PRELIMINARY IMPACT EVALUATION OF XCEL ENERGY'S IN-HOME SMART DEVICES PILOT BASED ON SUMMER 2012

IHSD Impact 2012 – 01

EnerNOC Utility Solutions Consulting
500 Ygnacio Valley Road
Suite 450
Walnut Creek, CA 94596
925.482.2000
www.enernoc.com

Prepared for:
Xcel Energy

Presented on:
March 8, 2013

This report was prepared by

EnerNOC, Incorporated
500 Ygnacio Valley Blvd., Suite 450
Walnut Creek, CA 94596

Project Director: C. Williamson

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION

This report covers the impact evaluation of Xcel Energy's In-Home Smart Devices (IHSD) pilot program. The program is part of Xcel's SmartGridCity^{TM} (SGC) efforts. Limited customer interest and installation issues within the city of Boulder limited the total number of installations within Boulder to just over 100. The scope of the pilot was thus modified through a 60 day notice to increase the number of installations planned outside of Boulder. Customers outside of Boulder were recruited to participate and just over 1,000 additional devices were installed in the two cities of Westminster and Centennial. Of the 100 installed in Boulder, about two thirds were installed at the homes of customers who were participating in the Pricing Pilot, and 33 were installed in homes on the standard residential rate. These 33 are included in this analysis. The devices associated with the pricing pilot are discussed in the Pricing Pilot Annual Report and are not included in this analysis. The devices were installed over the course of several months in the spring and summer of 2012.

As described in more detail in the analysis section below, the only interval data available for those participants in Centennial and Westminster was the data collected and stored by the devices. Because of this we used a split design, where about half of the customers were notified of the event each time and the other half were not. Event notification was alternated between these two groups. In this way, the two groups could serve as a control group for each other by comparing the load shapes of the group receiving notice on that day with the load shape of the group that didn't receive notice.

To implement this approach, once all devices were installed, they were assigned to one of two groups, labeled group A and group B. The goal was to split the participants randomly to ensure an unbiased comparison, but to do so in such a way that the two groups had about the same number of participants in each city and were similar demographically. For demographics, PRIZM codes were used in the split so that there was approximately the same number of customers with each of these demographic codes in each group. Within the PRIZM codes, the customers were split between the two groups randomly. This resulted in two groups that were similar, with no systematic bias inherent in the assignments.

Once the devices were in place and operating, and the participants were assigned to the two groups, events were called. For each event, only one of the groups was notified of the event, on an alternating basis. Group A was notified for the first event, but Group B was not, then Group B was notified for the second event. In this way, the group not notified in each event could be used as a control group to estimate what the load during the event would have been for the group that was notified.

In order to estimate any changes in energy use resulting from participants having devices installed in their homes, the energy use of participants had to be compared to a group of non-participants. A matched control group was selected from among the general Xcel Energy customer population in the cities with participants. For each participant, a non-participating control group customer with similar energy use in the same city was selected. This control group was then used to estimate the overall energy impacts across the whole summer. Of course, since these control group customers did not have interval data available, they could not be used to estimate the DR event impacts.

Table 1-1 shows the events that were called during the summer of 2012, between June 1 and September 30. Events began being called after the majority of devices were installed, with the first event called on July 13. The criteria for calling events was the same used for the SGC Pricing Pilot. If the forecasted temperature was above a certain threshold, the next day was considered a candidate event day. Across the summer, the candidate event days were called as actual event days two thirds of the time, following a "call-don't call-call" sequence that was repeated throughout the summer. This resulted in event days called on many but not all of the hot days. There was also one mild event day called in September, to see how customers responded on days that were not as hot. The table shows the date and day of week of each IHSD DR event, the high and low temperatures on those days, and whether the events were called for group A or group B.

Table 1-1 IHSD Event Days June 1, 2012 - September 30, 2012

Date	Day of Week	High Temperature	Low Temperature	Group Notified
7/13/2012	Friday	96	67	A
7/20/2012	Friday	101	67	B
7/23/2012	Monday	100	64	A
8/1/2012	Wednesday	94	65	B
8/8/2012	Wednesday	96	61	A
9/14/2012	Friday	77	49	B

Events start at 2:00 pm and end at 8:00 pm.

IMPACT RESULTS

This analysis of the In-Home Smart Device (IHSD) pilot covers impacts for the period June 1, 2012 through September 30, 2012. It includes analysis of the customers with IHSDs outside of Boulder (in Centennial and Westminster, collectively referred to as "Denver") and those 33 customers with IHSDs in Boulder that are not participants in the pricing Pilot, but are on the standard residential rate.

Analysis Method

This analysis employed a "difference of differences" analysis method to estimate energy savings. This method provides a savings estimation that is particularly useful for situations where there may be preexisting differences between the participants and the customers in the control group. In this method, participating customers are compared with a control group of similar but non-participating customers both during the participation period (also referred to as the treatment period) and for a time before participation started (referred to as the pre-treatment period). This allows for the calculation of the change in energy use, corrected for any preexisting differences between the participant and control groups.

