Demand Response Opportunities and Enabling Technologies for Data Centers: Findings From Field Studies

Authors:
Girish Ghatikar, Venkata Ganti, Nance Matson, and Mary Ann Piette
Lawrence Berkeley National Laboratory

August 2012

DISCLAIMER

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

Acknowledgements

This work described in this report was coordinated by the Lawrence Berkeley National Laboratory (LBNL) Demand Response Research Center (DRRC) and funded by the California Energy Commission (Energy Commission), Public Interest Energy Research (PIER) Program, under Work for Others Contract No. 500-03-026; Pacific Gas and Electric Company (PG&E), and San Diego Gas and Electric (SDG&E) under Work for Others Contracts POPG03-L01 and POSD01-L01; and by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

The authors would like to thank the following for supporting the study:

- The subcontractors, Megawatt Consulting and SCG for their support with recruitment, technical assistance, and design of DR strategies and conducting the tests.
- Janie Page, Roozbeh Mahdavi, Henry Coles, Steve Greenberg, Bill Tschudi, and Dale Sartor of LBNL for their data center expertise, assistance with the strategies and tests, and review.
- Greg Kurtzer, Gary Jung, Karen Fernsler, and Yong Qin of the LBNL High Performance Computing Services Group and the use of LBNL's Lawrencium cluster and review.
- Ed Ritenour of the LBNL Building 50 1275 (50B) data center team for assistance with the DR strategies and tests.
- Steve Aguirre of University of California, Berkeley and The University of California Office of the President for the use of their Shared Research Computing Services clusters.
- Stan Cox, Dave Shroyer, Victor Ifediora, and Pridhvi Appineni of Network Appliance (NetApp) data center for assistance with the strategies and tests.
- Matthew Campbell and Todor Milkov of San Diego Supercomputer Center (SDSC) for assistance with the strategies and tests.
- Clemens Pfeiffer, Tony Montalban, and Jim McCray from Power Assure for providing the resources and enabling technologies support and data for NetApp and SDSC.
- Patricia Nealon and Garret Smith from SynapSense Corporation for providing the resources and the enabling technologies support and data for LBNL 50B

The authors also want to acknowledge all others who assisted this project, and for their ongoing support, including California Energy Commission Public Interest Energy Research (CEC/PIER) Anish Gautam and Paul Roggensack, PG&E's Albert Chiu and Jonathan Burrows, and SDG&E's Eric Martinez. We would also like to thank the members of our Technical Advisory Group for their guidance, comments, and review.

DISCLAIMER

This report was prepared as the result of work sponsored by the California Energy Commission. It does not necessarily represent the views of the Energy Commission, its employees or the State of California. The Energy Commission, the State of California, its employees, contractors and subcontractors make no warrant, express or implied, and assume no legal liability for the information in this report; nor does any party represent that the uses of this information will not infringe upon privately owned rights. This report has not been approved or disapproved by the California Energy Commission nor has the California Energy Commission passed upon the accuracy or adequacy of the information in this report.

Table of Contents

Executive Summary **1**

1. Project Goals and Background **2**
- **1.1. Goals and Objectives** **2**
- **1.2. Background and Data Center Energy Use** **2**
- **1.3. Report Organization** **3**

2. Methodologies **3**

3. Data Center Characterization **4**
- **3.1. Participating Data Centers and Functions** **5**
- **3.2. Enabling Technologies** **5**
- **3.3. Load Profile Characterization** **6**

4. Demand Response Strategies **7**
- DR Strategy 1: Server and CRAC unit shutdown 8
- DR Strategy 2: Load shifting or queuing IT Jobs – Server idling 8
- DR Strategy 3: Temperature set point adjustment 9
- DR Strategies 4 and 5: Cooling relative to IT equipment load reduction 10
- DR Strategy 6, 7: Load migration between homogeneous and heterogeneous systems 11

5. Field Test Analysis and Results **14**
- **5.1. LBNL Building 50 Data Center** **15**
  - Field Tests: Baseline and Demand Savings Analysis 15
- **5.2. NetApp Java-1 Data Center** **17**
  - Field Tests: Baseline and Demand Savings Analysis 17
  - Economic Analysis 20
- **5.3. University of California Berkeley and San Diego Super Computer Center** **20**
- **5.4. San Diego Super Computer Center and LBNL 50B** **21**

