LBNL-1004415

Lawrence Berkeley National Laboratory

Field Testing of Telemetry for Demand Response Control of Small Loads

Steven Lanziera, Andrew Weber, Anna Liao, Oren Schetrit, Sila Kiliccote and Mary Ann Piette

Lawrence Berkeley National Laboratory

November 2015

The work described in this report was coordinated by the Demand Response Research Center and funded by the California Energy Commission (Energy Commission), 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

Energy Research and Development Division
Final Report

FIELD TESTING OF TELEMETRY FOR DEMAND RESPONSE CONTROL OF SMALL LOADS

Prepared for: California Energy Commission
Prepared by: Lawrence Berkeley National Laboratory

NOVEMBER 2015
CEC-500-YYYY-XXX

PREPARED BY:

Primary Author(s):
Steven Lanzisera
Andrew Weber
Anna Liao
Oren Schetrit
Sila Kiliccote
Mazry Ann Piette

Lawrence Berkeley National Laborator
1 Cyclotron Road
Berkeley, CA 94720

Contract Number: 500-26-026

Prepared for:

California Energy Commission

David Hungerford
Contract Manager

Virginia Lew
Office Manager
Energy Efficiency Research Office

Laurie ten Hope
Deputy Director
ENERGY RESEARCH AND DEVELOPMENT DIVISION

Robert P. Oglesby
Executive Director

Acknowledgements

The work described in this report was fund by the California Energy Commission (Energy Commission), 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. The authors would like to thank Chris Scruton, David Hungerford and Mike Gravely from the Energy Commission for their support in this and other Demand Response Research Center Projects.

ABSTRACT

The electricity system in California, from generation through loads, must be prepared for high renewable penetration and increased electrification of end uses while providing increased resilience and lower operating cost. California has an aggressive renewable portfolio standard that is complemented by world-leading greenhouse gas goals. The goal of this project was to evaluate methods of enabling fast demand response (DR) signaling to small loads for low-cost site enablement. We used OpenADR 2.0 to meet telemetry requirements for providing ancillary services, and we used a variety of low-cost devices coupled with open-source software to enable an end-to-end fast DR. The devices, architecture, implementation, and testing of the system is discussed in this report. We demonstrate that the emerging Internet of Things (IoT) and Smart Home movements provide an opportunity for diverse small loads to provide fast, low-cost demand response. We used Internet-connected lights, thermostats, load interruption devices, and water heaters to demonstrate an ecosystem of controllable devices. The system demonstrated is capable of providing fast load shed for between $20 and $300 per kilowatt (kW) of available load. The wide range results from some loads may have very low cost but also very little shed capability (a 10 watt [W] LED light can only shed a maximum of 10 W) while some loads (e.g., water heaters or air conditioners) can shed several kilowatts but have a higher initial cost. These costs, however, compare well with other fast demand response costs, with typically are over $100/kilowatt of shed. We contend these loads are even more attractive than their price suggests because many of them will be installed for energy efficiency or non-energy benefits (e.g., improved lighting quality or controllability), and the ability to use them for fast DR is a secondary benefit. Therefore the cost of enabling them for DR may approach zero if a software-only solution can be deployed to enable fast DR after devices are installed for other reasons. We recommend that the DR research community continue to engage with the IoT community to encourage the use of documented and open development interfaces. A library of device drivers and machine-readable interface specifications would significantly reduce the burden on users or system integrators for deploying systems in large numbers of buildings in California.

Keywords: demand response, electricity loads, telemetry, automation, renewable integration, end-use load control

Please use the following citation for this report:

Lanzisera, Lanziera, Weber, Andrew, Liao , Anna, Schetrit, Oren, Killicote, Sila, Piette, Mary Ann. (Lawrence Berkeley National Laboratory). 2015. ***Field Testing of Telemetry for Demand Response Cotnrol of Small Loads.*** California Energy Commission. Publication number: CEC-500-YYYY-XXX.

For more information about the Energy Research and Development Division, please visit the Energy Commission's website at www.energy.ca.gov/research / or contact the Energy Comission at 916-327-1551

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

EXECUTIVE SUMMARY

Introduction:

The electricity system in California, from generation through loads, must be prepared for high renewable penetration and increased electrification of end uses while providing increased resilience and lower operating cost. California has an aggressive renewable portfolio standard that is complemented by world-leading greenhouse gas goals. Taken together, it is clear that all elements of the electricity ecosystem will need to be smarter and more interactive to ensure grid reliability and minimize overall system cost.

Project Purpose:

The purpose of this project was to evaluate methods of enabling fast demand response (DR) signaling to small loads for low-cost site enablement. The term "fast DR" is defined as demand-side resources that respond without advanced notification and with fast response time (within minutes to seconds). We used OpenADR 2.0 to meet telemetry requirements for providing ancillary services, and we used a variety of low-cost devices coupled with open-source software to enable an end-to-end fast DR. The devices, architecture, implementation, and testing of the system is discussed in this report.

Project Results:

We demonstrate that the emerging Internet of Things (IoT) and Smart Home movements provide an opportunity for diverse small loads to provide fast, low-cost demand response. We used Internet-connected lights, thermostats, load interruption devices, and water heaters to demonstrate an ecosystem of controllable devices. The utility-installed smart meter with a home area network (HAN) radio provides near real-time power data for telemetry feedback to the OpenADR 2.0 virtual top node (VTN, also commonly called the *server*). The system demonstrated is capable of providing fast load shed for between $20 and $300 per kilowatt (kW) of available load. The wide range results from some loads may have very low cost but also very little shed capability (a 10 watt [W] LED light can only shed a maximum of 10 W) while some loads (e.g., water heaters or air conditioners) can shed several kilowatts but have a higher initial cost. These costs, however, compare well with other fast demand response costs, with typically are over $100/kilowatt of shed. We contend these loads are more attractive than their price suggests because many of them will be installed for energy efficiency or non-energy benefits (e.g., improved lighting quality or controllability), and the ability to use them for fast DR is a secondary benefit. Therefore the cost of enabling DR may approach zero if a software-only solution can be deployed to enable fast DR after devices are installed for other reasons.

Project Benefits:

As mentioned, we demonstrate that the cost of enabling DR may approach zero if a software-only solution can be deployed to enable fast DR. One barrier to widespread deployment of small loads for fast DR is the availability and documentation of open network interfaces for the devices under control and for the smart meter HAN interface. Today devices use a custom communication protocol, and the level of protocol documentation varies widely from device to device. OpenADR does not naturally fill the role of providing specific control to individual devices. We recommend that the demand response research community continues to engage with the IoT community to encourage the use of documented and open development interfaces. A library of device drivers and machine-readable interface specifications would significantly reduce the burden on users or system integrators for deploying systems in large numbers of buildings in California.

CHAPTER 1: Introduction

Deployment of renewable portfolio standards in 29 states in the U.S. increased the focus on the use of demand response (DR) to address four major challenges related to renewable generation penetration: (1) over-generation during low-load hours, (2) steep and unpredictable ramps, (3) forecast errors associated with renewable generation, (4) and intra-hour variability of these resources (Kiliccote et al. 2010a). Increased flexibility of demand-side resources and availability of real-time signals from the electricity grid are the key ingredients for successful supply and demand interactions. Automated Demand Response (AutoDR) is machine-to-machine enablement of DR in response to grid signals without a human in the loop. While AutoDR is seen as a critical part of managing the smart grid with high renewable penetration, today's wholesale DR systems seem experimental, and retail DR systems typically work on slow time scales as open-loop systems to address peak load reduction (Kiliccote et al. 2010b). In contrast, generation resources continually report generator conditions to and accept commands from grid operators to provide real-time closed-loop control and status information of the grid. Any demand-side resource that transacts with the electricity grid by providing a bid, price, and duration with short (minutes) or no notification is required to adhere to the same communication requirements as a generator. These transactions today happen in the wholesale ancillary services markets.

