Internet Protocol Architecture for the Smart Grid

Executive Summary

The smart grid is a data communications network integrated with the electrical grid that collects and analyzes data captured in near-real-time about power transmission, distribution and consumption. Based on these data, smart grid technology then provides predictive information and recommendations to utilities, their suppliers, and their customers on how best to manage power. Integration of the generation, transmission, distribution, and end-user components of the electrical grid is crucial to the smart grid. For largely historic reasons, the existing electrical grid consists by-and-large of isolated "islands" of transmission and distribution capability with limited communications among them, as these various systems "speak" using different communications methods.

To achieve the level of interoperability and security that will meet the technical and security goals of the smart grid, its data communications network architecture must be built using standard, open protocols. The standard suite of protocols best-suited for the smart grid is the Internet Protocol (IP). IP-based communications networks deliver the reliability, scalability, extensibility, interoperability, manageability, security, future-proofing and cost-effectiveness necessary to meet tomorrow's energy needs.

While IP made the Internet possible, IP is not the Internet. Rather, the Internet is a public network of networks that uses the Internet Protocol (IP) and the Internet Architecture. It is not the only network that does so; corporate networks are typically private networks that also use IP, but connect to the Internet only at controlled points. Various militaries and private entities also run private IP networks, some of which are layered over the Internet and some, which do not connect to the Internet at all.

Given the security, flexibility, and interoperability of IP, Cisco proposes using the Internet Architecture, based on the suite of Internet Protocols, as the basis for the architecture of the Smart Grid, implemented via a migration strategy that avoids market disruption. Therefore, to enable a secure and innovative smart grid system, we respectfully submit that the suite of IP standards should be included on the final list of Smart Grid Interoperability Standards that NIST publishes later this year.

The Challenge of the Current Electrical Grid Architecture

Most of the nation's electricity system was built when primary energy was relatively inexpensive. Grid reliability was ensured mainly by having excess capacity in the system, with unidirectional electricity flow to consumers from centrally-dispatched, coal-fired power plants. Investments in the electric system were made to expand the existing system to meet increasing demand—not to fundamentally change the way the system works. While innovation and technology have dramatically transformed other industrial sectors, the electric system has generally continued to operate the same way for decades. This lack of investment, coupled with increased demand, has resulted in an inefficient and increasingly unstable system.[1]

This outdated grid has both environmental and economic consequences. As seen in Figure 1, electricity generation is the dominant source of growth in CO_2 emissions.[2] As well, power outages and power quality issues cost American businesses and estimated $100 billion each year, while growth in peak demand for electricity exceeds growth in power transmission in the United States by almost 25% each year. [3]

Further increasing the economic cost is the fact that the current electric grid cannot supply consumers and businesses with information necessary to allow them to take advantage of time-based fluctuations in power costs that would enable them to adjust consumption to lower both their energy bills and strain on the power grid during peaks demand. As shown in Figure 2, overwhelming demand on the power grid increasingly results in brownouts and blackouts with massive economic impact. For example, the Northeast blackout of 2003 resulted in a $6 billion economic loss to the region.[4]

---

[1] "Cisco's Smart Grid Vision for America, Cisco Systems, Inc., April 2009.

[2] Energy Information Administration, *Annual Energy Outlook 2009 (AEO2009)—Early Release*, Dec. 2008. For more information, visit: http://www.eia.doe.gov/oiaf/aeo/index.html

[3] "The Smart Grid: An Introduction" U.S. Department of Energy. http://www.oe.energy.gov/1165.htm

[4] Ibid.

Figure 1. Electricity Generation: The Dominant Source of CO_2 Emissions Growth

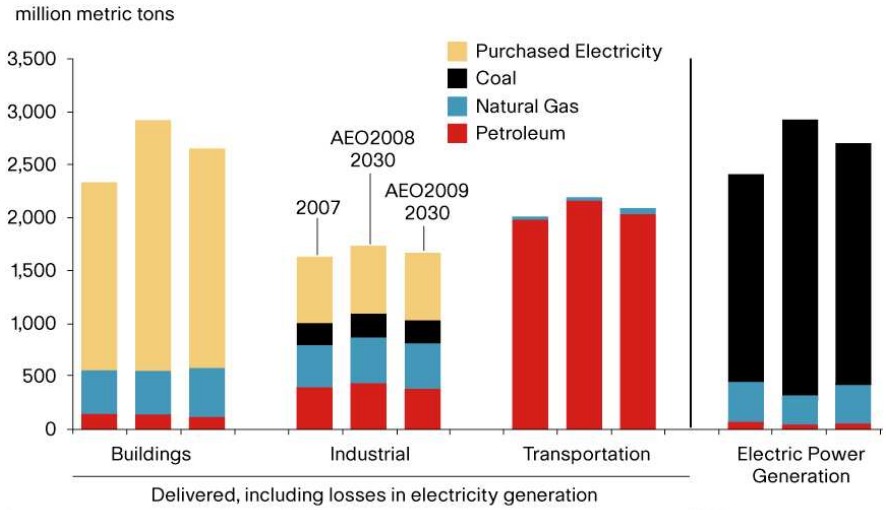

Source: EIA Annual Energy Outlook 2009 Reference Case Presentation—December 17, 2008

The effects of climate change, the rapid rise of fuel costs, the investment in renewable power sources, and the outdated and fragmented grid infrastructure are forcing change upon those who own, manage, regulate, sell into, and use the system. These changes on both the demand and supply sides require a new, more intelligent system that can manage the increasingly complex electric grid.

