Creating an Energy Intelligent Campus: Data Integration Challenges and Solutions at a Large Research Campus

Dylan Cutler, Stephen Frank, Michelle Slovensky, Michael Sheppy, Anya Petersen
National Renewable Energy Laboratory

ABSTRACT

Rich, well-organized building performance and energy consumption data enable a host of analytic capabilities for building owners and operators, from basic energy benchmarking to detailed fault detection and system optimization. Unfortunately, data integration for building control systems is challenging and costly in any setting. Large portfolios of buildings—campuses, cities, and corporate portfolios—experience these integration challenges most acutely. These large portfolios often have a wide array of control systems, including multiple vendors and nonstandard communication protocols. They typically have complex information technology (IT) networks and cybersecurity requirements and may integrate distributed energy resources into their infrastructure. Although the challenges are significant, the integration of control system data has the potential to provide proportionally greater value for these organizations through portfolio-scale analytics, comprehensive demand management, and asset performance visibility.

As a large research campus, the National Renewable Energy Laboratory (NREL) experiences significant data integration challenges. To meet them, NREL has developed an architecture for effective data collection, integration, and analysis, providing a comprehensive view of data integration based on functional layers. The architecture is being evaluated on the NREL campus through deployment of three pilot implementations.

Introduction

NREL, one of the U.S. Department of Energy's 17 national laboratories, develops and validates new clean-energy science, technologies, and practices for sustainable energy systems integration. NREL's sustainable campus and clean-energy research embody the living-laboratory concept, as one of the laboratory's goals is to demonstrate and deploy energy-efficient solutions. NREL considers energy informatics innovation an important path to fulfilling its mission to advance building energy infrastructure and grid modernization.

Many programs and projects at NREL require access to high-quality campus energy data, but the specific nature of data collection, storage, and analysis needs vary widely. To address individual needs, a patchwork of energy information systems emerged, often with significant duplication of functionality. In the final quarter of 2014, NREL embarked on the unification of these disparate systems into a single, well-defined Energy Management and Information System (EMIS), with a view toward more actionable and integrated energy management, demonstration, and replication.

An EMIS—as we define it—is a capability rather than an individual product. While this capability could be provided by a suite of different software/hardware products or by a single vendor, it is the attributes of an EMIS that are the focus of this paper. At NREL, the primary purpose of the new campus EMIS is to collect real-time data resources and support analytics to enhance operational awareness and decision-making with respect to site energy use. Enabling advanced, holistic control of campus systems is a secondary goal.

To develop a design and the requirements for a new EMIS, we conducted an analysis of use cases, benefits, and barriers for the laboratory. From this analysis, we defined a conceptual system architecture that organizes EMIS capability into functional layers and developed

requirements for each layer. We then selected three separate pilot implementations, installed each on the NREL campus, and evaluated them through a series of functional tests.

Benefits and Barriers

An effective EMIS delivers significant benefits to any facility. This is particularly true at NREL, where we have a diverse portfolio of facilities, significant renewables penetration, and a mission to reduce our energy consumption. Like the buildings industry in general, NREL has experienced significant barriers to successful EMIS deployment. In this section, we examine key benefits and barriers and how they relate to NREL's experience.

Benefits

Combined with appropriate analytic tools, the benefits that an EMIS can provide are wide ranging. They include: fault detection and diagnostics (FDD), advanced control techniques, improved facility operations, enhanced maintenance programs, portfolio management, and tenant engagement. Many case studies have demonstrated EMIS benefits, including commissioning, benchmarking, and FDD (Mills 2009, Granderson et al. 2013). In addition, the Federal Energy Management Program (FEMP) has created a comprehensive guide to achieving operational efficiency in facilities, including EMIS deployment recommendations (Sullivan et al. 2010).

NREL has identified several of these benefits as deployment goals for its own EMIS, which we describe below. Although each of the benefits may be delivered via a number of different products or solutions, the common thread is that they are unlocked or significantly enabled via access to integrated and well-organized data.

Fault detection and diagnostics. FDD provides the ability to continuously monitor and commission a building by alerting the building operator to operational problems (faults) in building systems or equipment. Most commercial FDD tools apply a set of pre-programmed rules to data collected from the building control system (BCS) to determine whether equipment and controls are operating properly. Faults identified by FDD tools are reported to building operators for investigation and correction, ensuring that the building continues to perform as expected, and that energy conservation measures are functioning as intended. The effectiveness of FDD relies on an EMIS to provide accurate and well-described building performance data. At NREL, FDD capability supplements the required 4-year retro-commissioning and energy auditing cycle stipulated by EISA (2007) and assists facility energy managers in identifying, tracking, and retaining additional savings.