Ancillary services are support services in the power system and are essential in maintaining power quality and reliability. There are typically two types of ancillary services products with which DR participates: contingency and operating reserves. Table 1 summarizes the specifics of these products. Regulation, the product with the fastest communication requirements in ancillary services markets, allows the system operator to request upward or downward changes in output. It is used to track and balance system wide generator output with system wide load on a sub-minute by sub-minute basis (Kiliccote et al. 2010b). In California and in the Electricity Reliability Council of Texas, regulation is separated into two products: regulation up and regulation down. In all other markets, regulation products are symmetric, meaning the generator signs up to deliver equal amounts of up and down product.

Resources providing regulation products are certified by the independent system operators (ISOs). The ISOs currently have a number of requirements for generators seeking to provide the ancillary service of regulation, including a direct communication and control system that communicates key parameters and allows the generator to respond without operator intervention. These generators must respond to automatic generation control (AGC) signals to change their operating levels depending upon the service they provide, regulation up or regulation down.

Each system operator states specific telemetry requirements, which include remote measurement and communication specifics, that generators must use to participate in regulation products. These range from two- to six-second (four-second in California) granularity in measurement and communication (CAISO 2012).

This project's goal was to evaluate methods of enabling fast DR signaling to small loads for low-cost site enablement. The term "fast DR" is defined as demand-side resources that respond without advanced notification and with fast response time (within minutes to seconds depending on the ancillary services products' requirements outlined in Table 1). In this project we used OpenADR to meet current telemetry requirements for providing ancillary services. We considered a variety of devices and architectures for Internet-based fast demand response. Each of these areas is addressed in a separate section of this report.

Table 1. Summary of ancillary services for fast DR participation

Service	**Service Description**		
	Response Speed	Duration	Market Cycle
Operating Reserves			
Regulating Reserves, or Regulation up/down; Automatic Generation Control (AGC); Frequency responsive reserves	Online/Spinning reserve, immediately responsive to AGC, to allow the Balancing Authority to meet the NERC Real Power Balancing Control Performance.		
	<1 minute; must be able to reach max amount of Reg within 10–30 minutes	30 minutes (Real Time); 60 minutes (Day Ahead)	Hourly; every 15 minutes looking ahead 2 hours
Load Following or Fast Energy Markets	Similar to regulation but slower. Bridges between the regulation service and the hourly energy markets.		
	~10 minutes	10 minutes to hours	5 minutes
Contingency Reserves			
Spinning Reserves	Online generation, synchronized to the grid, that can increase output immediately in response to a major generator or transmission outage and can reach full output within 10 minutes		
	Instantaneous response; <10 minutes for full output	30 minutes	10 minutes
Non-Spinning Reserve	Same as spinning reserve, but need not respond immediately; resources can be offline but still must be capable of reaching full bid within the required 10 minutes		
	< 10 minutes	30 minutes	10 minutes

1.1 Fast DR and the Need for Telemetry (Telemetry Requirements for Ancillary Services)

The California Independent System Operator (CAISO) provides telemetry specifications for communications between generators and the ISO (CAISO 2012). These specifications provide a reference case for any demand-side resources that aim to provide generation-like services. The CAISO specifies that a specific gateway system and communication must take place over the

CAISO's dedicated Energy Communications Network. Both voice and data communications are required between the generators and CAISO, which is typically done over a single data link with Voice-over-Internet Protocol (VoIP).

There are two relevant DR cases to consider: (1) a large load directly participates, or (2) two or more smaller loads participate through an aggregation service. The aggregation service can be a traditional aggregator or a cloud-based system that aggregates and optimizes resources. Figure 1 shows a conceptual schematic of the different portions of the current telemetry system for real-time demand response using information from the CAISO for direct participation of a large load. Structured wholesale markets around the country have similar ancillary services products and requirements that vary primarily in nomenclature and in timing. This figure shows a load directly connected through the dedicated Energy Communications Network (ECN), as well as loads connected via an intermediary—the aggregator. The timing requirements shown reflect those required for resources providing the ancillary service spinning reserve. The 1-minute rate is increased to 4 seconds for the regulation ancillary service and decreased to 5 minutes or 10 minutes for other services or for wholesale energy generation (e.g., load following).

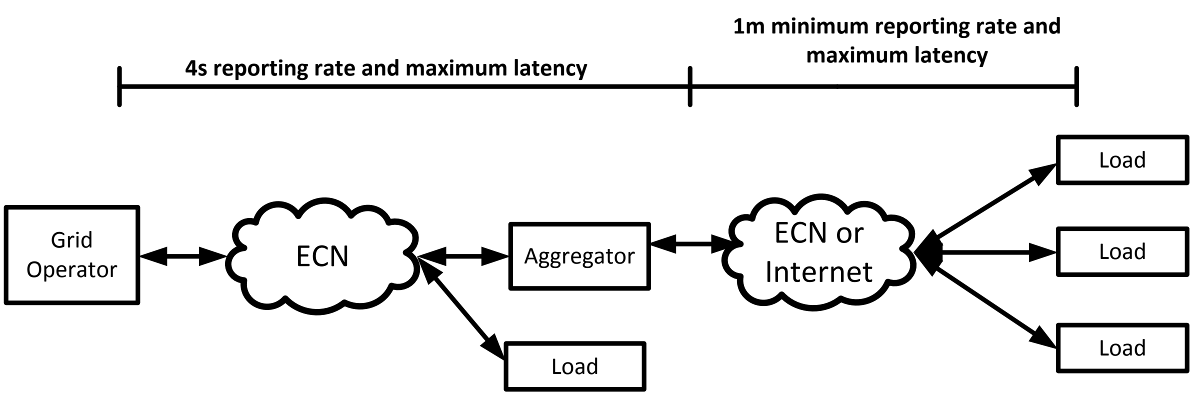

Figure 1. An overview of the architecture showing communication and latency for DR resources with and without aggregation

For our project, we served as in intermediary between the load being shed and the ISO to provide specific data formats, forecasts, capabilities, etc., regardless of the load size. Case 2 (aggregation) is similar to our field studies and implementation. The CAISO requires meter data to be reported every four seconds to the CAISO. These data should be four-second interval data if it is a single load, but aggregated loads can provide 1-minute interval data instead of four-second data. The value must still be reported every four seconds, and the aggregating load meter data server (ALMDS) should update its load estimate with every report it receives (i.e., if there are two loads that report data 20 seconds apart, the ALMDS will report one value for the first 5 samples or 20 seconds, then provide an updated value for the last 10 samples or 40 seconds). An example situation is shown in Figure 2, where an ALMDS is used.

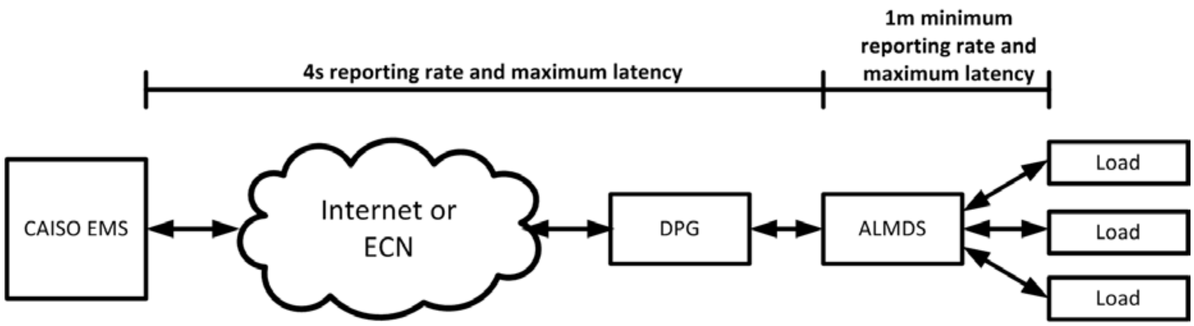

CAISO EMS: The energy management system used to manage the supply demand balance
ECN: Energy Communications Network, a Internet like entity dedicate dto CAISO communication with generators and loads
DPG: Data Processing Gateway. A device the provides protocol and data support in compliance with CAISO standards.
ALMDS: Aggregating Load Meter Data Server. A server that aggregates load information from more than one load

Figure 2. Schematic showing communication latency and reporting rate requirements for demand side aggregated loads, specific to CAISO.