Figure 2. Yearly U.S. Power Outages in Minutes Lost per Year

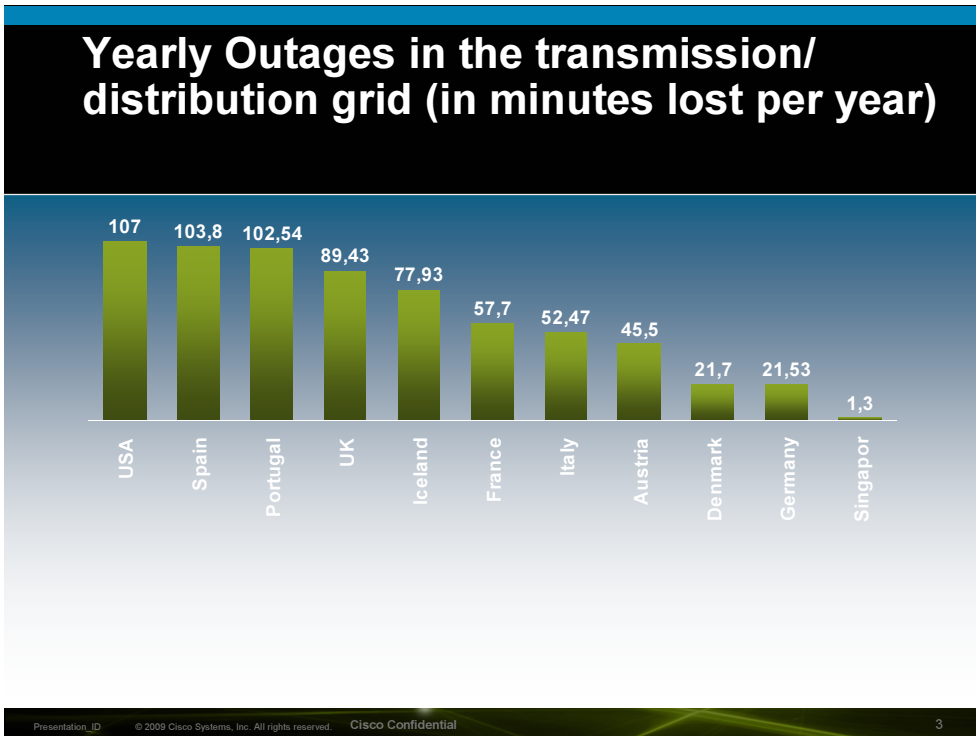

The Smart Grid

Recognizing these challenges, the energy community is starting to combine information and communications technology (ICT) with electricity infrastructure. According to former U.S. Vice President Al Gore, "Just as a robust information economy was triggered by the introduction of the Internet, a dynamic, new, renewable energy economy can be stimulated by the development of an electranet or Smart Grid."

The "smart grid," as it is known, is a data communications network integrated with the power grid that collects and analyzes data captured in near-real-time about power transmission, distribution and consumption. Based on this data, smart grid technology then provides predictive information and recommendations to utilities, their suppliers, and their customers on how best to manage power. Having near-real-time information and recommendations allows utilities to manage the entire electricity system as an integrated framework, actively sensing and responding to changes in power demand, supply, costs, quality, and emissions across various locations and devices. Having better information enables consumers to manage energy use to meet their needs.[5]

The smart grid will enable an ecosystem of interconnected and interoperating functions and services that optimize energy generation, distribution, and consumption. It will enable millions of distribution field devices, thousands of transmission substation devices, millions of customer premises devices (for example, smart meters, home energy controllers, and electric vehicles), data center applications, and customer service and support apparatuses to interoperate and communicate transparently. Facilitating this communication requires interconnection of the disparate components of the electric-grid networks—including home energy management, business energy management, data centers, and substation automation—as well as the interconnections among operators, federal and state agencies, and others.

Modernization of the various components of the electric grid will create millions of endpoints and petabytes of data that will need to communicate in a highly secure, low-latency, predictable, and standardized way. In addition, the electric utility data center and customer support operations that are evolving require more flexible and sophisticated solutions.[6]

---

[5] "Security for the Smart Grid," Cisco Systems, Inc., May 2009. http://www.cisco.com/web/strategy/energy/smart_grid_solutions.html

[6] "Cisco's Smart Grid Vision for America," Cisco Systems, Inc., April 2009.

The Smart Grid: Benefits and Challenges

The smart grid will not only conserve energy and reduce energy bills, but also benefit the environment. The Global e-Sustainability (GeSI) Smart 2020 Report for the United States[7] estimates that, "Smart Grid built on better information and communication could reduce CO_2 emissions by 230–480 MMT of CO_2, and save $15–35 billion in energy and fuel costs." See Figure 3.

