WHITE
PAPER

Application Domain Partitioning for the Smart Grid

Deploying a Multi-Tier Network for Security and QoS

Executive Summary

Today's electric grid employs a variety of different monitoring and control applications, many of which use different forms of communications. Utilities also operate a number of enterprise applications on their corporate networks. The advent of the Smart Grid affords an opportunity to create an integrated communications infrastructure capable of supporting these and other applications in a secure, efficient, and cost-effective way.

The US Department of Energy asserts that just such an integrated communications infrastructure is absolutely essential to deploying the Smart Grid. Without one, it would be difficult or cost-prohibitive to support the many applications required to integrate sources of renewable energy or implement demand response programs to control loads during peak periods.

Achieving the maximum return on the investment in a fully integrated communications infrastructure requires that it support the full spectrum of an electric utility's applications. Each application will place different demands on the network, and in the aggregate, these demands will require a network that is secure, capable, and scalable enough to satisfy the past, present, and future needs of the grid.

What the future holds is always unknown, but one thing is certain: Every utility wants to remain in control of its destiny. This is why a growing number of utilities are choosing to deploy their own end-to-end network infrastructure. Having control over the entire network infrastructure is the only way to ensure that every application will perform well and be protected from attack or tampering.

Most utilities now realize that the best way to implement an integrated communications infrastructure end-to-end is with multiple tiers—beginning in the subscribers' homes and businesses, extending through the meters and distribution network, and ultimately encompassing the utility's entire service area. A multi-tier network provides the operator control over its devices and a secure domain environment for its applications.

While likely to be new to most utilities, the ability of modern data networks to support multiple applications well and securely is both standardized and proven. Supporting multiple applications with an integrated, multi-tier Smart Grid network has two fundamental requirements: virtualizing the physical network infrastructure to securely separate or segment each application from all others; and managing the aggregate traffic to satisfy each application's specific needs for performance and Quality of Service (QoS).

Virtual networking is a proven technique that is used in both the Public Switched Telephone Network and the Internet—the ultimate multi-application, multi-tier network. The two widely used standards for virtual networking are Virtual LANs or VLANs, and Virtual Private Networks or VPNs intended for use with the Internet Protocol (IP). Managing traffic to deliver satisfactory performance to each and every application also enjoys standardized and proven techniques.

This white paper, intended for business decision-makers, is organized into two remaining sections followed by a brief conclusion. The first section, *The Networking Challenge,* identifies many Smart Grid applications and assesses the individual and collective demands these applications place on a utility's communications infrastructure. The second section, *Application Domain Partitioning for the Smart Grid,* explains how utilities can employ various virtual networking and traffic management techniques to achieve a successful implementation.

The Networking Challenge

The Smart Grid promises to bring us a multitude of important benefits, including improved reliability, reduced costs, and greater security. The benefits are delivered by a number of key applications, both existing and envisioned, that are the foundation of the need for the Smart Grid network.

Applications specific to an electric utility's Smart Grid include:

- *Distribution/Feeder Automation* to improve voltage/VAR control and better manage fault and outage conditions
- *Distributed Generation and Distributed Storage* that are becoming increasingly necessary to meet Renewable Portfolio Standards for wind, solar and other sources of renewable energy
- *Power Quality Monitoring and Control* across the utility's entire grid, normally employing systems that use Supervisory Control and Data Acquisition (SCADA)
- *Advanced Metering Infrastructure (AMI)* that includes automatic meter reading, potentially with Time-of-Use pricing and remote (dis)connect
- *Demand Response and Demand-Side Management* that leverage the AMI Neighborhood Area Network (NAN) to shed load during peak periods in residences and businesses
- *Wide Area Measurement System (WAMS)* that requires taking precisely synchronized readings from Phasor Measurement Units (PMUs)

Additional enterprise applications employed by the utility include:

- *Client/Server and Host/Terminal* systems used by various departments, including billing, engineering, operations, finance, regulatory compliance, human resources, etc.
- *Voice over IP (VoIP)* to supplement or replace use of the Public Switched Telephone Network (PSTN)
- *Site Security with Video Surveillance*, particularly at substations
- *Wi-Fi Access in a Field Area Network (FAN)* around substations to provide work crews with access to enterprise client/server applications and, optionally, VoIP

Achieving the benefits delivered by the above applications requires a scalable, high-performance network that can accommodate the different needs of the many different applications. Two key requirements that characterize the needs of an individual application are its demand for bandwidth and its tolerance of latency. Some applications do not require very much bandwidth, but what little bandwidth they do consume must be delivered consistently in real-time. Some will require a considerable amount of bandwidth or throughput, but may be quite tolerant of latency or delays in transmission. Table 1 provides a summary of these two requirements for the applications identified above.

