LBNL-6013E

ERNEST ORLANDO LAWRENCE
BERKELEY NATIONAL LABORATORY

HAN Attack Surface and the Open Smart Energy Gateway Project

Justin Searle

UtiliSec

Chuck McParland

Lawrence Berkeley National Laboratory

May 2013

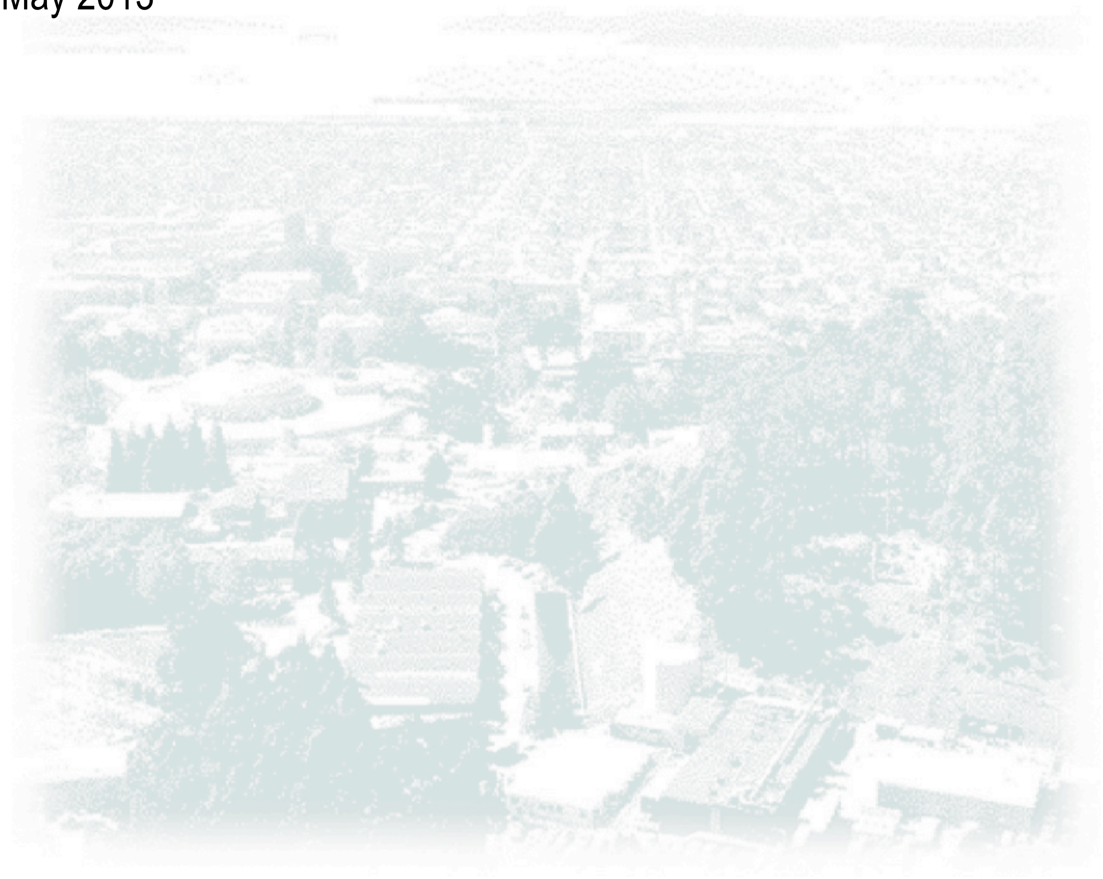

Disclaimer

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.

BERKELEY LAB
Lawrence Berkeley
National Laboratory

HAN Attack Surface and the Open Smart Energy Gateway Project

Justin Searle, UtiliSec
Chuck McParland, Lawrence Berkeley National Laboratory

Smart Meter Security (AMI and HAN)

The cost of deploying smart meters throughout many of California's utility service areas has been justified by a combination of benefits to both utilities and consumers. Utilities would receive operational benefits from the use of modern Smart Meter communications capabilities (i.e. Advanced Metering Infrastructure – or AMI) for both automated meter reading and enhanced monitoring of the power distribution grid. Consumers would benefit from newly available services that would allow near real-time readout of energy usage – both power and price – and enable, through ubiquitous Demand Response (DR) signaling, cost-saving automatic responses to changing energy price conditions. At this point in time, some of the utility goals related to the "back end" or AMI communications systems have been achieved. However, many of the benefits promised to consumers, such as enhanced control over their energy consumption and related bills, have yet to materialize.

Although the installed systems are technically capable of utility-to-residence communications, California utilities have not yet enabled smart meter communications into the home. The reluctance on the part of utilities to enable wireless communication between smart meters and residential devices (e.g. thermostats, energy displays, etc.) has been the primary factor in limiting the availability of these new consumer services. While some of this reluctance has been based on technical shortcomings of the currently selected communications technology (ZigBee PRO and ZigBee SEP 1.0), the overarching issue has been concern about the level of security provided by this particular set of network and application-level protocols, Utilities remain uncertain about the ultimate, system-wide risk entailed by allowing customers to directly interact, via a wireless network, with their smart meters. As a result, the proposed consumer benefits that depend on such communications have not been achieved.

For obvious and justifiable reasons, utilities consider the security details of their power distribution and metering operations to be highly sensitive. As such, they are rarely discussed in the absence of stringent non-disclosure agreements. These same concerns surround the security aspects of the recent AMI/Smart Meter upgrades. Consequently, public discussions of the status and schedule of smart meter installation programs have been unable to address – or even describe – the security concerns surrounding utility reluctance to enable smart meter communications into the home. This paper attempts to provide, on a non-proprietary basis, insight into the major design elements of both smart meters and the larger AMI systems within which they operate and, within this context, enumerate known security issues. It discusses the particular architectural areas that might be vulnerable to cyber attacks (i.e. "hacking") and sheds some light on the potential risks of such attacks at both the local smart meter level and also at the wide-area distribution grid level. And lastly, it describes and evaluates a technique for reducing the overall level of risk accompanying wide-scale smart meter communications into the home – *utilizing presently installed and operational meter platforms and software*. It is our hope that an increased understanding of both system components and the larger architecture within which they operate will promote a more meaningful – and open - understanding of the security risks incurred by enabling smart meter communication into the home.