Figure 3. Smart Grid: Potential Impact in United States in 2020

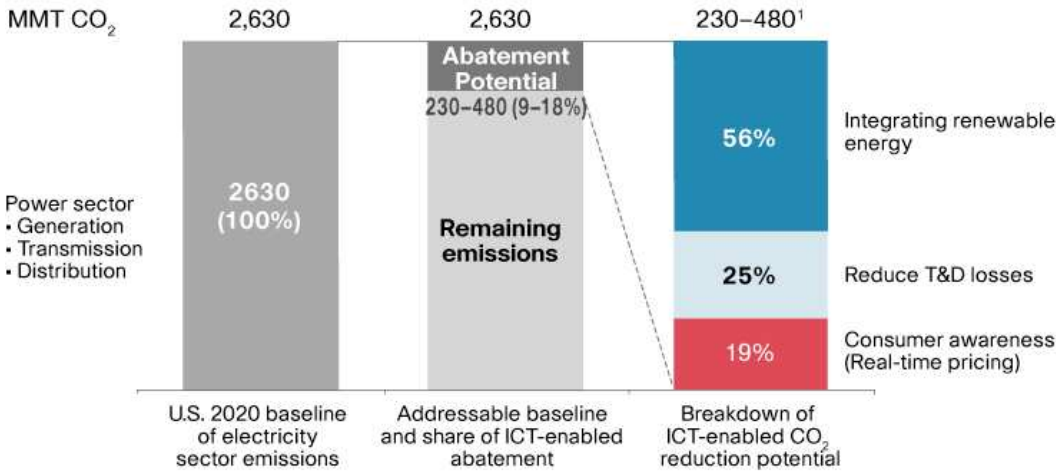

1 Multiple levers contribute to the reduction potentials. The mid-point was used to obtain the percentage breakdowns.
Note: See appendices for baseline, adoption, and reduction assumptions.
Source: Global e-Sustainability Initiative (GeSI), 2008.

To integrate renewable sources of energy, ICT can help balance the unpredictable supply of renewable sources with fluctuations in demand. To reduce transmission and distribution losses, ICT can enable the remote monitoring of grid performance and balance resource usage. Increasing consumer awareness by providing information about usage, prices, and sources can be achieved by installing smart meters with time-of-use prices, intelligent thermostats and appliances that adjust usage based on prices, and web-based interfaces to control and analyze usage. Taken together, these changes are predicted to help lower the GHG emissions associated with electricity generation, distribution, and usage.[8]

To provide the benefits described above, system architects, power industry companies, and policymakers have set a series of goals for the smart grid. Specifically, the smart grid architecture must address the following critical issues:

---

[7] http://www.theclimategroup.org/assets/resources/publications/Smart2020UnitedStatesReportAddendum.pdf

[8] "Cisco's Smart Grid Vision for America," Cisco Systems, Inc., April 2009.

Transmitting Data over Multiple Media – Smart grid data must be able to travel rapidly and reliably over a variety of different network media, from copper cables to fiber infrastructure to wireless networks.

Collecting and analyzing massive amounts of data rapidly – The smart grid must be able to capture massive amounts of data, analyze it, and provide power management recommendations based on that data—all in near-real-time.

Changing and Growing with the Industry – The smart grid must be able to evolve as technological advances yield new hardware, applications and devices. At the same time, it must incorporate such advances into the network with minimal cost and difficulty.

Connecting large numbers of devices – The smart grid architecture must enable communication and correlation of data from approximately 61,000 substations, 180 million transformers and meters, and millions of new smart grid sensors – with the number of devices connected to the smart grid rapidly growing every year.

Maintaining Reliability – High network availability is absolutely critical. As seen in Figure 2, network outages are costly and debilitating—and currently all too frequent. In fact, 41% more outages affected 50,000 or more consumers in the second half of the 1990s than in the first half of the decade.[9] Ensuring uninterrupted electrical service to ratepayers is a prime challenge for the smart grid. Therefore, ensuring that the smart grid data network is reliable, so that it, in turn, can ensure uninterrupted electrical service to ratepayers, is crucial.

Connecting Multiple Types of Systems – The smart grid must connect and exchange data freely with many different types of hardware, ranging from smart sensors in home appliances to home energy meters to transformers and beyond.

Ensuring Security – The unfortunate reality is that because of the critical nature of the technology and the services it provides, the grid becomes a prime target for acts of terrorism and cyberattacks. The transformation of traditional energy networks to smart grids requires an intrinsic security strategy and specific security mechanisms to safeguard this critical infrastructure.[10]

Maximizing Return on Investment – Utilities and regulators must ensure that the estimated $42 billion investment in smart grid technology will reap an acceptable return on investment for ratepayers. [11] This includes minimizing smart grid system deployment expense, operating costs, and risk.

---

[9] "The Smart Grid: An Introduction," U.S. Department of Energy.

[10] "Security for the Smart Grid," Cisco Systems, Inc., May 2009.

[11] Hendricks, Bracken, "Wired for Progress 2.0: Building a National Clean-Energy Smart Grid," Center for American Progress, April 2009.

Underlying each of these goals is the challenge of exchanging data freely and securely among all components of the power grid. To achieve these goals, the smart grid architecture must be open and distributed, consisting of many interconnected, addressable and autonomously-administered entities. These goals mirror the initial goals set when the first Internet architecture was developed, and similar approaches will need to be used in the context of standards that make up the smart grid.

Data Communications Challenges

The existing electrical grid consists by and large of isolated "islands" of transmission and distribution capability with limited data communications among them. Unfortunately, the success of the smart grid depends upon fast and free exchange of data among all components of the smart grid from generation plants to substations home and business meters.

At the highest level, today's grid is operationally and functionally divided into two large systems, the transmission system and the distribution system, each of which has multiple subsystems. In general, each of these subsystem has its own specialized rules for exchanging data within the subsystem. These data exchange rules are known as *communications protocols* or simply *protocols*. Because these different devices "speak" using different communications protocols that are not designed to communicate with each other, they essentially function as information "islands," and thus it is very difficult to integrate communications across systems or to correlate data from different systems and devices.

Current power grid systems now in place do transmit and receive data, but almost always via specialized communications protocols that cannot exchange data broadly with different types of subsystems. For example, the Supervisory Control and Data Acquisition system, or SCADA, communicates via its own specialized communications system and protocol, which is designed only to send and receive SCADA data. So, while the devices on the existing electric grid may exchange data, they generally communicate only within their subsystem and using only their own specialized protocols.