Demand management. Demand management can take two different forms: *customer-initiated* (customer scheduling and operation of assets to achieve demand reduction and reduce utility bills) and *utility-initiated* (control signals or requests to the customer to reduce demand when the utility requires). Utility-initiated demand management can be accomplished through dedicated protocols such as OpenADR (2012) that communicate directly with pre-selected devices; this can be achieved through direct interaction with a BCS. Customer-initiated demand management, on the other hand, depends on integrated, organized, and accessible data to achieve the coordinated controls required for reliable demand reduction. Although both of these demand-reduction methods can be facilitated by an EMIS, customer-initiated demand management in particular benefits enormously from the visibility and control of the building or campus loads and assets that an EMIS provides. NREL is especially interested in customer-initiated demand management because of the large number of unique buildings and onsite generation present on its research

campuses. NREL's experience is typical of campuses with a common owner and geographical location; these customers often have a single billing meter, have large demand components in their utility bills, and may have controllable onsite generation or energy storage that can be dispatched to perform demand management.

Predictive/condition-based maintenance. The insight into real-time performance of building equipment enables maintenance programs to be conducted based on the current or predicted performance of that equipment. Predictive and condition-based maintenance can reduce the number of faults that FDD must identify, equipment failure rates, and work-hours required to check on equipment status. It can also more accurately match replacement schedules with the required equipment performance. NREL is pursuing this type of enhanced maintenance program and is interested in automating the approach using EMIS-enabled analytics.

Optimized controls. Most BCS are not designed to integrate large numbers of data points into control sequences; run advanced algorithms and optimization routines using integrated data; or coordinate control among multiple buildings, renewable generators, and energy storage devices. Beyond BCS, however, forward-thinking companies are now leveraging ever-increasing computing capabilities and innovative optimization algorithms to improve building performance and minimize energy consumption. An EMIS provides the access point to integrated BCS data, enabling many different technology providers to deliver optimized controls that increase the efficiency of building or campus operation. As with demand management, NREL desires to capture both the energy-savings opportunities associated with optimized control tools available in the marketplace and to pilot the laboratory's research on this topic.

Dashboarding and tenant engagement. Dissemination of energy information through well-designed dashboards can enhance situational awareness for a building control team, enable portfolio management via a single interface, extend energy awareness to building occupants, and educate the general public. The suitable solution to achieve these benefits will vary with the particular goals of an organization, but all are enabled by a high-quality integration point and depend on the well-curated data that an EMIS provides. NREL recognizes the value in fostering awareness of energy consumption for both employees and visitors to the laboratory and has active research in this area (Schott et al. 2012). Compelling dashboard graphics enabled by the EMIS catalyze educated and nuanced discussions that otherwise would not take place.

Barriers

The benefits described above present a compelling case for the deployment of an EMIS, yet there are numerous barriers to effective deployment that must be addressed and overcome to achieve the type of EMIS operation that delivers the identified benefits. The barriers listed below are indicative of the wide range of site components and personnel that an EMIS interacts with, including: building managers and technicians, IT and networking, cybersecurity, energy engineers, and sustainability managers. This wide-ranging impact is possibly the largest barrier to adoption, as communication and coordination between all of these stakeholders is critical to successful deployment.

Data integration. A functional EMIS requires integration of data from many devices via many communications protocols. BCS communication protocols include BACnet, LonWorks, oBIX, and a vast number of proprietary protocols offered by individual controls providers (Livingood et

al. 2016). Most of these protocols support communication over physical Master-Slave Token Passing (MSTP) and Transmission Control Protocol/Internet Protocol (TCP/IP) connections. Additional protocols (Modbus, DNP3, IEC 61850, and others) are widespread in the electrical generation and distribution industry and are also commonly found in building equipment. The wide variety of protocols in use at any given site—and associated software drivers required to translate into standardized data formats—presents one of the most significant integration challenges facing EMIS. Related data-quality and coherency issues include misaligned timestamps, dropped data, and mismatched precision as the data are transferred from the BCS into the EMIS platform.

Unlike many sites, NREL is fortunate that its BCS is limited only to Modbus and BACnet, with BACnet providing the dominant control network. The NREL campus BCS contains a mix of BACnet MSTP, BACnet Ethernet, and BACnet IP devices, which had impact on the evaluation of EMIS offerings and integration troubleshooting. NREL also sought to integrate systems from vendors of submeters, batteries, etc., whose data are only accessible via web-based application programming interfaces (APIs). The internet access required to pull data via these APIs (along with web-based data sources such as weather feeds or real-time pricing in the future) has implications for both cybersecurity and integration cost.

Data Context. Data integration converts all of the time-series data into a single format, yet gleaning value from those data requires rich context that describes what the data mean. This is typically accomplished through metadata, that is, data that describe relationships between sensors or equipment and describe attributes such as units, type of measurement, and sign convention of the measurement. BCS in buildings today do not typically contain sufficient metadata to enable the benefits described above; information is often limited to what is contained in the point name itself and knowledge of the control panel on which the point resides. The point name rarely contains sufficient information about the measurement. Even if it does provide a good description, it is unlikely that the naming is consistent and machine-friendly across the site's portfolio such that automatic parsing of metadata is possible.