Figure **2** shows the schematics of communication latency and reporting rate requirements for a demand side spinning reserve ancillary services system using several aggregated loads. The telemetry links have some general physical requirements. The loads must communicate with the Data Processing Gateway (DPG) or ALMDS (for single or multiple loads, respectively) using a broadband communication medium. There are no requirements on the data rate of this link. The protocol to the right of the DPG in Figure 2 is unspecified and can be selected by the implementer, but the communication from the DPG to the CAISO EMS must use DNP3 (Distributed Network Protocol Version 3). DNP3 is a Supervisory Control and Data Acquisition (SCADA) protocol that is commonly used in the electricity utility industry. The connection between the DPG and the Internet or ECN must be an always-available connection (i.e., not dial up), but it does not need to be dedicated to only providing ancillary services.

A variety of data types must be reported, depending on the size of the load and the needs of the ISO at that time. For loads less than 10 MW, that is: real power in MW, reactive power in megavolt-ampere reactive (MVAR), Voltage in kilovolts (kV), and connectivity status for each load (if the source is an ALMDS). For loads over 10 MW, gross values for real power and reactive power are also required. A gross value is one measured at the terminals of the load itself rather than one measured at the connection point to the grid. The DPG must also provide a "heartbeat" that changes from 0 to 100 at one increment per second. The CAISO specifies a target data reliability of 99.0 percent. Billing settlement is handled using five-minute interval data, and these data are not provided in real-time to the ISO. Discrepancies between telemetry data and settlement data are reviewed by the ISO, but there is no defined standard for the maximum discrepancy allowed between the settlement and telemetry data. The ISO specifies meters that are allowable for settlement, and it reserves the right to inspect facilities to review how the metering is conducted.

In addition to CAISO, there are other relevant grid operators in the United States. These entities have comparable requirements to the CAISO with changes in the data reporting time (e.g., 2s or

6s telemetry for regulation in ERCOT or PJM, respectively). Networks such as ECN are also widely used, although PJM is starting to use the Internet for some services today. PJM estimates that the cost of meeting their Internet-connected system requirements is roughly $10,000 per site (PJM). Even in the cases where the Internet is allowed, the overall site enablement costs are quite high. In this work we strive to use the Internet for telemetry and keep costs low.

1.2 Smart Meters and the Home Area Network (HAN)

The California utilities have deployed an extensive network of smart meters. The primary value of these meters in the vast majority of installed situations is that the automated meter reading services allow utilities to avoid manual meter reading for billing, while also providing relatively fast, automated information on outages. These features are supported over the utility automated metering infrastructure (AMI) network, which is utility owned and operated. The overall AMI network capacity is limited, and this prevents the utility from receiving real-time power values from more than a handful of meters in a small geographical area at one time. To mitigate this limitation, California smart meters also incorporate an additional data network link. This link, known as the *home area* or *building area* network link provides access to near real-time utility meter data over a low-power, wireless link using the ZigBee Smart Energy Profile 1.0 standard. With proper configuration, this link can provide four-second power data that can be used to provide telemetry information as part of a fast demand response or ancillary services system served through an aggregator of many small loads.

1.3 State of Standards (OpenADR)

The OpenADR Alliance has developed a solution that enables fast demand response and fills the gaps left by OpenADR 1.0. It solves the communications issues and provides for scalability of demand response for fast DR without locking the resource into proprietary communications networks. It employs a hierarchical architecture in which virtual top nodes (VTN) push or allow polling of information to virtual end nodes (VEN). Typically, the VTN is a demand response automation server (DRAS) that is receiving reliability, price, or power instruction signals from grid operators that are then passed down to client VENs that are either at the load directly or at load aggregators who can then disaggregate the signal to load sites below them. The 2.0 data format has become a national standard through National Institute of Standards and Technology's Smart Grid Interoperability Panel. There are currently two profiles of the standard available: 2.0a and 2.0b. OpenADR 2.0b was specifically designed to be able to handle advanced DR capability such as ancillary services communication, both signal receipt and reporting telemetry.

OpenADR 2.0 was released in two phases, with OpenADR 2.0a released in August 2012 and OpenADR2.0b released in July 2013. OpenADR 2.0a is a slightly expanded version of OpenADR 1.0 that is conceptually similar to its predecessor. OpenADR 2.0 is the only existing open data model to exchange messages for DR events between a service provider, aggregator, and end-user in commercial, industrial, and residential markets that has international acceptance as a standard. OpenADR 2.0 is one of a suite of Smart Grid communication profiles that provides dynamic pricing, grid reliability, and transactional signals using a server-client model to convey information rather than a network-based control structure.

It took over three years and thousands of person-hours by a broad group of stakeholders to update OpenADR version 1 to OpenADR version 2. This effort changed OpenADR from a de-facto standard to an internationally recognized standard by the Organization for the Advancement of Structured Information Standards (OASIS), Energy Interop Technical Committee. The recognition of OpenADR as an official standard has several positive implications, including: (1) increased interest and adoption internationally, and (2) creation of a defined compliance process. LBNL is on the board of directors of the OpenADR Alliance and committed to its success. OpenADR 2.0 was accepted by the International Electrotechnical Commission (IEC) as IEC/PAS 62746-10-1 ed1.0 2, and by the U.S. Smart Grid Interoperability panel into its catalog of standards.

OpenADR 2.0b contains the features required to perform closed-loop control, as required to use fast DR for a variety of ancillary services. Features include the following:

1. Set point (in kilowatts [kW] or percentage)
2. Reporting services (i.e., feedback, in kW or percentage)
3. Price and scheduled price
4. Historical energy profiles, real-time power, forecast
5. Opt in/Opt out service
6. Registration service (defines where in the system hierarchy the software client resides)

The features in OpenADR 2b are designed to support the requirements to use fast DR in existing and future ancillary services markets. To that end, the OpenADR alliance that developed the standard had active stakeholder involvement from 10 ISOs and regional transmission organizations (RTOs) in North America. These organizations represent two-thirds of the electricity load, and their contributions were critical to ensure the technical requirements for wholesale markets, such as ancillary services, were met by the specification.

1.4 Summary of Work in this Task

In this task we developed a low-cost system which contains all of the elements required for providing fast demand response. We collect four-second power data using the HAN radio, and used the open source OpenSEG (Open Secure Energy Gateway) software for this service. We used an implementation of OpenADR 2.0 to provide event notification and report telemetry data back to the OpenADR virtual top node (VTN, sometimes called the “server”). We selected low-cost loads that can be actuated using available, open application programming interfaces (APIs), and we demonstrated the system in field test sites.

CHAPTER 2: Related Work

2.1 Overview of Demand Response in California

In the early 2000s (2000–2002) California experienced an energy crisis that resulted in, among other things, rolling blackouts and brownouts for many customers. Poor demand forecasting and energy market speculation were identified as the major contributing factors to these service disruptions. During that time, only large industrial customers were exposed to time-of-day usage rates. Also, only those same large industrial customers had the information necessary to respond with load-shedding procedures during a high-demand event.

