriptions and eliminating duplication in broadcasts, would significantly reduce the amount of traffic between any two Hubs. It will further reduce signal transmission by any given PGW, as each signal needs only to be transmitted to the Hub once and then sent to each subscriber in an efficient manner. More reductions could be achieved by bundling disparate signals into one packet on the sending Hub, and un-bundling them on the receiving Hub.

Similar aggregation models can be utilized in a distributed storage mechanism. We propose that each Hub maintain a storage array, being the primary archival point for locally attached PGWs. This archived data would then be available for historical requests. Such requests could then be cached by the receiving Hub, reducing redundancy in transfer. Further, the data could be opportunistically replicated to other nodes, creating a network of backups that is geographically diverse. Finally large arrays of disks require processing power for data access and retrieval, this processing power could be used to host the other services, and run other distributed computing jobs that we will discuss next. Interestingly, the community is seriously considering systems, such as OpenPDC (http://openpdc.codeplex.com/), that can provide distributed storage. These solutions show great promise for future development and can be leveraged and adapted to our proposed architecture.

C. Distributed Computing

Distributed computing refers to computing that occurs on multiple autonomous systems. These systems communicate via a network, and divide the computational load amongst themselves. Such systems can divide the work by task, by data segments, or by some combination of the two. The systems can be spread over a large geographic area, or can be located in the same data center rack, however there is a tendency to refer to the latter cases by more specific terms, such as grid, cluster, or cloud, and using the term distributed to imply a geographic distribution. With a distributed system approach, the system should support multiple different paradigms of computation if possible to provide the most flexibility for future applications. A very strong approach may be to use a domain specific query language to address this with a common syntax for all applications.

There are many reasons to use distributed computing in a NASPInet context. The first benefit is maximizing the use of computational resources. The various distributed services will already require each Hub to house computers to act as local servers for the Hub network. Further, many large-scale storage schemes are built around arrays of smaller file servers. Due to the nature of storage arrays these servers are frequently CPU idle. This is the type of situation faced by large companies like Google and lead them to creating Map-Reduce, Google's internal distributed computing solution. Using a distributed computing paradigm allows these otherwise idle processors to be put to use.

The second benefit is minimizing network transit via the use of in-network data processing. Transmitting data across the WAN is expensive, the links are fairly slow, and Ambrust *et. al.* [15] show bandwidth per dollar to grow slower than other computing resources. Since NASPInet is expected to handle hundreds of terabytes a year in data to store, bandwidth costs to transport such a load are a serious consideration. A further consideration is that querying large data sets for requested signals already requires processing power at the point of data storage, to select it and prepare for WAN transit. Such limits suggest instead of transmitting raw data, supporting in-network data processing and then transmission of full or partially processed query results would be significantly more efficient.

In our scheme we put pools of compute nodes close to the edge of the network. These pools will maintain data storage, distributed services, and general purpose distributed computing on the NASPInet transited data. The obvious place to maintain a pool of compute nodes will be the Hub, taking advantage of the physical infrastructure and networking resources already present. Forward looking, it is noteworthy that the edge devices could one day be the actual power devices themselves. Imagine a PMU communicating up the hierarchy to get information about any calculations the system would like it to perform and then performing those calculations (perhaps one stage of a map-reduce equation) and sending those results for further use. When there are multiple data points needed from varied devices, the hierarchy would naturally provide the structure to operate on those data sets as well.

This core computation pool would be available to NASPInet members for the purpose of running computations against the collected signal data. It can be made available via a common API or via a domain specific query language. The computations need not be complete, they may just aggregate or filter signals, but even a partial reduction would significantly alter bandwidth requirements. Of course a computation that could complete within the computation pool would significantly reduce the bandwidth requirements by sending a single value in place of a large number of signals.

The type of computations, and their usefulness to the Power Grid are varied, and will certainly change in nature as the applications of PMU data mature. This is another reason for providing a general purpose computing node, instead of specialized services. One example is distributed state estimation [16]. Distributed state estimation would allow for more data points to be considered in the state-estimation, providing a more complete picture or a wider area picture, and do so in the same or less time than the centralized version. Further, due to the distribution of the Hubs, a two level [17] state estimation approach could be used. The first level would compute the state for the Hub or (even PGW-local) region. Results would then be transmitted to other Hubs for the overall state computation. The result of this computation could then be accessed by the appropriate entities. Such a wide area distributed state estimator would provide several efficiencies. First it would reduce the bandwidth used by raw signal transmission. Second it would take full advantage of a common resource, allowing for a backup state estimation approach, without the expense of a redundant system at each utility. Finally, it would allow for multiple models to be run simultaneously, taking advantage of the large computational resources, providing better situational awareness.