At the beginning of this project, no consistent point naming convention existed at NREL. The campus BCS contains a variety of naming practices adopted by the controls design team for each building. This has made it challenging to know exactly what data are being measured without direct, prolonged discussion with the building manager or the contractor who installed the system. A limited amount of metadata (units of measure) is available within the BCS and within a legacy data acquisition system, but neither data source has a fully defined taxonomy for the metadata or a well-developed meta-model (model for how the metadata are applied). In addition, our team found many instances in which the existing metadata were completely incorrect. The level of effort required to correct existing metadata and assign accurate new metadata to existing data points is substantial and remains a key hurdle to full EMIS deployment at NREL.

IT Infrastructure. Within the general category of "IT infrastructure", the main items of concern are: types of communications networks installed in the building or campus, configuration of those networks, and performance of connected hardware such as supervisory control panels. Historically, BCS networks and general IT infrastructure are often designed and implemented separately. BCS networks in particular may combine serial, Ethernet, and IP communication, complicating integration with enterprise data networks. Installed BCS infrastructure may be insufficient to support the volume and frequency of data transmission required for effective

EMIS integration, and network hardware may not be configured in such a way that BCS and enterprise networks can be easily joined. It is important to understand the installed infrastructure and its potential flexibility in order to design the suitable solution for pulling data off that transport network.

Our team has faced IT hurdles associated with integrating the data from NREL's BCS due to the configuration of the BACnet network (relatively flat with multiple translation points to avoid excessive broadcasting) and the need to convert between BACnet Ethernet to BACnet IP used by the gateways that communicate with some of the pilot installations. This BCS network structure complicated network communications and device discovery during NREL's EMIS pilot implementations. Another challenge was in establishing a server environment to support operation of EMIS software with appropriately configured networking, and then enabling access to this environment for contractor staff. In several cases, these challenges were not a technical barrier but rather an organizational obstacle associated with the separation of IT from the site operations team. Addressing this issue required coordinating the correct personnel to implement requirements and to facilitate sharing of expertise that sometimes remains compartmentalized at NREL, as in any large organization.

Cybersecurity. In today's connected world, cybersecurity is an all-encompassing and critical topic when working with BCS. Health and safety systems are often integrated with or operate on the same networks as these control systems, requiring continuous and effective operation of the common network. At NREL, this includes many laboratory buildings with safety-critical ventilation systems. There are also a large number of physical assets that could cause harm to the building and/or its occupants if a malicious act or human error were introduced into the control system. BCS have a number of physical interfaces and may operate using outdated systems, firmware, and protocols, which makes them especially vulnerable to cyber-physical attacks. This paper does not address the cybersecurity considerations associated with BCS; these are addressed elsewhere in the literature (for instance, Zito 2014). Germane to the present discussion is the impact of cybersecurity concerns on the interoperability and integration of an EMIS. To reduce risk, control systems often operate on "moderate-security" networks, disconnected or firewalled from enterprise networks. Such dedicated control networks often have little or no Internet connectivity.

NREL faced challenges ensuring that the correct components of the EMIS were located at an appropriate security level in the network, that the firewall had the correct ports opened, and that the EMIS supported the protocols required by the pilot systems. Similar to the IT infrastructure discussion above, the barriers were both technical (such as hardware not having the appropriate number of Ethernet ports to support crossing multiple networks) and organizational (such as obtaining approval for bidirectional communication between enterprise and BCS networks). Resolving the organizational issues via effective communication between building managers, the IT team, cybersecurity team, and our team has been possible, but laborious.

Physical meters and sensors. Over time, meters and sensors may fall out of calibration, reducing the accuracy of measurements and increasing possibility of BCS malfunction. An EMIS system relies on functioning and calibrated meters and sensors to correctly identify high energy-consuming equipment, track the performance of energy-conservation projects, and make predictions about future performance/conditions.

At NREL, meters and sensors are calibrated at regular intervals according to their relative impact on research and operations. However, calibration intervals have not always been aligned

with the priorities of all data users, and meter malfunctions have often persisted undiscovered for weeks or months. Going forward, NREL intends to organize meter calibration through a system of tiered priority assignment.

Policies, procedures, and personnel. EMIS benefits do not materialize simply by the installation of an EMIS; an EMIS by itself does not provide any energy or cost savings (except in certain automated controls applications). The vast majority of the benefits identified rely on people to interface with the EMIS and then act based on recommendations for actual savings to accrue. The policies, procedures, and personnel put in place to support the EMIS determine the success of an energy intelligent campus. Appropriate staffing, applicable training, standard procedures, and adequate finances for the ongoing maintenance of the system are all critical to the success of a deployment.

For NREL, like most large organizations, the topic of personnel required to sustain the operations of the EMIS precipitated discussion of the up-front and annual financial obligations. Because costs are certain but benefits are uncertain, decision makers are often hesitant to fund EMIS support at a level sufficient to realize savings. Additionally, success is not only defined by the creation of the system but also communicating and facilitating the utilization of its capabilities to all stakeholders.