Because of the need for participants and to avoid customer dissatisfaction, we did not us a classic randomized experimental design, which would have necessitated recruiting twice the number of customers to participate, and then telling half of them that they could not participate and would not receive a device. In addition, for those IHSD customers outside of Boulder, the only interval data available for the analysis was the interval data collected and stored by the devices. As a result, there was no pre-treatment interval data for participants, and no interval data at all for any representative control group customers. We dealt with this situation in two ways, using two variations on the difference of differences.

In order to estimate any overall change in energy use (generally referred to as a conservation effect), we created a matched control group for the participants. We used a pool of potential control group candidates that was much larger than the number of participant customers (approximately 10 times the number) for each city and PRIZM Lifestage code. For each participant, we selected the customer from the pool of candidates in the same city with the same PRIZM code that had energy use (billing data) that was most similar to the participant's energy use during the pre-treatment period. Savers Switch (SS) participation was also included in the match, so that IHSD customers that were SS participants were matched to control group customers who were also SS participants, and non-SS IHSD customers were matched with Non-SS control group customers. This resulted in a control group of non-participants that was the same size as the participant group, and was very similar based on geography, demographics, SS participation, and energy use. We then used a traditional difference of differences based on customer billing energy as described above, with June-September of 2011 serving as the pre-treatment period, and June-September of 2012 serving as the treatment period.

The lack of interval data for non-participants, and for the participants before the devices were installed meant that selecting a matched control group would not provide the ability to estimate the demand impacts. For this reason, we used a split design. The participants were divided into two groups, A and B. When an event was called, only one of the two groups was notified, and the other group was treated as if the day was a non-event day. The participant group was divided in order to create two groups that were as similar as possible, so that each could function as a control group for the other. First, the participants were grouped into cells based on the city in which they were located (Westminster or Centennial) and their Lifestage code. Then within each cell, the participants were split randomly into the two groups. On non-event days, the two groups have similar (but not identical) load shapes. For each event day, the event day is considered the treatment period, and the average of several non-event days with about the same temperatures is considered the pre-treatment period. The group that received notification is the treatment group, and the group not receiving notification is the control group. Then the next event day, the two groups are switched. For example, for July 20, a Group B event day, Group B was the treatment group, and Group A was the control group. July 20 was the treatment period, and the average of days with similar temperatures made up the pre-treatment period.

There is potential for confusion with terminology. During events, devices are notified of the event and controls are put in place, including changing the thermostat set points and disconnecting loads on remote plugs. This can be referred to as a customer being "controlled" during an event. Because of the potential confusion with the idea of a statistical control group, we refer to customers being "notified" of an event. When we refer to control group customers, we are talking about customers not receiving notice of the event. This also reflects the fact that not all who are notified of an event will participate – those who opt out of an event do not have their energy use controlled in any way.

The average savings across all the hot event days was calculated based on the first five event days. There was no need for an adjustment, since both the participant (notified) group and the control group included data from both groups, averaged together with appropriate weights. Because of this, any pre-existing differences between the two groups cancel out. This estimate also does not depend on finding a similar non-event day, and so is a more robust estimate. We also include the one mild day as typical of what can be expected if events are called on mild days. Because this was based on only one day, September 14 when Group B was notified, we included the correction for the difference between the two groups on similar mild days.

Data Used in the Analysis

The DR analysis was conducted using the 15-minute interval energy usage data collected from the devices installed at customers' homes, and billing data extracted from the Xcel Energy billing system. Table 2-1 below shows how many customers were included in the DR analysis for each event. Nine Group A and four Group B customers had no data at all for the summer. Not all customers were used for all events – Some customers were missing data for only part of the summer, and those were used for the event days when their data was available, but excluded when their data was missing.

Table 2-1 IHSD Customers Used for DR Analysis

Event Day	Group A			Group B		
	All customers	At least some data available	Used for event day	All customers	At least some data available	Used for event day
July 13, 2012	509	500	399	513	509	407
July 20, 2012	509	500	405	513	509	423
July 23, 2012	509	500	406	513	509	427
August 1, 2012	509	500	430	513	509	444
August 8, 2012	509	500	444	513	509	452
September 14, 2012	509	500	452	513	509	459

The energy analysis was conducted using billing data for the summer of 2011 and the summer of 2012. The 909 customers with complete billing data available for both summers were used in the analysis. A matched control group of 909 customers, based on summer 2011 (pre-treatment) billing data was selected from a bigger pool. The difference of differences analysis was then done based on the summer billing data for the two summers.

2012 Preliminary Results

There are two types of impacts that are estimated for the IHSD pilot. During events, customers may reduce their energy use in response to notification. Customers may also change their behavior in general, reducing their overall energy consumption. We estimate these two impacts separately, and this results section is split into two parts, with the overall energy (consumption) impact shown first, and then the demand response impacts. As described above, the energy impacts are estimated based on a matched control group, and the demand response impacts are estimated using a split participant group.