In response to the systematic shortcomings related to demand response capabilities, the California Energy Commission (CEC) brought forth a series of measures that would move California away from the flat rate system, to a more modern system where more utility customers would be exposed to time-of-day use rates and be given the ability to respond appropriately by shedding load during periods of high demand. In 2001, the California legislature allocated the budget to install high-resolution (15-minute interval) meters at all customer sites that consumed 200 kW or greater. Those customers could subsequently be exposed to time-of-use rates.

In 2002, in response to the growing need for research in the area of demand response, Lawrence Berkeley National Lab (LBNL) established the Demand Response Research Center (DRRC). Since its establishment, the DRRC has worked with investor-owned utilities (IOUs), policymakers, and researchers to promote methods for improving California's capability to respond to high-demand periods.

More recently, in 2009, the DRRC, in collaboration with California IOUs, released version 1 of the Open Automated Demand Response (OpenADR) specification. This specification provides a standard template for utilities to communicate real-time price information to customers. Now in version 2.0b, OpenADR allows utilities to communicate real-time demand response events to customers. Version 2.0b also allows customers to return telemetry data back to utilities, which confirms load shed at that customer's site.

The work currently being conducted at LBNL, and reported here, seeks to pair the OpenADR 2.0b specification with a system that automatically sheds loads and reports that load shed back to the utility. The pairing of the OpenADR specification with this sort of automated response and reporting system represents a significant step forward—one that could allow for widespread automated response to DR events combined with precise information regarding the amount of load shed for any given event.

2.1.1 ARPA-E project on Fast-DR

In 2012 through early 2014, AutoGrid, Inc., in conjunction with LBNL and Columbia University, designed and demonstrated a highly distributed Demand Response Optimization and Management System for Real-Time (DROMS-RT). The project intended to show that "personalized" DR signals could be sent to millions of customers in extremely short time frames. This capability would enable customers to reduce or increase their demand, depending on grid needs. The work focused on the cloud-computing side of the enablement, but it also included field test sites. These sites typically utilized traditional, electrician-installed electricity meters and low-cost load actuation methods such as thermostats or relay-based on/off controls. However, OpenADR 2.0 was still fairly new, and the ecosystem for supporting OpenADR 2 signaling was not widely available. This project showed, however, that overall low enablement costs could be achieved with simple installations on relatively small loads. This work builds upon earlier work that used the smart meter HAN radio as a source for telemetry data and OpenADR 2.0 for all DR signaling and telemetry reporting.

CHAPTER 3:
Deployed System Architecture

The demonstration system (Figure 3) consisted of a cloud-based server and client sites. The relationship between the server and clients was one-to-many. A single server can be paired with an arbitrary number of client sites.

3.1 Hardware Overview

The cloud-based server ran Ubuntu Linux, and acted as an OpenADR virtual top node. The VTN posts the current DR signal in the case of an active event. In addition, this server received telemetry data sent from client sites. In our demonstration, the system ran on a virtual machine (VM) housed in a data center at LBNL.

Each client site consisted of an embedded system, an Open Smart Energy Gateway (OpenSEG), and the controlled loads. Together, this hardware allowed the controlled loads to respond to DR signals, and for current site-level power consumption data to be sent back to the VTN server. The VTN does not store historical data on its own; it only handles direct interactions with VENs. We used the open-source sMAP data historian to store telemetry and DR event data.

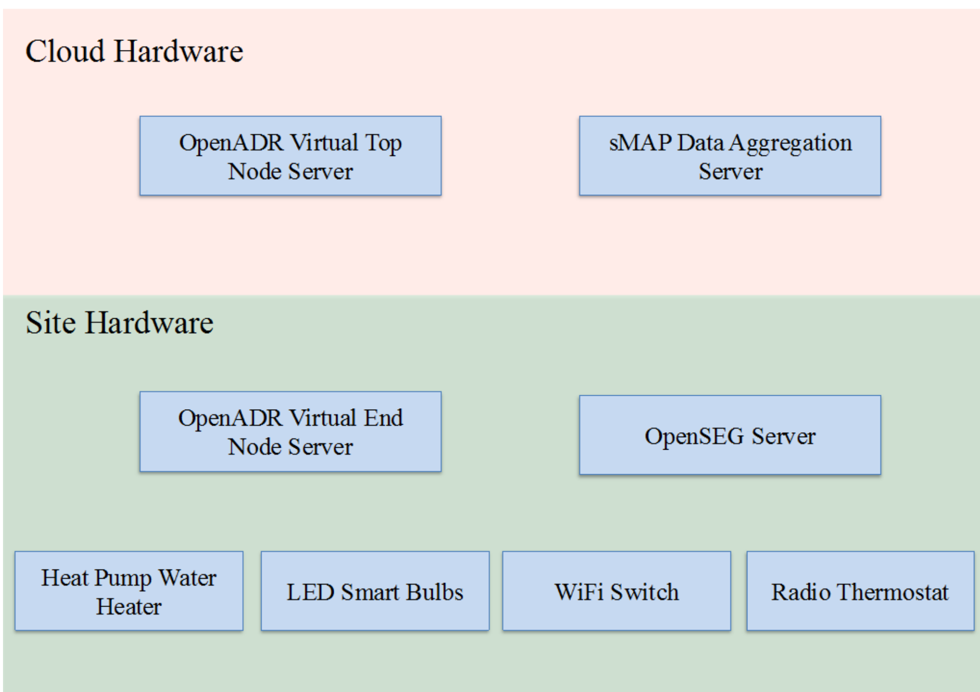

Figure 3. The demonstration system

BeagleBone Black[1]-embedded Linux computers (Figure 4) were selected to act as the central control entity for each client site. These systems act as OpenADR virtual end nodes (VENs),

[1] BeagleBone Black. http://beagleboard.org/BLACK

responsible for receiving DR signals from the VTN and replying with telemetry data gathered from the OpenSEG. A Rainforest RAVEn USB Radio Adapter (Figure 5)[2] was connected to the BeagleBone Black, and provided the wireless interface to the OpenSEG, allowing the central entities to gather telemetry data.

Figure 4. BeagleBone Black embedded Linux computer

Figure 5. Rainforest RAVEn USB radio adapter

All end use loads receive actuation signals from the BeagleBone Black in response to DR events.

[2] RAVEn™ Radio Adapter. http://rainforestautomation.com/rfa-z106-raven/

3.2 Software Overview

3.2.1 Virtual Top Node Software

The server acting as the VTN ran a web server based on the Grails Framework.[3] The primary purpose of this web server was to exchange extensible markup language (XML) documents with each client site. These documents are properly formatted to meet the specifications detailed by OpenADR 2.0b. The web server also provided a graphical user interface (GUI), which allowed users to create, modify, and remove DR events.

In addition to providing the necessary services to exchange OpenADR 2.0b documents, the VTN server sent all received telemetry data to a data archiver, which allowed tracking of the effects of the DR programs over time.

3.2.2 Site-Level Software

The software running on the central control entity at each client site (Figure 6) performed three primary duties:

1. Received DR signals from the VTN
2. Actuated connected loads appropriately at the start and conclusion of a DR event
3. Sent telemetry data from the OpenSEG back to the VTN

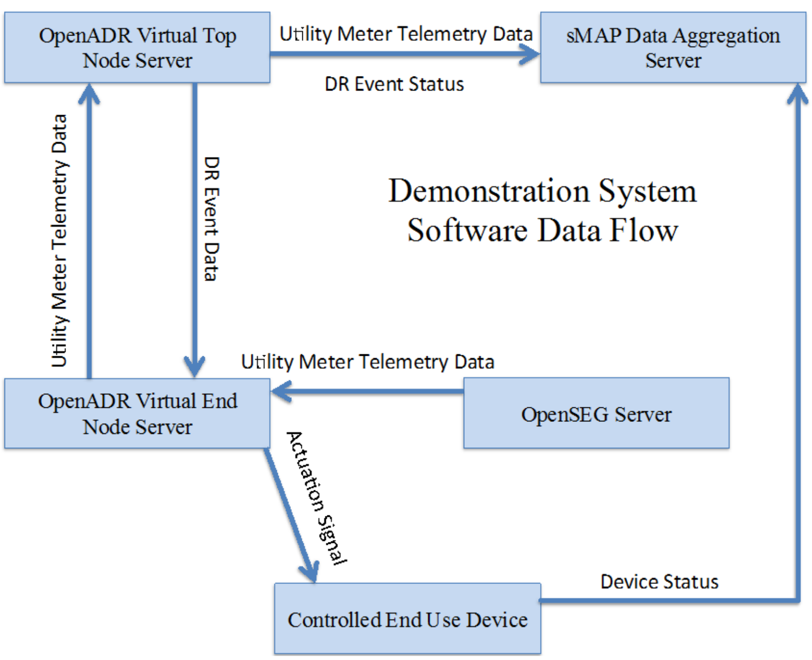

Figure 6. Software diagram

[3] Grails Framework. https://grails.org/.