Internet Protocol and the Smart Grid

Smart grid technology must connect the islands of data found in the current electrical grid. Connecting all this data requires adoption of a common communication protocol. Internet Protocol (IP) has been the protocol of choice in businesses and the home for decades. Using the same protocol end-to-end is undoubtedly the most scalable and flexible approach for decades to come.

Overall IP is the glue that allows multiple types of physical equipment to operate seamlessly from end to end over a wide variety of media without adding the burden

of conversion of protocols. Smart grid deployments need the kind of transparent interconnection that IP provides to connect the various types of equipment and sensors that will be deployed to make the grid smart.

Technical Advantages of IP Architecture for the Smart Grid

IP made the Internet possible, but IP is not the Internet – The Internet is a public network of networks that use the Internet Protocol (IP) and the Internet Architecture. It is not the only network that does so; corporate networks are typically private networks that also use IP, but connect to the Internet only at controlled points and are often layered on top of it. Various militaries run private IP networks, such as US DoD's Secret Internet Protocol Router Network (SIPRNet), which do not connect to the Internet at all. Industrial automation is frequently based on the IP protocol with a similar air gap between itself and the public Internet. DOE's own NASPInet, a low latency "real time" phase synchronizer application, is likewise designed to use IP in a private IP network.

Because of the inherent design of the Internet Protocol, the Internet Architecture addresses each of the goals set for the Smart Grid:

Transmitting data over multiple media – The Internet Protocol layers atop underlying link layer networks, such as Ethernet, wireless radio networks, and serial lines, providing a common and flexible way to use and manage a network composed of disparate parts.

Collecting and analyzing massive amounts of data rapidly – Congestion management is a shared responsibility. Performance guarantees can only be given if both the behavior of the transport and the behavior of the network elements are known; they must both behave predictably to manage congestion.

Changing and growing with the industry – One of the principal benefits of the Internet Architecture is its ability to add a capability such as a new transport, a new link layer, or a new application without having to redesign the entire stack.

Connecting large numbers of devices – One of the main challenges with connecting large numbers of devices is providing a unique identifier, or address, for each device. Unlike the many architectures that went before it, IPv6 offers straightforward addressing and routing for a huge network such as the Smart Grid.

Maintaining Reliability – End-to-end reliability is generally managed at a "transport" layer above IP. Different applications may use different transports to meet different requirements.

Connecting multiple types of systems – The primary role of IP is to identify the interface, and thus the system, to which data is going and the interface from which it

came. IP routing delivers the data to its destination, and enables the receiver to respond back to the sender.

Ensuring Security – "Security" is the general practice of avoiding or managing threats, and every communications layer has threats. Hence, security is a shared responsibility. In general, security concerns are managed by a combination of technologies that ensure one is communicating with appropriate systems, ensure the confidentiality of conversations, and ensure that applications behave in a manner consistent with their roles.

Maximizing return on investment – As mentioned earlier, currently the development of standards for the smart grid is highly fragmented, and solutions for its various needs are specified in discreet, isolated standards. The implication for the utility that uses a given solution is that it has to run a network for that solution. Vendors use this as a lock-in strategy. Imagine that a utility buys a set of products that implement a given standard because they get a good price on a specific package on a given date. Part of the purchasing decision is the expected cost of operation over the lifetime of that build-out. When the utility makes the next purchasing decision, the operational and training costs of building a new network for a competitive solution will greatly exceed those of expanding the existing infrastructure, resulting in the utility being biased in favor of the same vendor or sent of vendors *even though a competing solution may be superior.*

In general, the operational costs to a utility will be much lower, and its options far more flexible, if the utility can deploy the best solution for a given requirement at the time it makes the purchasing decision. This argues that the utility should be able to run its existing applications using components from a different manufacturer that implements a common end to end infrastructure with the possible exception of a new physical layer, or be able to run a new application using updated firmware in existing equipment.

A Smart Grid IP Reference Architecture

A design framework for the smart grid based on the Internet Architecture will avoid the long-term costs associated with solutions built upon proprietary standards that create a fragmented electric system, as described above. To this end, Cisco proposes an architecture based on the Internet Reference Model, as described below.

Internet Reference Model

The Internet Reference Model[12] is a close relative of the ISO Open System Interconnect (OSI) Reference Model. Unlike the OSI Model, the Internet Model assumes the application protocol, the way the application encodes its data, and the

[12] In this section, many of the architectural comments apply to both IPv4 and IPv6, but the specifics of IPv6 addressing and its expected place in the market make it a far superior solution for smart grid. As such, we will only discuss IPv6.

management of application sessions are not layers that can be specified with application programming interfaces (APIs) and protocol state machines, but are functions of a single application layer. The Internet Architecture therefore has five layers:

Application Layer	User functionality
Transport Layer	End to end reliable delivery
Network Layer	Delivery of data across a network
Link Layer	Delivery of data across a physical medium
Physical layer	Encoding and management on a physical medium

The Network Layer, which runs the Internet Protocol, traverses a wide array of underlying network protocols to create a seamless addressing context. This independent set of lower layer protocols may include IP, recursively, when building Virtual Private Networks (e.g., by overlaying private networks on the public network, by overlaying an IPv6 network over an IPv4 network to facilitate early deployment, etc). This lower layer may alternatively use other link protocols, such as switched Ethernets, Institute of Electrical and Electronics Engineers (IEEE) 802.11, IEEE 802.15.4, Frame Relay, Asynchronous Transfer Mode (ATM), etc. The primary advantage of this is flexibility: **An application that runs on IP may use any lower level protocol or any combination of lower level protocols without adaptation, because the network decouples the application from the lower layers**.