6. Conclusions and Future Research **22**
- Conclusions 22
- Future Research 23

7. Glossary of Key Terms **24**

8. References **26**

9. Appendices **27**
- **9.1. Research Methodology** **27**
  - TAG Members and Roles: 27
  - Site Selection Criteria and Framework 27
  - Technology Selection Criteria and Framework 28
  - Assessment, Outreach, and Site Recruitment 28
  - Develop Custom Field Tests and Monitoring Plans 30
- **9.2. Data Center Characterization** **31**
  - Equipment description 31
- **9.3. Results** **34**
  - Baseline Analysis 34
  - LBNL Building 50B 1275 Data Center 34
  - NetApp 41
  - Field Tests: Baseline and Demand Savings Analysis 41
  - Economic Analysis 43

Executive Summary

The overall goals of the study are to improve the understanding of demand response (DR) opportunities in data centers and evaluate an initial set of control strategies in field tests. The study is intended to accelerate the applicability and adoption of DR in data centers and relate the findings to similar data centers. This is a phase two study; a follow-up to the phase one study completed by Lawrence Berkeley National Laboratory (LBNL) in 2010. The driving factor for these studies are (1) increasing data center energy use and cost, and (2) the impact of data center energy use during peak periods and high prices. According to a U.S. Environmental Protection Agency (EPA) forecast for 2011, 20% of U.S. data center energy use is in the Pacific region alone. In the California Pacific Gas and Electric Company utility territory, data centers are estimated to consume 500 megawatts of annual peak electricity. The recent 2011 data show that there is an increase in U.S. data center energy use, although it is slightly lower than the EPA forecast.

A technical advisory group (TAG) was recruited to review findings and provide guidance. The TAG members represented the data center industry, utilities, regulatory bodies, and participating data centers. The TAG provided insights to data center operations and real-world strategies to enable DR in data centers. The site and technology selection criteria and framework were prepared to recruit data centers. A comprehensive outreach was conducted to recruit data centers that met the criteria. A total of four data centers were recruited for DR tests—LBNL; NetApp, Inc.; San Diego Super Computer Center; and the University of California Berkeley. A site survey questionnaire and monitoring and test plans were prepared to characterize the data centers, identify DR strategies, and conduct the field tests. The field test results were analyzed and presented.

The results from the implementation of manual DR strategies for both IT equipment (server, storage) and site infrastructure (cooling) identified that similar data centers can participate in DR programs with no impact to operations or service-level agreements, which were set by the data center operators. In certain instances, a small IT equipment load shed strategy produced a demand savings of 25% at the data center level or 10% to 12% at the whole building level. As data center site infrastructure becomes more efficient with lower power usage effectiveness (PUE), a large DR opportunity is within the IT equipment, with larger savings when combined with DR automation and integration with the cooling systems. With the growth in cloud computing, the load migration strategies that are unique to data centers can enable participation in DR and price-responsive programs such as ancillary services and address intermittency issues with renewables and other issues with minimal or no impact to data center's day-today operation.

However, some key barriers remain for data centers to participate in DR programs, and future studies must consider them. Although the results are encouraging, it represents a small dataset. More field-tests of component-level performances, data center types, and automation is necessary for wide scale DR adoption. The enabling technologies need to link data center operational requirements with the supply side systems, and provide aggregation to visualize metered DR information. The evolving DR market design must describe the value proposition, the measurement and verification models, and how it can benefit a data center's social, economic, and efficiency goals.

1. Project Goals and Background

1.1. Goals and Objectives

The overall goals of the study are to improve the understanding of the demand response (DR) opportunities in data centers and evaluate an initial set of control strategies in field tests.[1] Demand response is a set of actions taken to reduce electrical loads when contingencies, such as power grid emergencies or congestion, threaten the electricity supply-demand balance and/or when market conditions cause the cost of electricity to increase. The study is intended to accelerate the applicability and adoption of DR in data centers and relate the findings to similar data centers. Through field tests, the specific study objectives are to evaluate and improve the understanding of:

- the feasibility and adoption of DR in data centers by exploring practical opportunities, and perceived versus actual risks and methods to overcome risks;
- a set of potential DR strategies for site infrastructure (e.g., cooling, lighting), and information technology (IT) equipment (e.g., servers, storage), including the use of load migration strategies, used for disaster recovery, for DR; and
- enabling technologies for IT equipment and control systems, and the needs to enable DR automation.