A Python script polled the VTN server at regular 20-second intervals to check for active DR events. This script listened for event start and end times in addition to signal level. At each time step, the current DR signal was sent to a data archiver.

In the case that an event is detected, the signal is forward via TCP to a Python process responsible for actuating the end-use loads. This process changes the state of each controlled device in response to the DR event. The pre-event state of each device is saved, so that the device can revert to the pre-event state when the event expires. The state of each device was sent to the data archiver at each poll interval.

Finally, an HTTP server running on the control entity provided an interface to the OpenSEG. A Python process also running on the control entity polled the server via an API every 20 seconds. The server returned the current power consumption measured by the utility meter. Those data were, in turn, packaged into the correct OpenADR 2.0b format and sent to the VTN server.

3.3 Loads

Each controlled load used OpenADR 2.0 for event notification and telemetry responses. We demonstrated the ability to control four different load types:

1. WiFi Thermostat
2. Dimmable LED Bulbs
3. Binary Switches
4. Heat Pump Water Heater

3.3.1 WiFi Thermostat

This demonstration used the Radio Thermostat of America CT-80 thermostat (Figure 7).[4] These thermostats connect to networks via WiFi, and are accessible via an API over HTTP.

Figure 7. CT-80 radio thermostat

During DR events, the central control entity sent a signal to the CT-80 to either increase the temperature set point by 3 degrees if it was in the "cool" mode, or decrease the temperature set

[4] Radio Thermostat. http://www.radiothermostat.com/wifi/.

point by 3 degrees if it was in the "heat" mode. Allowing the temperature to drift up or down during DR periods decreases the use of heating and cooling systems.

3.3.2 Dimmable LED Bulbs

This demonstration used Belkin's WeMo LED Smart Bulbs (Figure 8)[5] to show how domestic lighting could respond to a DR event. The lights were paired with a Belkin WeMo Link device, which communicates with the bulbs using the ZigBee communication protocol. The WeMo Link in turn communicates over WiFi to the central control entity.

Figure 8. WeMo LED bulb

During a DR event, the central control entity dimmed all bulbs by 50 percent. Dim settings can be changed on a per-bulb basis, allowing users to further dim low-priority lighting areas, or choose not to dim bulbs where lighting is always needed at a certain level. At the conclusion of a DR event, the bulbs returned to their dim level set point prior to the beginning of the event.

3.4 Binary Switches

3.4.1 Belkin WeMo Switch

The demonstration used Belkin WeMo switches (Figure 9) to control binary (on/off) 120v loads. The switch communicates via WiFi to the central control entity. In this demonstration, we showed the switch controlling an ice maker.

[5] WeMo® Smart LED Bulb. http://www.belkin.com/us/p/P-F7C033/.

Figure 9. Belkin WeMo switch

During DR events the switch toggles off, and then toggles on at the conclusion of the event.

3.4.2 Insteon 220/240V 30A Load Controller

The Insteon 220/240v load controller (Figure 10) is well suited to be connected to a large home load such as a pool pump. In terms of behavior, it is identical to the WeMo switch. The controller toggles the load off during a DR event, and returns it to its previous state after the event's conclusion.

Figure 10. Insteon load controller

3.4.3 Heat Pump Water Heater

To demonstrate appliances responding to a DR event, a GE GeoSpring Water Heater (Figure 11) was paired with a FirstBuild Green Bean Maker Module (Figure 12). The maker module interfaced with the GeoSpring via an RJ45 cable. In turn, the maker module exposed an API to the central controller via a USB connection.

Figure 11. GeoSpring water heater

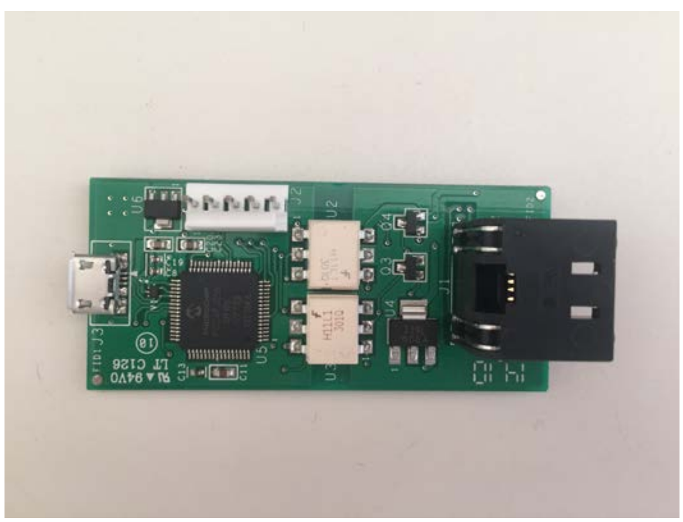

Figure 12. Green Bean maker module

Using this API, the water heater lowers its set point, and enters heat-pump-only mode during a DR event. At the conclusion of the event, the water heater is returned to its previous settings.

CHAPTER 4:
Site Installation s

4.1 Site Selection Criteria

The goal of this task was to show that open-source, low-cost software and hardware could provide fast demand response and ancillary services. Today, several large utilities offer a number of DR programs for peak capacity management and ancillary services. However, implementing new DR programs requires a substantial investment in IT equipment and personnel on the utility side, as well as telemetry and control equipment expense on the customer end. Most small utilities are unable to make such investments without knowing how successful the programs could be.

This study aimed to show that low-cost, scalable IT-based demand response is possible. This method requires a load-selection strategy that minimizes friction and costs of installation and configuration. Many factors contributed to identification of appropriate loads, and the following key factors were given the highest weighting: scalability, low-cost, security, low-latency, accuracy, and repeatability.

We selected sites that were accessible and provided access to a controllable load and a smart meter with HAN capability. Both residential and small commercial buildings in Pacific Gas and Electric (PG&E) territory have equivalent HAN (or Building Area Network, BAN) capability, and both of these site types are representative of the opportunity for low enablement cost. We selected a publically accessible demonstration site at the Berkeley Lab Guest House, a small hotel on the LBNL site. This site has a PG&E smart meter that meters only the guest laundry area, and this site was attractive because it allowed for a demonstration site that remained after the project's conclusion. We also enabled a small commercial site and a residential site that met our criteria.

4.2 Site 1: Berkeley Lab Guest House

4.2.1 Site overview

The Berkeley Lab Guest House (Figure 13) is a hotel available for guests who have business with Berkeley Lab. The Guest House is the main demonstration site for this project, and it is a test bed for the Demand to Grid (D2G) Lab in the Demand Response Research Center (DRRC).

Figure 13. Berkeley Lab Guest House

4.2.2 Installation overview

The D2G lab in the Guest House was augmented to contain many of the elements described in the system architecture section of this report (Section 4). This includes the WiFi-controlled lights, WiFi load control switch, thermostat, and heat pump water heater (and associated elements), as well as the OpenSEG, to provide real-time power data for telemetry feedback.