By comparison, many International Electrotechnical Commission (IEC) standards series, including IEC 14908, 61780, 61334, and 62056, appear to specify the entire communication architecture from the application API to the way bits are encoded in the physical layer. With such monolithic standards, deployment of a new physical layer (such as IEEE 802.15.4g) requires a re-specification of the entire silo, if only to say "use the same one as this other standard." Whereas, layered standards only need to specify the revisions. As a result, the Internet Architecture – which was initially deployed using a purpose-built network specified by BBN TR1822 and called the ARPANET – has migrated to support X.25, the Ethernet, many IEEE 802 series protocols, SONET/Synchronous Digital Hierarchy (SDH), Frame Relay, Link Access Procedure/Protocol Balanced (LAPB), Point-to-Point Protocol (PPP), and ATM. By extension, should another MAC/PHY be developed in the future, IP could be directly migrated to it, and all existing applications be usable on it with little to no architectural change.

Communicating Applications

The systems in a network that initiate connections and run applications are generically called "hosts," or following OSI, "end systems." Applications come in many varieties. In the smart grid, hosts in the home include the power meter and other computers that control power circuits and as a result loads placed on the electrical grid. Related applications at the utility communicate with the meter and

control applications, report readings, issue pricing signals to the home, etc. Other hosts in the power generation and transmission networks include sensors that report telemetry and enable automated control of those systems, and the management systems that interact with the sensors and control points. Hosts in the smart grid back office include more familiar laptops, desktop computers, and servers that support them.

In general, hosts (and the applications that run on them) are identified using human friendly strings that are translated into network addresses by the Domain Name Service. Domain names provide a level of indirection, enabling users to access services like www.example.com without having to know the underlying Internet addresses or choose among possible alternatives addresses. It is possible for applications to refer to other systems by their IP address directly, but this generally results in management problems; it is difficult to locate and change an IP address used explicitly across a large set of distributed applications. Besides providing a human-friendly translation service, the Domain Name Service provides a centralized database to manage any evolutionary address/name binding changes.

Applications Are Located In Places Which Have Addresses

IPv6 addresses identify topological locations in a network, and host interfaces within. For example, if a pole-top router is serving 5,000 neighboring homes, it might have a single subnet (IPv6 prefix), and the homes connected to it choose addresses within that subnet. Datagrams travel to homes by finding the router serving the home's subnet (identified by the prefix in the address); the router in turn forwards the datagrams to the addressed home (identified by a 64 bit interface identifier in the address).

Prior to any communication, an application will typically translate a domain name to an IPv6 address. When a home connects to the utility, it creates an address in that Local Area Network (LAN), translates the utility's service name to an IPv6 address, and opens a session between itself and the utility host running the desired service.

Consider the implications of the address design on the pole-top router. Xerox and Novell allowed for 2^{32} LANs in their architecture, but gave no way to aggregate LANs. Hence, route tables were large and complex. OSI Connectionless Network Service (CLNS) used a 20-byte address that could potentially identify 2^{16} LANs within a single network. Route aggregation required inter-domain routing, with its complexities. The design of IPv6, on the other hand, targeted the ability to address 10^{12} LANs and 10^{15} hosts (RFC 1726), and is in fact capable of 2^{64} LANs (RFC 4291) readily aggregatable into manageable administrative chunks, each containing an arbitrarily large number of hosts. Large networks are allocated prefixes addressing 2^{32} LANs, may obtain larger or multiple prefixes if their businesses require it, and within themselves can aggregate as they deem appropriate (ARIN). Technologies like Cryptographically Generated Addresses enable the network to authenticate the use of those addresses, something no architecture has in the past addressed.

In addition to topological interface addresses, two forms of group addresses are available in IPv6: multicast and anycast. An IP multicast enables an application to send a single datagram to an arbitrary set of receivers. In services that are supported by a large number of interchangeable servers, anycast addressing enables a client to communicate with whichever server is nearest. Group addressing may be useful in demand response applications such as load shedding, and for service discovery.

Support for Application Communication Requirements

Application quality of service (QoS) requirements are a side effect of the application design. *Real Time* applications, such as voice and video, have very specific bandwidth and delay requirements due to the nature of human communications; *Elastic* applications, such as the World Wide Web, adapt themselves to ambient characteristics, using whatever bandwidth is available and delivering service consistent with the network it uses. The Internet Architecture enables many applications to multiplex use of the network, which has both positive and negative implications; applications each have the capabilities they require, but the network may have to allocate capacity intelligently to provide that consistently.

Smart grid applications are in many cases properly described as *Process Control* applications. Such applications are not generally rate dependent, like voice or video, but generally require timely service and can be both delay- and loss-sensitive. Hence, appropriate transport protocols should be used to ensure appropriate delivery behavior, and may require specific configuration for smart grid networks.

The lower layer networks serving each function also need adequate bandwidth for their applications. IEEE 802.15.4 provides 20-250 KBPS over a limited range, depending on environmental factors and encoding, which is ample for simple telemetry and control applications. Backhaul requirements may call for pole-top routers to use cellular radio, IEEE 802.11n, or IEEE 802.16, or wireline technologies when appropriate. Local issues will determine the most appropriate backhaul technology. In the network core, higher-capacity fiber technologies are more appropriate, and traffic engineering may be necessary to manage costs.

Below the network layer, there is a question of burst behavior. Intermixing delay/loss sensitive process control exchanges with more complex file exchanges implies either providing enough bandwidth that bursts are not an issue, or managing traffic rate and priority to provide for the needs of each application. This is accomplished using the Differentiated Services Architecture, which has applications marking traffic and the network providing needed services to each class of traffic.