Application	Bandwidth	Latency
Distribution/Feeder Automation	Low	Low
Distributed Generation and Distributed Storage	Low	Tolerant
Power Quality Monitoring & Control	Low	Low
Advanced Metering Infrastructure	Moderate	Tolerant
Demand Response and Demand-side Management	Low	Tolerant
Wide Area Measurement System	Low	Low
Water/Gas Infrastructure Monitoring & Control	Low	Low
Client/Server and Host/Terminal	Low-Moderate	Tolerant
Voice over IP	Moderate	Low
Site Security with Video Surveillance	High	Tolerant (Low Jitter)
Wi-Fi Access in a Field Area Network (without VoIP) - With Support for VoIP	Low Low-Moderate	Tolerant Low
Rural Broadband Access (without VoIP) - With Support for VoIP	Moderate-High High	Tolerant Low

Table 1 – Bandwidth and Latency Requirements of Smart Grid Applications

As shown in the table, the bandwidth requirement for most individual applications is relatively low. In their aggregate, however, the demand can be high. For this reason, the WAN portion of the integrated, multi-tier Smart Grid network should support genuine broadband data rates, especially in deployments that will need to accommodate video surveillance, rural broadband access and/or VoIP applications for mobile workforce management. Numerous applications do require a relatively low latency, however, particularly those that are used for real-time or near-real-time monitoring and control. Voice communications also demands low latency, as anyone who has attempted to carry on a conversation via satellite can attest.

In any network there is a natural contention between bandwidth and latency as the network approaches 100% capacity. Once the network becomes congested, at least some traffic will need to be stored temporarily, or buffered, until there is sufficient capacity to forward it. Normally these periods of congestion are intermittent and brief, but they can wreak havoc on those applications that are intolerant of latency. The challenge for the network, therefore, is to manage traffic in such a way that latency-sensitive applications are treated with a higher priority.

Another challenge when supporting so many different applications on a common network infrastructure is security. Two important components of security are privacy of communications and protection from problems caused by other applications. The latter component can actually present a greater challenge in some networks where a single application going awry could have an adverse impact on the entire, end-to-end infrastructure. The reason privacy is not normally a major issue is the wealth of proven security mechanisms (e.g. access control, authentication, encryption, tunneling, firewalls, etc.) that have long been incorporated into most network solutions.

The collective requirements for both copious bandwidth and intelligent traffic control (to manage latency and protect applications) make the public cellular network unsuitable as an integrated communications infrastructure for electric utilities. Because they are designed primarily for the voice communications and Internet access needs of subscribers, public wireless networks lack the ability to accommodate applications that require high, sustained bandwidth and/or consistently low latency. When these networks become congested, all users and applications are affected equally. Additional limitations of cellular networks include security exposures, incomplete coverage and high ongoing costs. Perhaps the most significant limitation, though, is the utility's inability to control the network infrastructure. While cellular service may be acceptable in trial or pilot projects, these limitations make it an unsuitable, long-term choice. Additional information on public vs. private WANs can be found in another Trilliant white paper titled *Wireless WAN for the Smart Grid*.

Since a public WAN does not allow for network partitioning, Quality of Service (QoS) for applications cannot be easily accomplished, and a critical application may not be provided width sufficient network bandwidth. Communications to one device can potentially affect all applications on the network. For example, a meter firmware download could disrupt grid control. In addition, the security is a major concern, as any device on the entire public WAN can be accessed. For example, hacking an in-home device could provide access to grid control.

Application Domain Partitioning for the Smart Grid

The creation of secure, isolated Application Domains on the Smart Grid provides an effective solution to the types of issues described above. Network partitioning on an integrated, multi-tier Smart Grid network has two fundamental objectives: separating or segmenting each application from all others to provide security; and managing all traffic to satisfy each application's specific needs for Quality of Service (QoS).