4.2.3 Overall cost

Table 2 assumes that either an electrician or a handyman installs all components. Many components of this system could be self-installed by a user with basic familiarity of home networking, and this would reduce system cost. The WeMo devices in particular were specifically designed to be installed and commissioned by a typical user.

Table 2. Summary of Installed System Costs

Product	Retail Cost	Installer	Install Hours	Install Cost/hr	Total Installed Cost
Insteon 220/240 Volt 30 Amp Controller	$100	Electrician	1.5	$150	$325

WeMo LED Bulb Starter Set (2 bulbs and a control hub)	$99	Handyman	0.5	$60	$129
WeMo Switch	$49	Handyman	0.5	$60	$79
Radio Thermostat CT-80	$190	Handyman	1	$60	$250
Green Bean Maker Module	$20	Handyman	0.5	$60	$50
BeagleBone Black	$45	Handyman	1.00	$60	$105
Total Cost	**$503**				**$938**

These modules can each control different size loads, but many of the elements are useful and desirable to building occupants for reasons unrelated to demand response. For example, the WeMo devices enable smart phone control of light level in a low-cost device. Each light only uses 10 W, resulting in a high cost-per-kilowatt of load shed ($3,000/kW), but the non-energy benefits of the system and potential widespread deployment of wirelessly connected, dimmable lights make them an attractive load. Heat pump water heaters that are installed for efficiency reasons can be augmented cheaply ($20) to provide demand response service provider and software-agnostic demand response capability because of its open interface. A WiFi connected version of the Green Bean interface device may increase costs to $30, but the overall cost per kilowatt remains low for this add-on (400 W to 1,200 W shed for $30, corresponding to $25 to $75 per kilowatt of shed). WiFi thermostats and other devices provide similar cost per kilowatt capability, but higher cost per kilowatt devices have increased non-energy-related benefits to consumers.

The Internet of Things and Smart Home movements are driving increased adoption of devices capable of providing fast, low-cost demand response. Open software interfaces and low-cost energy metering of controlled devices are important to leverage these new technologies for grid benefits.

4.3 Site 2: Small Commercial

We selected a 6,250 square foot (ft^2) community center in San Jose (Figure 14) as a small commercial test site. This building contains office space, a community theatre, and a community meeting space. The building is cooled by three split-system type heating, ventilation, and air conditioning (HVAC) units and has a peak electrical load of approximately 18 kW. We installed two WiFi-controllable thermostats covering approximately 65 percent of the building by floor area (Figure 15).

Figure 14. Aerial view of the community center and controlled zones

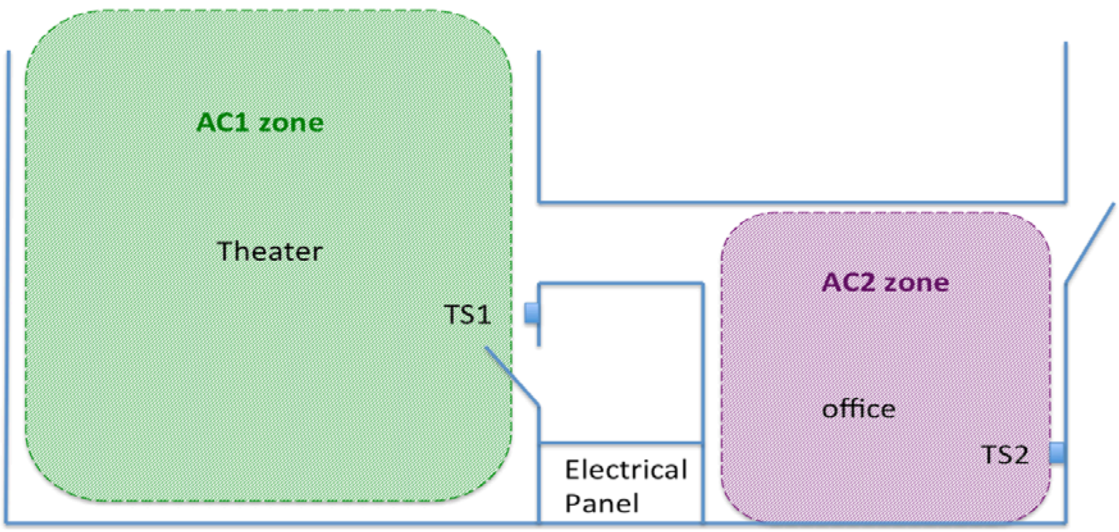

Figure 15. Community center floor schematic and thermostat placement

4.4 Site 3: Residential

Our third site was a 1,400 ft^2 residential building constructed in 1912 (Figure 16) that contains three bedrooms. In this building, we used a WiFi-connected thermostat and a WeMo switch to control loads. Telemetry data were provided using the OpenSEG.

Figure 16. The three-bedroom, 1,400 ft^2 residential site

CHAPTER 5: Test Events

5.1 Test Events

5.1.1 Test Event Setup

Test events were simulated using the devices described in Section 4. A heat pump water heater, radio thermostat, WiFi LED bulbs, and a WiFi switch connected to a fan were used as sample loads. A server at LBNL acted as an OpenADR virtual top node, and an embedded system (BeagleBone Black) acted as an OpenADR virtual end node at the test site.

The VTN broadcast DR event signals, which when received by the VEN triggered a response by the controlled loads. In turn, real-time power consumption information was collected via an OpenSEG system at the test site, packed in OpenADR 2.0b-compatible format, and sent from the VEN to the VTN. At the conclusion of each DR event, each controlled device was returned to the last seen state before the event.

5.1.2 Latency

The VEN polled the VTN at 60-second intervals. When a DR event was detected by the VEN, the process of changing the state of all connected devices took approximately 20 seconds during test events. The speed of this state change could be improved by multi-threaded software that allows for the device actuation in parallel. However, that software development was outside the scope of this demonstration. Latency could also be improved by changing the embedded system at the test site. The BeagleBone was chosen for its ease of use for rapid prototyping systems and reliability. Alternative systems developed specifically for this application would likely improve performance.

Despite room to improve the software and hardware elements of this system, it still represents a significant step forward. The time between an event being issued and the response being observed at the VTN in the telemetry data is typically under 90 seconds

5.1.3 Reliability

The primary drivers of reliability in the test system are the connected devices and the APIs used to access them. For this demonstration, custom interfaces were programmed to translate received DR signals into appropriate device behaviors. The radio thermostat and the heat pump water heater were designed by the manufacturer to allow for open access to the devices by developers. This made those two devices extremely easy to interface with, and highly reliable in their response to DR events. The WiFi switches and LED bulbs were not designed with open access in mind. This translated to slightly less reliability during testing, but improvements were continuously made during the development process. The final test setup responded as expected to DR events the vast majority of the time. No data were collected to characterize overall system reliability; however, the final system worked reliably for demonstration purposes. Reliability would be further improved by using only devices intended to be open to developers.

The exchanges between the VTN and the VEN were highly reliable. The primary driver of reliability in these exchanges is network reliability, and in testing, the DR signals and telemetry data transactions were extremely reliable.

5.1.4 Load Sheds

Several sample events were run to demonstrate example load sheds using this system. The following charts show the results of these runs. The results can be extrapolated to estimate the effect of this type of system if deployed in a residential setting.

During DR events, the LED bulbs were dimmed to 50 percent of full brightness. Figure 17 shows the two lamps in the test setup responding to a DR event.

