

The Evolution of the Internet Community and the “Yet-to-evolve” Smart Grid Community: Parallels and Lessons-to-be-learned

Charles McParland, Lawrence Berkeley National Laboratory

November 2009

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

The Evolution of the Internet Community and the “Yet-to-evolve” Smart Grid Community: Parallels and Lessons-to-be-learned

Charles McParland

Lawrence Berkley National Laboratory
1 Cyclotron Rd., 50B-2270B
Berkeley, CA 94720
mcparland@lbl.gov

Keywords: Smart Grid, Internet protocols, Internet ecosphere

Abstract

The Smart Grid envisions a transformed US power distribution grid that enables communicating devices, under human supervision, to moderate loads and increase overall system stability and security. This vision explicitly promotes increased participation from a community that, in the past, has had little involvement in power grid operations – the consumer. The potential size of this new community and its member’s extensive experience with the public Internet prompts an analysis of the evolution and current state of the Internet as a predictor for best practices in the architectural design of certain portions of the Smart Grid network.

1. INTRODUCTION

Although still evolving, the vision of the Smart Grid is that of a community of communicating and cooperating energy-related devices that can be directed to route power and modulate loads in pursuit of an integrated, efficient and secure electrical power grid. The remaking of the present power grid into the Smart Grid is considered as fundamentally transformative as previous developments such as modern computing technology and high bandwidth data communications. However, unlike these earlier developments, which relied on the discovery of critical new technologies (e.g. the transistor or optical fiber transmission lines), the technologies required for the Smart Grid currently exist and, in many cases, are already widely deployed. In contrast to other examples of technical transformations, the path (and success) of the Smart Grid will be determined not by its technology, but by its *system architecture*. Fortunately, we have a recent example of a transformative force of similar scope that shares a fundamental dependence on our existing communications infrastructure – namely, the

Internet. We will explore several ways in which the scale of the Internet and expectations of its users have shaped the present Internet environment. As the presence of consumers within the Smart Grid increases, some experiences from the early growth of the Internet are expected to be informative and pertinent.

2. A VERY BRIEF HISTORY OF THE INTERNET

The term “internet” has multiple meanings and contexts. But, references to the “Internet” usually focus on two areas. One is the set of low level network protocols used to communicate between processors, switches and components of the internet. The other is the collection of data and web services that form the environment users experience when surfing and interacting through web browsers and specialized programs. The development of both of these aspects of the Internet are relevant to the Smart Grid and will be examined separately.

2.1. The Internet Protocols

Like many other protocols, the internet protocols (IP) began development in the early 1960’s. Under DOD funding, the IP protocols were originally developed to provide a robust, self-healing computer network for defense-critical computers and applications. The result was an elegant and fairly simple design that required a minimal amount of centralized administrative effort to support the exchange of messages between participating computers. True to their fundamentally simple design, the internet protocols concentrated solely on the task of reliably exchanging messages over communications links that were, at least in the early 1960’s, potentially unreliable themselves. These protocols made few demands on the internal design of the computer applications that utilized them and, thus, helped create a “layered” application programming model that improved overall reliability and application design freedom. This successful approach of isolating program responsibilities inspired the development of the well known

OSI 7 layer protocol “stack” model [1] that is shown in figure 1.




Figure 1. OSI 7 layer Model

The internet protocols were only one of many competing network designs that vied for mind – and market – share during the 1970’s and 1980’s. In the ranks of network designers and programmers, this period was known, literally, as the “great protocol wars” and strongly held technical opinions, on all sides, were the rule [2]. By the early 1990’s, it became clear that the internet protocols had prevailed and were quickly becoming the primary networking mechanism for inter-computer communications. Increased investment and deployment of IP communications infrastructure rapidly followed. While this outcome was based on a number of factors, the clean separation between IP networking and application program layers were a good match for an expanding variety of new networked computer applications, such as the world wide web. The result was the foundation of the “Internet” and the development of a unique “ecosphere” shared by millions of individuals.

2.2. “Map vs. Territory”, The Internet Ecosystem

While the IP protocols make the Internet possible, the actual Internet that we access and search is a very different thing. The basic IP protocols require a small number of specialized programs that assist in maintaining and diagnosing network operations and integrity. These functions include name servers that help translate numeric IP addresses into human-readable forms and programs like “ping” that allows verifying the presence of a particular computer on the network at any given time. With the exception of the staff that manages the issuance of Internet numbers and names, little additional effort is required for network governance. Network service providers (ISPs) are able to independently configure their equipment and service their users with little interaction required from higher network administration levels. This is essentially the “map” portion of the analogy and it’s worth noting that the lack of a large centralized administrative and support staff is both the genius and

primary strength of the original internet protocol design. Given its ability to support the continually-evolving Internet on a basically unchanged protocol design, the IP protocols represent a success story worthy of study.