The Open Smart Energy Gateway (OpenSEG) Project

The goals of the Open Smart Energy Gateway project (OpenSEG)[1] are two-fold. Overall, the project sought to engage consumers in the smart meter initiative by promoting the flow and display of energy consumption data from recently deployed smart meters into the residential environment. And, once demonstrated, measure the acceptance and usefulness of displaying smart meter near real-time energy consumption in the home. While the logistics of selecting display devices and enabling smart meter HAN communications were initially considered a minor part of the project, utility reticence to allowing such communications quickly became a major impediment to achieving its goals. Although several California utilities were running small pilot programs involving smart meter HAN communications in residential settings, the clear sense across all utilities was to forgo wide-scale use of smart meter HAN capabilities until a revised set of protocols was available. There is general agreement that the existing set of protocols will be re-architected to achieve a strict layering between the network and application portions of the communications software "stack". The existing Zigbee PRO network protocol will be replaced by an industry standard IP stack and the Zigbee SEP 1.0 application protocol will be totally re-written and become SEP 2.0. Although the revised version of SEP was initially expected to be specified, tested and generally available by mid 2011, its progress was hampered by both design disagreements and by implementation issues resulting from minimal computing resources on the current generation of deployed meter platforms. And, while the specification has recently been released (Aug. 2012) for public review and comment, the timeline for its appearance in currently deployed Smart Meters remains unclear.

It should also be pointed out that, while the SEP revision process has been moving forward, both vendors and utilities realized that additional functionality and specification clarity was required in the currently deployed SEP 1.0 application-level protocol. As a result, the original SEP 1.0 specification was revised to include a new SEP Cluster to support "over the air" HAN device firmware upgrades and portions of the testing and certification plan were improved. This revision, SEP 1.1, is backward-compatible with the original SEP 1.0 and no changes were made to any of the underlying security design or implementation. In fact, it is not unusual to see any pre 2.0 SEP software referred to as "SEP 1.x" in technical literature. It is ironic that, given the insistence that HAN communications be delayed until the availability of SEP 2.0, some utilities are currently upgrading their existing SEP 1.0 stacks to SEP 1.1. Given the fact that, from a security perspective, both SEP 1.0 and 1.1 behave identically, they will be used interchangeably in the remainder of this paper.

Within this environment, it became clear that utilities would only entertain wide-scale enablement of Smart Meter HAN communications if the system-level risks were better understood and, to the extent possible, minimized. Given the uncertain delays in realizing the "ultimate" resolution of utility HAN security problem, the OpenSEG project designed and proposed an interim solution that, while allowing

[1] Although originally referred to as the "REDS gateway", this design is now known as the "Open Smart Energy Gateway" or OpenSEG.

the dissemination of near real-time power consumption data into the home, constrained the remaining system features in a way that significantly reduced the real and perceived risk involved in enabling smart meter HAN radio communication. While OpenSEG has received favorable responses from both utilities and vendors, it became clear that a more formal security review of both its capabilities and the current status of smart meter HAN security was needed. In addition to promoting a more open and fruitful discussion of smart meter HAN security, this report is specifically intended to evaluate the effectiveness of the OpenSEG design in addressing known security issues present in the utility HAN and the wider smart meter/AMI environment. A more detailed discussion of OpenSEG specifics will follow the general discussion of HAN security concerns that appears below.

What Can Attackers Attack?

Any computer system that accepts and interprets external inputs has the potential to be attacked. By sending inputs that a system does not expect or that are intentionally malformed or malicious in nature, an attacker can attempt to push a computer system into a state that the developer or engineer may not have anticipated and, thus, properly planned for. Such inputs can affect target systems in a variety of ways from decreased or unstable performance to complete system "lock up" or operating system crashes. In extreme cases, which unfortunately are not uncommon in today's complex computer systems, an attacker can gain some degree of control over the target system and have it perform actions on its behalf. The program elements capable of accepting input messages – messages that attackers can leverage in their attempt to compromise a system - are broadly considered the "attack surface" of a system. Without any inputs, there is little to no potential for attack. While all program elements that process inputs do not present the same degree of risk, a useful rule of thumb is that the greater number of inputs, the greater the attack surface of any given system.

When we consider the attack surface for Home Area Networks (HANs), there are several different levels we can discuss. At the highest level we can consider each HAN device as a possible input point that an attacker can attack. Consider the following diagram:

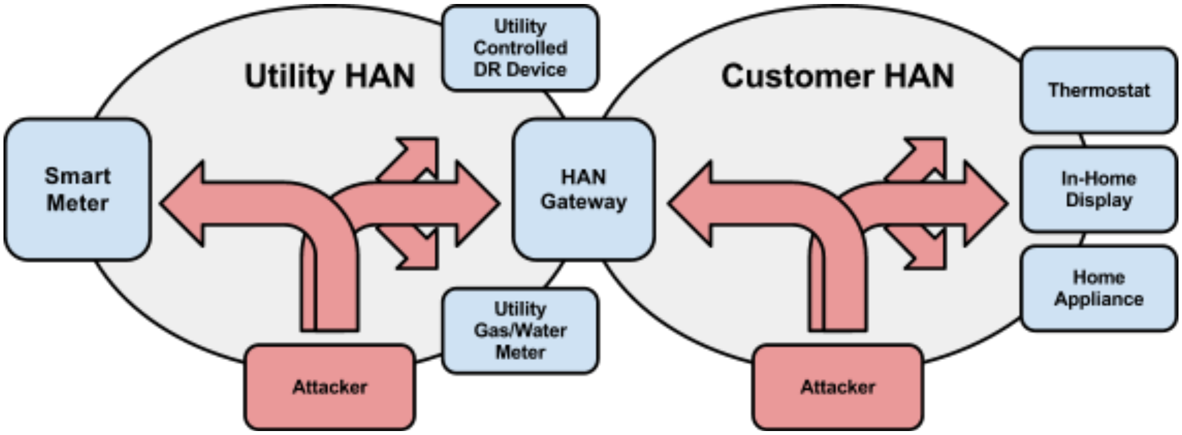

Figure 1: High Level Attack Surface

This diagram depicts a superset of common architectures used in the United States where a electric utility have deployed a smart meter to service a customer's home. It also depicts possible attacker entry paths for one or both networks. In some cases, deployed systems match the left side of the diagram with no "HAN gateway" present in the network. Other deployments include multiple networks with some form of intervening gateway.