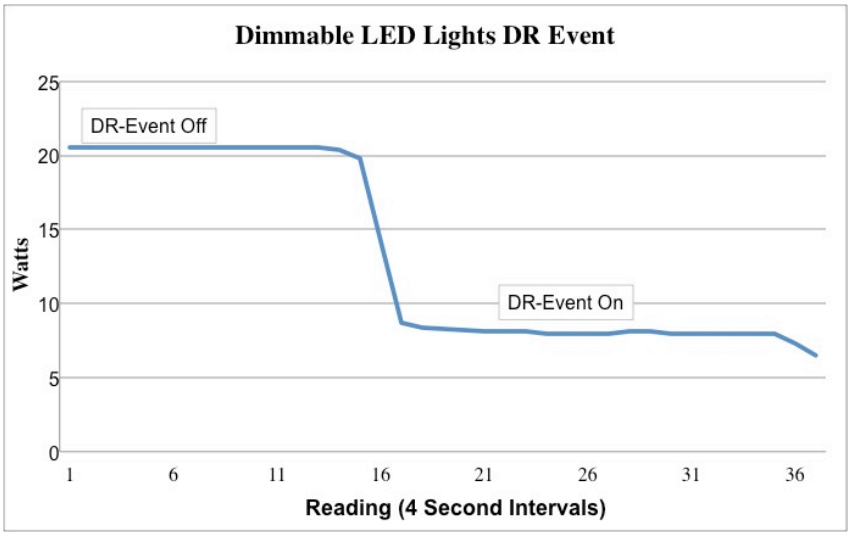

Figure 17. Dimmable LED bulbs sample DR event

Figure 18 shows a WiFi-enabled switch responding to a DR event. In the test setup, this switch was paired with a 50 W fan. In actual deployments of the system, it would be appropriate to pair this switch with any non-essential plug loads in the home. During a DR event, the switch toggles off if it was on prior to the event. After the event conclusion, it returns on only if it was on prior to the start of the event.

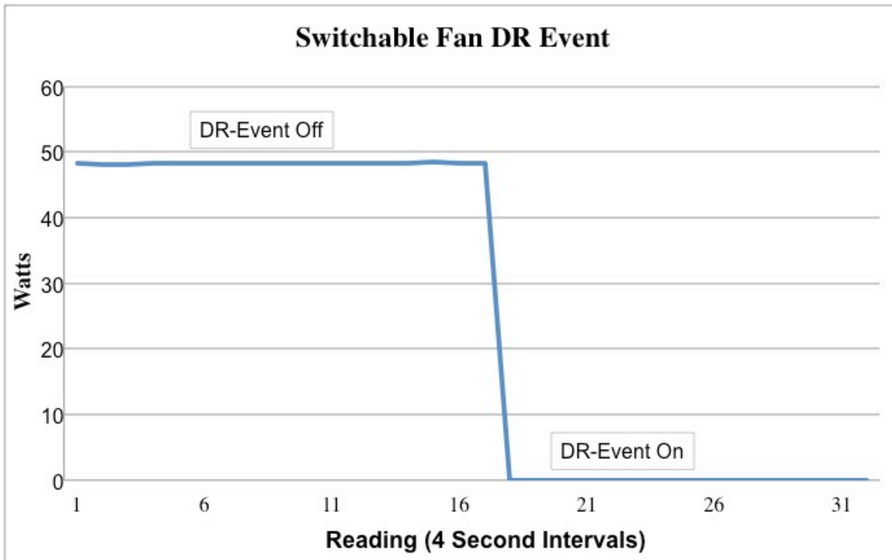

Figure 18. WiFi switch connected to 50 W fan sample DR event

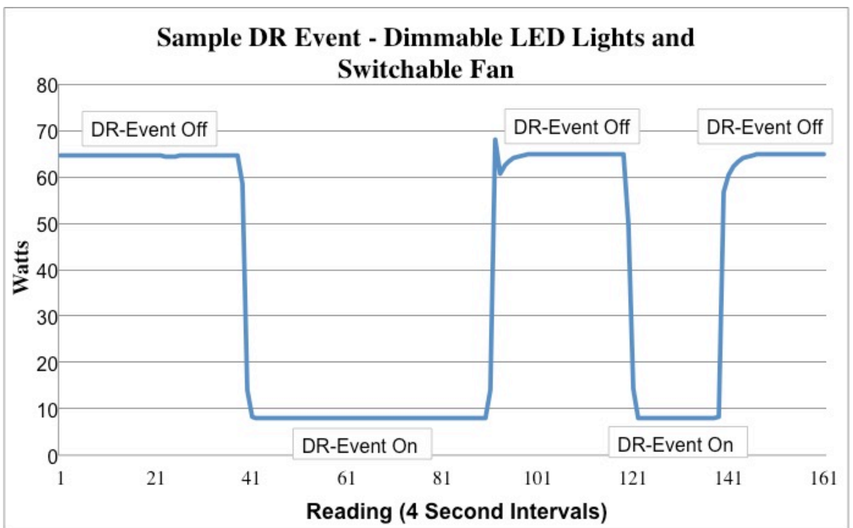

Figure 19. WiFi switch connected to 50 W fan and LED bulbs sample DR event

During a DR event, the heat pump water heater lowers its set point by 5°F. This will prevent the water heater from turning on during a DR event unless absolutely necessary. Figure 20 and Figure 21 show a normal water heater event, followed by an example where the event is interrupted by a DR signal, which results in the water heater interrupting its normal cycle and entering standby mode. At the conclusion of the DR event, the water heater returns to its previous set point.