The actual “territory” of the Internet is the tangible environment within which Internet users interact. The Internet, as we know it, is a set of well known services that, over time, have evolved to occupy this environment. They create a recognizable and evolving “commons” that is available to anyone capable of accessing the Internet. As the Internet evolved, programs, such as bill board systems - BBS), were developed to allow users to directly interact with each other. With increasing program sophistication and increasing network bandwidth, these services ultimately changed into today’s Facebook and Twitter. Search services such as Google that are widely used to conveniently navigate the Internet and locate services and web sites evolved from earlier, now little used, efforts such as Gopher. In all these cases, the stability and flexibility of the underlying internet protocols enabled the continual evolution and improvement of the Internet.

While it would be a stretch to think of all Internet users as a special “community” with similar interests, the Internet has evolved an identifiable environment on which users rely. The Internet is essentially a “built” environment that receives little guidance from its administrative staff on behavioral values or standards. However, the Internet ecosystem does tend to express the values and agendas of its users. So, for example, automated port scanning of computers on the internet, an internet-legal exercise frequently used to probe for a “hackable” security weakness, is forbidden by most ISPs. And Internet sites that allow users to post feedback and review products all require responsible behavior from participants. The Internet ecosystem is beginning to display some of the common behavioral values associated with the concept of a “community”.

For some users, these values are simply a consistent “look and feel” that permeates almost all applications with which they interact. For others, it’s an expectation of consistent and reliable back-end financial settlement services that promote e-commerce. But, for all Internet users, it’s a relatively consistent and supportive environment that allows effective interactions with both other users and automated services. We will refer to this difficult to define quality as the “Internet ecosystem”.

Both the size and the economic impact of this community have prompted studies of user expectations and preferences. These studies have helped formalize the current “look and

feel” of the Internet and, in many ways, determine the kinds of interactions that are offered. [3]




Figure 2. Evolution of Web Site User Expectations

2.3. Size and Scale

Present Internet usage within the United States is surprisingly high, with actual user numbers only exceeded by China. Out of an estimated population of 307 million, roughly 227 million are estimated to regularly use the Internet [4]. At approximately 75% of the population, US usage is consistent with that of Japan, Canada and much of Europe. The explosive growth of web sites available to visit is equally interesting – increasing almost exponentially over the past three years to almost 160 million sites in 2008 [5]. The continued effectiveness of the Internet, particularly in the face of these increasing usage numbers, indicate an ongoing enterprise that is successfully serving its user community – a goal of great importance to the Smart Grid deployment effort.




Figure 3. Growth of Internet websites 1990 - 2008

3. INTERSECTION OF SMART GRID AND INTERNET

As is the case in the present power grid environment, the Smart Grid will make substantial use of IP-based networks

for power distribution system data communications and normal business-related IT functions. With the exception of real-time constrained communications and specialized device-specific protocols, there is general acceptance of the utility and cost-effectiveness of IP networks within the power generation and distribution domain. And, while data security issues remain a serious concern for many portions of the power grid infrastructure, existing IP technology has already been leveraged by creating private, physically isolated networks, by implementing encrypted virtual private networks (VPN) that utilize the public network while maintaining a degree of advanced security, or by encapsulating and tunneling non-IP protocols over IP routed infrastructure.

3.1. Smart Grid and Internet already share a context

Having said this, it is interesting to examine the likely context for IP usage within the emerging Smart Grid. As pointed out, IP networks will be used in application areas where their economic and performance strengths permit and not used where security and latency issues dictate other solutions [6]. The specialized requirements of distribution and sub-station monitoring and control have already promoted the development of specialized non-IP data communications protocols. Given latency demand constraints for these applications and their large installed base, it is unlikely that these applications will move to IP networks in the near future. In fact, given the high level of engineering and administrative effort needed to maintain fixed power distribution system resources, it is not clear that these application areas would substantially benefit from the inherent flexibility and reconfiguration capabilities found in IP networks. While this topic is the subject of on-going research, it could be argued that the design of distribution control systems that are capable of the rapid and seamless reconfiguration (i.e. facilitated by the IP protocols) is not yet a mature discipline and should be approached with some care [7]. So, within the power generation and distribution portions of the Smart Grid, IP protocols and, to some extent, the Internet itself, will be used in piecemeal fashion within the existing framework of acceptable grid engineering practices.