The smart meter plays a central role in all of these systems. If properly provisioned by the electric utility, the smart meter has the capability of forming a HAN network that permits communication between the electric utility and devices deployed in the customer's home, such as thermostats or energy usage displays. This "Utility HAN" is most often a ZigBee PRO network using the ZigBee Smart Energy Profile (SEP 1.0) application-level protocol to communicate between application software running in multiple HAN devices. This ZigBee/SEP network is controlled by the smart meter, which acts as the coordinator for the network. In order for a particular device to join this utility HAN, it must be individually registered with the utility and, through a technique known as "white listing", have its identification information forwarded to a specific customer's smart meter over the utility AMI network. It should be pointed out that the process of interacting with the utility, authenticating yourself and exchanging device-specific information varies from utility to utility and has its own set of security risks and requirements. If this "out of band" exchange between utilities and customers is somehow compromised, it would be possible to instruct customer smart meters to allow foreign HAN devices, unknown to the customer, to join the HAN and access the smart meter over the utility HAN. While such compromises do not exploit any security weakness of the Zigbee PRO/SEP 1.0 stack and are outside the scope of this paper, these concerns should be included in any assessment of overall system-level security.

Once utilities have enabled the HAN radio interface on the smart meter, they can maintain control over what devices can join a specific meter's HAN through the registration and white listing mechanism described above. In practice, utilities intentionally limit network participation to specific device types that have been tested and approved by their own staff, such as those deployed as part of a Demand Response (DR) program. The limitation of devices which can join the utility HAN is an effective means of limiting the attack surface in the utility HAN. The fewer HAN devices they choose to communicate with, the smaller the chance one of those devices will have been compromised by attackers and have unforeseen consequences for the operation of their network.

While limiting which devices are permitted to join the utility HAN provides some benefit to the electric utility through a decrease in risk, this limitation is not always beneficial for utility customers. Customers may want to purchase and use HAN devices not yet tested by utilities – devices which require data from their smart meters. For this reason many vendors have created HAN gateway devices to allow a second, customer controlled HAN to be created. This second HAN, often called the "Customer HAN", allows customers to make their own decisions on which HAN devices to permit on their network. This solution allows customers and utilities to each control their own set of HAN devices on the customer premises, allowing for limited information sharing through the intermediate HAN gateway. Information transmitted between the utility and customer HAN, though limited, typically involves sharing energy usage data between smart meters and customer HAN devices.

In practice, attacks that can be launched at the utility HAN can also be directed at the customer HAN. Attackers can attempt to discover and exploit vulnerabilities in any of the connected HAN devices, regardless of which HAN network they belong to. However by dividing the HAN devices into two separate networks, utilities can effectively create an additional defensive barrier between the HAN devices they control and the HAN devices the customer controls. This defensive barrier is created and controlled by the HAN Gateway which handles most of the customer HAN traffic and only interacts with the utility HAN in a well-defined and controlled manner. The more traffic this HAN gateway can limit, the greater its ability to defend against attacks from compromised devices in the customer HAN. The OpenSEG design, which will be discussed in greater detail below, is based on these ideas and derives its enhanced security characteristics from them.

The protection provided by the HAN gateway extends in both directions. From a utility's perspective, compromised customer HAN devices have very limited or no control over utility HAN traffic, such as messages sent from the gateway to the smart meter and beyond into the utility data center. Likewise, compromised utility HAN devices, such as energy displays and DR devices, have no control over what information they can pass to the customer HAN network, as the HAN gateway controls what information is forwarded and retrieved from the smart meter. A properly configured HAN gateway can effectively limit the attack surface between devices each HAN network.

Of course the greatest security risk to both the electric utility and the customer is the potential to compromise the network ZigBee coordinator and its SEP trust center. The SEP trust center manages the security environment for devices participating in the HAN and is considered the primary mechanism for supporting security on the HAN. Attackers deem these areas to be high value targets. For the customer HAN, the HAN gateway assumes both of these roles. A compromised HAN gateway can be used to attack any HAN device on the customer or the utility HAN networks. However of greater concern are the attacks from the HAN gateway towards the devices installed in the customer HAN. Since the HAN Gateway is acting as the SEP trust center for the customer HAN, it contains the security keys for communication on the customer HAN, and is inherently trusted by connected customer HAN devices. Attackers can also leverage a compromised HAN gateway to manipulate information from the utility smart meter network being forwarded to customer HAN devices, such as falsifying energy consumption information, sending malicious messages to in-home-displays, and even send spurious demand response and pricing signals to customer HAN devices causing potential harm such as rapidly cycle between no-power, low-power, and high-power consumption states. While a compromised HAN gateway can create all sorts of trouble for a consumer HAN, luckily its risk to electric utilities is roughly the same as the compromise of any other device on the utility HAN.

For the utility HAN, the smart meter becomes the high value target since it plays the role of the ZigBee coordinator and the SEP trust center. A compromised smart meter can be used not only to attack the utility HAN devices as described above, but it can potentially attack the electric utility's smart meter network and the servers which accept data from the smart meter network in the datacenter. While potential attacks on this upstream utility infrastructure is a grave concern, there is only a moderate increase in attack surface between a smart meter with its HAN interface enabled and one with its HAN interface disabled. The simple presence of a smart meter that uses bi-directional communication with the electric utility's AMI infrastructure creates a potential attack surface on upstream utility infrastructure. All inputs exposed to the smart meter by the upstream infrastructure makes attacks on the utility infrastructure potentially possible, even if those meters do not contain external interfaces such as HAN interfaces and administrative optical ports.

ZigBee Smart Energy Profile: How are networks formed?

The Zigbee protocol, or more accurately, the Zigbee PRO protocol, is a wireless protocol layered on top of the IEEE 802.15.4 RF signaling layer standard. It implements a partially co-mingled network and application layer on top of this 802.15.4 signal. Since wireless networks are based on a shared media (i.e. RF), the Zigbee network layer needs to differentiate its messages from other 802.15.4-complient networks using the same RF channels (e.g. proprietary protocols), it must identify messages that are considered to be a part of its network. It does this by creating its own “logical” network identified by a unique Personal Area Network Identifier (PAN ID) which is used as part of all messages transmitted within that logical network. It is the job of the ZigBee network coordinator to select an appropriate PAN ID and advertise it to potential ZigBee network members. For clarification, note that, in the example above, the utility and customer HANs are, in fact, different Zigbee networks (w/different PAN IDs) each managed by their own coordinator – the smart meter and HAN gateway respectively. Prior to petitioning a coordinator to join a particular network (i.e. a PAN), a ZigBee device is capable of having limited communication with other Zigbee devices which may already be members of co-located Zigbee network. This is accomplished by using the ZigBee Inter-PAN messaging service that permits a device to offer selected services to ZigBee devices outside their network. The role of such ad hoc communication is poorly defined and therefore an area of some security concern.