The analysis shows that while there appears to be some overall energy savings of about 3% across the four summer months, the savings are not statistically significant based on the analysis done. On the DR side, however, there are clearly significant savings during hot events, with small but significant pre-cooling and significant snapback after the event.

It is difficult to compare the DR savings found here to other pilots or programs in the industry, because to our knowledge, there are not other programs which, like this pilot, notify customers of events, but allow customers to opt out of events, and have no financial incentive or penalty associated with the use of the device. That said, the shape of the DR impacts, with the pre-cooling and snapback, are consistent with other programs using a similar technology to raise thermostat set points. The magnitude of the load reductions is somewhat smaller than other programs, but that could be due to the lack of financial incentives or penalties, or could be due to lower AC usage in Colorado.

The level of the estimated energy savings is comparable to other programs and technologies that involve providing the customer with information about their energy use.

Energy Consumption Impacts

In this section, we present the analysis of the energy impacts resulting from access to the devices. We first show the information graphically, then in a table. As discussed above, this is based on a matched control group, and the pretreatment period of the summer of 2011 and a treatment period of the summer of 2012. While it appears that there may be energy savings resulting from the presence of devices, the differences are not statistically significant. Figure 2-1 shows the four monthly savings estimates and their 90% confidence intervals. The 90% confidence intervals for each month include zero, meaning the savings are not statistically significant.

Figure 2-1 Energy Savings Estimates for Summer 2012

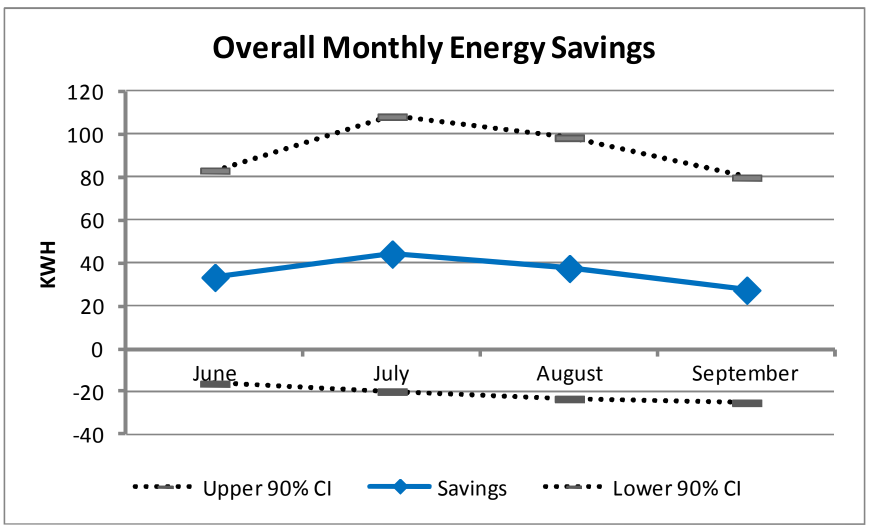

Table 2-2 shows the adjusted control group, the savings, and the percent savings for each month. The savings estimates are about 3% of the adjusted control group energy use. Again, these are not statistically significant, so should not be considered as valid savings estimates. There is not sufficient evidence from this analysis to say that customers with IHSDs are conserving energy as a result of having the devices. It is important to note that this does not necessarily mean there were no overall energy savings – it just means that this analysis was not able to detect statistically significant savings. During the final impact analysis of the IHSD pilot, we will have more data and plan to look into potential alternative analysis methods that may show significant savings.

Table 2-2 Energy Savings Estimates

	Adjusted Control Group Billing Energy	Estimated Savings	Percent Savings	Significant Savings?
June	924	33	3.6%	No
July	1,269	44	3.5%	No
August	1,157	38	3.2%	No
September	959	27	2.9%	No

Demand Response Impacts

In this section, we present the analysis of the demand response impacts resulting from calling events. These results are for the six event days from Table 1-1. Table 2-3 shows the estimated average demand reductions across the event period for each event, along with the average of the five hot event days. The average demand reduction during hot events ranged from a low of 7.6% and a high of 19.2%. With only five event days, it is difficult to characterize any relationship between temperature and percent savings – the highest percent savings did not occur on either the warmest or the mildest of the hot days. With more data for the summer of 2013, it may be possible in the future to see a relationship. The load reductions were statistically significant on all but the last event day, which was the one mild day with a high of only 77 degrees. The August 1 event day also exhibited some unusual characteristics, with a control group load and corresponding savings estimates lower than expected, which may warrant further investigation.