An Intelligent Energy Data Architecture

Effective requirements development for EMIS acquisition must address the barriers identified in the previous section. To bring structure to the requirements conversation, we developed a conceptual EMIS architecture based on distinct functional layers. Using the architecture as a guide, we then developed functional requirements that addressed technical, informational, and organizational usability and interoperability. These, in turn, drove selection of technologies for NREL's EMIS pilot implementations.

System Architecture

A conceptual system architecture consisting of five distinct functional layers—devices, drivers, data historian, applications, and supervisory controls—is shown in Figure 1. Although commercially available EMIS products often combine multiple layers into a single, seamless package, separating the functions conceptually assists in the identification of requirements that address specific barriers.

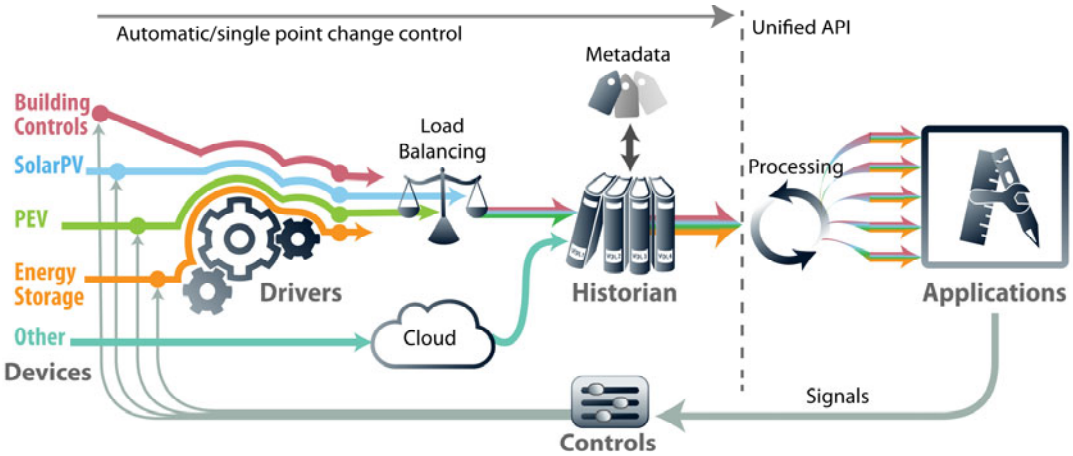

Figure 1. NREL's conceptual EMIS architecture, including functional layers and key features.

Data acquisition begins with sensors and actuators within the Device Layer. The Driver Layer transfers data from devices to the data historian and upstream applications; the drivers provide protocol translation, sampling, and load balancing. The Historian Layer stores historical data in a time series database. Applications in the Application Layer consume those data for a variety of use cases via a common API. Finally, working in concert with applications, the Control Layer provides feedback or control signals to devices to execute supervisory control actions.[1]

To date, we have focused on requirements for the first three layers: devices, drivers, and the historian. Many of the key data integration challenges occur at the nexus of these layers. Crafting requirements such that all communication to higher-level applications and supervisory control occurs via a standard, unified API mitigates data integration challenges at the Applications Layer.

The Device Layer is the collection of equipment, instrumentation, and local control devices that comprise the physical hardware of a building. It may encompass entire low-level communication networks, such as BCS networks and associated equipment. Time series building performance data originates in this layer.

The Driver Layer is responsible for managing communication between the Device Layer and the Historian Layer, including protocol translation. The performance of this layer is strongly influenced by the responsiveness of the devices and by the quality of the low-level communication networks. Many BCS networks are not designed for heavy data acquisition; excessive data requests can overload the underlying control systems and compromise their ability to perform basic control functions. Therefore, the Driver Layer should include a load-balancing capability that restricts network traffic as needed to preserve performance.

The Historian Layer stores time series data and associated metadata in one or more databases, providing those data on request to applications. Of primary concern with the Historian Layer is database performance, as the volume, velocity, and concurrency of input/output can challenge certain database configurations. Historian metadata management is also a major challenge, and the value of accurate, well-organized metadata cannot be overstated. Incorrect or nonstandard metadata directly affect accuracy: they may result in inaccurate interpretation of measured values, cascade into incorrect operation of upstream energy analysis and FDD tools, or create difficulty in locating the equipment for calibration or repair.

The Application Layer consists of all high-level analysis tools that rely on collected building performance data. Examples include dashboards, benchmarking and reporting software, and FDD tools. Applications query historical data from the historian and may also read real-time data directly from the drivers. Rather than the specifics of each application, NREL's requirements development targeted two main sources of overhead in application development: locating relevant energy data resources and integrating them with application code. Metadata improvements at the Historian Layer accelerate resource location while standardized data access

[1] In certain configurations, the Control Layer could communicate back through to devices through the intermediate API and the Driver Layer (perhaps represented by two-way arrows in the diagram), but we have kept the communication pathways separate for clarity.

enables a wide suite of applications to interface easily and effectively with the lower-level layers via an open, documented API that serves the time series data with rich context surrounding it.