3.2. Relationship of the Internet and the GridWise Architectural Framework (GWAC Stack)

The GridWise Architectural Framework, also known as the GWAC stack, represents an effort to codify, at a high level, the key elements of the power grid enterprise and describe the relationships and interfaces between these elements. [8] When accompanied by additional constructs that enumerate issues that cut across the multiple areas of technical, informational and organizational grid activities, the GWAC stack methodology creates a reasonably complete picture of

how the grid operates. Given the scope of grid activities and interactions that are represented within this diagram, it is reasonable to ask if such a methodology could be used to describe the Internet as well. The short answer is yes. The actual internet protocols would comfortably fit within the first two layers of the technical portion and the application level protocols, including both syntactic representation and semantic meaning functions would properly occupy the remainder of the diagram.




Figure 4. GridWise Architectural Framework

However, while the GWAC stack may be useful in describing the relationship of elements of any particular Internet application, it does little to describe the entirety of the present Internet environment. At its core, the GWAC stack enumerates functions, relationships and interfaces for a single enterprise (i.e. the Smart Grid). In this context, the success of the Smart Grid enterprise will depend on effective and binding governance of all of the functional interfaces depicted in this diagram. Acceptable operation of this large, integrated system will require more than simply the development of suitable standards. It will require their mandatory adoption before any element can participate effectively within the Smart Grid. This fact is widely appreciated and organizational efforts to create a suitable form of standards and behavioral governance for the Smart Grid are now under way.

While these issues of standards governance are applicable to any enterprise that combines technical, informational and organizational elements to accomplish its task, the Internet, unlike the Smart Grid, is not a single enterprise. Throughout its history, multiple and varied enterprises have

simultaneously inhabited the Internet and, over time evolved with little coordinated governance. It is true that, as some applications (e.g. eCommerce) gained a large and successful presence on the internet, important cross cutting issues (e.g. financial settlement, information security) started being treated in a common and standardized manner. Clearly, the almost complete acceptance and support of standards in some areas gives the appearance of Internet behavior being governed in a coordinated way. However, there has not – and, currently, is not – any high level governance that requires particular approaches on any given issue. At the highest levels of the Internet stack, acceptance of any application-level standard is an individual and voluntary decision.

In the Internet ecosphere, change is more the rule than the exception. Even the long established financial “back end” credit card services have seen new service models (i.e. PayPal) appear and take hold. This level of freedom has fostered the development of entirely new enterprises like uTube and Google Map/Earth. And, because of the minimal governance style of the Internet, they were free to evolve alongside more conservative applications, such as online banking, with little interaction or interference.

While the GWAC stack depicts the development of a single, stable Smart Grid enterprise, within the Internet ecosphere, this same model describes the almost continual – and simultaneous - formation of many new enterprises. At the highest level, the model terminologies - vocabularies and task descriptions - of these two worlds are similar. But, at the behavioral level, they are very different. The necessarily “closely” coordinated governance of the emerging Smart Grid stands in stark contrast to the minimal governance style – “open” – seen in the Internet at large. While this culture clash may have little adverse effect on the core design and evolution of most elements of the Smart Grid enterprise, at the point where these two domains intersect, i.e. the Internet-dominated world of the Smart Grid end user, the difference between these two worlds will become very apparent.

3.3. Smart Grid Population Boom – the Consumer

As stated earlier, one of the major motivations for constructing the Smart Grid is to monitor and, to the extent possible, moderate end use loads that are connected to the national power generation and distribution grid. While future power storage techniques may make it possible to store electrical power generated during off-peak times and distribute it when needed, the primary presently available path leading to greater stability and control of the national power grid is the real-time monitoring and control of increasingly larger numbers of attached power loads. This requirement directly motivates the implementation of an expansive data communications network that allows

instrumented power loads to communicate their status with other power grid components and permit some degree of grid-wide control. Given that the residential and commercial sectors are responsible for approximately 39% [9] of current US energy use, the attached loads (both aggregate and individual) within this sector are important candidates for inclusion in the set of “devices” that will be required to communicate within the emerging Smart Grid. Furthermore, it is expected that the development and growth of the electric automobile industry will move personal auto power consumption from the oil-based transportation sector to the predominately electricity-based residential/commercial sector – thus further increasing the motivation for integrating residential load monitoring and control into the Smart Grid. Therefore, the single area where new participants will enter the Smart Grid community – in large numbers – is the power-consuming end user.