Table 2-3 Demand Response Impact Estimates

	Adjusted Control Group Average kW	Estimated Average kW Reduction	Percent Reduction	Event Day High Temperature	Significant Savings?
July 13, 2012	2.85	0.37	13.0%	96	Yes
July 20, 2012	3.40	0.29	8.7%	101	Yes
July 23, 2012	2.93	0.29	9.9%	100	Yes
August 1, 2012	2.01	0.15	7.6%	94	Yes
August 8, 2012	2.68	0.52	19.2%	96	Yes
September 14, 2012	0.98	0.02	1.9%	77	No
Average of hot events	2.75	0.31	11.3%		Yes

Figure 2-2 and Figure 2-3 below show the average load shapes and savings for an average hot event day, based on the first five events. The load shape graph shows the hourly load shapes for the control and treatment groups. The control group load shape is an average of the non-notified customers across the five hot event days (for some events, this is Group B, for others it is Group A). The treatment group load shape is the average of the notified customers across the five hot event days. The savings graph shows the estimated savings, along with the associated 90% confidence intervals as dotted lines on either side of the estimate. When the confidence interval includes zero, the savings are not statistically significant, but where the confidence interval does not include zero, we can say that there are statistically significant savings during those hours. Because the values on the graph represent savings, a negative value means that the participating customer used more energy.

Note that the two loads are nearly identical overnight, but then there is a slight difference in the hours leading up to the event which might be pre-cooling, which is small but statistically significant. Then during the event there is a clear reduction in load that lasts throughout the event, peaking near the middle, and dropping off somewhat in the last three hours of the event. These event savings are all highly significant, with little doubt that the savings are different from zero. After the event, there is a period of higher usage, which shows up as negative savings, which represents the "snapback" effect when the customers' AC units use more energy to "catch up" after the event. The snapback is statistically significant as well.