Secure, Reliable Connectivity between Applications

As noted, security is an issue at every communication layer, and requires solutions appropriate to the threats. Fiber, for example, is often carried in pressurized bundles to make detection of fiber cuts and attacks instant and efficient. Radio networks are subject to signal degradation and interference. In the Internet, attacks

of various kinds make the news. Many attacks are best described as applications misbehaving; a virus is an application distributed in mail or other applications that attack neighboring computers.

Security design starts with a threat analysis, and formulates specific responses to those threats. A likely model for power management layers communication on the power line itself or on communication lines bundled with it. However, that cannot be the only avenue, as power lines get cut. Hence, in addition to power line communications, physical connectivity in the smart grid needs to also be able to take advantage of residential broadband if it exists, cellular radio, and other technologies – and be able to advance to future technologies as they are developed without having to redesign applications.

Common threat models share the issue of the trust one places in a communicating process. *Identification* of those processes, proving that the identity is correct (*authentication*), and enabling those processes to perform only a limited set of actions (*authorization*), are critical to establishing and maintaining limited trust. *Confidentiality,* often required for business information security, can be accomplished using related technologies. The IP Security Architecture is built into IPv6, and provides tools for cryptographic authentication and privacy of communications between hosts.

In addition, network security architectures generally use some form of *network access control*, and active intrusion detection and management. Tools here vary, but are generally about protecting critical regions with prophylactic firewalls, ensuring source address validity, preventing inappropriate access, and monitoring the network.

Network Management

Network management is the weak point of the Internet Architecture; Simple Network Management Protocol (SNMP) is widely-used for monitoring network behavior, but configuration management of the network core is generally accomplished using proprietary configuration languages and scripts. Access to these is generally via Secure Shell (SSH), providing authenticated access. The Internet Engineering Task Force (IETF), the standardization body in charge of IP, is moving in the direction of the NetConf standard, which uses XML-encoded exchanges to parameterize network devices such as switches, routers, and middleware.

At the edge, however, Cable Labs and Digital Subscriber Line (DSL) operators have demonstrated an ability to rapidly deploy and manage tens of millions of consumer Customer Premise Equipments (CPEs).

Choice of Transport Protocol

Concerns have been raised in NASPInet about Transmission Control Protocol (TCP) as a transport. TR-GS-009 reports that:

> Computer science researchers have long understood that TCP/IP, while "reliable" and quite useful for general purpose web applications, is inadequate for many

> mission-critical wide-area applications, because it has narrow coverage of failures and a high and unpredictable latency (for a recent study, see [BCH+05])…
>
> …Fortunately, distributed computing technologies in recent decades have shown a number of ways to provide broader coverage of failures with latencies that are lower and more predictable, using advanced protocols and communication services built, for example, over UDP/IP, as we show later in this report.

It is true that TCP is designed primarily for file transfers and elastic transactions, such as the World Wide Web, and it is true that researchers have written many papers on minor modifications of its congestion avoidance algorithms. An important point to note is that the congestion avoidance algorithms it uses were originally designed for DECNET, and are the same as those used by OSI TP4; the issue has to do with the behavior of windowed transports, not the specifics of the TCP protocol. The lack of papers on the topic in other protocol suites speaks to their lack of widespread use, not their efficacy.

Running an application on User Datagram Protocol (UDP), which has no OSI counterpart, does allow the application to be in control of time variance issues. The nearest OSI relative to UDP is Transport Protocol (TP) 0, which performs segmentation and reassembly, but presumes an underlying circuit with the reliability and ordering characteristics of X.25, and permits no multiplexing. Applications using UDP can, for example, directly control the timing of retransmissions, or simply count losses as losses and move on to the next event. Note that UDP does not do this itself – applications using UDP are forced to implement transport services, like session creation and management, congestion control, and security in the application itself, ultimately making the applications more complex and potentially less robust.

One standard Internet transport that may be worth considering in the smart grid is Datagram Congestion Control Protocol (DCCP) [RFC 4340]. It enables common transport services for normal transport operations including session open and close, congestion management services, etc,, and supports Transport Layer Security [RFC 5238] for authentication, authorization, and confidentiality. It, however, implements an *unreliable* datagram transmission service, leaving acknowledgement, if required, to the application. This is very appropriate for unacknowledged telemetry or command/response applications in which the response is implicitly the acknowledgement of a command and transmission, and retransmission of commands needs to be directly controlled by the application. DCCP also measures the capacity available in the network, advising the application if the application is exceeding the network's capabilities.

Concerns about the Internet Architecture

The following points have been raised as objections or concerns. In the interest of clarity, they are addressed directly in FAQ fashion.

Relationship between the Smart Grid and the Public Internet
The public Internet is a set of cooperating commercial networks that connect a large number of private networks at its edge. These private networks generally connect to the Internet through firewalls, which apply restrictive rules to prevent unauthorized use of their networks. These private networks (assuming IPv4) frequently use private addressing with address translation. Examples of such private networks include most corporate networks and residences attached via residential broadband.

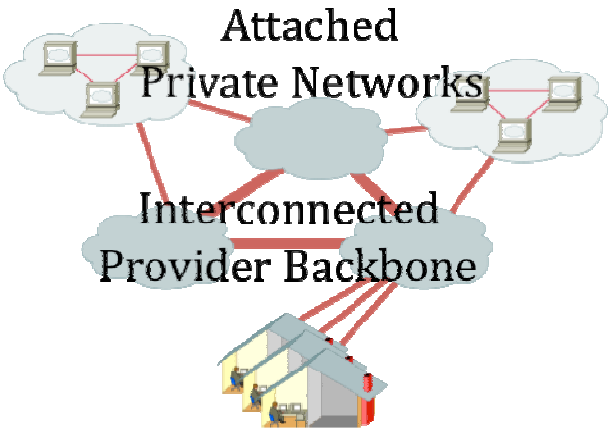

Attached private networks often communicate using separate, non-internet bandwidth, or might use a tunnel overlay called "Virtual Private Networking" to connect privately using the public Internet. Like SIPRNet, the smart grid might use existing infrastructure, such as residential broadband services, to augment its primary paths, but in general, use private bandwidth.