1.2. Background and Data Center Energy Use

This is the second phase of a two-phase study, following on the phase one study completed in 2010 (Ghatikar et al. 2010). The increasing data center energy use and cost, in particular, the impact during peak periods and high prices are driving factors for these studies. According to a 2007 U.S. Environmental Protection Agency (EPA) report, 20% of the data center energy use is in the Pacific region alone (EPA 2007). In Pacific Gas and Electric Company (PG&E) territory, data centers are estimated to consume 500 megawatts (MW) of peak electricity annually A recent study has shown that there was a significant increase in data center energy use in 2011, by 85.6 billion kilowatt-hours (kWh). This was slightly lower than the EPA forecast of more than 100 billion kWh by 2011 (Koomey 2011).

The phase one study examined data center characteristics, loads, control systems, and technologies to identify DR and Automated DR (Auto-DR) opportunities and challenges. The phase one study was coordinated with experts, technology vendors, and data center operators. The ongoing research on commercial and industrial DR was collected and analyzed (Motegi et al. 2007). The results from the phase one study suggested that data centers, with increasing and rapidly growing energy use, have significant DR potential. Because they are highly automated, they are excellent candidates for Auto-DR. Of the data center types, the "non-mission-critical" data centers are the most likely candidates for early adoption of DR, although "mission critical" data centers can also participate if the data center operators find it feasible.[2] Data center site infrastructure DR strategies

[1] Typically, in data center terminology *DR* refers to Disaster Recovery. In this report, we use the acronym DR in reference to Demand Response.

[2] "Non-mission-critical" data centers do not provide production-level services. These data centers are typically used for back-ups, research and development, high-performance supercomputing, etc.

have been well studied for other commercial buildings (Motegi et al. 2007); however, DR strategies for IT equipment have not been studied extensively. The largest opportunity for DR or load reduction in data centers is in the use of virtualization to reduce IT equipment energy use, which correspondingly reduces a facility's cooling loads. Demand response strategies could also be deployed for data center lighting, and heating, ventilation, and air conditioning (HVAC). The study concluded that additional studies and demonstrations are needed to quantify benefits to data centers of participating in DR and to address concerns about the possible impact of DR on data center performance or quality of service and equipment life span. This phase-two study is the first step in that direction.

1.3. Report Organization

The remainder of this report is organized as follows: Section 2 describes the research methodology for the study; Section 3 provides data center characterization to apply the field test study results to similar data centers; Section 4 details DR strategies for both IT equipment and cooling systems; Section 5 shows results of field tests using baseline and demand savings analyses; Section 6 presents conclusions and next steps; and Section 7 and 8 list references, followed by glossary of key terminologies used in this report and appendices that contain further details from the main report.

2. Methodologies

For the study, a Technical Advisory Group (TAG) of experts was recruited to review findings and guide the study. The TAG consisted of members from the data center industry, utilities, and regulatory fields, who reviewed the ongoing work continuously. The field test participating data center representatives provided pragmatic insights to data center operations and real-world strategies for data center participation in DR programs. In coordination with the TAG, the site and technology selection criteria and framework were prepared for data center recruitments for field tests. A comprehensive exercise was conducted to focus the outreach and recruitment of the data centers that met these criteria. A data center site survey questionnaire and monitoring and test plans were prepared to better understand the recruited data centers. The field test results were analyzed and presented. This methodology is described in detail below. Appendix 9.1 lists the TAG members and their affiliations.

Site and Technology Selection Criteria and Framework: The purpose of this activity was to identify site and technology requirements based on study goals and objectives, and for data center recruitment. The site and technology selection identified control systems and enabling technologies for site infrastructure (cooling) and IT equipment (server and storage). Technologies that are interoperable, standards-based, and are vendor neutral were considered to recruit data centers. Appendix 9.1 lists this activity in detail.

Assessment, Outreach, and Site Recruitment: The purpose of this activity was to select data centers and evaluate them against the study objectives and site and technology selection criteria. The recruitment was pursued using the network from LBNL, the California investor-owned utilities, the TAG members, and others. The recruitments focused on the data centers within the PG&E and San Diego Gas and Electric (SDG&E) utility territories and others within California due to sponsor requirements. The recruited data centers signed a non-binding memorandum of understanding to outline project

expectations. The site survey questionnaire included data center characteristics and enabling technologies for the field tests. Appendix 9.1 lists this activity in detail.