Figure 2-2 Average Hot Event Day Load Shapes

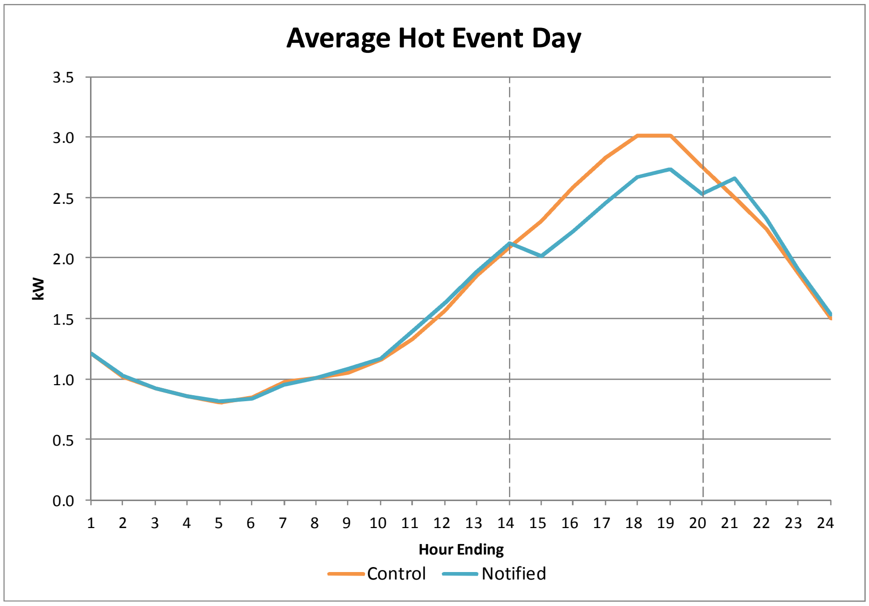

Figure 2-3 Average Hot Event Day Savings

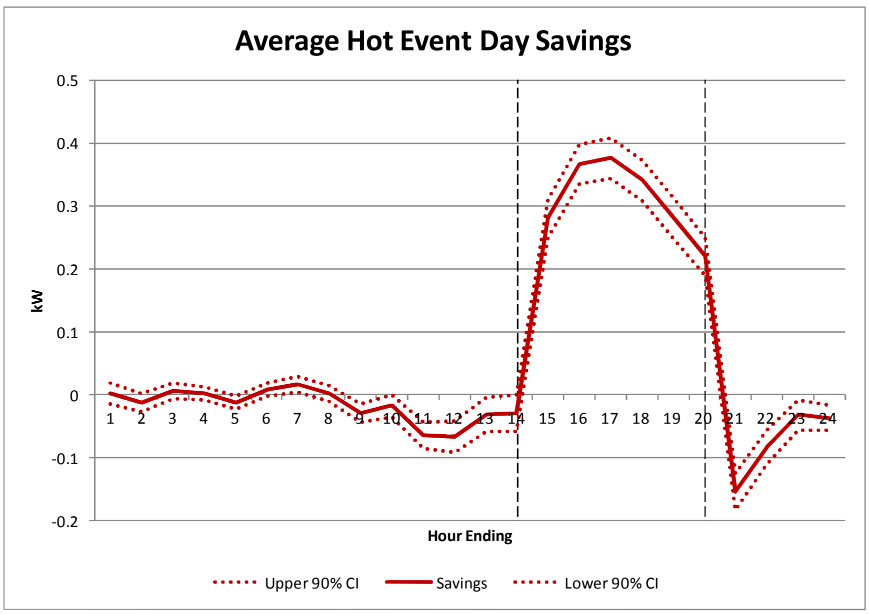

Figure 2-4 and Figure 2-5 below show the average load shapes and savings for the mild event day. Because there was only one event called on a mild day, this load reflects that single event day, September 14. The control group load shape is the Group A load, adjusted for similar non-event day differences between the two groups. The treatment group load shape is the average Group B shape, since for this event, Group B was notified. The savings graph shows the estimated savings, along with the associated 90% confidence intervals as dotted lines on either side of the estimate. The scale on the savings graph here is the same as the scale on the hot event day savings graph above, to facilitate comparison. When the confidence interval includes zero, the savings are not statistically significant, but where the confidence interval does not include zero, we can say that there are statistically significant savings during those hours. Because the values on the graph represent savings, a negative value means that the participating customer used more energy.

None of the savings on this mild event day are statistically significant. This is not surprising, as there is little or no AC usage because of the mild temperature, so there is not much available load to reduce. It appears that there might be savings during the event, particularly later in the event, but it is hard to say if that is meaningful or just random variation.

Figure 2-4 Mild Event Day Load Shapes

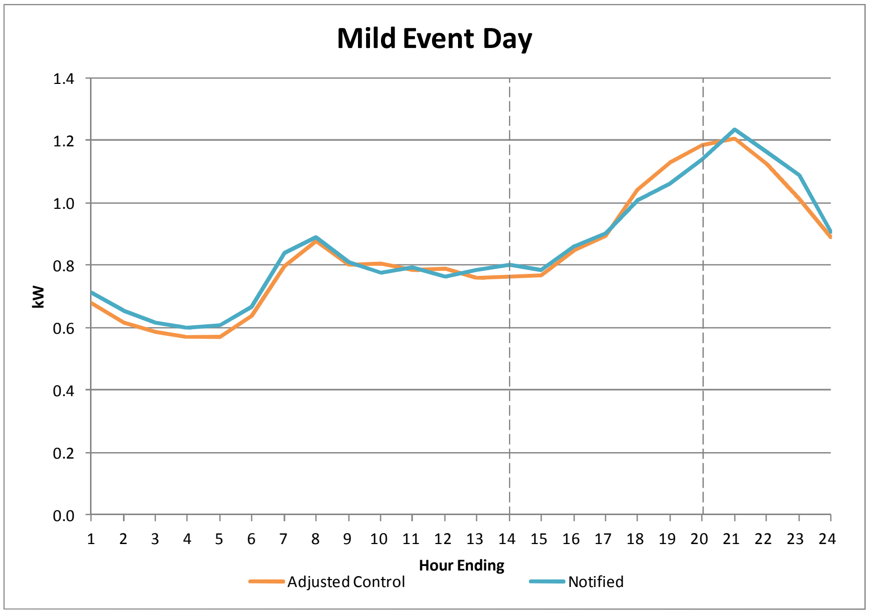

Figure 2-5 Mild Event Day Savings

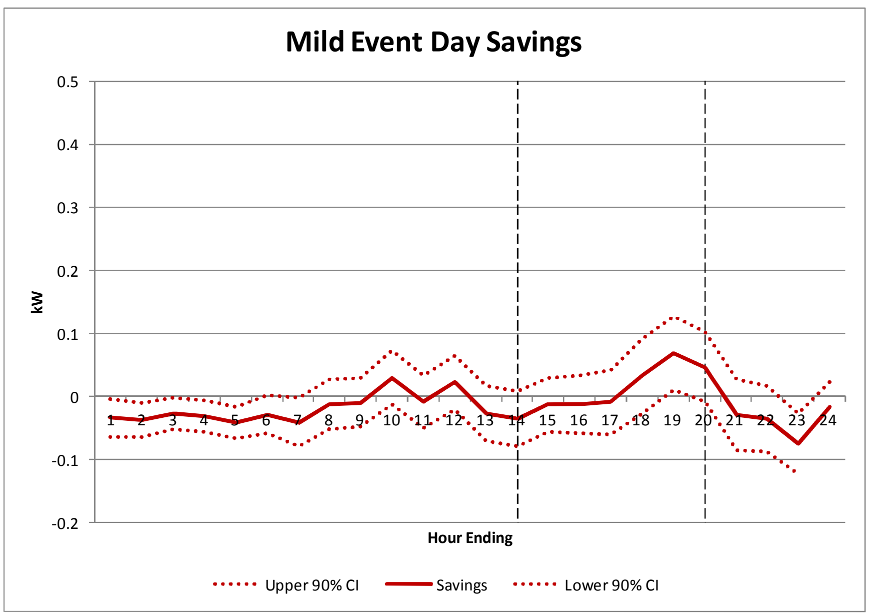

Load shapes and savings graphs for the individual events are included in Appendix A.