It has been said that the Smart Grid will, in a tangible way, eventually touch and affect every household in the US [10]. Given the level of energy consumed within the residential sector, there is little doubt that the Smart Grid will be motivated to instrument large numbers of residential loads. As the Smart Grid moves forwards, installation of residential Advanced Metering Infrastructure (AMI) [11] and Smart Grid-enabled appliances will bring both residences and, more importantly, their occupants into the Smart Grid world. These residential households will, either through their utilities or local power aggregators, become active participants within the Smart Grid domain. As the pervasiveness of the Smart Grid expands, so will the number of households affected – numbers far in excess of any other organizational entity within the Smart Grid and, ultimately, will become the largest identifiable stakeholder within the Smart Grid.

To the extent that participation remains passive, perhaps because Smart Grid demands on household behaviors are initially minimal, these large numbers of end users may remain unimportant. However, as power grid constraints begin to impact their daily lives or as their active participation is encouraged through ancillary energy service providers, they will interact more closely with the grid. And, as described above, given the pervasiveness of the Internet as the standard mode of interacting with external organizations and entities, end users will expect and, eventually demand, that their Smart Grid interactions take place through familiar, web-based Internet mechanisms – those that they already use for activities varying from checking the weather and paying bills to downloading movies.

3.4. A Word on Security

Communications security is central to many aspects of Smart Grid design and operation. While security is best

treated as an integral part of any overall communications design, a detailed discussion of Smart Grid security is beyond the scope of this paper. In particular, portions of the Smart Grid architecture that place the highest constraints on network latency and security are typically far removed, architecturally, from the consumer and are specifically not discussed here.

However, since the nexus of the Smart Grid and the Internet is being discussed, the question of overlapping security domains is an important issue. To the extent that some Smart Grid functions, notably those involving the end user, are performed via the Internet, they will take place within the security domain of Internet service providers. This is a domain with an existing, and evolving, collection of security services already in use to service areas such as e-commerce. While security implementation on the Internet can take on a number of different forms (e.g. SSL, Public Key Systems, time-synchronized key generation cards, etc.), the particular mechanism chosen is ultimately administered by service providers within the Internet domain. Providing a seamless security environment that spans both the Internet and Smart Grid domain may be problematic. While it may be possible to layer security technologies in such a way that both Internet and Smart Grid security domains co-operate to control user interactions, in practice, multiple, layered security mechanisms – particularly when separately administered within different network domains – are difficult to incorporate into existing Internet frameworks and present awkward user interfaces. In the end, the suitability of any transaction over the Internet will need to be evaluated on the basis of best practice Internet security mechanisms and those Smart Grid transactions that are deemed too critical for the available level of Internet security should be disallowed.

4. INTERNET LESSONS AT THE PROTOCOL LEVEL

At the “source” end of the Smart Grid, we have power generation and transmission systems. The cost and importance of these facilities require that their control systems be highly engineered and carefully configured to maximize their safe and secure operation. As noted earlier, while IP protocols and, perhaps, the public Internet may have some role in their control and monitoring systems, their inherent and operational value dictates conservative control system design practice.

However, as one looks at the consumer end of the Smart Grid architecture, the networking terrain changes. As the Smart Grid grows towards the benchmark of touching every US household, roughly 114 million in 2010 [12], with potentially multiple network addressable devices in each household, we approach the scale of the present Internet. In terms of both numbers of nodes and overall network

stability due to intentional and accidental network outages, the consumer end of the Smart Grid shares many characteristics of today's Internet.

The issue of how to reliably find nodes within the Internet and, once found, how to properly "phrase" digital communications with them has been focus of multiple software engineering efforts within the networking community. Smart Grid applications that promote grid-related messaging between "smart devices" will share many of the design and operational problems addressed by these internet engineering efforts. A brief examination of the history of these efforts may prove fruitful when applied to the development of the Smart Grid.

4.1. How Do You Find Things on the Internet?

As with all data networks, each node within the Internet is assigned an individual numerical identifier or network address. Ideally, each node's address is unique. In practice, the present Internet has grown beyond its original ability to assign unique addresses to every device and, as a result, some portions of the Internet remain partially "hidden" private subnets. This shortcoming is being addressed by modifications to the basic, underlying IP protocols and is referred to as version 6 of the IP protocols (IPv6). Fortunately, Smart Grid transactions that transit the Internet will have the benefit of these improvements.