Develop Custom Monitoring and Field-Test Plans: The purpose of activity was to prepare monitoring and DR test plans in coordination with the data center operators. The monitoring plan included metering, procurement, installation, and technology requirements. The field test plans listed DR strategies, sequence of operations, data points, and anticipated outcomes. Appendix 9.1 lists this activity in detail.

Conduct Field Tests: The purpose of this activity was to conduct field tests within the recruited data centers, as outlined in the field test and monitoring plans. The DR strategies were refined based on the test results. The relevance of results to similar data centers and their economics for participation in DR programs were important considerations for field tests.

Evaluate Results: Following field-tests, the results were evaluated through baseline and demand-savings analyses to understand if and how data centers could participate in DR programs. The analysis was conducted relative to the whole building level and, where applicable, at the system or equipment level, including actual and perceived risks.

3. Data Center Characterization

This section summarizes the data center sites that participated in field tests and characterizes enabling technologies, equipment, systems, and load profiles to identify their DR strategies. This characterization provides a guide to which similar data centers the study findings are applicable. Table 1 summarizes the participating data centers, their function, and enabling technologies for both IT equipment and site infrastructure. The Power Usage Effectiveness (PUE), which is the ratio of total data center power and IT equipment power, is not included in this table, as it cannot be calculated from this data.

Table 1: Summary of the Recruited Data Centers

	LBNL 50B-1275	**NetApp**	**SDSC**	**UC Berkeley**
Floor area (%) relative to whole building.	5,000 ft^2 (100%)	1,100 ft^2 (37%)	19,000 ft^2 (54%)	10,000 ft^2 (22%)
Function	Storage Systems & HPC	Storage Systems	Storage Systems & HPC	Storage Systems & HPC
Utility territory	WAPA Power	PG&E	SDG&E	PG&E
Whole building power	550 kW	816 kW	2.3 MW	1 MW
IT equipment average power	350 kW	145 kW	1.6 MW	550 kW
Enabling technologies (IT and site)	SynapSense, Power Assure	Power Assure, Automated Logic Control	Power Assure, Opto22, Johnson Controls	Pulse Energy, Emerson Controls

SDSC = San Diego Supercomputer Center; UC = University of California; HPC = High Performance Computing; WAPA = Western Area Power Administration; PG&E = Pacific Gas and Electric Company; SDG&E = San Diego Gas and Electric

3.1. Participating Data Centers and Functions

This section summarizes the data centers recruited for DR field tests. A total of four California data centers were recruited for the study, including the NetApp Java 1 data center, University of San Diego's (UCSD) San Diego Supercomputer Center (SDSC), University of California, Berkeley (UCB), and the LBNL Building 50B-1275 (LBNL 50B) data center. The NetApp is a storage back-up data center; the SDSC, UCB, and LBNL 50B are a combination of high-performance computing (HPC), research and development (R&D), and production data centers. The SDSC is similar to LBNL 50B, except that one part of the data centers is a newer facility. All these data centers are self-owned and operated. Thus the operators determine the service level agreements (SLA).

An initial assessment of these data centers identified potential DR strategies for the site infrastructure, IT equipment (server and storage), and load migration. A few of the strategies (e.g., temperature set point adjustment) were tested at multiple data centers, with repeat tests to compare factors that influence the results. The assessment of these data centers included in-depth analysis of enabling technologies, equipment, and load profile characterization. This assessment was critical in determining the above DR strategies and also how relevant these strategies are for similar data centers.

3.2. Enabling Technologies

The study evaluated enabling technologies that existed within the recruited data centers and played a role in DR field tests.[3] These technologies, specific to data centers, can provide information to the data centers to facilitate DR program participation. These technologies provide real-time management and control of IT equipment, cooling, and monitoring of temperature and humidity conditions for cooling and air management.[4] The technologies that manage computing loads also provided data to better characterize the field test results. Table 2 provides an overview of these enabling technologies.