Of course the creation and maintenance of such a computing pool brings with it a set of challenges, both technical and social. Problems such as resource provisioning and accounting must be solved. Fortunately there are many answers in the field of Cloud Computing. This is a fairly new field, and the definition of cloud computing is still a bit fluid [18] however the basic premise is that cloud computing transforms large scale computing into a service. This service dynamically provisions and deprovisions resources for users, based on their current needs. The overall result of such a system is an efficient and well accounted use of resources [15].

VIII. CONCLUSION

In this work, we explored a tiered design approach to the NASPInet architecture, which reflects the hierarchy among power grid operators, monitors and regulators in order to provide a realistic deployment path for NASPInet. We discussed how this architecture favorably impacts the implementation of Quality of Service (QoS), cyber security and network management services. We also showed that this architecture provides distributed storage and computing opportunities that have the potential to reduce the bandwidth and latency requirements on the communication infrastructure.

ACKNOWLEDGEMENTS

REFERENCES

[1] Federal Energy Regulatory Commission. (2009, July) Smart Grid Policy. [Online]. Available: http://www.ferc.gov/whats-new/comm-meet/2009/071609/E-3.pdf

[2] National Institute for Standards and Technology. (2009, September) Framework and Roadmap for Smart Grid Interoperability Standards Release 1.0 (Draft). [Online]. Available: http://www.nist.gov/public_affairs/releases/smartgrid_interoperability.pdf

[3] "www.naspi.org."

[4] North American SynchroPhasor Initiative. (2009, May) Data Bus Technical Specifications for North American Synchro-Phasor Initiative Network. [Online]. Available: http://www.naspi.org/resources/dnmtt/naspinet/naspinet_databus_final_spec_20090529.pdf

[5] North American SynchroPhasor Initiative . (2009, May) Phasor Gateway Technical Specifications for North American Synchro-Phasor Initiative Network. [Online]. Available: http://www.naspi.org/resources/dnmtt/naspinet/naspinet_phasor_gateway__final_spec_20090529.pdf

[6] "U.S. Departmet of Energy Selections for Smart Grid Investment Grant Awards," October 2009. [Online]. Available: http://www.energy.gov/recovery/smartgrid_maps/SGIGSelections_Category.pdf

[7] S. Chakrabarti, E. Kyriakides, T. Bi, D. Cai, and V. Terzija, "Measurements get together," *Power and Energy Magazine, IEEE*, vol. 7, no. 1, pp. 41–49, January-February 2009.

[8] K. Mandl, W. Simons, W. Crawford, and J. Abbett, "Indivo: a personally controlled health record for health information exchange and communication," *BMC Medical Informatics and Decision Making*, vol. 7, no. 1, p. 25, 2007.

[9] U.S.-Canada Power System Outage Task Force, "Final Report on the August 14, 2003 Blackout in the United States and Canada: Causes and Recommendations," April 2004.

[10] J. Dagle, "Postmortem analysis of power grid blackouts - The role of measurement systems," *Power and Energy Magazine, IEEE*, vol. 4, no. 5, pp. 30–35, Sept.-Oct. 2006.

[11] R. E. Schantz, J. P. Loyall, C. Rodrigues, D. C. Schmidt, Y. Krishnamurthy, and I. Pyarali, "Flexible and adaptive qos control for distributed real-time and embedded middleware," in *Middleware '03: Proceedings of the ACM/IFIP/USENIX 2003 International Conference on Middleware*. New York, NY, USA: Springer-Verlag New York, Inc., 2003, pp. 374–393.

[12] David E. Bakken and Carl H. Hauser and Harald Gjermundrd and Anjan Bose, "Towards More Flexible and Robust Data Delivery for Monitoring and Control of the Electric Power Grid," Technical Report EECS-GS-009, School of Electrical Engineering and Computer Science, Washington State University, May 2007.

[13] D. E. Bakken, R. E. Schantz, and R. D. Tucker, "Smart Grid Communications: QoS Stovepipes or QoS Interoperability?" Technical Report TR-GS-013, School of Electrical Engineering and Computer Science, Washington State University, 2009.

[14] A. Tanenbaum and M. Van Steen, *Distributed systems*. Citeseer, 2002.

[15] M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica *et al.*, "Above the clouds: A berkeley view of cloud computing," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28*, 2009.

[16] R. Ebrahimian and R. Baldick, "State estimation distributed processing," *IEEE Trans. Power Syst*, vol. 15, no. 4, pp. 1240–1246, 2000.

[17] T. Yang, H. Sun, and A. Bose, "Two-level pmu-based linear state estimator," in *Proceedings of the IEEE PES Power Systems Conference and Exposition*, Seatle, WA, 2009, pp. 1–6.

[18] L. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, "A break in the clouds: towards a cloud definition," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 50–55, 2008.