Wireless WAN for the Smart Grid

WHITE PAPER

Trilliant helps leading utilities and energy retailers achieve their smart grid visions through the Trilliant Communications Platform, the only communications platform purpose-built for the energy industry that integrates disparate systems of systems into a unified whole. The Trilliant Platform is deployed with more than 200 utilities worldwide to enhance energy efficiency, improve grid reliability, lower operating costs, integrate renewable energy resources and electric vehicles, and empower consumers to better manage their energy consumption.

1100 Island Drive
Redwood City, CA 94065
t 650-204-5050
f 650-508-8096
www.trilliantinc.com

Introduction

A Smart Grid communications infrastructure allows utilities to communicate with one another in regional grids, as well as with customers and distributed power generation and storage facilities. To achieve the full vision of the Smart Grid, individual utilities will need to support multiple networks: the Home Area Networks (HANs) for consumer energy efficiency; the Neighborhood Area Network (NAN) for advanced metering applications; and the Wide Area Network (WAN) for distribution automation and the backbone of the Smart Grid. The exclusive focus of this paper is the WAN.

The material, intended for a business decision-maker audience, is organized into three sections followed by a brief conclusion. The first section describes the role of the WAN in a utility's own Smart Grid. The second section compares the public and private alternatives available today for implementing a WAN, concluding that a private wireless WAN is the most prudent choice. The third section outlines the critical requirements for a private wireless WAN that are of particular importance to a utility. Although some of the topics are technical by their very nature, every attempt is made to cover each in a non-technical way.

The WAN's Role in a Utility's Smart Grid

The US Department of Energy's (DOE) National Energy Technology Laboratory (NETL) prepared a report titled "A Systems View of the Modern Grid" for the Office of Electricity Delivery and Energy Reliability. This comprehensive, systems-level view of the Smart Grid identifies the following five key technology areas:

- Integrated Communications
- Sensing and Measuring
- Advanced Components
- Advanced Control Methods
- Improved Interfaces and Decision Support

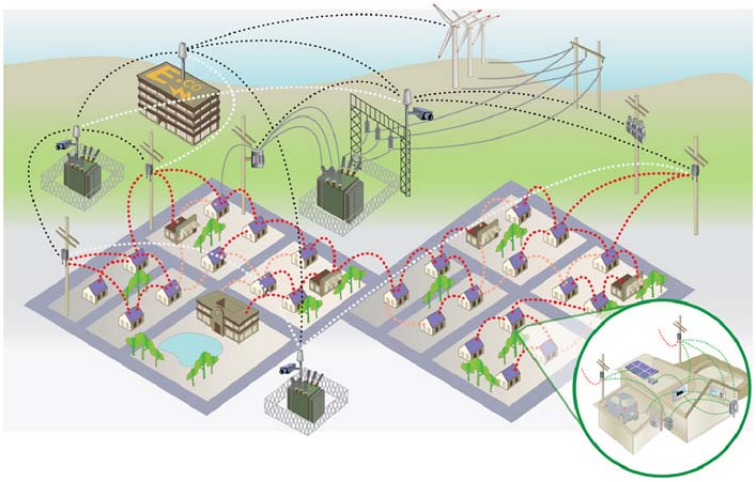

A "Typical" Private Utility Network

Although all utilities are different, it is possible to imagine what a "typical" utility's private network might look like. The backbone of the network exists along the transmission and distribution lines, making it capable of reaching the utility's entire infrastructure. This backbone would normally consist of wireless point-to-point and/or fiber optics links over any long-distance transmission lines, and point-to-multipoint wireless nodes spaced as required along shorter transmission lines and distribution lines. The substation, as a point of demarcation between transmission and distribution lines, is an ideal location to integrate wireless network segments with any existing fiber optic or other wired links. Each substation would also have either a wired and/or wireless LAN, the latter being in the form of a Wi-Fi network capable of covering both indoor and outdoor equipment. A utility's own Distributed Generation facilities would normally be integrated as edge nodes on the wireless portion of the backbone.