The Control Layer represents supervisory control systems that have a need to affect the operation of building devices in an automated or semi-automated manner. Examples include optimization of HVAC operation and automated demand response—two applications for which products currently exist in the marketplace. For closed-loop control, the Control Layer specifically requires a communication path from applications to device hardware, either by allowing bidirectional communication through the Driver Layer or via a separate feedback path. The Controls Layer is therefore conceptually separate from the Applications Layer because closed-loop control introduces a host of integration and security challenges that data-consuming applications with one-way data transfer do not require. Without closed-loop control, it is still possible for human operators to take action based on the analytics performed in the Application Layer.

Requirements

NREL's requirements development shadowed the barriers identified above, with a particular focus on interoperability: enabling seamless communication and reducing the burden associated with maintaining accurate and up-to-date metadata. To date, we have developed detailed requirements for the first three layers (device, driver, and historian), which is consistent with our data collection and interoperability focus. Hardin et al. (2015) provide a useful framework for buildings interoperability by dividing requirements into three categories: technical, informational, and organization. Figure 2 maps several existing industry standards into these three columns of interoperability.

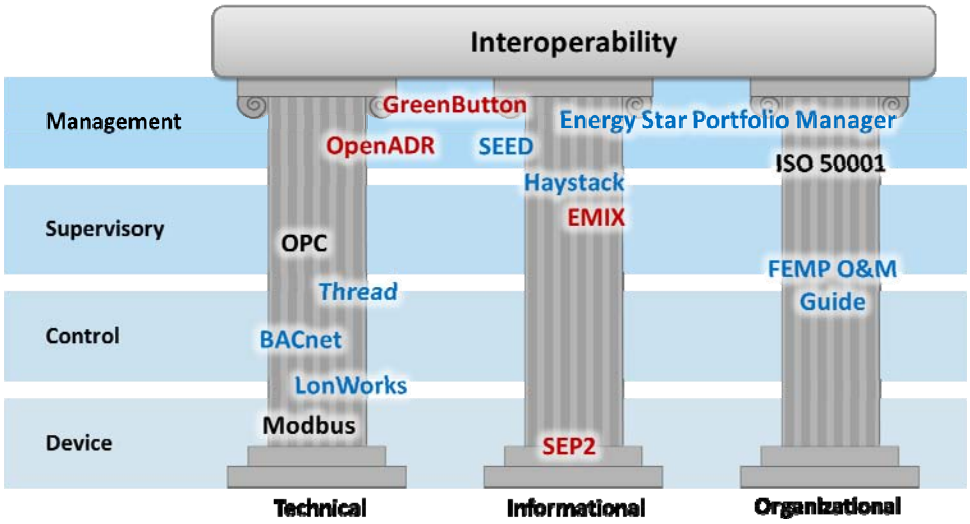

Figure 2. Industry standards pertaining to three pillars of interoperability in buildings. Red text indicates utility standards; blue text indicates buildings standards; standards in black text are broadly applicable.

Technical Requirements. Technical requirements for interoperability represent hardware, communication protocols, and network specifications that collectively ensure functional, efficient data collection. We focused on technical requirements pertaining to physical sensors, data integration, and IT infrastructure barriers. The key technical requirements are:

1. Sensors used for billing, internal recharges, or federally mandated reporting must comply with ANSI C12.20 and C57.13 Class 0.3 standards for accuracy (ANSI 2008).
2. The Driver Layer must support BACnet and Modbus communication at minimum; import capabilities for comma-separated values (CSV), JavaScript object notation (JSON), and extensible markup language (XML) data are highly desirable.
3. To conserve network bandwidth, the drivers must support change-of-value (COV) and change-of-state data collection, when enabled by the underlying communication protocol.
4. Drivers must perform load balancing via automatic or user-configurable traffic throttling.
5. During data collection, timestamps between the Device, Driver, and Historian Layers must be aligned, that is, drivers must correct any time-shift errors that result from data caching or unsynchronized clocks.
6. To avoid information loss, data type and precision must be compatible across all layers.
7. The driver and historian must mitigate issues associated with dropped data, e.g., by automatic back-collection of historical data following outages.
8. The historian must scale to support 1,000,000 simultaneous time series, 10 simultaneous users, and 5+ years of data retention for 1-minute interval data.

Informational Requirements. Informational requirements govern the organization and interpretation of data. Successful collection of time-series data is only one prerequisite for successful analytics; intuitive, well-organized, and searchable metadata is another. Of the several available standards that address the organization and exchange of building performance data (Figure 2), we selected Project Haystack (2016). Among the standards we considered, Haystack is not only the most comprehensive, but also the simplest and most flexible.

The Haystack standard defines three main elements: a flexible and extensible tag model, a list of standard tags with accepted definitions, and a representational state transfer (REST) API. Haystack tags are a collection of name/value pairs applied to objects that describe both intrinsic characteristics of the objects and relationships between objects. Some tags, called "markers," do not have an associated value. For example, an air handling unit might have the tags *equip* (designating equipment), *ahu*, and *dis:"Air Handler 1"* (which provides a user-friendly display name/description). Because there is no fixed schema that constrains which tags are used, extending Haystack is as simple as adding custom tag definitions for specific needs. However, for maximum compatibility with other Haystack-compliant software, custom tags should be used only when the standard tag definitions are insufficient to fully characterize the data.