Table 2: Summaries of Enabling Technologies from Recruited Data Centers

<table>
<tr><th rowspan="2">Data center</th><th colspan="3">IT equipment</th><th colspan="2">Site Infrastructure</th></tr>
<tr><th>CPU Utilization</th><th>Memory I/O</th><th>IT power</th><th>Temperature</th><th>HVAC Power</th></tr>
<tr><td>LBNL 50B</td><td colspan="2" rowspan="2">Warewulf Cluster Toolkit Adaptive Computing® Moab™ Job scheduler.</td><td colspan="3">SynapSense® Optimization Platform™</td></tr>
<tr><td>SDSC</td><td>Power Assure® EM/4™</td><td colspan="2">Johnson Controls® Metasys™* Opto 22® systems;</td></tr>
<tr><td>NetApp</td><td colspan="3">Power Assure® EM/4™</td><td colspan="2">Automated Logic Corporation®*</td></tr>
<tr><td>UCB</td><td colspan="3">Warewulf Cluster Toolkit Adaptive Computing® Moab™ Job scheduler</td><td colspan="2">Emerson Controls*, and Pulse Energy*</td></tr>
</table>

Integrated monitoring and data from IT equipment and site infrastructure, necessary for baseline and demand savings analyses, was collected and aggregated using the EM/4

[3] Other technologies were used within recruited data centers (e.g., IBM® Active Energy Manager™ in UCB data center), which were not part of the field tests and not detailed in this study.
[4] *Real-time* refers to the on-screen refresh frequency of up to one-minute.
* These mature technologies, used in commercial buildings, are included, as they played a role in DR tests.

software from Power Assure. The EM/4 – P100 data collection appliance was installed in NetApp and SDSC data centers for the study. The SynapSense® Data Center Optimization Platform was used for real-time environmental monitoring (temperature, humidity, subfloor pressure differential) and power metering (PUE, branch circuit monitoring and cooling, cabinet and server outlet level power). For load migration tests between LBNL 50B and SDSC, the Warewulf toolkit was used for operating system management and provide data from large High Performance Computing (HPC) Linux computing clusters. It was also used for efficient management of physical and virtual HPC systems within local and distributed data centers. For load migration tests between LBNL 50B, SDSC, and SDSC, the Adaptive Computing® Moab™ HPC and cloud suites were used to manage and optimize computing load transfers, workload management of clusters, and automated job scheduling of HPC systems in real time during a DR event. A HPC system performs large computations using a number of processors operating in parallel. The HPC system performance is measured in term of teraflops or petaflops, which is the number of floating point operations (flop) per second. An HPC system comprises tens, hundreds, or thousands of servers (each with its own processing, memory, and disk storage) connected by a high-speed network and shared storage. An HPC system can be used as a stand-alone system or as a grid using a high-speed, low-latency network connection with another HPC system. Appendix 9.2 lists hardware specification for HPC systems in SDSC, LBNL, and UCB, including the network topology of SDSC and UCB clusters used for load migration tests.

The IT equipment consisted of storage hard drives, high-performance computing servers or nodes, and networking systems. The study evaluated cooling systems such as the computer room air conditioning (CRAC) units commonly used across data centers to maintain a zone temperature, air distribution, and humidity. The operational behavior of these equipment and systems provided the load patterns and the right DR strategy. The study did not look at the IT jobs or component-level details to understand their impact on DR participation. This section describes the IT equipment used in the DR field tests.

The NetApp Java-1 data center contains storage hard drives clustered together in racks. These storage clusters are primarily used to back up data from production data centers. A storage cluster includes hard disks, filer heads, and networking equipment. The filer head is the computing brain, which is used to manage the hard disks in a cluster. All the data storage drives utilized in the NetApp DR tests are rated at 7,200 rpm, which determines the equipment's total power. Higher speeds result in increased power demand. Of the eight rows of storage racks, only two rows participated in the DR tests. Table 14 in Appendix 9.2 provides information on the row configuration of these racks.

3.3. Load Profile Characterization

End-use loads in a data center environment can be categorized in to three broad categories: (1) IT equipment, (2) cooling or site infrastructure, and (3) support loads, which consist of an uninterruptible power supply (UPS) and lighting. Table 3 summarizes the end-use loads at the four data center sites considered in this project.