An alternate model might be to offer broadband Internet service as a tunnel overlay over the smart grid. As the smart grid would already have a service delivery point into each premise receiving power, the home gateway could be provisioned to tunnel Internet service through a delivery interface on the far side of the smart grid.

IP Addressing Combining Location and Identity
The Zigbee Architecture, which is itself moving to IPv6, uses an application identifier that is also embedded in the IPv6 address. While companding these is unnecessary (there may be many applications with different identifiers running on the same host, but they need only one IP address), it can be made to work for simple applications.

QoS over IP Networks
As noted, QoS is about application requirements and their implications for the underlying network, and may have a cost/benefit trade-off. Commercial networks are routinely designed that deliver peak delay variation less than ten milliseconds Post Office Protocol (POP)-to-POP and median delay variation on the order of a millisecond, and negligible loss rates (http://www.ieee-infocom.org/2004/Papers/37_4.PDF).

Real World Viability of IPv6
As recently as 2008, the world had not taken IPv6 seriously, although some networks have had it deployed for a decade. Recent announcements from the Regional Internet Registries (ARIN, for example) indicate that the IPv4 address pool is nearing exhaustion. Service providers are in the process of moving to IPv6. We fully expect IPv6 to be deployed worldwide within the next few years.

Managing Mass Deployments of IP Endpoints
Very large numbers of IP endpoints are now deployed in public and private networks, and the Internet handles them gracefully. Examples include the rollout of Cable Modem and DSL broadband networks, which support tens of millions of customers in various markets.

What is different with the smart grid is that large numbers of IP endpoints will be deployed en masse by the network operator, as opposed to being added piecemeal by its customers. The IP community has many tools and techniques available to readily manage this massive rollout of meters, etc. DHCPv6 enables central administration of endpoint addressing and configuration. IPv6 Stateless Address Autoconfiguration can be used for places where central administration is unnecessary. Techniques learned from the Cable and DSL industries are clearly applicable to the smart grid space.

Securing Mass Deployments of IP Endpoints
Securing mass deployments is also doable with planning. In Cable Modem and DSL markets, attacks on the network infrastructure have not generally been successful. Consumer hosts, which are not under the control of the operator, are notorious for their insecurity. CPE routers, such as Linksys and D-Link have, however, proven generally reliable in preventing certain classes of attacks.

Securing large numbers of smart grid endpoints is in large part dependent on deploying them with appropriate certificates for themselves and the systems they will communicate with, and enabling those to be updated in a secure manner. This can be handled as a manufacturing process for initial deployment, and updated as a network management action.

Key to this, however, is a proper notion of "security;" regardless of architecture, the key issues in security are not an application of a few "silver bullet" technologies, but critical analysis of the threats that the application and its network will be subject to. Proper solutions will include notions of identity and the verification of identity, authority and the denial of inappropriate requests, confidentiality of information exchanges, hot redundancy to work around accidental or deliberate infrastructure attacks, and the ability to deploy new security technologies as needs change without rolling trucks.

Application Migration to IP
Numerous non-IP IEC or International Telecommunications Union (ITU) standard protocols have been deployed in the grid, which if IP is used, will need to be ported

to the Internet Architecture. IP has a long history of successful and straightforward application conversion, with OSI applications running on TCP, Novell and Apple porting their proprietary applications to IP to replace their proprietary architectures, and IBM first porting Systems Network Architecture (SNA) applications to Advanced Peer-to-Peer Networking (APPN)/Data Link Switching (DLSw), and then to native IP support.

One of the most noteworthy applications to migrate to IP is voice. While the core major interconnecting circuits have always used the same technology as the legacy telephone system, many of the early feeder circuits, and virtually all of the consumer access, starting in the mid-1990's, ran the Internet Protocol over the voice system as modulated audio. Over time, the economics of delivering voice as an application drove it toward digital compression techniques, resulting in a flip where voice is now just another application running over the Internet Protocol. While analog components of the voice system remain in use, the smaller and smaller islands continue to interoperate with the growing Voice over IP (VoIP) service. For many consumers with direct VoIP today, that is a service running over their connection to the public Internet, although the decision for separation or integration often is driven by the relationship between the provider of Internet service vs. the provider of VoIP service.

Grid applications need not be an exception. Some, such as IEC 60870-5-104, are already hosted on IP; some of the protocols for smart metering are designed atop IP. The remaining protocols could also be adapted to the Internet Architecture, as well.

Network Availability

Network reliability is undoubtedly a key requirement for the networks that will support smart grid applications and tolerance to network failures (software, hardware) is extremely limited, to say the least.

With the increasing number of sensitive applications such as voice, video and other real-time application supported by IP networks, high reliability has become crucial to present service providers. This was in part due to their existing service requirements, but became particularly important when Service Provides started to migrate their Public Switched Telephone Network (PSTN) over IP network infrastructures. PSTNs had been notoriously reliable, and IP-based networks had to provide similar reliability levels.

To that end, IP protocols have been greatly enhanced in many ways:

- New protocols have been developed to provide fast failure detection and mechanisms have been designed to receive failure notifications from lower layers with extremely short delays.
- IP routing protocols in charge of routing packets have been augmented so as to find alternate paths (what is referred to as "routing convergence") in a matter of tens to hundreds of milliseconds.