However, since smart meter SEP services can only be requested by other SEP-authenticated nodes, HAN devices must execute several additional authentication steps in order to become fully capable SEP HAN devices. The device must first join a ZigBee network, which requires authentication using a network key. The ZigBee device can then communicate with all devices in the ZigBee network and interact with the services those devices offer. It should be noted that there are serious concerns about the level of encryption used in network key exchange during the network join operation. However, it has been argued that HAN devices that have joined a ZigBee network are still required to successfully complete an additional, more rigorous certificate-based authentication procedure before being able to interact with SEP services on other devices. All utility HAN devices perform this second authentication procedure by exchanging certificate-derived information between themselves and the SEP trust center – located, by common practice, in the smart meter. The successful completion of this process results in the generation of authenticated SEP keys that can be used for encrypted SEP data exchange between HAN devices.

Smart Energy Profile’s Attack Surface

When evaluating the attack surface created by ZigBee PRO and Smart Energy Profile (SEP), we should consider the new inputs created by these interfaces and protocols. The following diagram maps out the various levels in which these inputs exist in the ZigBee/SEP stack.

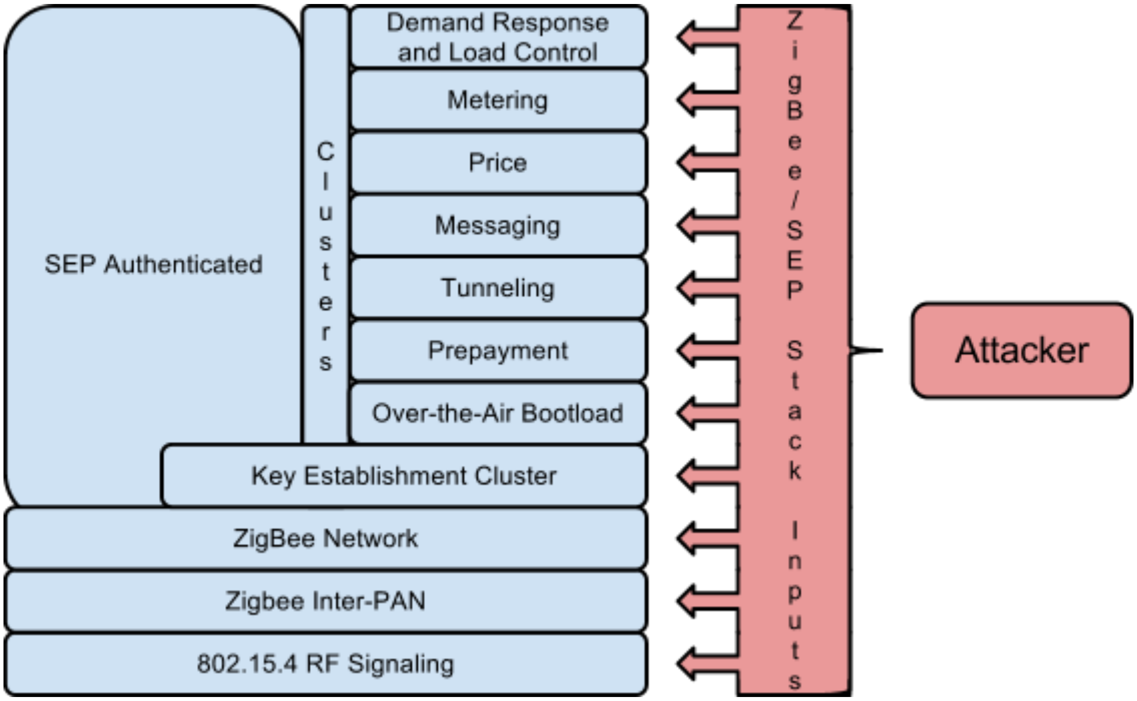

Figure 2: ZigBee/SEP Attack Surface

Upon establishment of an authenticated SEP key, the SEP enabled device can access any of the offered SEP Clusters on any of the participating HAN devices (see diagram above). Clusters can be thought of as the software components that provide specific services (e.g. meter reading) on a given HAN device. Or, when viewed from the perspective of the ZigBee network "stack", they can be thought of as distinct "applications" present on a HAN device to which incoming messages can be directed for processing. Each Cluster has its own set of functions such as retrieving energy consumption data in the Metering Cluster or acknowledging demand response signals in the Demand Response and Load Control Cluster. Each of these layers, though the operation of the intervening ZigBee network stacks, accepts various inputs from other ZigBee devices. Thus each input of each Cluster represents a potential attack surface where, if not handled properly, malformed incoming network messages could cause unanticipated and unpredictable results within the HAN device.

While many people do not consider the 802.15.4 MAC (Media Access and Control) layer as a potential attack surface, they are wrong. At the most basic level, this is a most readily available layer in which to perform denial of service (DNS) attacks which interrupt RF communications between HAN devices. While much of the HAN security discussion above has focused on the potential misinterpretation of intentionally malformed messages by HAN device SEP Clusters, a denial of service attack can, in the proper circumstances, be extremely disruptive. Although less targeted and specific than the message-based attacks described above, the complete loss of communications between a group of coordinated HAN devices could result in damage to physical devices being controlled. Furthermore, intermittent denial of service attacks can dramatically effect network stability and induce erratic communications latencies that can have unpredictable effects on system-level behavior and performance – particularly when individual Cluster services attempt to recover from communications outages. In the case of spoofing attacks, which target higher levels of device behavior, attackers will need to compromise this layer to manipulate device MAC addresses to match that of the ZigBee device being imitating. Any weaknesses or implementation shortcomings at the MAC layer of the network stack will potentially facilitate such attacks. And finally, buffer overflow-like attacks can effectively target this layer if the 802.15.4 stack does not properly handle the MAC header inputs.