The Trilliant Multi-Tier Smart Grid Architecture

In the appendix on Integrated Communications, NETL makes the following important point: "Of these five key technology areas, the implementation of integrated communications is a foundational need, required by the other key technologies and essential to the modern power grid. Due to its dependency on data acquisition, protection, and control, the modern grid cannot exist without an effective integrated communications infrastructure. Establishing these communications must be of highest priority since it is the first step in building the modern grid."

A multi-tier network integrates communications throughout the distribution grid and uses an infrastructure-wide network or wide area network (WAN). To be fully effective, the utility's WAN will need to span its entire distribution footprint, including all substations, and interface with both distributed power generation and storage facilities as well as with other distribution assets such as capacitor banks, transformers, and reclosers. The utility's WAN will also provide the two-way network needed for substation communication, distribution automation (DA), and power quality monitoring while also supporting aggregation and backhaul for the advanced metering infrastructure (AMI) and any demand response and demand-side management applications. And many utilities will want to take full advantage of the investment in this WAN infrastructure to run other enterprise networking applications, including wireless communications for work crews in the field, site security with video surveillance, Voice over IP (VoIP), asset management, and more.

Each application running on the utility's WAN has its own set of requirements. Some applications like Supervisory Control And Data Acquisition (SCADA), automatic restoration and protection, and VoIP will require prioritization for real-time or near-real-time response and satisfactory Quality of Service (QoS). Some applications like AMI backhaul and video surveillance will consume considerable bandwidth, requiring broadband data rates end-to-end. And others like substation load management and crew communications will require both high bandwidth and fast response times.

In the aggregate, the many applications are expected to demand much from the WAN. But as NETL points out, the benefits will be significant: "Integrated communications will enable the grid to become a dynamic, interactive medium for real-time information and power exchange. When integrated communications are fully deployed, they will optimize system reliability and asset utilization, enable energy markets, increase the resistance of the grid to attack, and generally improve the value proposition for electricity."

Smart Grid WAN Alternatives

Utilities face a daunting array of options today when attempting to choose the optimal WAN solution. Network vendors and service providers alike are vying for a share of the Smart Grid market, particularly with the availability of billions of dollars in stimulus funding. And the regulatory agencies and standards bodies are struggling to get ahead of the situation to provide better guidance to utilities.

A contextual comparison of the various alternatives, however, reveals that some networking solutions are simply not viable as a Smart Grid WAN for one or more reasons. Of those that are viable, the choices quickly distill to a select few when related factors, such as investment protection and risk mitigation, are considered. This section first divides the choice facing utilities into two fundamental options—public vs. private networks—then examines the viable choices for private WANs.

Public vs. Private

This fundamental choice comes down to a classic "build vs. buy" decision. The "build" option in this case involves deploying a private WAN (covered next). The primary "buy" option considered here is the extensive public cellular data/phone network now being used in typical AMI projects. The same considerations also apply, however, to the other public network services available, including dial-up modems, leased lines, digital subscriber lines (DSL) and even cable.

The National Institute of Standards and Technology (NIST) is tasked with establishing the full set of standards that will be needed to make the Smart Grid fully interoperable. In its Framework and Roadmap for Smart Grid Interoperability Standards (Release 1.0), NIST cites an important role for public networks: "Examples of where [grid] communications may go through the public networks include: customers to third-

party providers, bulk generators to grid operators, markets to grid operators, [and] third-party providers to utilities."

Each of these examples involves communications between different parties, making public networks a natural choice. Indeed, in some situations, it may be necessary to utilize a public network, including the Internet, for such multi-party communications. The real question confronting the utility is whether or not public networks are sufficient for its own WAN infrastructure?

For a growing number of utilities, the answer is "No." The most significant limitation is the utility's inability to control the network infrastructure. This fact alone has convinced most utilities that they simply cannot risk making the Smart Grid's many operational applications—some of which are mission-critical—dependent on one or more service providers. Consider what would happen in a natural disaster when power lines are down and the cellular network is flooded with emergency calls and families trying desperately to reach one another. Even under normal circumstances periodic traffic congestion, occasional tower outages and other problems inevitably occur in such large-scale and often overloaded networks.