While it is expected that the lesson of running short of network addresses is already well appreciated by engineers everywhere, it is worth noting that running out of numbers is not uncommon in the history of the computing and networking worlds (e.g. growth of PC addresses from 16 to 32 and, ultimately, to 64 bits in length). If fully realized, the vision of the Smart Grid includes a vast number of communicating devices and services. While Internet lessons of network addressability may have been universally learned, there may well be other areas within the Smart Grid architecture that suffer from artificially constrained addresses or identifiers. Scaling systems to accommodate large numbers of components can create problems in unexpected ways – many of which have been encountered during the development of the Internet. When conceptualizing a large, geographically-dispersed system with a very large number of devices and services, a detailed examination of Internet best practices will be of benefit.

4.2. Third Party Web Search Services

The advent of large, highly capable search engine services demonstrates the tremendous system-level flexibility made possible by the original IP protocols. While services such as Google, Yahoo and MetaCrawler were never envisioned when the IP architecture was being formulated, each of these successful search services have been designed and deployed without alterations to the original underlying

protocol. For many network users, these applications represent the only tool necessary for navigating throughout the Internet and, as far as their experience is concerned, they "are" the Internet. While some Internet URLs (e.g. www.ebay.com) are easily remembered, often the simplest way to find a particular site is to simply search for it by providing a partial name or inclusion of a key word likely to be found as part of a web site's content. Anecdotal evidence of the through, highly detailed cataloging of web sites accomplished by these services is common. With the growing scale of the Internet, these services may prove to be the *only* effective way to navigate and locate services on the web in the future.

Although the use of search services has become common for users within the Internet ecosphere, what possible roles can such services have in the Smart Grid domain? Much will depend on the level and scale of innovation at the consumer end of the Smart Grid. The highly engineered control systems within the power generation and transmission portion of the grid will continue to require careful and well understood network communications patterns. Since they will consist of well known resources, there will be little value in locating them on the web by searching for them. However, in a future world of truly distributed power generation and storage, it may be reasonable to apply these powerful tools to the end user portion of the Smart Grid. For example, it may be useful, within a limited geographic area, to search for customers that have power available for injection into the grid. It may also be effective to search, within a particular power distribution area, for facilities advertising unusual or ad-hoc load curtailment opportunities. Effective application of these search services implies that some elements of the Smart Grid architecture need to properly "advertise" themselves on the Internet in order to be seen and properly cataloged by these services. Future leveraging of these services will only be possible if accommodated by the underlying Smart Grid architectural design.

Third party search services are not being suggested as replacements for the dedicated communications protocols and systems (e.g. OpenADR, SEP 2.0) [13] that are now being defined and engineered for operational demand response applications. By their very operational nature, these systems will need to be carefully administered and configured and, therefore, not benefit from the "web search" paradigm. However, it is not unreasonable to expect that future refinements to Smart Grid operations or ancillary services could make effective use of the capabilities offered by a Google or Yahoo in locating sites offering various power-related services.

This entire scenario implies that, conceptually, a portion of the Smart Grid ecosphere exists within the Internet. In other words, the Smart Grid architecture is expansive enough to

accommodate some role for the innovative and powerful services that are presently available within the Internet ecosphere. These functions and services may be limited to searches for power related service offerings and market-supporting financial functions. Internet lessons at the “ecosphere” level

4.3. Building Complex Applications on the Web

As the reliability and scale of the Internet grew, it became possible to design single applications that were composed of many programs executing on computers at various locations within the network – all co-operating to accomplish a single, complex task. Instead of users logging onto multiple computers and overseeing the coordination of multiple tasks, programs were being designed to find and communicate with their designated “peers” to accomplish these operations with little or no user assistance. However the development of these “peer-to-peer” applications, such as e-commerce, inventory control and financial transaction processing, required communicating components to carefully synchronize the contents of their messages to match program communications expectations at each end. Successfully managing this process in a way that was reliable, flexible and easy to maintain proved to be an extremely difficult task and has dogged the IT and networking world since the early 1990’s. It was widely held that something else besides the two communicating peer programs was required to ensure their correct and functional rendezvous within the network. Whereas the traditional high-level diagram of the Internet has a network “cloud” at its center, there was a growing concern that something more was needed to make these distributed enterprises work. This assisting entity was termed “middleware”.




Figure 5. Relationship of Middleware to Web Applications

4.3.1. The Birth and Death of Middleware

The design, development and deployment of appropriate middleware programs throughout the Internet became the battle cry for IT and network engineers in the early 1990’s. Since this coincided with the growth of a new programming model referred to as “Object oriented programming”,

middleware was thought of as an enabling software technology that allowed programs to cooperate over the network by finding and executing small program elements (i.e. objects) on remote computers. As with Internet-wide engineering efforts, competing designs were pitted against each other in demonstration battles and in heated and highly opinionated discussion forums. The “winning” technology was a standard known as CORBA (Common Object Request Broker Architecture) [14].