The ability to send and receive Inter-PAN messages is a required component for ZigBee/SEP-certified devices, and was introduced to ZigBee because of a requirement in the Smart Energy Profile. Although these Inter-PAN services are not commonly exposed, they can, when enabled, allow ZigBee devices to communicate in an anonymous, insecure manner, without the need to join the ZigBee network controlled by the coordinator. This becomes a concern for security, not only for data and functionally exposed in these Inter-PAN services, but also because a buffer overflow exploits here can cause the same damage as any other input in the ZigBee/SEP stack – without the effort needed to achieve any degree of successful authentication. The existence of Inter-PAN communications services is not widely understood and, as a result, its exploitation, albeit difficult and unlikely, will probably go undetected.

Once a ZigBee device has joined a particular PAN, it can attempt to access whatever ZigBee profiles and services that are being offered within the network, or in the case of SEP, access to the secondary SEP authentication process. To gain access to the PAN, an attacker must first gain access to the ZigBee network key. Occasionally these network keys can be discovered through clear text transfers between ZigBee devices during a network join operation. In SEP enabled networks, where additional certificate-based authentication processes are required, the join process typically offers more effective protection against the risk of network key exposure. However a poorly implemented or non-standard initialization process can still expose the network key. Once the key has been obtained, attackers cannot only communicate with non-SEP services, but also can access second layer SEP authentication processes. Thus, they can attempt to identify and exploit SEP authentication vulnerabilities, including potential buffer overflow issues in the authentication inputs themselves.

The secondary SEP authentication process is officially called the Certification Based Key Establishment (CBKE) process. This is a strong authentication mechanism leveraging ECC asymmetric Elliptic Curve Cryptography (ECC). Successful authentication here provides access to all the SEP Cluster functionality. However obtaining this CBKE authentication requires a valid ECC certificate that has been cryptographically signed by a trusted source. While these ECC certificates are not possible for an attacker to create, they are possible for an attacker to steal from a valid and permitted ZigBee device. By compromising a permitted ZigBee device or by stealing a ZigBee device’s ECC certificates and cloning its IDs, an attacker can successfully complete this CBKE authentication.

If an attacker can successfully obtain access to the SEP Clusters, they will have access to the greatest number of services and functionality in the HAN network. All layers below this point are usually implemented directly by a small handful of ZigBee chip manufacturers who typically provide working and tested ZigBee network stack software. However, the inputs being accepted in the SEP Cluster functions are being passed to program elements written by the exponentially larger number device manufactures. Where the chip manufacturer’s code is being seen and vetted by numerous manufactures, the device manufacturer’s code is usually not seen nor vetted beyond the small team of developers who have written and presumably tested the code. Furthermore, within the ZigBee SEP software environment, this is the area that contains and potentially exposes the greatest number of inputs (i.e. attack surfaces). These conditions provide fertile ground for attacks.

Through simple request and response interactions with the appropriate SEP Cluster after CBKE authentication has been successfully completed, the attacker can interact and potentially manipulate functionality in the connected HAN devices. Every function in every SEP Cluster accepts inputs to accomplish its task. By fuzzing messages (i.e. methodically creating malformed messages) and directing them to these inputs, an attacker can potentially find buffer overflow vulnerabilities and gain remote code execution and system level control of the target device. It is at these SEP cluster layers where most of a device's custom code is written, thus providing the most inputs to be vulnerable to buffer overflow or malformed message attacks. Between the manipulation of the intended functionality in the SEP Clusters and the increased potential in finding buffer overflows in custom code, the SEP Cluster functionality is the layer where most of the damage can be done.

While the inputs at each of the various layers in the ZigBee/SEP stack adds to the attack surface, the majority of these inputs are usually handled by the ZigBee chip and not by the main microprocessor in the ZigBee device. Attackers that exploit buffer overflows vulnerabilities in those inputs processed entirely in the ZigBee chip may gain greater access to the RF traffic but cannot normally extend that control to the main system functionality of the ZigBee device. However, a small number of those inputs are pushed beyond the ZigBee chip and sometimes even beyond the ZigBee device itself.

Attacks Beyond the Immediate Target

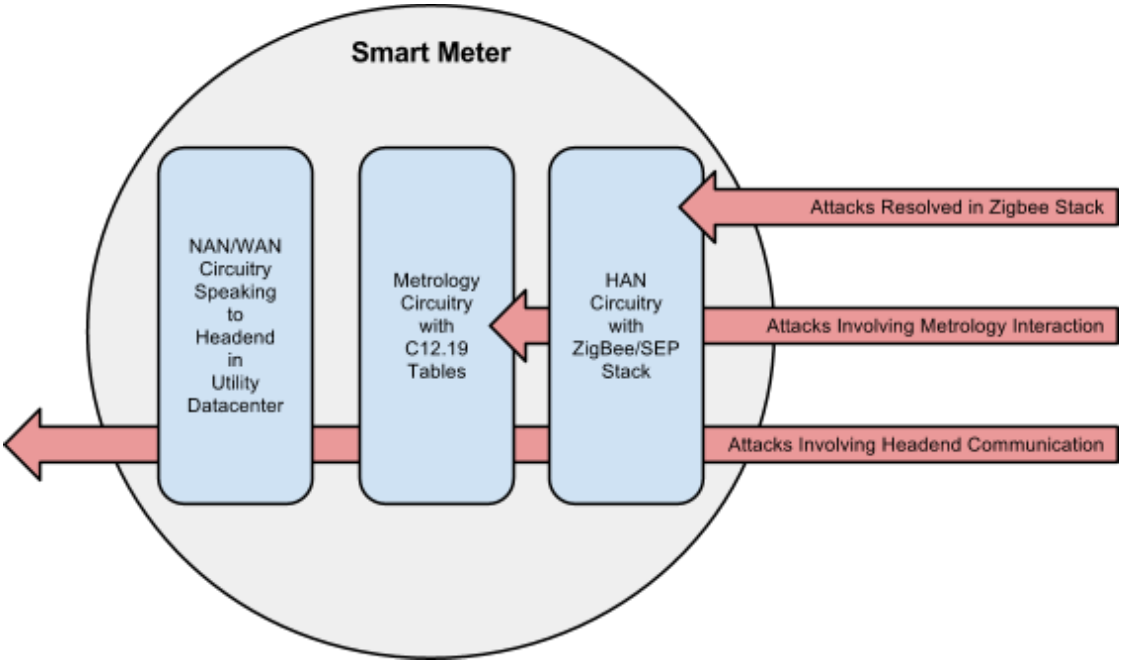
Smart Meter
NAN/WAN Circuitry Speaking to Headend in Utility Datacenter
Metrology Circuitry with C12.19 Tables
HAN Circuitry with ZigBee/SEP Stack
Attacks Resolved in Zigbee Stack
Attacks Involving Metrology Interaction
Attacks Involving Headend Communication

Figure 3: Meter Components and SEP Clusters

The above diagram shows the relationship between the outward-facing SEP network interface and other internal components within the smart meter itself. The advent of smart meters has allowed meter manufacturers to include a host of new features into the previously "simple" electric meter form factor. Implementation details for these new meter designs – from improved metrology to AMI network interfaces – remains vendor proprietary and can, within this study, only be treated in the abstract.