There are number of other limitations inherent in public cellular networks that make them unsuitable as a utility-wide WAN for Smart Grid applications:

- Lack of security is a major concern being raised not only by the Department of Energy and NIST but by many other organizations, as well. Consider how vulnerable the public Internet continues to remain despite the unprecedented effort to secure it over nearly two decades.
- High latency can be an insurmountable problem for some command and control grid applications that require real-time or near-real-time response.
- Insufficient bandwidth is an issue for some other individual applications, such as substation video surveillance, and for all applications in the aggregate.
- Excessive cost of operation is a problem that does not occur in a pilot project, but the relatively high ongoing monthly fees quickly exceed both the capital and operational costs of deploying a private network.
- Stranded assets can cause costs to increase even more as changes in cellular technology force upgrades in the utility's infrastructure. Consider the migration to date in the public cellular network, which began using analog technology, and now uses Second and Third Generation (2G and 3G) digital communications with plans to migrate to 4G and beyond.
- Incomplete coverage (including poor reception) remains an issue in some areas, especially in rural environments and the remote settings that are ideal for substations. Where coverage is unavailable or inadequate, the utility is forced to implement some form of private network.

Private WAN Alternatives

Although the public cellular network has been used to provide two-way communications in AMI pilot projects, the limitations outlined above make it an unacceptable choice for the full suite of applications the Smart Grid must support. Which is why a growing number of utilities are now implementing private networks.

A private network consists of an infrastructure owned and operated—in its entirety—by an organization for its own, exclusive use. For this reason private networks have some major advantages but few enterprise organizations are large enough to have a real need for one. Instead, most organizations achieve some of these advantages at a lower cost with Virtual Private Networks. A VPN is a hybrid of a public and private network that employs the public network infrastructure, but with special traffic segmentation and other security provisions that make it operate in ways similar to a private network.

Advances in technology, especially in the wireless arena, now make private networks quite affordable, and in many situations, even more affordable than using public network services, including as part of a VPN. Utilities also have another advantage favoring private networks: rights-of-way along transmission and distribution lines. Of course, with wireless technologies, there is no need for a physical right-of-way. But a utility's existing rights-of-way affords the option of laying ultra-high bandwidth fiber optic cabling—something very few other organizations would even consider.

The combination of demanding applications over a large footprint with rights-of-way all along the way makes the private network an obvious choice of a utility's Smart Grid WAN. Very few other organizations (with the possible exception of railroads and municipalities) are in such an enviable position to fully control their networking destiny.

While the significance of full control should not be understated, there are a number of other advantages inherent in a properly-designed private network, including:

- Inherent security that derives from exclusive use of the infrastructure—end-to-end
- Genuine broadband capacity to support all Smart Grid and other applications
- Ubiquitous coverage for all facilities with no voids
- Guaranteed performance with the ability to set priorities without any dependence on any other entity
- Lowest possible total cost of ownership based on today's much lower capital and operational expenditures (CapEx and OpEx), and elimination of ongoing monthly fees for public network services
- The private wireless or hybrid fiber/wireless network is far more affordable and versatile than any other network solution available now or in the foreseeable future. Indeed, the ability to fully control one's networking destiny virtually eliminates both short- and long-term risk by ensuring that the infrastructure will always be able to handle the demands being placed on it. Applications and networking technologies will change over time, of course, but with a properly-designed private network, the utility will be able to accommodate these changes cost-effectively and without disruption.

Rural Broadband Access

Utilities in rural areas often find that their investment in a private wireless WAN is capable of generating incremental income from a completely unrelated application: rural broadband access. Rural areas that are underserved or even unserved with broadband access from the local telephone company or other Internet service provider (ISP) are a lucrative market for some utilities and cooperatives. The most common business model is to offer the utility's wireless WAN and 'last mile' access infrastructure wholesale to one or more independent ISPs. Under this model, the ISP(s) handle(s) all of the arrangements with the customer, including support, enabling the utility is to receive an ongoing revenue stream without actually getting into the Internet access business. The utility utilizes a portion of the WAN for broadband access and can securely partition the network to virtualize access without effecting reliability.