CHAPTER | 3

KEY FINDINGS AND RECOMMENDATIONS

The following were key findings of the IHSD impact analysis:

- There was a statistically significant load reduction during DR events on hot days, averaging 0.31 kW or an 11.3% reduction across the six hours of the five hot event days.
- The average demand reduction during hot events ranged from a low of about 8% and a high of over 19%; there was no clear relationship between the percent savings and temperature.
- During the one mild event day, the estimated average reduction was only 1.9%, and was not statistically significant. This is most likely due to the lack of AC load to reduce during this event, and if so, implies that most of the savings result from the change to the thermostat set point.
- While there was some evidence of overall energy savings of about 3% for the four summer billing months of 2012, the estimates of those savings were not statistically significant. Adding data for summer 2013 and investigating alternative analysis methods may well result in estimates that are statistically significant, but we won't know that until the final analysis period.

We recommend the following change to the operation of the IHSD pilot:

- Call DR events on all pricing pilot candidate event days, including the "don't call" days. This will result in more DR event days, which will result in more precise savings estimates across the multiple event days. Because groups A and B act as control groups for each other, there is less of a need for comparable non-event days.

We recommend the following additional analysis for the final impact analysis of the IHSD pilot:

- Test regression models and energy stratification for the energy analysis to see if more precise estimates of savings might result in statistically significant savings.
- Include consideration of the participants' actions in the analysis, specifically separating the participants based on whether customers opted out of events or not, and whether systems were online during the event. This would separate those customers that are likely to reduce demand during an event from those that are not, making the estimation of the reduction more precise.

INDIVIDUAL EVENT DAY GRAPHS

This Appendix includes the load shape graphs and savings graphs for the six individual event days. The first graph for each event day is a comparison of the load shape of the notified group with the adjusted load shapes for the group that was not notified. The second is the estimated savings based on a difference of differences using an average of comparable non-event days as the pre-treatment period.