However, there are high-level similarities between smart meter designs that are based on current "best practices" and the requirement that they provide metrology data in a standardized format. So, to a limited extent, we can discuss known internal smart meter attack surfaces beyond those associated with SEP Cluster inputs.

While an attacker may be able to deliver SEP messages to one or more HAN device SEP Clusters, the level of damage they can inflict on both the meter and upstream AMI components remains an open question. The smart meter attack surface exposed to the home is based on the ZigBee SEP protocol. Although the SEP messages, as processed by the various Clusters within the HAN device, ultimately interact with a number of critical smart meter components, the contents of these messages, as well as any implied commands they may convey, are abstractly expressed within the constraints of the SEP protocol. In other words, SEP messages are interpreted by meter vendor firmware, they do not directly act on any internal smart meter components. Without detailed knowledge of actual smart meter hardware and firmware implementation, absolute statements about what is possible – or impossible – cannot be made. However, in any security analysis of smart meter architectures and their interaction with external SEP HAN devices, this area would be considered the primary attack vector and, as a result, be implemented with some care.

The problem of buffer overflows, while possibly addressed within the smart meter MAC layer, presents a less certain situation. Since successful buffer overflow attacks can allow an attacker to execute unanticipated portions of the firmware, some of the care taken in safely interpreting SEP messages may be effectively bypassed. Having said this, execution of a successful, targeted attack through this mechanism would require extensive knowledge of the meter's code design and implementation – which will be difficult to obtain. Furthermore, in many of the current smart meter implementations, the entire SEP network stack and initial message interpretation software are executed in a separate computational environment – a single-chip device that shares none of the basic smart meter computing resources (e.g. memory). As a result, attempts to subvert the SEP Cluster inputs, even if partially successful, remain contained within the SEP subsystem and are not propagated to other smart meter resources. Having said this, it must be stated that, without detailed knowledge of both smart meter hardware design and its software implementation, categorical statements about the potential for a successful attack are difficult to articulate.

One particular area of smart meter design deserves mention – that of the "C 12.19" tables. The ANSI C 12.19 standard describes the table structure used for conveying meter information to other devices, be they handheld meter readers or utility AMI communications systems. These tables contain, among others, stored meter metrology values such as instantaneous energy and aggregated power consumed. Given the ubiquitous requirement to provide metrology data in the order and format proscribed by this standard, most meters implement some form of these tables internally. For example, a request sent to the smart meter Simple Metering Cluster for the aggregated power value consumed would, after suitable interpretation by meter firmware, result in accessing the appropriate internal C 12.19 table and returning the stored value. So, SEP Cluster requests can, in some cases, cause meter firmware to access these tables – presumably in a safe and well-protected manner.

These C 12.19 tables can also be extended in manufacturer-specific ways. And, given their required presence for meter data transfer purposes, have proven to be a convenient "organizational" element used for storing a variety of general and proprietary meter data. In some smart meter designs, these tables may contain, in manufacturer-specific locations, critical meter configuration information like the status and control of the smart meter's power disconnect switch or, perhaps, whether the HAN interface is enabled and powered. A reasonable attack methodology would explore the feasibility of accessing areas of the C 12.19 table that control such critical resources. While these sorts of attacks also rise to the top of any reasonable security analysis and would, therefore, be given great attention by vendors and utilities, little is known about the current level of protection in this area and what measures have been implemented to prohibit access to critical configuration elements.

Perhaps the greatest risk to an electric utility when enabling HAN interfaces on an AMI meter lies in the data passed from the HAN, through the meter, to servers located in the utility's datacenter. Often this data is passed to the AMI Head-end and then to whatever servers are managing device and data in the customer HANs such as demand response (DR) management servers and distributed energy resource (DER) management servers. While this is a very small attack surface due to the very limited number of SEP functions that pass data through meter firmware to these backend servers and is further limited by the short data field lengths these inputs permit, this is a real threat and one of the few ways an attacker could gain wider access to the utility's infrastructure via the HAN interface. If an attacker were able to find a SEP function that passed inputs directly to the backend servers and was able to discover a buffer overflow or similar vulnerability in how that backend server handles those inputs, an attacker could feasibly pass place a malicious payload into those inputs to exploit that vulnerability and execute the code of his choice on the backend server.

Overall, by enabling a HAN interface on the smart meter, we are allowing some additional inputs into the smart meter and thus increasing a smart meter's attack surface and relative risk profile. However, the vast majority of these additional inputs (i.e. those exposed by the ZigBee/SEP interface on the smart meter) are resolved within the ZigBee/SEP stack and dedicated HAN hardware in the smart meter. Very few of these inputs are passed to other internal meter elements or beyond the meter into the utility data center. If we were to disable that small number of inputs going into the meter and the utility datacenter servers, we could greatly decrease the risk of attack on utility backend infrastructure created by enabling the HAN interface. While we can, at an abstract level, say that some connections between the SEP interface and critical smart meter resources do exist, the risk – if any - associated with these connections lies in their implementation and not with architectural design.