While the details of this system are not germane to this discussion, one element is worth stressing. A great deal of effort was invested in the design and implementation of a critical piece of network software that was deemed necessary for the continued growth and evolution of the Internet into new and beneficial application areas. While this effort required additional complexity within the network, most software designers agreed that the lack of such a service in the original set of IP protocols was a serious shortcoming.

But, as time passed and major applications began depending on the ubiquitous availability of these new middleware services, network administration costs for these applications began to soar and increasing programming effort was required to properly manage the middleware that governed how applications were deployed on the Internet. It soon became apparent that in the network world, as in most other areas, everything has a cost. The additional cost and complexity of maintaining and updating these newly-required services was not anticipated and, lacking sufficient support for these middleware components, the Internet became a far less flexible and accommodating environment for these applications.

Eventually, a group of network and IT software engineers crafted an alternate technology that eliminated the need for additional Internet services (i.e. middleware). This new software approach, known as SOAP (Simple Object Access Protocol), allowed communicating programs to share tasks previously relegated to middleware components. In the end, by sacrificing a bit more communications bandwidth, they were able to eliminate the added complexity of additional Internet service components. In keeping with the battlefield-like analogy that permeated the original “great protocol war”, several influential network experts confidently announced “the death of middleware”. [15] And, while a number of large stable applications continue to use the CORBA middleware architecture, the mind share of the SOAP efforts has grown into the web services and service oriented architecture (SOA) applications of the present Internet.

What is the relevance of this tale to the evolution of the Smart Grid? First, issues related to distributed, cooperating programs are typically uncovered within the Internet

community and addressed there first – its history is worth careful examination. In the end, all complexity has a cost and solutions within the Internet are judged solely on their effectiveness and simplicity – proving the adage that the original Internet design is surprisingly “lean and agile”. Efforts, even major ones as described above, to address a particular shortcoming can, and in many cases have been, discarded because they have been found to add unnecessary complexity to an otherwise conceptually clean design. The design practice of adding complexity only when absolutely necessary and being constantly wary of the network’s scale has served the Internet development (and indirectly, end user) community well. So, as the large-scale, end user portion of Smart Grid evolves, close attention should be paid to Internet engineering efforts that have likely encountered and solved (in some cases, on multiple occasions) similar problems.

5. LESSONS FROM THE END USER COMMUNITY

5.1. Home Automation

While most home automation systems and devices do not directly rely on IP protocols, they are, in increasing numbers, interfaced to a variety of Internet services. The growth of the home automation market and its integration into the Internet ecosphere presents both opportunities and architectural issues for the expansion of the Smart Grid into the home environment.

5.1.1. The State of Home Automation

The home automation market has developed over the past two plus decades and represents a wide range of devices, systems and services. In its infancy, the market was targeted at the convenience factor stemming from the ability to remotely control major household lighting and HVAC systems. Since its “value proposition” was based on convenience, hardware components and systems were highly cost-constrained. This cost sensitivity resulted in a dependence on communications technologies that, while relatively affordable, were not as reliable as their more expensive IP protocol- capable counterparts. The result was an industry with a plethora of communications technologies and private standards with varying degrees of reliability and little or no interoperability.

Eventually, as the explosive growth of the Internet drove the cost of reliable, network-capable data communications technologies down to affordable levels, modern communications technologies and design practices entered the home automation domain. In particular, the availability of inexpensive wireless communications technologies, with their low deployment and installation costs, has driven the installation of home automation systems. Installation of whole house systems is expected to 4 million households by 2013 [16].

5.1.2. Home Automation and the Internet

While not explicitly linked, strong growth in residential Internet subscribers has coincided with growth in home automation system installations. One result has been increasing interest in joining these domains – the nexus of the convenience proposition of home automation and the ubiquitous access proposition of the Internet. Examples include internet-based home HVAC thermostats, residential irrigation systems that access weather forecasts over the Internet and, incredibly, even a TV remote control that automatically updates your Facebook web page as you change channels.

With broadband Internet subscriptions expected to reach 77% of US households by 2012 [17] and the continued integration of home automation and internet technologies, the home will increasingly take on characteristics of the larger Internet. Besides remote access capability, homes on the Internet will be inevitably searched by services, such as Google, and some portions of their internal state made available on the web. And, it will become common for household systems to automatically and transparently access Internet-based web services for everything from weather and environmental forecasts to energy costs for water, gas and electricity.