Requirements for a Private Wireless WAN

A solid design for a private wireless WAN requires a full understanding of the fundamental requirements. Knowing both the current and likely future requirements is essential when designing a network that will be able to function properly initially, and then migrate gracefully as needed to accommodate new locations and/or applications. This section outlines a set of five fundamental requirements that utilities can (and should) use when evaluating the myriad wireless solutions available today.

NIST recognized the importance of thoughtful evaluation in its Framework and Roadmap for Smart Grid Interoperability Standards (Release 1.0), where the organization warns that great care must be taken when choosing a wireless solution: "Wireless technologies are candidate media for meeting Smart Grid requirements, especially those for which alternative media are too costly or not workable. However, different types of wireless technologies also have different availability, time sensitivity, and security characteristics that may limit their suitability for certain applications. Therefore, the capabilities and weaknesses of specific wireless technologies must be assessed in all possible conditions of Smart Grid operations."

Another important consideration for understanding the requirements of a private wireless WAN is the strategic significance of the substation. Substations are a logical and even inevitable location for nodes in the utility's WAN. Because the utility already owns these facilities, which are conveniently spread throughout the service area, there is little or no need to enter into leasing agreements for separate towers on hilltops or high-rise buildings. Having each substation as a node on the private wireless WAN (or hybrid fiber/wireless network) also ensures support for all distribution automation, AMI, power quality monitoring and control, and demand response and demand-side management applications.

Coverage/Capacity/Cost

These three requirements are treated here as one because most solutions require that tradeoffs be made among the three. Nearly all wireless solutions can satisfy two of these requirements well—but at the expense of the third. For utilities that require both coverage (large footprint) and capacity (with broadband and real-time applications), cost becomes the tradeoff; that is, cost can rise dramatically to satisfy the coverage and capacity required.

The discussion below of some of the other requirements will help explain this relationship among coverage, cost and capacity. The important point to remember is that most wireless solutions force this compromise, and the need to tradeoff one of these critical factors early on can mean going to great lengths to overcome it later.

The least painful compromise in a network's initially deployment is usually capacity. In the early stages, when only a few applications are up and running, and only pilots or trials may be involved, it is tempting to deploy a network that is good enough for now. It covers the geography at a reasonable cost, and delivers adequate capacity—all seemingly without compromise. But as the network grows, it becomes obvious that it fails to scale. In a worst case scenario, the entire infrastructure may need to be replaced.

This problem can be avoided by having a long-term plan for the network that considers the evolution of each tier (WAN, NAN, and HAN). Look out five years, at a minimum, and consider all of the possible applications that the network might need to support. And be conservative by estimating on the high side what each application might demand from the network's resources, especially its bandwidth. As the aggregate demand for a finite amount of bandwidth grows, performance can suffer in networks that fail to scale. And when performance degrades, some applications may fail to run well—or at all.

Range

The transmission range of a wireless node is an important factor in the tradeoff between coverage and cost, because shorter range requires more nodes. But it is also worthy of separate consideration for other reasons.

The primary factor that determines the range of a wireless node is the regulatory restriction on transmit power levels, and all regulatory agencies (such as the Federal Communications Commission, or FCC, in the U.S.) utilize the potential for radio frequency (RF) interference in making this determination. Because directional antennas transmit only along a narrow path, thereby minimizing the potential for RF interference, they are permitted to transmit at higher power levels for longer range.

The major disadvantage of using fixed, directional antennas is the need to manually and carefully aim each antenna during its installation. This problem can be mitigated or even eliminated, however, with the use of beam-forming or switched directional antennas that are aligned electronically, rather than physically. These antennas transmit along the same narrow, directional path as fixed antennas do, but can do so in different directions automatically as needed. For this reason, switched or beam-formed antennas afford another advantage: the ability to support the long-distance point-to-point communications needed to create the alternate, redundant paths employed in a resilient mesh network topology.