Table 3: Summary of End-Use Loads of Participating Data Centers

End Use Loads	LBNL 50B	NetApp	SDSC	UC Berkeley
IT equipment	58%	52%	71%	55%
Cooling systems	38%	42%	28%	40%
Lighting, UPS, other	4%	6%	1%	5%

The IT equipment governs the data center energy use. Depending on the data center, load profiles are based on IT equipment characteristics (load, utilization, storage capacity, vintage, and access) and site infrastructure loads. The load profile characterization showed no sensitivity to the weather. There is a slight variation in energy consumption of a cooling system, which is in the noise of a load profile curve, particularly in data centers with cooling systems consuming much lower energy than IT equipment. A cooling plant with air- or water-side economizer will have significant variation in energy use. The study used these end-uses for data center load characterization attributes to determine DR strategies for the participating data centers and their relevance to similar data centers.

4. Demand Response Strategies

Building from the previous studies, this section provides in-depth understanding of engineering and operational considerations for the DR strategies, their sequence of operations, and other relevant details using field tests. The strategies were developed based on the data center characterization described earlier. Multiple field tests were conducted for these DR strategies in different data centers, which lead to further refining of these DR strategies; in particular, the sequence of operations and identifying potential challenges. Table 4 summarizes the DR strategies and their field-test data centers. Results and findings for these DR strategies are described in Section 5.

Table 4: DR Strategies for Field Tests of IT Equipment and Cooling Systems

	Demand Response Strategy	LBNL 50B	NetApp Java-1	SDSC, UCB, LBNL 50B
1	Server and CRAC unit shutdown	■		
2	Load Shifting or Queuing IT jobs – Server idling	■	■	
3	Temperature set point adjustment	■	■	
4	Shutdown and idling of IT storage clusters	■	■	
5	Cooling relative to IT equipment load reduction		■	
6	Load migration between heterogeneous systems.[5]			■
7	Load migration between homogeneous systems[6]			■

IT = IT equipment; CRAC = Computer Room Air Conditioners; LBNL = Lawrence Berkeley National Laboratory; SDSC = San Diego Supercomputer Center, UCB = University of California Berkley.

The field tests for these sites included load shed, load shift, and load migration strategies. *Load shed* refers to a temporary decrease in power or processing load of the IT equipment by shutting it down, and/or by raising the temperature set points. The *load shift* refers to rescheduling IT equipment power by shifting the IT jobs to a time outside of the DR event window, followed by idling or shutting down the corresponding IT

[5] Heterogeneous systems consist of different IT equipment and processing capabilities in both data centers.
[6] Homogeneous systems consist of identical IT equipment and processing capabilities in both data centers.

equipment. The *load migration* refers to geographic shifting of computing load to another data center from a data center that is participating in DR event.

The DR strategies below are based on the collective findings from field tests, with emphasis on their relevancy within similar data centers and considering operational requirements. The field-test experience helped finalize the sequence of operations.

DR Strategy 1: Server and CRAC unit shutdown

In this strategy, the data center demand was reduced by graceful shutdown of IT equipment and site infrastructure loads.[7] The tests were conducted and results were validated multiple times at the LBNL 50B data center. Table 5 provides additional information about this strategy.

Table 5: Server and CRAC Unit Shutdown

Definition	Shut down servers and compressors in CRAC units to shed load
Applicability	IT equipment and site infrastructure
End-Use Type	Servers, networking, HVAC, UPS
Target Loads	Servers, processors, HVAC pumps, chillers, and fans
Category	Load shed
Sequence of Operations for Load Reduction	▪ Clear the jobs queues of IT equipment to stop all processing jobs ▪ Gracefully shut down the servers and other IT equipment ▪ Shut down all HVAC equipment
Sequence of Operations for Recovery	▪ Following the DR event, turn on the chiller and CRAC units to bring zone temperatures to a standard operating level. ▪ Stage the IT equipment restart to avoid higher rebound power. ▪ Monitor power and zone temperatures to ensure failure-free restoration from the DR event. ▪ Restart IT job queues.
Notes	Temperatures must be within standard ranges before servers restart.

DR Strategy 2: Load shifting or queuing IT Jobs – Server idling

The IT load shifting or queuing strategy entails reducing the utilization of IT equipment to reduce its overall power draw. The strategy is conceptually the same for both computing and storage clusters, however, they differ in the load types shifted. DR field tests were conducted for both system types. Table 6 provides information on this strategy.