Goal of the OpenSEG HAN Gateway

As initially discussed, the specific goal of the OpenSEG design is to reduce security-related risks presently associated with the operation of the smart meter utility HAN and, thereby, encourage California utilities to enable these HANs and allow more direct consumer engagement with smart meter programs. Specifically, the OpenSEG interposes itself between the smart meter utility HAN and a residential WiFi/IP network. By providing the only available communications channel for interacting with the smart meter, it significantly constrains the types of interactions permitted with the smart meter. Early analysis, similar to that shown above, indicated that, by constraining the available smart meter feature set to only the minimal number of SEP Clusters needed (i.e. Simple Metering Cluster), thus removing all SEP inputs that get passed to the internal meter and those that get passed all the way to the utility datacenter, the overall level of risk would be reduced. And, while the attendant constraints on SEP functionality may not be acceptable in the long term, the application of such constraints will allow utilities to achieve some of the promised "meter to home" communications and promote positive consumer engagement with the smart meter concept.

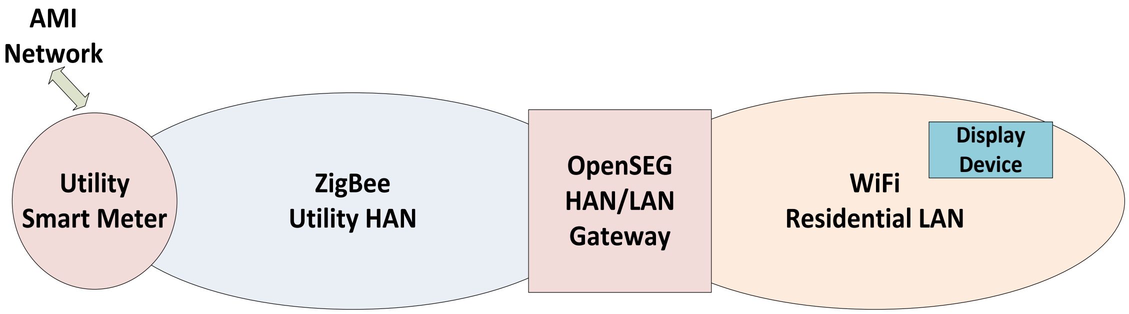

Figure 4: OpenSEG Architecture

As shown in this simplified diagram, the utility HAN is constrained to allow only a single additional HAN device, namely the OpenSEG device. This device acts as a gateway for smart meter information requests flowing to it from the residential local area network (LAN). This network is intentionally specified as a 802.11 compatible WiFi or wired IP network to further constrain the number of SEP HAN devices within a single residential smart meter network environment. The OpenSEG specification document is currently (Aug. 2012) being updated. For further information, see references. An analysis of the security aspects of this design follows.

The OpenSEG Security Environment

As mentioned above, the goal of the OpenSEG architecture is to minimize the system-level risk encountered when enabling residential smart meter communications on currently deployed SEP 1.0 smart meter platforms. This is accomplished by minimizing the number of existing smart meter SEP attack surfaces that can be trivially accessed and reducing the number of devices whose function needs to be fully trusted in order to assure the safety of utility assets located upstream of the smart meter. In light of the above analysis of the current smart meter security environment, specific implementation aspects of the OpenSEG architecture will be discussed and evaluated.

In terms of hardware, the OpenSEG design requires a modest, embedded computing platform with sufficient memory and computational resources to support both a ZigBee SEP 1.0 and 802.11 compliant WiFi network stack. Many suitable hardware designs are already available. Some of these designs incorporate ZigBee and/or WiFi “single chip solutions” as part of their communications interface design. As a result, multiple hardware platforms with well-tested, embedded ZigBee and WiFi network subsystems can serve as suitable OpenSEG development platforms. Since some of these hardware platforms have already been designed with the intention of being used as HAN devices, some level of attention has been paid to the requirements of physical security as well (e.g. minimal electrical exposure of security data passing between circuit board chips).

In addition to the required underlying hardware and communications interface platform, the OpenSEG architecture consists of three software elements. A tested, ZigBee Alliance-certified ZigBee SEP 1.0 (or greater) network stack (to obtain signals from the Zigbee radio embedded in the utility meter), a WiFi Alliance approved IEEE 802.11 family IP network stack (to communicate with the receiver typically found in consumer devices) and, a simple gateway application that, upon reception of an external WIFi request, presents a single SEP request to the smart meter and conveys the resulting response back to the original requestor. This design focuses on several security shortcomings discussed in the above utility HAN analysis.

First, it limits the number of utility HAN ZigBee devices requiring binding to the smart meter itself to one. The only device allowed, by prior arrangement with the utility's registration process, to be white listed within a customer's smart meter will be their OpenSEG device. While this does not eliminate the potential for physical or input-based compromise of the OpenSEG device by attackers, it does dramatically reduce the variety of devices to be tested and certified as safe against such attacks.

As discussed earlier, the largest number of attack surfaces found inside a ZigBee HAN device are those associated with SEP Cluster processing. Furthermore, the code implementing these services, by virtue of its being unique to each device, is extremely customized and, when compared to that found in the ZigBee stack itself, minimally tested. By limiting the number and the variety of HAN devices present on the smart meter network to one, we can both limit possible compromise attacks and promote more robust and thoroughly tested implementations for the small number of HAN gateway variants designed.

While we are clearly reducing the number of attack surfaces in the system, we are also introducing one or more new surfaces – namely those associated with the HAN gateway application that connects the ZigBee SEP and WiFi network stacks. This application will require careful design and testing. However, given its functional requirements, it can be designed and implemented in a transparent manner that simplifies both security analysis and functional testing. Given its function as an in-line processor for WiFi generated smart meter requests, this application will allow only a minimum number of internal states with no overlapping or parallel operations. Requests seeking access to non-supported SEP Clusters will be rejected by a small (essentially trivial in terms of memory requirements) message inspection filter. Therefore its correct behavior can be verified in a straightforward manner. Since the primary function of this application is to constrict the visible set of smart meter services and access patterns, its implementation should be, in practice, relatively straightforward. In fact, if its physical design incorporates the use of a vendor-based single-chip ZigBee communications interface, the SEP traffic can be constricted (i.e. filtered) in both the application and the ZigBee communications engine – thus providing two independent mechanisms to prevent unwanted access to smart meter SEP Clusters. Having said this, ample care will be needed to provide protection from buffer overflows and call frame manipulation. But, the scope of its behavior, when compared to that of the full set of SEP Cluster functions it is blocking, is relatively simple.