- IP software and hardware have been enhanced so as to allow for software and hardware upgrades with no traffic interruption.

Implementation Recommendations

To achieve the level of interoperability and security that will achieve the technical goals stated earlier in this paper, the smart grid data communications network architecture must built using standard, open protocols. We propose using the Internet Architecture as a basis for the architecture of the smart grid.

One of the concerns of smart grid equipment vendors, particularly Advanced Meter Infrastructure (AMI) vendors, has been that a mandated fundamental change to their product offerings could result in an "Osborne Effect." Osborne Corporation, the company to which the name refers, made a marketing error by announcing a product that was significantly superior to their current offering and which they were not ready to deliver. As a result, all sales stopped, and they ran out of money before they could deliver the new product. Vendors see a very real possibility of this happening, and have understandable concerns.

To avoid a market-crushing Osborne Effect, between now and the time advanced products are ready to ship, existing products must continue to be sold, and utilities must be able to continue to purchase and build networks to support them with gateways to the more general network. The specification and introduction of smart grid technologies must allow an evolutionary approach that may begin with the deployment with "islands" to be interconnected later and that may re-use existing technology to fill temporary gaps. In short, a seamlessly interoperable, end-to-end smart grid won't happen with just a few targeted investments overnight.

In the meantime, AMI vendors, for example, must build an end-to-end architecture that enables more cost-effective specification, manufacture, deployment, and management of products. If an 802.15.4-based Home Area Network (HAN) is appropriate to a home, for example, one could use the electrical connection to the utility as a primary link, and a Virtual Private Network running over the home's broadband connection (the means by which it receives TV and telephone connectivity, as well as general Internet access) as a backup link, and run the management applications over a common network architecture to the devices in the home.

Deployment of this new architecture is done the same way one deploys any of the proprietary architectures, with one important difference – the others all have gateways to the new network. When the new architecture is initially deployed, the total build-out may in fact be smaller than the existing proprietary networks and do nothing but interconnect them. The legacy proprietary networks continue to operate as long as the utility considers them useful. However, new build-outs use the interoperable architecture, and all vendors are able to sell into them, saving

costs for the utility due to competition and given the vendors market share due to interoperability and differentiation based on functionality. In time, the interoperable architecture therefore dwarfs the legacy proprietary networks, and they become less important.

This same model was followed in building out the IP Internet itself, and the fact of its existence and widespread use is an existence proof that the model works. If the proprietary technologies that were in use prior to the widespread deployment of the Internet had been as much better as they claimed to be, they would still be in use; for both the vendors and the network operators, interoperability and interconnectivity have proven to be of such value that the networks the Internet replaced are barely even remembered at this point.

This implementation will be done by an entire ecosystem of smart grid vendors and integrators. The ecosystem Cisco sees necessary has as its hallmarks interoperability, adaptability, and innovative flexibility. In our experience, lower layer technologies, such as IEEE 802.15.4, come and go as technology progresses and market requirements change. As such, it is in the utility's interest to be able to use the best set of products for a job without having to completely retrain and retool.

For example, AMI networks that are currently using IEEE 802.15.4 or IEC 14908 should be able to adapt by adding new underlying technologies without having to change their applications or see new stovepipe standards to describe the environment.

As such, we expect the market to include sets of vendors that provide network communications, Advanced Metering technology and services, and may include intelligent control services that offload the consumer, while enabling the market. In general, we find that markets that devolve to a single vendor are not well-served; we would recommend that each facet of the problem have a selection of vendors addressing it in a capitalist fashion, with the innovation and economies that competition brings. We would also avoid regulation in this space; in our experience innovation and regulation don't cohabit.

Timeframe of Recommendations

In the near term, 2009-2010, AMI vendors should not be mandated by government to change their products. While the current specifications have obstacles, the market and any mandates from NIST would be best served by introductory tests of the technology. However, those networks that are currently using IPv4 must deploy IPv6 as well, because address space will become an issue in 2011-2012, and deploying IPv6 now will provide a basis to deal with those issues as they arise. Consumers will find things to like and dislike in this transition phase, and both the products and the service offerings must be designed to be adapted accordingly.

In the medium term, 2011-2012, early-deployment networks must be interconnected to their utility's networks using gateway technology, and utilities must deploy technology that uses an IPv6 end-to-end architecture. This will help to promote the goals of interoperability, adaptability, and innovative flexibility. While this should not be a regulatory mandate, utilities should use the power of the purse to encourage this transition.

At this point, the industry does not know the long-term future of smart grid technology. Therefore, utilities and regulators must remain flexible, allow the market to evolve and allow vendors to innovate.

Conclusion

The Internet Architecture is the proven, scalable, secure, cost effective, and interoperable foundation for communications, information and commercial networks around the world. Cisco believes that the Internet Architecture should similarly serve as the foundation for the smart grid. As such, Cisco respectfully suggests that NIST uses the Internet Architecture as the framework for smart grid interoperability, and adds the suite of IP standards and protocols to its list of Smart Grid Interoperability Standards.

Cisco recognizes that the transition of existing electrical systems to an IP-based smart grid will not happen overnight. Rather, the transition to a truly interoperable smart grid will take time.

Cisco is committed to partnering with industry, standards bodies and government to develop a migration path for utilities and government that optimizes existing assets, and over time, migrates new investments into an IP-based smart grid that delivers the reliability, scalability, interoperability, manageability, security and cost-effectiveness necessary to meet tomorrow’s energy needs.