5.1.3. Juncture of Home Automation and Smart Grid

As this scenario evolves, the interaction and overlap between the Smart Grid and the internet will become more important – and potentially – problematic. For example, it is anticipated that one of the critical services provided by the Smart Grid will be the delivery of energy price rates and individual, real-time consumption measurements to many, if not all, households on the grid. With increasing levels of automation within the home, this information will need to be distributed to various devices throughout the house. As the house becomes increasingly Internet-like, it would be convenient for this information to flow into the home over the public Internet. However, if this information originates on a private network partitioned from other home automation devices, some mechanism (e.g. gateway) will be needed to bridge between private and household networks or route information out to the public Internet and back into the home.

There are several approaches to resolving this data path problem. And the interaction and relationship between the Smart Grid and home automation systems is being actively addressed by working groups such as OpenHAN. The following figure gives a high-level architectural view, within the home, of the co-located Smart Grid and home automation network functions.

But, at this point in time, US homes are not a completely “blank slate” waiting to be automated. They are, to varying degrees, being instrumented and accessible as part of the

AMI-HAN Interface

Slide #1

April 21, 2007 rev 1.0

The diagram illustrates the AMI-HAN Interface, showing the connection between Utility Owned and Consumer Owned systems.

Utility Owned

- interval energy
- time
- billing start time
- peak power
- messages
- acknowledgments
- price signals
- reliability signals

Consumer Owned

Utility Owned (Left side):

- Private Fixed Networks WAN/LAN
- 2-way connection to Any interval meter or pole-top collector
- 2-way connection to Any gateway (protocol xfr)
- 2-way connection to Broadband TV, music

Consumer Owned (Right side):

- Any gateway (protocol xfr)
 - Special box
 - Internet modem
 - Router
 - Media PC
 - Security panel
 -
- Broadband TV, music
- 2-way connection to T24 PCT
- 2-way connection to Pool Pump
- 2-way connection to Appliances
- 2-way connection to Display Devices

Connections:

- Utility Owned to Consumer Owned: 2-way connection (Interval energy, time, billing start time, peak power, messages, acknowledgments, price signals, reliability signals)
- Utility Owned to Consumer Owned: 2-way connection (Any interval meter or pole-top collector)
- Utility Owned to Consumer Owned: 2-way connection (Any gateway (protocol xfr))
- Utility Owned to Consumer Owned: 2-way connection (Broadband TV, music)
- Consumer Owned to Utility Owned: 2-way connection (Any gateway (protocol xfr))
- Consumer Owned to Utility Owned: 2-way connection (Broadband TV, music)
- Consumer Owned to Utility Owned: 2-way connection (T24 PCT)
- Consumer Owned to Utility Owned: 2-way connection (Pool Pump)
- Consumer Owned to Utility Owned: 2-way connection (Appliances)
- Consumer Owned to Utility Owned: 2-way connection (Display Devices)

Protocols:

- HAN Protocols¹
 - Zigbee
 - Z-wave
 - Insteon
 - Wi-Fi
 - EIA/485
 - HomePlug
 - Bluetooth

Other Connections:

- Utility Owned to Consumer Owned: 1-way connection (T24 PCT)
- Utility Owned to Consumer Owned: 1-way connection (Pool Pump)
- Utility Owned to Consumer Owned: 1-way connection (Appliances)
- Utility Owned to Consumer Owned: 1-way connection (Display Devices)

Notes:

- e.g., 802.11b, proven mesh LAN protocol, etc.
- To be determined
- Up to 45 active protocols worldwide

Ron Hofmann

It is worth making one final point on end user expectations viz. interactions with the Smart Grid. The convenience and control proposition that has driven the development of the home automation market place has also helped crystallize user expectations of how well these household systems should perform. Control over household systems is expected to be effective and reliable with monitored information, such as electricity or water usage being accurately displayed. Recent anecdotal information gleaned from Smart Meter installations in California indicate that some Smart Grid applications have not yet achieve home automation performance expectations [19]. At least one individual was surprised to see their Internet-displayed power meter indicate substantial power consumption during the middle of an extended power outage. Although this behavior was traced to a programming artifact that attempted to animate web-based displays when real meter telemetry was unavailable, such behavior would be considered unacceptable from home automation systems. The successful introduction of Smart Grid capabilities into the end user environment should take into account existing technologies and the user expectations they have already established.