Next, the choice of 802.11 WiFi networking protocols for the residential HAN interface is based on several considerations. It can be argued that the maturity level of WiFi standards, with respect to security, exceeds that of the ZigBee SEP 1.0 protocol found in currently deployed smart meters. And, given the increasing numbers of WiFi enabled devices being sold into the marketplace, the level of awareness WiFi security issues as well as the required knowledge of configuration practices (e.g. WEP, WPA1-2, etc.) is far higher for WiFi than that for "yet to be mass marketed" ZigBee HAN devices. Therefore, at this point in time, WiFi protocols offer both a more secure network and a more supportable deployment environment within the home. And, in keeping with the goal of reducing the number of system-wide attack surfaces, reducing the number of SEP HAN devices and promoting increased use of WiFi-based devices clearly removes a potentially large number of attack surfaces for the resulting system.

It should be pointed out that this design does not directly address problems associated with possible network key capture or InterPAN access. However, in both cases, by reducing the number of HAN devices, the attack "cross section" is minimized by virtue of having greater control of which utility HAN devices are allowed in the network. In the case of network key capture, since capture is only possible during the period when a device joins an existing network, limiting the configuration to a single device dramatically limits the time interval when key capture is possible. Furthermore, since some SEP implementations provide increased network key security when joining a Zigbee network, we can, by careful selection of the Zigbee chip manufacturer, provide increased protection in this area. Similarly, in considering the potential abuse of InterPAN services, it is anticipated that neither the existing smart meters nor the HAN gateway will respond to InterPAN requests, thus eliminating their effect on resulting system security. Since no additional HAN devices will be present, ZigBee InterPAN requests on the utility HAN will go unanswered and will have no effect.

And lastly, it should be kept in mind that, while utility risk analysis primarily looks at the potential for system threats and compromise from the AMI perspective, there is also concern over the utility's data custodial responsibilities with respect to disclosure of customer energy consumption data. In California, as in most states, utilities have well-defined custodial responsibilities with respect to customer data. Data held within the utility's "back office" environment must be kept private and only released to third parties by customer permission. Even under those circumstances, limitations on both data use and further dissemination remain, partially, a utility responsibility. At this point in time, with smart meter utility HAN radios currently disabled, the extent of a utility's responsibility for the protection of customer data privacy, when disseminated into the utility HAN, is not clear. However, it is likely that, since the smart meter is the utility-administered ZigBee coordinator of the utility HAN, utilities will retain some custodial and privacy responsibilities for customer data when accessed directly on its network. The insertion of a HAN gateway partitions the HAN into two separately administered networks (i.e. utility HAN and residential HAN). Thus, the gateway becomes a demarcation point where responsibility for data security clearly changes hands. Using different network protocols, administered by different parties on each network, reinforces the case that responsibility for data dissemination on each of these networks resides with different parties. Utilities clearly bear no responsibility for the security and privacy of residential WiFi networks. Therefore, the creation of a "meter data demarcation point", as provided by the OpenSEG architecture can reduce the scope of utility data custodial responsibilities thus limiting their liability and regulatory risk associated with distribution of customer data within the home.

Concluding Remarks

In the area of network security, particularly wireless network security, there are few absolutes. But there are choices that can be made to significantly limit vulnerabilities. One can only review the evolution of the security suites associated with WiFi networks to be reminded that network security is more of a "life-style choice" than an achievable goal. While the currently deployed SEP HAN protocol (SEP 1.0 and its derivatives) has known shortcomings, it is arguably suitable – and sufficiently secure – for some of its originally intended functions. Given the lengthy delays in specifying its replacement (SEP 2.0) and the expectation that, like all network designs, "it may be better but will not be perfect", it is prudent to utilize the existing, deployed SEP 1.0 infrastructure wherever feasible. The OpenSEG design, by significantly constraining the SEP 1.0 feature set, provides a relatively secure environment within which utility and consumer smart meter engagement can begin in earnest. The system-level risk analysis contained in this document, while making few concrete statements, indicates that this effort can move forward with acceptable risks. Given the commitment to provide smart meter communications into the home now, the risks outlined above can be further quantified by utilities and vendors with specific knowledge of particular smart meter implementations. We feel the time to move forward is now and the risks, while difficult to quantify exactly, can be minimized and, ultimately, deemed acceptable.

In summary:

- There are known security shortcomings in the currently deployed ZigBee PRO/SEP 1.x network stack. If properly exploited, these shortcomings could allow an unauthorized party (i.e. attacker) to deliver command messages to one or more of the devices on the utility HAN and request them to be processed. All devices on the utility HAN, including the smart meter itself, are vulnerable to these exploits.
- While many smart meter implementation details remain proprietary, we can, with reasonable accuracy, enumerate and describe the software elements within the smart meter that are most visible and therefore most vulnerable to receiving and processing such unauthorized command messages.
- In performing the above analysis, we have found that the potential number of security weaknesses found in a smart meter implementation is roughly proportional to the number of distinct services the design supports (i.e. attack surfaces). In the current utility deployment environment, meter designs support many more services than are currently used or planned for use in the near future.
- As part of the project, we have developed an Open Smart Energy Gateway (OpenSEG) design that reduces the number of smart meter services visible to residential and third party devices - thus substantially reducing the system-level security risk associated with enabling smart meter HAN communications. This reduction in risk is sufficient to allow California utilities to confidently enable smart meter wireless communications into the home and immediately give customers access to near real-time energy consumption information.

Acknowledgements

The work described in this report was coordinated by the Demand Response Research Center and funded by the California Energy Commission (CEC), Public Interest Energy Research (PIER) Program, under Work for Others Contract No. 500-03-026 and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Reference Material

American National Standard. 2009. “American National Standard for Utility Industry End Device Data Tables.” http://www.nema.org/Standards/ComplimentaryDocuments/ANSI-C1219-2008-contents-and-scope.pdf

“Published Specifications,” accessed on August 27, 2012. https://www.wi-fi.org/knowledge-center/published-specifications

“Residential Energy Display Survey (REDS) Pilot,” accessed on August 27, 2012. http://drrc.lbl.gov/news/residential-energy-display-survey-reds-pilot

ZigBee Alliance. “ZigBee SEP 1.0 – Profile Specification – 07536r15.” Available from http://www.Zigbee.org

ZigBee Alliance. “ZigBee SEP 1.1 – Profile Specification – 07536r16.” Available from http://www.Zigbee.org

ZigBee Alliance. “ZigBee SEP 2.0 Application Specification 0.9 (11-0167-32).” Available at http://www.Zigbee.org