In Haystack, each modeled object has an *id* tag with a unique identifier. Relationships between objects are modeled with references, which are tags that point to other objects' *id*s. While the primary relationship structure in Haystack is the site/equip/point hierarchy illustrated in Figure 3, the tag model can also easily accommodate peer-to-peer and other relationships that do not fit conveniently into traditional tree structures. Electrical distribution systems, communication networks, drivers, and functions are all readily accommodated within Haystack.

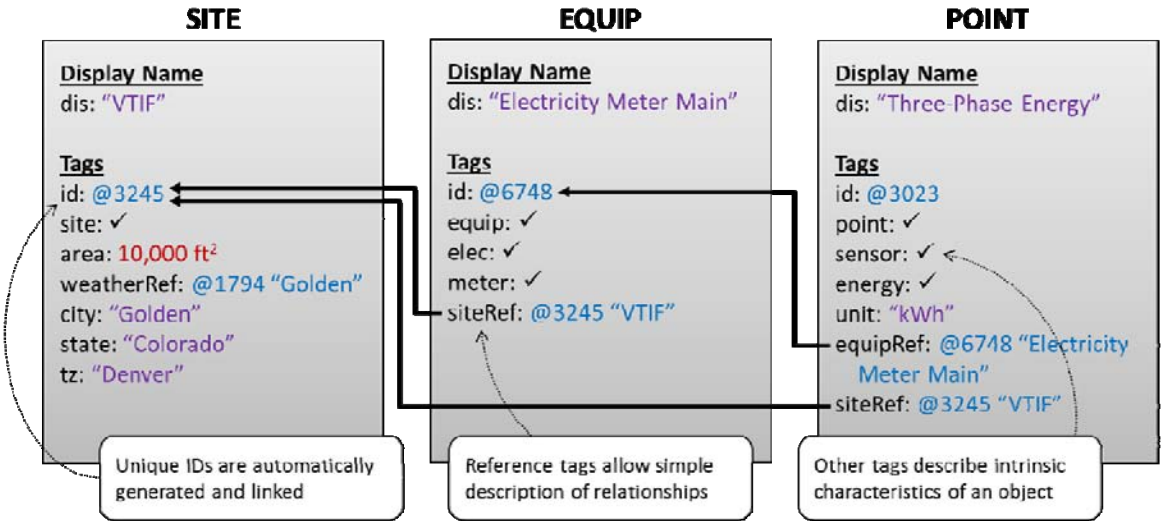

Figure 3. Example application of Project Haystack tags to describe an electricity meter, illustrating how references are used to construct the typical site/equip/point hierarchy. Marker tags are indicated with a checkmark symbol.

The Project Haystack standard addresses several key data integration and data context barriers: it provides standard unit notation (via the *unit* tag) and conversions via a built-in unit database; it provides a well-defined taxonomy of common building equipment and sensor tags; it allows rapid and flexible queries via arbitrary tag searches; and it efficiently models relationships between various systems and equipment.

Organizational Requirements. Organization requirements include policies, procedures, and personnel training requirements that are needed to ensure the long-term viability of an EMIS. Without organizational infrastructure to match, even the best EMIS technology will be ineffective. Standards such as ISO 50001 (ISO 2011) and FEMP's Operations and Maintenance Best Practices guide (Sullivan et al. 2010) provide high-level guidance, but ultimately, organizations must create procedures specific to their EMIS.

To address organizational requirements, we created a data-quality document that describes procedures that will standardize and maintain the EMIS infrastructure that we seek to adopt. These procedures address critical actions that still require human intervention given EMIS technology available today. These include:

- Regular calibration and maintenance of critical sensors and equipment.
- A comprehensive commissioning procedure for new equipment that adds metadata input and verification to the typical functional testing requirements.
- A standard workflow for responding to EMIS data-collection errors or outages, including clearly delineated roles and responsibilities.
- A recommendation to hire personnel dedicated to supervising the EMIS and responding to the issues it will inevitably uncover; this assures that NREL will follow through on energy-savings opportunities identified via the EMIS.

Pilot Implementations

Based on the requirements described above, we selected three pilot implementations for deployment and testing on the NREL campus. The intent of the pilots was to shed light on the ability of industry solutions to deliver on the requirements outlined above. In many cases, the

requirements could be met with multiple approaches—with varying levels of effectiveness in achieving the desired benefits—therefore, we wanted to evaluate multiple implementations before pursuing a larger deployment. All of the pilots were supported by commercial offerings, with some being contained within a single product and others consisting of multiple products and associated interfaces.

This section describes the salient characteristics of the three pilots and maps functional layers to features in the pilots. Although some of the products provided an Application Layer and/or Control Layer, this was not central to our evaluation. Emphasis was instead placed on the level of interoperability between the Historian and Application Layers.