***Strategy 2a: Computing clusters (IT)*:** The computing cluster energy consumption can be reduced by preventing new compute jobs and idling partial nodes. The tests conducted to evaluate load-shifting strategy exposed a clear correlation between CPU utilization and power draw of the cluster as a whole. The jobs, which were prevented from starting during the DR event, began after the DR event concluded. To achieve success in implementing this strategy, a capacity reservation block was created to modify the job

[7] A "graceful" shutdown, or restart refers to a consolidation, shutdown, or restart in which applications or equipment goes through the normal consolidation, shut down, or startup process, as opposed to an abrupt consolidation or shutdown, such as when a plug is pulled or an application process is forced to quit.

scheduler of the computing cluster. The duration of the capacity reservation block depends on notification period of the DR event, number of jobs in the scheduler queue, and their priority. Depending on the flexibility of system policies, the idle nodes can be shut down completely to achieve a more aggressive load shed. Field tests of this strategy were conducted at the LBNL 50B data center.

Strategy 2b: Storage clusters (IT): Load on the storage clusters can be reduced by rescheduling tape and data backup jobs to a time outside the DR event window. This is a classic application of the load shift strategy. The revised job schedule will free storage hard drives, filer heads, and other resources, which can be idled to reduce the system demand. By gracefully turning off the storage shelves and filer heads, significant energy can be reduced. Multiple scenarios of this strategy were evaluated at NetApp.

Table 6: Load Shifting or Queuing IT jobs

Definition	Use job scheduling techniques to remove compute load on IT equipment and have it go to idle state.
Applicability	IT equipment
End-Use Type	Server, storage systems, networking equipment
Target Loads	Server, processors, hard drives, filer heads, network switches
Category	Load shift or shed
Sequence of Operations for Load Reduction	**2a: High Performance Computing Systems** ▪ Set capacity reservation block on the job scheduler to prevent jobs from starting or running during the DR event period. ▪ Monitor the server status to ensure the load shed is achieved.
	2b: Storage Cluster ▪ Send halt command to filer heads in the storage cluster. ▪ Shutdown the filer head first, followed by the storage racks. ▪ Monitor the storage cluster status to ensure load shed is achieved.
Sequence of Operations Recovery	**2a: High Performance Computing Systems** ▪ Lift the capacity reservation block on the server after the DR event. ▪ Restart the jobs in the queue.
	2b: Storage Cluster ▪ Power on all the storage racks and hard drives. ▪ Once all the storage racks are powered ON; turn on the filer heads.
Notes	▪ After the DR test, a slight increase in the power demand of the servers and storage equipment were observed over the initial pre-test conditions. If this negatively impacts data centers, these equipment should be brought back to full availability incrementally so to reduce the immediate rebound impact. ▪ Applicable to non-mission critical applications.

DR Strategy 3: Temperature set point adjustment

Cooling systems have components such as chillers, chilled water pumps, computer room air conditioning (CRAC) units, computer room air handler (CRAH) units, and variable frequency drives (VFD). When global temperature set point adjustments are made in the control system, the components react in tandem to bring the temperature to new levels.

For DR, when the zone temperatures are increased over the standard operating set point, the overall demand of the cooling system decreases. Zone temperatures can be raised to a level not to exceed the temperature range specified by the IT equipment manufacturer or the American Society of Heating, Refrigeration, and Air-Conditioning Engineers (ASHRAE 2008). For example, in NetApp, the zone temperatures on two CRAC units were increased from 72°F to 78°F, resulting in a decrease of the cooling system power draw. Currently, the NetApp Java-1 data center has constant air volume fans (CAV) in the CRAC units. If VFDs were installed on all the CRAH units, the facility manager estimates an additional fan energy savings of approximately 17%. Enabling technologies with one-minute data granularity play a very important role in giving real-time feedback about the internal environmental conditions and overall data center thermal stability to the data center operators. Table 7 provides additional information about this strategy.

Table 7: Temperature Set Point Adjustment

Definition	Globally increase facility HVAC temperature set points in order to decrease HVAC power demand.
Applicability	Site infrastructure
End-Use Type	HVAC
Target Loads	HVAC pumps, chillers, and fans
Category	Load shed
Sequence of Operations for Load Reduction	▪ Increase temperature set points by two degrees. ▪ Monitor the zone temperatures and the IT equipment response in real time using the control system. If appropriate, increase the temperature set point incrementally. ▪ Hold temperature steady for allotted DR period, provided the data center space conditions do not exceed operating thresholds.
Sequence of Operations for Recovery	▪ After demand-response period, decrease temperature set points in incremental steps to avoid large power demand rebound.
Notes	Strategy should be tested before the full event to understand temperature profiles and prevent issues with hot spots causing server problems.