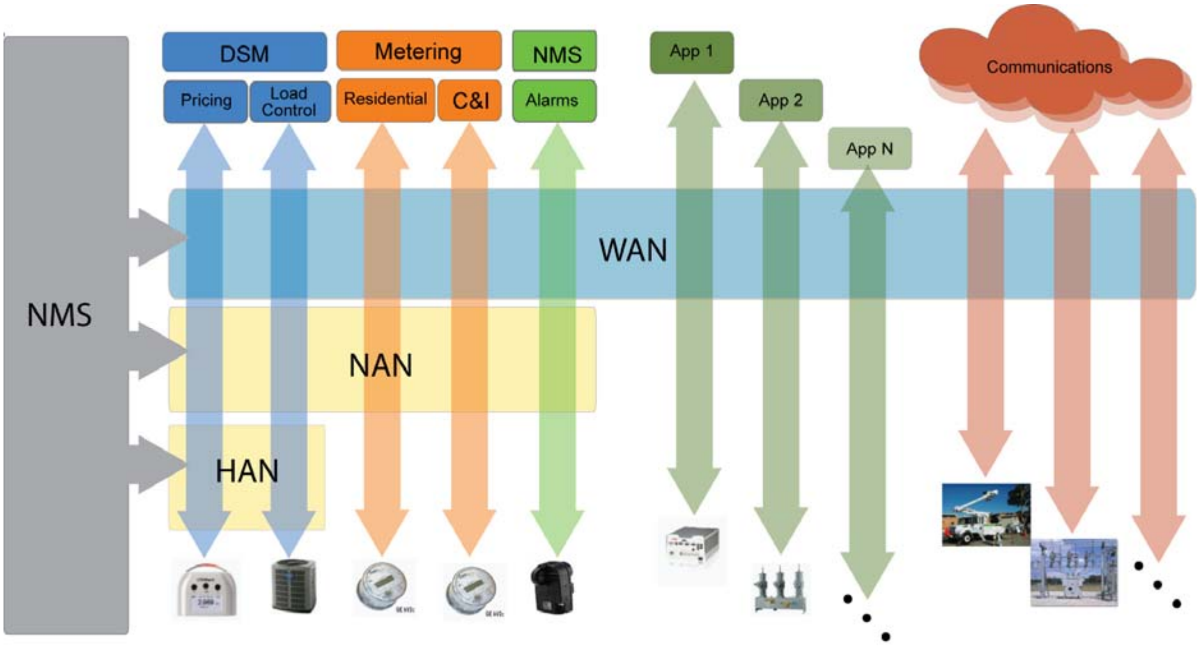

Figure 1 – multiple applications, each partitioned into their own virtual network domain

Application Domain Partitioning for the Smart Grid

Figure 2 illustrates secure application domains on a Smart Grid network. Network partitioning allows for the creation of secure isolated Application Domains. As a result, a strong level of security can be provided. Access to one application does not provide access to other applications or devices. Breach of one device does not compromise other devices. A compromise of a device used in one application does not compromise other applications.

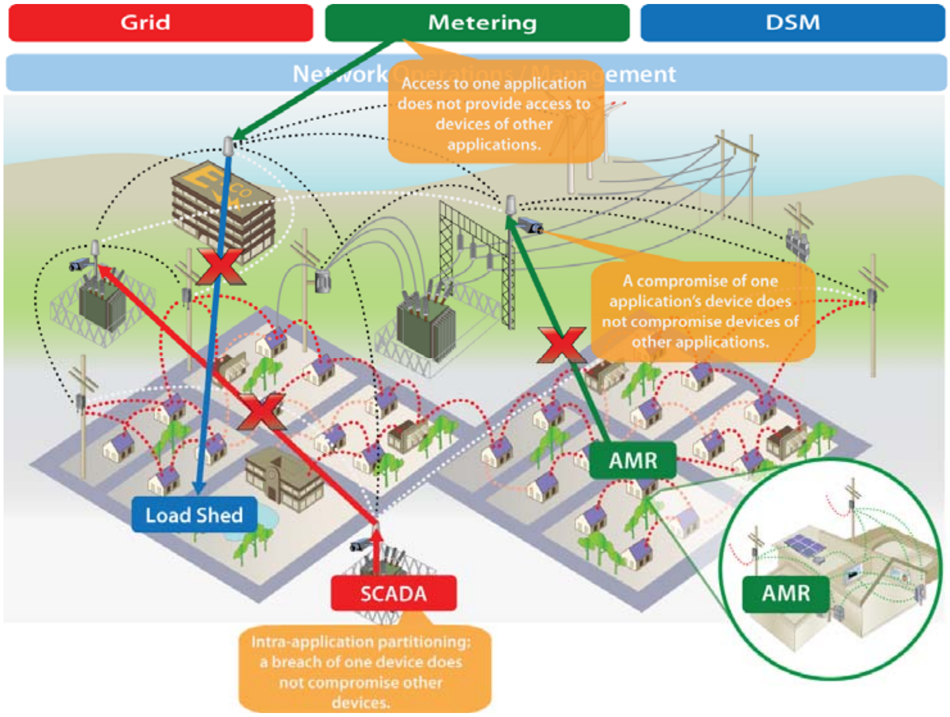
Grid
Metering
DSM
Network Operations/Management
Access to one application does not provide access to devices of other applications.
A compromise of one application's device does not compromise devices of other applications.
AMR
Load Shed
SCADA
AMR
Intra-application partitioning: a breach of one device does not compromise other devices.