Pilot A: Open Platform Communications (OPC). The first of the three pilots was based on Open Platform Communications (OPC) (OPC, 2016). OPC provided the Driver Layer in our system architecture, enabling communication with BACnet and Modbus. The OPC server was paired with a Haystack-compliant analytics platform with a scalable time-series database. The database associated with this analytics platform provided the Historian Layer and enabled Haystack-compliant metadata. The OPC server resided on the moderate network to enable full communication with the BACnet network and Modbus data. The analytics platform resided on the enterprise network with firewall exceptions to allow access to the OPC server through appropriate ports, enabling user-defined access for laboratory staff.

The OPC approach enabled the use of well-vetted, hardened driver software, but also created integration challenges for metadata. Data transfer from OPC to the Haystack-compliant historian required an intermediate MySQL database. This additional data link provided opportunity for translation error and lost data and was therefore a key weakness of the implementation. In the future, native OPC communication from the analytics platform may mitigate this weakness.

Pilot B: Dedicated Software on the Local Network. This pilot consisted of a single software solution providing Driver and Historian Layers. The software runs on a server in the same moderate security network enclave as the OPC server in Pilot A. The dedicated software provides the required drivers to access and translate the BACnet and Modbus networks and writes that data to a local historian. These data can then be accessed via the Haystack API by multiple analytics platforms. The dedicated software also provides certain analytics capabilities (supervisory control, FDD, others) as part of the software solution. This software solution, while not as time-tested as the OPC standard, was attractive due to its tightly integrated feature set.

Pilot C: Middleware Onsite/ Database in the Cloud. This pilot demonstrated a scalable, cloud-based data aggregation and warehousing capability. This implementation used a combination of hardware and software to push the data from the local control systems to a secure cloud-hosted database. The hardware component was a middleware panel that served as the Driver Layer and enabled a physical connection with the network. This middleware also created a secure bridge between the BCS network and the Internet, enabling data access from the enterprise network. This architecture used a combination of Apache Cassandra for time series data and MongoDB for metadata with a Haystack-compliant API.

Outcomes

Three pilot EMIS installations were deployed on NREL's South Table Mountain campus. Each implementation transferred data from devices to the historian and provided unified access

via the Haystack API. We evaluated how well each implementation met the requirements using a series of standard tests. Examples include “add a point,” “assign Haystack-compliant metadata,” and “induce a one-hour network interruption to historian with no data loss.”

All three systems met most requirements, but with varying degrees of usability. Pilot A offered the most stable communication between devices and drivers, but the MySQL table for data transfer between the OPC software and the Haystack-compliant historian is an undesirable link that may have high long-term maintenance cost. However, separating the historian from the drivers allowed the analytics software providing the historian to be located on NREL’s enterprise network, which increased accessibility.

Pilot B offered the most direct access to the underlying BCS and the best user interface for NREL staff to discover, tag, and organize points. However, because the same software served as both the Driver Layer and the Historian Layer, Pilot B also presented the greatest tradeoff between security and usability. Because the Pilot B software resides in a moderate security enclave alongside the BCS network, it is challenging to grant users on the enterprise network access to the data without a slow and cumbersome authentication process.

Surprisingly, despite the requirement for direct Internet access, Pilot C proved the simplest from the perspective of security integration. The dedicated hardware used two separate Ethernet ports, allowing physical separation of BCS network and Internet communication. Bypassing NREL’s moderate security network entirely allowed quicker approval of network interconnection. However, the user interface to the cloud database for Pilot C made data management difficult: point creation, renaming, and metadata tagging were cumbersome and batch processing was difficult. The Pilot C database is not intended as a full analytics platform but only for data storage, which may have contributed to the poor user experience.

All three implementations had difficulty satisfying two key requirements: automated device discovery and COV data collection. In theory, polling a BACnet network for devices and constructing device table is straightforward; but in practice, the sheer volume of devices, inconsistency in replies, and network fragmentation made effective device discovery nearly impossible. In all three cases, we had to resort to manual device discovery using BCS network and device information provided by NREL staff knowledgeable of the underlying connectivity. Similarly, although the software used in the pilots advertised COV data collection, in most cases the drivers relied on polling the underlying BCS (even if the data were *recorded* using a COV strategy). Improved COV implementation would reduce unnecessary network traffic and will be addressed in future deployments.

Lessons Learned

- *Selection of well-defined, smaller-scale test beds creates scalable solutions*. Beginning with a small, contained test enables understanding of the system’s construction, functionality, and expansion capabilities without an expensive implementation investment. Small tests also mitigate risk and the perception of risk, allowing experimentation without impacting the entire campus portfolio.
- *Incorporation of open standard products and protocols provides flexibility and interoperability.* Use of proprietary standards leads to vendor lock-in and limits scalability, adaptability, and the ability to readily incorporate future technologies. Open standards, such as Project Haystack, provide structured and standardized deployments and accelerate new application development. By constructing the EMIS with open standard components, owners can retain flexibility.