While some anecdotal examples may suggest that most consumers and end users are technically unskilled, the Internet market place argues otherwise. The evolution – and explosion – of the Internet has led to a greater understanding of data networking functions and related equipment installation issues on the part of the end user. A recent scan of major Internet ISP's revealed that all offered a do-it-yourself DSL installation program for new subscribers. The success of such programs indicates both a welcome standardization of data communications hardware and a growing competence and understanding on the part of the Internet community of basic IP protocol concepts - further evidence of a growing technical competence within the Internet community and public at large. It also indicates that vendors have invested heavily in standardizing a potentially complex network installation operation and have educated a large portion of the entire user community to perform it correctly.

6. CONCLUSION

By looking at the present user expectations within the Internet ecosphere, we can observe some qualities that are valued and, with some modification, attempt to integrate them into Smart Grid transactions. As the domains of the

Internet, home automation and the Smart Grid meet in the home, issues such as data transparency, platform independence, and ubiquitous local and remote access will be highly valued – and expected. And, above all, the continual evolution, openness and innovation experienced within Internet world will be expected for end user applications that interact with the Smart Grid.

6. ACKNOWLEDGEMENTS

This work was sponsored in part by the Demand Response Research Center which is funded by the California Energy Commission (Energy Commission), Public Interest Energy Research (PIER) Program, under Work for Others Contract No.150-99-003, Am #1 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

- [1] “Internetworking Technology Handbook”, <http://www.cisco.com/en/US/docs/internetworking/technology/handbook>
- [2]Holenstein, Bruce. *Breaking the Availability Barrier II: Achieving Century Uptimes with Active/Active Systems*. p. 213
- [3]Shaikh ,A and Lenz, K. “Where’s the Search? Re-examining User Expectations of Web Objects”, http://www.websitecriteria.com/Website_research/Userexpectations_online.html
- [4] “Top 20 Countries with the Highest Number of Internet Users”, <http://www.internetworldstats.com/top20.htm>
- [5] “How we got from 1 to 162 million websites on the Internet”, <http://royal.pingdom.com/2008/04/04/how-we-got-from-1-to-162-million-websites-on-the-internet/>
- [6]NIST, “Smart Grid Priority Action Plan Guidelines for the Use of IP Protocol Suite for Smart Grid”, http://www.nist.gov/smartgrid/paps/1-Guidelines_for_Use_of_IP_Protocol.pdf
- [7]”Advanced Control Methods: Appendix B4-A systems View of the Modern Grid”, National Energy Technology Laboratory, March 2007
- [8] <http://www.gridwise.org/>
- [9]”Intermediate Energy Infobook”, p. 46, The Need Project, 2008, http://www.need.org/needpdf/infobook_activities/IntInfo/ConsI.pdf
- [10]Staff, ”Wiser Wires”, The Economist, Oct. 8, 2009
- [11]Itron, “How AMI Enables the Smart Grid”, <http://www.nextgenpe.com/article/How-AMI-Enables-the-Smart-Grid/>
- [12]Bureau of the Census, “Projections of the Numbers of Households and Families in the United States:1995 to 2010”, US Dept. of Commerce
- [13] Piette, M.A., et. al., “CEC OpenADR-Version 1.0 Report”, <http://drrc.lbl.gov/openadr/pdf/cec-500-2009-063.pdf>

- [14]Siegel, Jon. *Corba, Fundamentals and Programming*, John Wiley and Sons
- [15] J. Schwartz, “The End of Middleware,” *Java Developer’s J.*, vol. 9, no. 2, 2004, pp. 16–18.
- [16]Knott, J. “Number of Home Automation Systems to Hit 4 Million in 2013”, http://www.cepro.com/article/number_of_home_automation_systems_to_hit_4_million_in_2013/
- [17]Reed, B. “Broadband to reach 77% of US households by 2012: Gartner Group”, <http://www.networkworld.com/news/2008/072408-broadband-penetration-gartner-study.html>
- [18]Kuniavsky, Mike. *Observing the User Experience: A Practitioner’s Guide to User Research*, Morgan Kaufmann, 2003
- [19]Baker, D. “Customers say new meters not always smart”, San Francisco Chronicle, Oct. 18, 2009

Biography

Charles McParland is a Computer Scientist at Lawrence Berkeley National Laboratory. He has been involved in the design and deployment of large-scale distributed data acquisition and remote network communications systems for more than 25 yrs. Recent research activities have centered on wireless (802.15.4) protocols and their integration into both residential energy management and demand response control systems. He has been a member of the IEEE for the past 12 years and has participated in both instrumentation and communications standards efforts