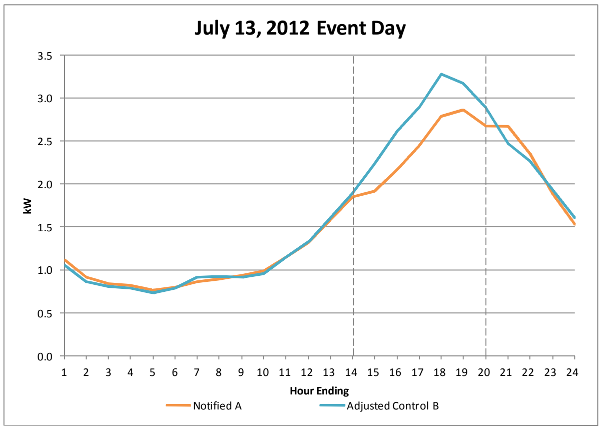

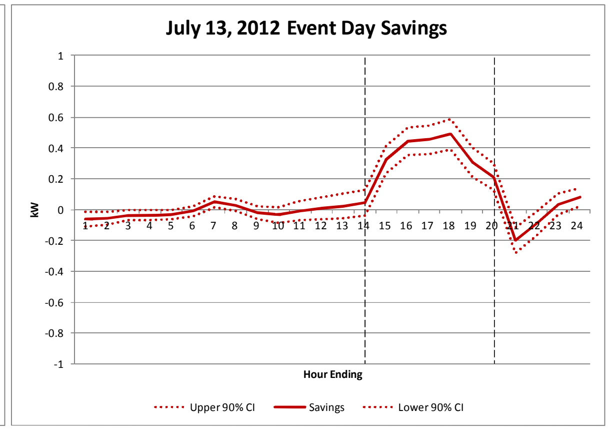

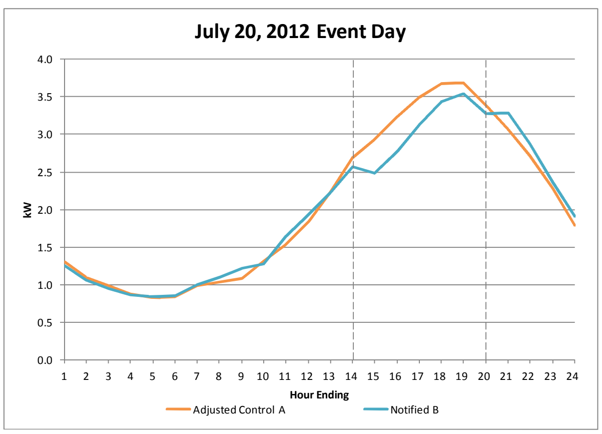

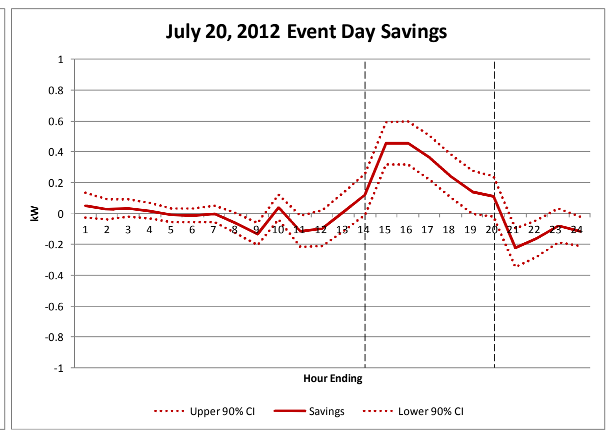

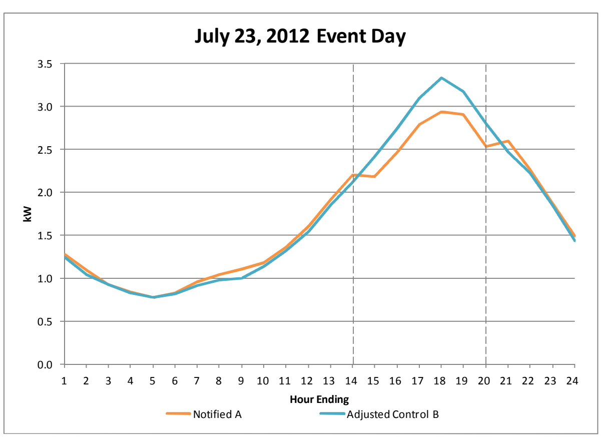

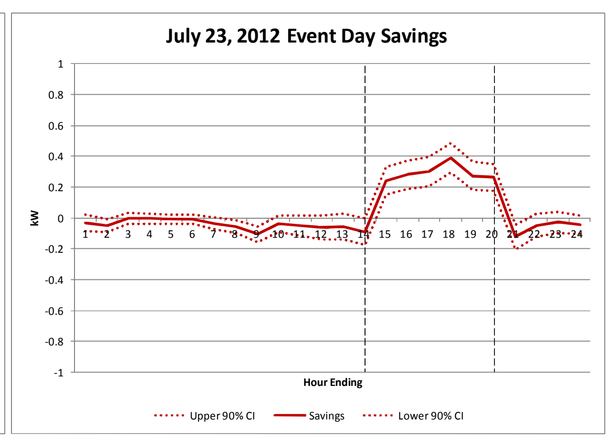

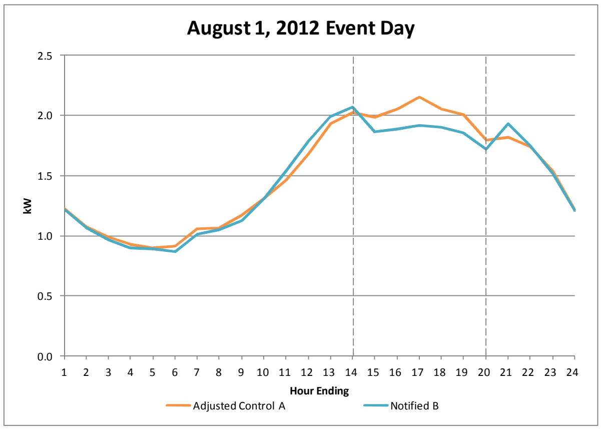
August 1, 2012 Event Day
kW
Hour Ending
Adjusted Control A
Notified B

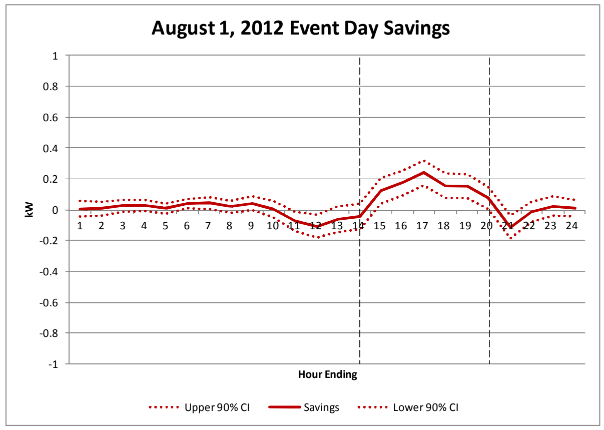
August 1, 2012 Event Day Savings
kW
Hour Ending
Upper 90% CI
Savings
Lower 90% CI