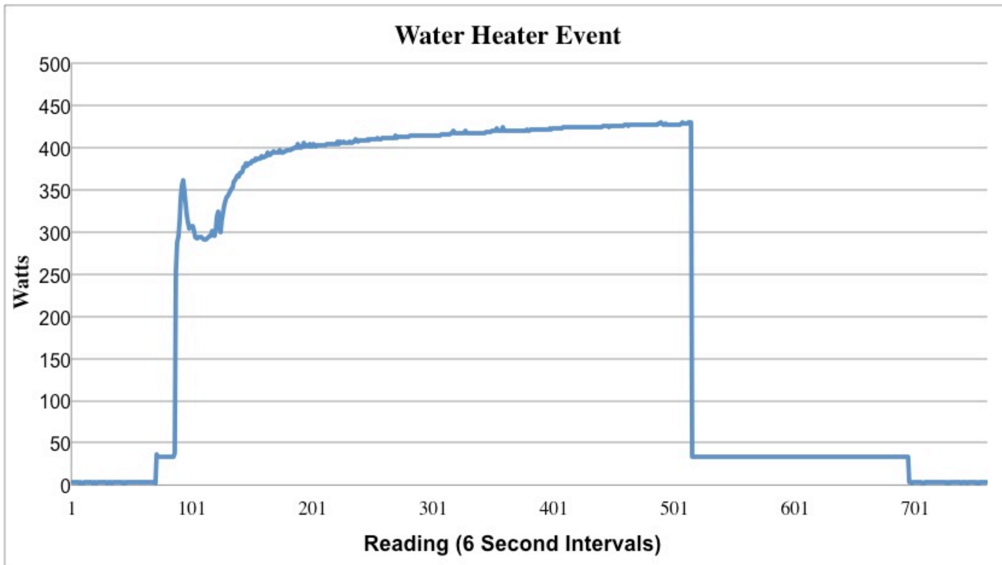

Figure 20. Heat pump water heater normal event cycle

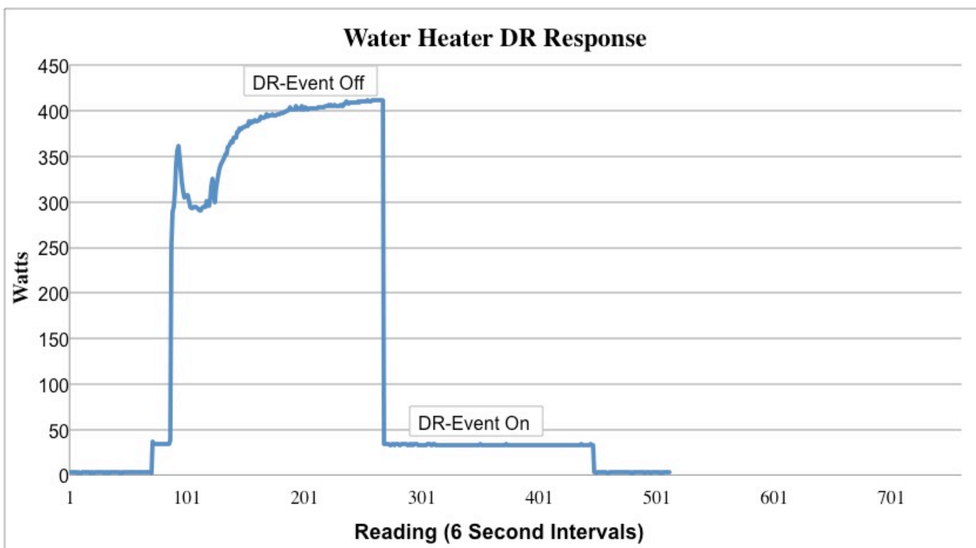

Figure 21. Heat pump water heater normal cycle interrupted by DR event

During a DR event, the radio thermostat increases its set point by 3 degrees Fahrenheit if it is in cooling mode, and decreases the set point by three degrees if it is in heating mode (Figure 22). This should result in the temporary shedding of the load associated with HVAC systems in the home. At the conclusion of the DR event, the thermostat reverts to the set point prior to the DR event.

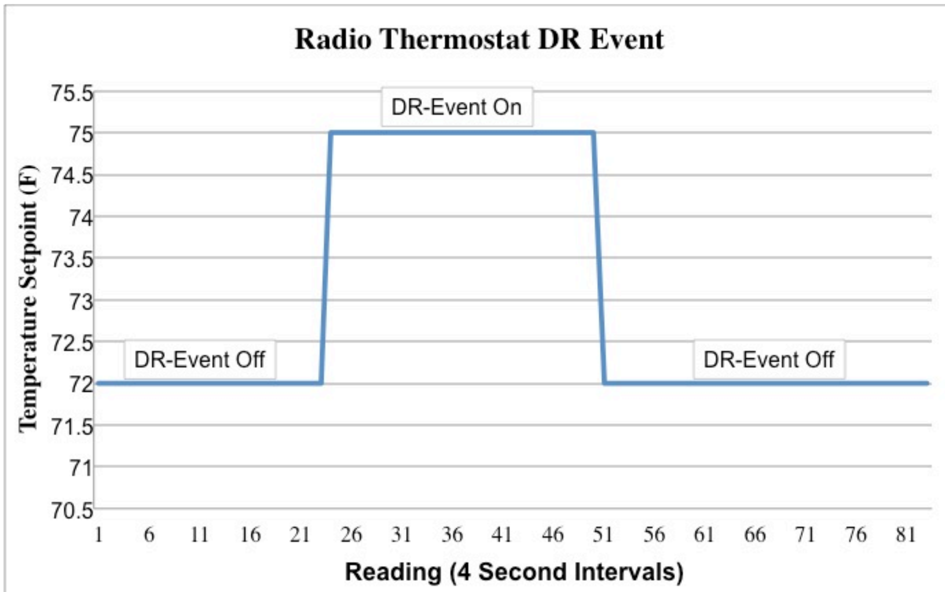

Figure 22. Radio thermostat sample DR event

Potential total shed attributable to this system is dependent on the number of controlled devices at the site and the total draw of those devices. HVAC systems and the water heater are the primary drivers of total shed potential in this setup. Table 3 summarizes the load-shedding potential of this system in a theoretical deployment in a residential setting. In the table, the small draws associated with the water heater and A/C unit represent standby power consumption.

Table 3. Potential load-shedding potential of demonstration system

Device	Normal Draw During Operation (W)	DR Event Draw (W)	Amount Shed (W)
Heat Pump Water Heater	300–400	2	298–398
LED Bulb x 2	21	8	13
WiFi Switch	1–1,400	0	1–1,400
Residential A/C	3,000–5,000	2	3,000–5,000
Total	3,400–6,800	12	3,400– 6,800

As shown by Table 3, approximately 3–7 kW could be shed by the as-built demonstration system. Increasing the number of controlled devices would increase shedding potential. However, HVAC system size dominates shedding potential in this system.

Table 4 summarizes the cost per watt shed for each controlled load.

Table 4. Cost per watt shed summary

Product	Total Installed Cost ($)	Potential Shed (W)	$/W of Potential Shed
Insteon 220/240 Volt 30 Amp Controller	328.00	2,000	0.16
WeMo LED Bulb Starter Set (2 bulbs and control hub)	129.00	18	7.20
WeMo Switch	79.00	1–1,400	79.00–0.06
Radio Thermostat CT-80	250.00	3,000–5,000	0.08–0.05
Green Bean Maker Module	49.95	400–500	0.12–0.10

Predictably, the larger controlled loads represent significantly higher cost effectiveness for load-shedding applications. A factor not included in this table is the incremental cost of these products as compared to non-controlled counterparts. The cost gap between network-enabled devices and their standard counterparts is narrowing. This is particularly applicable to the LED bulbs. When cost effectiveness is considered as a ratio between the incremental device cost and potential shed, the higher-cost items in the table above may become more attractive.

CHAPTER 6:
Conclusions and Future Work

This project's goal was to evaluate methods and the potential costs and barriers for small loads to perform fast DR in California. The term "fast DR" is defined as demand-side resources that respond without advanced notification and with fast response time (within minutes to seconds, depending on the ancillary services products' requirements). In this project we used OpenADR 2.0 for these services, deployed a smart meter home area network interface for real-time power data, and integrated multiple load control devices with our system.

We found that it is possible and potentially low-cost to use the emerging Internet of Things (IoT) and smart home devices as dispatchable loads for fast DR. Consumers are paying for these devices and deploying them in the field for quality of life or energy efficiency, but they also have potential to provide dynamic grid services. Because consumers install these devices for reasons other than demand response, relatively small incentives could increase the installed base and provide significant resources for advanced, fast DR in California.

One of the main barriers to widespread deployment of a system like that demonstrated in this project is the availability and documentation of open network interfaces for the individual devices under control. Each individual component uses a unique communication protocol, and the level of documentation of this protocol (or API, application programming interface) varies greatly from product to product. OpenADR does not fill the role of providing specific control to individual devices (although each device could support OpenADR natively, user optimization of this behavior is likely to be frustrating), and open, well-documented APIs that utilize widely available Internet standards will result in low-cost and easy device "driver" development.

We recommend that the demand response research community continues to engage with the IoT community to encourage the use of documented and open development interfaces. A library of device drivers and machine-readable interface specifications would significantly reduce the burden on users or system integrators for deploying systems in large numbers of homes in California.

REFERENCES

Kiliccote, S., P. Sporborg, I. Sheikh, E. Huffaker, and M. A. Piette, "Integrating Renewable Resources in California and the Role of Automated Demand Response," *Lawrence Berkeley National Laboratory, Berkeley, CA*, 2010. LBNL 4189E. November 2010.

Kiliccote, S., M. A. Piette, J. Mathieu, and K. Parrish, "Findings from seven years of field performance data for automated demand response in commercial buildings," *Proc. of ACEEE summer study on energy efficiency in buildings. Pacific Grove, CA*, 2010.

CAISO, "ISO Generation Monitoring and Control requirements for AGC/Regulation Units V4.6", http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.194.1264&rep=rep1&type=pdf accessed April 2012.

Lanzisera, Lanziera, Weber, Andrew, Liao , Anna, Schetrit, Oren, Killicote, Sila, Piette, Mary Ann. (Lawrence Berkeley National Laboratory). 2015.