Real-Time, Two-way Communications

Many wired and wireless networks share a dirty little secret: frequent "collisions" of two-way traffic. As with any shared medium protocol, including Ethernet and Wi-Fi, only one node can transmit at a time. If any two nodes attempt to communicate simultaneously a "collision" occurs halting all transmissions until the affected nodes have time to recover. Special protocols are used to avoid, detect and recover from these collisions, but they are terribly inefficient. In shared Ethernet networks, the problem can be mitigated with brute force bandwidth. But in a wireless environment, where bandwidth is not as copious, these inevitable and often frequent collisions can result in crippling limitations.

Overcoming a wireless collision-based system requires synchronization and coordination between devices. With scheduled transmissions, communications can be assured between devices.

The utility's need to support two-way traffic in real-time is perhaps even more critical than it is in the public cellular network, where users are accustomed to the occasional (and sometimes frequent) service disruption. Consider, for example, a grid reliability application that utilizes Phasor Measurement Units (PMUs) or synchrophasors, potentially in a full Wide Area Measurement System (WAMS). Each PMU needs to transmit up to 30 synchronous measurements every second to a centralized Phasor Data Concentrator (PDC). Control traffic, normally employing SCADA, must be able to travel simultaneously and in real-time in the opposite direction, outward from the central facility. Such mission-critical applications simply cannot tolerate poor performance in the wireless WAN. The same traffic management provisions that enable real-time, two-way communications also allow the network to support other latency-sensitive and/or bandwidth-intensive applications, such as VoIP and video surveillance.

Security

A private wireless WAN is inherently secure, especially for those that employ inter-node link authentication and strong encryption of backbone traffic. Special security provisions remain a requirement, however, for two reasons. One is that the RF signals from a wireless WAN can "leak" beyond the utility's facilities, making it accessible to anyone nearby without proper authentication and access control provisions. Fortunately, wireless access can be protected in several ways, including industrial-strength authentication and encryption techniques. The better solutions support all of these techniques to afford the best protection.

The second reason is the need to securely segment traffic when connecting with external parties, such as may be required for a utility's normal business operations or Smart Grid applications. The two most common and proven ways to segment this virtual traffic is with VPNs and Virtual LANs or VLANs. Most solutions support IP Security (IPsec) VPNs as an integral part of the Internet Protocol; the better solutions also offer robust support for VLANs to provide additional flexibility and security.

Reliability

Last but certainly not least is the reliability of the private wireless WAN. Most modern electronic equipment is quite reliable with high mean-time between failures (MTBFs). And most equipment today is also designed with a sophisticated diagnostic capability and a low mean-time to repair (MTTR).

High reliability in a wireless network, however, requires more than a low failure rate in the equipment used. Any outdoor wireless network, regardless of the technology employed, will experience period disruptions caused by severe weather conditions (that attenuate or weaken RF signals) and external RF interference in the unlicensed bands.

A well-designed mesh network topology can mitigate against and even overcome these and other periodic problems that adversely affect network uptime. Most wireless mesh network solutions are self-forming and self-healing, the latter requiring the presence of alternate paths to reroute around a problematic primary path. Some solutions do a better job, however, at solving another problem that affects reliability: periodic traffic congestion that can prevent mission-critical applications from running well or at all. To avoid problems caused by congestion, the better solutions offer advanced traffic management features like traffic filtering and shaping, QoS-based traffic prioritization, mesh route optimization and dynamic load-balancing.

Conclusion

Establishing integrated communications is of the highest priority because, as DoE's National Energy Technology Laboratory points out, "it is the first step in building the modern grid." Whether driven by mandate or the market—or both—a growing number of utilities are choosing to implement a private wireless WAN as the backbone for their own Smart Grid. They are finding that with a prudent choice in wireless WAN technology, they can not only satisfy today's demanding requirements, they also gain the versatility needed to deal with the uncertainty surrounding the future of the Smart Grid.

To learn more about Trilliant's multi-tier Smart Grid architecture, visit our online library at http://info.trilliantinc.com/library.