- *Organizational requirements facilitate EMIS operation.* The creation of a procedure to standardize point naming and metadata assignment ensures consistent data across the campus. The standardized naming ensures consistency between the historian and the point names in the control system, enabling efficient correction of problems that are identified via advanced analytics applications.
- *Investment in providing robust data quality serves current and future needs*. Performance failure during data collection undermines all applications that interface with the EMIS. Data accuracy is fundamental to the ability of researchers, building engineers, and site operations to confidently leverage the wealth of sensor data generated on campus.
- *Creation of integrated project teams (IPTs) facilitates product development.* An IPT leverages financial and human capital resources, and results in overall project effectiveness. To productively develop, operate, and support the EMIS, it is imperative to accurately understand the various levels of user needs and establish realistic requirements early on. Identifying team members in all functional disciplines is crucial because they influence the EMIS development throughout its evolution. Another aspect to the IPT's purpose is the identification of resource requirements such as staffing, equipment, and funding.

Conclusion

The deployment of an effective EMIS requires a clear understanding of desired benefits and anticipated barriers for the organization. The dissection of the EMIS into functional layers provides a clearer understanding of the requirements for each layer and creates opportunities for integrating different products to create a system that meets site-specific requirements. The requirements should address both the various functional layers, and the different types of interoperability (technical, informational, and organizational). Often, the technical requirements are the best understood and easiest to identify, but specifying and addressing the informational and organizational requirements can have a large impact on EMIS success.

In NREL's case, the key deployment barriers have been related to networking, security, and IT infrastructure. These barriers are organizational as much as technical, and they can be overcome by clear and frequent communication between site operations, research, and IT staff. Poor performance of the vendor solutions with respect to device discovery, efficient BCS network communication, and user interaction presented secondary barriers. These can be mitigated by working with vendors to improve software functionality; but effective vendor engagement requires staff time and financial resources to pay for improvements.

Future work will involve continued testing of the pilot implementations, assisting in deployment of a full-scale EMIS for NREL, and research into methods for increased interoperability between the functional layers of an EMIS. Our team is also actively engaged with NREL's cybersecurity team, and we are working on solutions for bidirectional control through the EMIS. This will facilitate onsite demonstration of demand management and controls optimization, and it will provide a venue for application of research activities in both these areas.

References

ANSI (American National Standards Institute). 2008. ANSI C12.1-2008: American National Standard for Electric Meters - Code for Electricity Metering.

EISA (Energy Independence and Security Act of 2007): Section 432. 2007.

Granderson J., Lin G., and Piette M. A. 2013. *Energy Information Systems: Technology Costs, Benefit, and Best Practice Uses*. LBNL-6476E. Lawrence Berkeley National Laboratory, Berkeley, CA. Accessed March 1, 2016: http://eis.lbl.gov/pubs/lbnl-6476e.pdf

Hardin D.B., Corbin C.D., Stephan E.G., Widergren S.E., Wang W. 2015. *Buildings Interoperability Landscape*. Pacific Northwest National Laboratory. Accessed March 2, 2016. http://energy.gov/sites/prod/files/2016/01/f28/BuildingLandscapeReport.pdf

IEEE (Institute of Electrical Engineers Power Engineering Society). 2008. IEEE Std C57.13-2008: IEEE Standard Requirements for Instrument Transformers. Revision of IEEE Std C57.13-1993.

ISO (International Organization for Standardization) 2011. ISO 50001: Energy Management Systems – Requirements with Guidance for Use. First edition. Reference number ISO 50001:2011 (E). Downloaded January 5, 2016.

Livingood W., Stein J., Considine T., Sloup C. *Review of Current Data Exchange Practices: Providing Descriptive Data to Assist with Building Operations.* National Renewable Energy Laboratory, Golden CO. Accessed March 10, 2016. http://www.nrel.gov/docs/fy11osti/50073.pdf

Mills, E. 2009. *Building Commissioning: A Golden Opportunity for Reducing Energy Costs and Greenhouse Gas Emissions*. Lawrence Berkeley National Laboratory. http://cx.lbl.gov/documents/2009-assessment/lbnl-cx-cost-benefit.pdf

OpenADR (OpenADR Alliance). 2011-12. *OpenADR 2.0 Profile Specification.*

Project Haystack. Accessed March 10, 2016. http://project-haystack.org/

Schott M., Long N., Scheib J., Fleming K., Benne K., and Brackney L. 2012. "Progress on Enabling an Interactive Conversation between Commercial Building Occupants and Their Building to Improve Comfort and Energy Efficiency." In *Proceedings of the ACEEE 2012 Study on Energy Efficiency in Buildings,* 7:250–265. Washington, D.C.: ACEEE.

Sullivan G.P., Pugh R., Melendez A.P., Hunt. 2010 W.D. *Operations and Maintenance Best Practices*. Federal Energy Management Program. http://energy.gov/eere/femp/downloads/operations-and-maintenance-best-practices-guide

Zito P. 2014. Decoding the World of BAS Security. ASHRAE Journal. Accessed March 10, 2016. https://www.ashrae.org/File%20Library/docLib/eNewsletters/Zito-092014--01142016feature.pdf