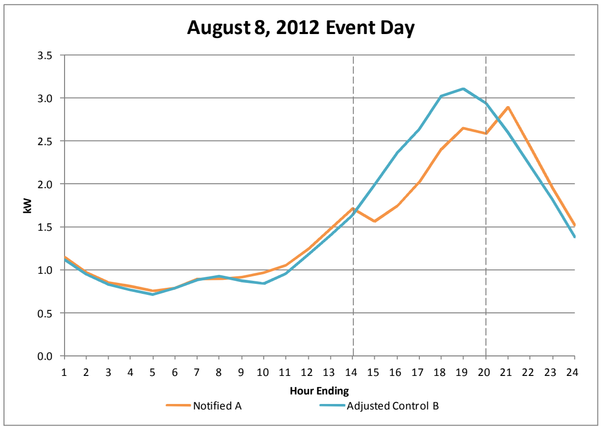
August 8, 2012 Event Day
kW
Hour Ending
Notified A
Adjusted Control B

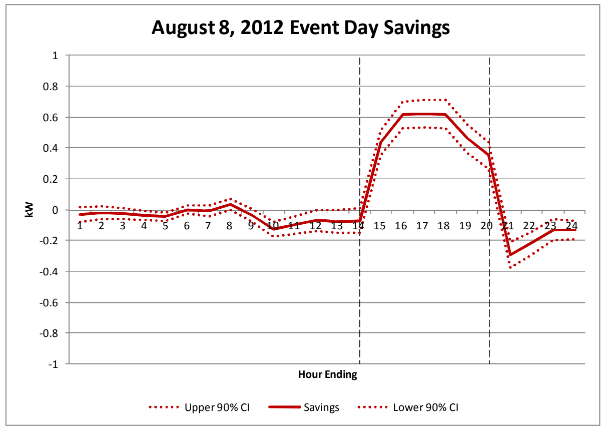
August 8, 2012 Event Day Savings
kW
Hour Ending
Upper 90% CI
Savings
Lower 90% CI

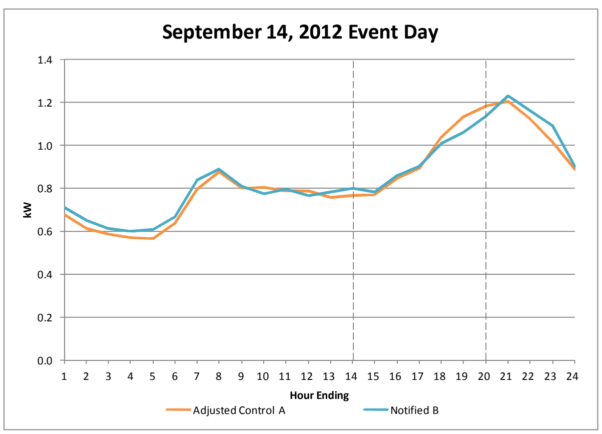
September 14, 2012 Event Day
kW
Hour Ending
Adjusted Control A
Notified B

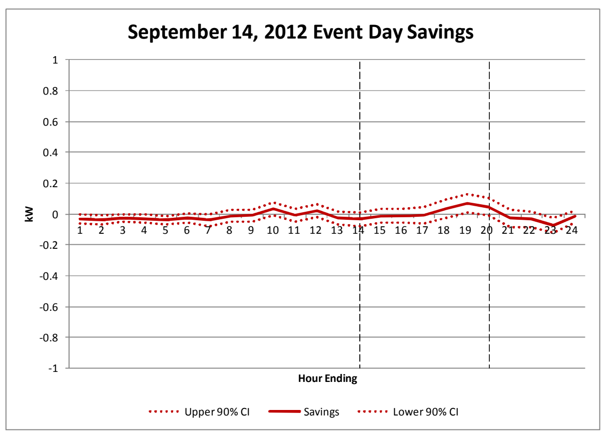
September 14, 2012 Event Day Savings
kW
Hour Ending
Upper 90% CI
Savings
Lower 90% CI

About EnerNOC Utility Solutions Consulting Team

Our Consulting team is part of EnerNOC's Utility Solutions, which provides a comprehensive suite of demand-side management (DSM) services to utilities and grid operators worldwide. Hundreds of utilities have leveraged our technology, our people, and our proven processes to make their energy efficiency (EE) and demand response (DR) initiatives a success. Utilities trust EnerNOC to work with them at every stage of the DSM program lifecycle – assessing market potential, designing effective programs, supporting the implementation of the programs, and measuring program results.

The Consulting team provides expertise and analysis to support a broad range of utility DSM activities, including: potential assessments; end-use forecasts; integrated resource planning; EE, DR, and smart grid pilot and program design and administration; load research; technology assessments and demonstrations; EE project reviews; EE and DR program evaluation; and regulatory support.

The team has decades of combined experience in the utility DSM industry. The staff is comprised of professional electrical, mechanical, chemical, civil, industrial, and environmental engineers as well as economists, business planners, project managers, market researchers, load research professionals, and statisticians. Utilities view EnerNOC's experts as trusted advisors, and we work together collaboratively to make any DSM initiative a success.

EnerNOC Utility Solutions Consulting
500 Ygnacio Valley Road, Suite 450
Walnut Creek, CA 94596

P: 925.482.2000
F: 925.284.3147