Figure 2 – applications communications secured from breaches

In addition to addressing potential security challenges, Application Domains also enforce QoS for each application. *Rate Shaping* is used to limit the capacity each device is allowed to consume, thereby protecting the minimum bandwidth required for each application. In addition, whenever WAN capacity is reached, *Traffic Prioritization* minimizes latency and ensures reliability of higher-priority applications by restricting lower priority traffic. Figure 3 shows how *Fault Recovery* is achieved on the network. A fault condition may cause traffic to follow alternate paths through the mesh network. This increases the load in certain regions of the network and may cause network capacity to be exceeded. Application Domain architecture enables the network to recover from this type of fault. *Priority Queuing* ensures that higher priority traffic is forwarded first, selectively dropping low priority traffic if capacity is exceeded. Rate Shaping protects against possible traffic flooding that may be caused by network misconfiguration, Denial of Service (DoS) attack, or another fault condition. Traffic rates can be shaped differently for each application, achieving a high level of fault tolerance.

Application Domain Partitioning for the Smart Grid

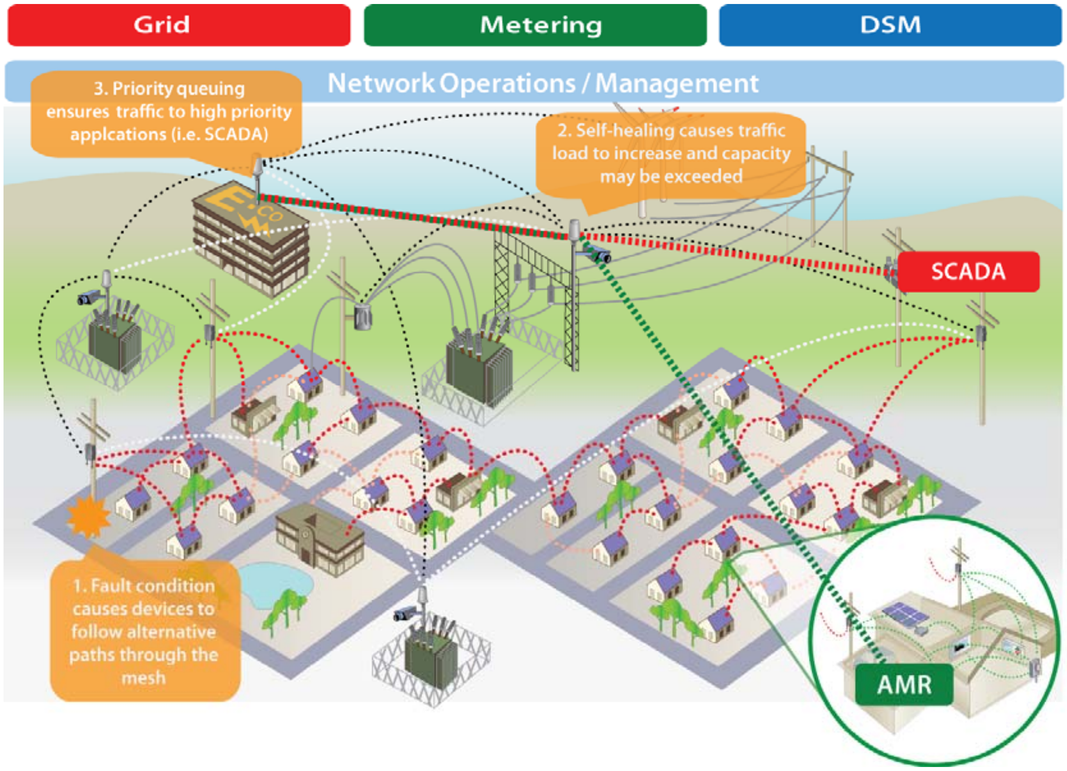
Grid
Metering
DSM
Network Operations / Management
3. Priority queuing ensures traffic to high priority applcations (i.e. SCADA)
2. Self-healing causes traffic load to increase and capacity may be exceeded
SCADA
1. Fault condition causes devices to follow alternative paths through the mesh
AMR

Figure 3 – Fault Recovery with QoS for each application

Segmentation through Virtualization

All networks are comprised of physical elements, including wires, fiber optic cabling and radio transmitters/receivers at the physical layer, and internetworking equipment at the higher layers to direct the flow of traffic. The flow of traffic can also be virtual; that is, the traffic can be segmented or partitioned in a variety of ways that make the physical network appear differently to different applications or users. Virtual networks have long been used in both the PSTN and the Internet to ensure security or QoS—or both. Virtual networking is ideally suited to isolating and securing the many applications that operate on a utility's physical network infrastructure. Two proven techniques provide both the necessary and sufficient virtual networking capabilities required by any utility: the Virtual Local Area Network or VLAN; and the Virtual Private Network (VPN) based on IP Security or IPsec. Both VLANs and VPNs are fully standardized with robust feature sets and solid security.

VLANs specify a means for partitioning and isolating traffic to control or deny communication between applications in the shared physical network infrastructure. Similarly, IPSec VPNs are broadly deployed on data networks to provide secure communications. Because the Internet Protocol (IP) was created originally without security, the Internet Engineering Task Force (IETF) formed the IP Security Working Group to create a general-purpose means of overlaying security to facilitate confidential communications via the Internet or any other public IP backbone. The result is IP Security, or simply IPsec, that consists of a fundamental architecture and a collection of IETF Request for Comment (RFC) standards. IPsec provides three separate forms of protection—authentication, integrity and confidentiality.

Virtual networking affords virtually unlimited possibilities for optimizing communications in an integrated, multi-tier Smart Grid network. Here are just a few examples:

- Substation SCADA using a VLAN to carry legacy Modbus data back to the head-end
- Substation Video Surveillance via a VPN with rate limiting and QoS
- Distribution Automation using a VLAN to carry Ethernet and/or serial DNP3 traffic from grid devices back to the head-end
- Smart Meter communications via ANSI C12.22 traffic on a dedicated VLAN
- Demand-Mide Management with consumers communicating through a VPN to a deregulated retailer

VLANs can also be combined with VPNs in a variety of ways. One example is the use of a single VLAN to segment all VPN traffic. This would allow any VPN to terminate anywhere within the utility's entire multi-tier Smart Grid network, thereby giving users at different locations access to the same VPN from their PCs. Embedding a VPN within a VLAN also enables the utility to apply Rate Shaping or other traffic management provisions more easily and consistently.

Virtual networking is indeed quite versatile. And that versatility ensures that a utility will get a high return on its investment in the physical integrated communications infrastructure.

QoS through Traffic Management

There are many techniques for managing traffic to assure adequate QoS. All traffic management techniques operate in essentially the same way by establishing a hierarchy of priorities. At the low end of the hierarchy is “Best Effort” treatment that essentially gives the packet no priority at all; it is forwarded only to the extent adequate capacity exists. Such a low priority is perfectly acceptable for many types of traffic, including file downloads. The high end is generally reserved for network management applications, which consume relatively little bandwidth but are mission-critical to the network's very operation. In between, network managers are free to assign whatever traffic priorities might be needed. Normally, applications like real-time monitoring and control such as SCADA are given higher priorities.

It is important to note that all traffic management provisions are usually implemented to take full advantage of any performance-enhancing capabilities in the underlying physical network infrastructure itself. For a private wireless mesh WAN, for example, these capabilities might include multi-path route optimization and load balancing with end-to-end synchronization of all traffic flows that together help ensure the network's ability to sustain peak levels of performance. Path vector-based routing is a particularly effective means for ensuring optimal performance and dependability throughout a wireless mesh WAN.

Some solutions offer another powerful traffic management capability: Rate Shaping. Rate Shaping provides a means for ensuring a degree of “fairness” during periods of congestion by placing a limit on the amount of bandwidth any particular application can consume. The algorithm employed also normally ensures that high priority traffic takes precedence, but that no single application's high priority traffic ever takes bandwidth allocated (proportionally) to another application. Rate Shaping even helps mitigate against certain forms of attack designed to flood networks with traffic, enabling the network to remain operational until the cause is found and fixed.

Conclusion

The future holds tremendous promise for the Smart Grid. But the Smart Grid's future also presents much uncertainty. The best way for a utility to remain fully in control of its networking destiny is to implement a private, integrated communications infrastructure capable of supporting any and all of its applications—securely, and with satisfactory performance. And the best way to make this infrastructure end-to-end is in a multi-tier hierarchy that spans the HAN, NAN and WAN. Such a private, integrated multi-tier Smart Grid network gives utilities the versatility and scalability they will need to accommodate whatever the future holds.

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