DR Strategies 4 and 5: Cooling relative to IT equipment load reduction

When the IT equipment load is reduced for DR, the net heat generated by the equipment reduces. Because of this decease in the generated heat, the cooling systems may respond automatically to increase the rack inlet air temperatures. However, this response period depends on the type of HVAC system (e.g., CAV or VFD). Alternatively, the inlet air temperature can be raised manually by a few degrees to lower the Delta T (ΔT) for a faster response. Depending on the IT equipment operating temperature rating and the ASHRAE specification for data center environments, zone inlet temperatures can be raised to prevent any equipment failure or negative impact on its life span. Field tests were conducted as part of this project to understand the interactive effects between the IT loads and cooling loads—both automatically and by manual adjustment. Section 5 presents detailed results for multiple test sites where this strategy has been tested. Table 8 provides additional information about this strategy.

4: Automatic response of cooling load relative to IT equipment load: This strategy applies to data centers, which have feedback loop from zone outlet temperatures to the damper positions of the CRAC units. When the indoor temperature falls due to reduced IT equipment operation, the outlet air temperature also decreases. The damper positions of the CRAC units will be adjusted proportionally to maintain the zone temperatures at set point values. This type of setup is ideal for fully automated or semi-automated DR.

5: Manual set point adjustment: In the absence of a control loop between outlet temperatures and CRAC units, manual set point adjustment can be made in the data center cooling management system to achieve faster demand savings. The adjustments depend on the comfort of the facility manager, availability of monitoring technologies, and the health of the IT equipment. In a semi-automatic control environment, manual set point adjustments can be pre-programmed.

Table 8: Cooling Relative to IT Equipment Load Reduction

Definition	Intelligent coordination of site infrastructure controls to respond automatically to IT infrastructure load reductions
Applicability	Site Infrastructure
End-Use type	HVAC
Target loads	HVAC pumps, chillers, and fans
Category	Load shed
Sequence of Operations for Load Reduction	**4. Automatic Response** ▪ Cooling load reacts automatically in response to IT load. ▪ Monitor the status of the IT and cooling load until end of DR event.
	5. Manual set point adjustment ▪ After reducing IT load, monitor the zone temperatures. ▪ Based on the decrease in the zone temperatures, increase the zone set points incrementally without adverse effects on any equipment. ▪ Hold at this set point until the end of the DR event.
Sequence of Operations for Recovery	**4. Automatic Response** ▪ Cooling system restore to normal operation automatically.
	5. Manual set point adjustment ▪ After the event, roll back the temperature set point to original value. ▪ Turn on IT equipment after the temperature becomes stable.
Notes	Staged turn on sequence can prevent any spikes after the DR event. IT equipment should be turned on only after the zone temperatures are within the operable conditions of the equipment (usually under ASHRAE temperature guidelines).

DR Strategy 6, 7: Load migration between homogeneous and heterogeneous systems

The homogeneous and heterogeneous systems with a dedicated high-speed network were used for load migration strategies. ShaRCS is a homogeneous system with two identical Linux HPC clusters: the North cluster, Mako, located at UCB; and the South cluster, Thresher, located 446 miles away at SDSC. Lawrencium (LR) is a heterogeneous system and consists of the LR-1 cluster located in SDSC and is connected to LR-2 cluster in LBNL 50B through a dedicated 10 gigabits per second (Gbps) low latency network. The LR system has shared central file storage at LBNL 50B and uses a job scheduler to

process LR-1/LR-2 jobs. The job scheduler information does not include the job nature, size of input/output operations, size of data stored or transferred, which is also true of ShaRCS system. While the ShaRCS system was dedicated for DR tests, the LR tests were conducted using the production systems. Table 9 provides details about this strategy.

6: Homogeneous system (ShaRCS): For DR tests, new and current IT jobs running on Mako clusters were migrated to Thresher clusters. A certain percent capacity reservation block was placed on clusters in one data center for the duration of the DR event. A capacity reservation block is a policy that is created in the job scheduler to block a certain number or percentage of nodes from running IT jobs. When this capacity reservation block is placed, IT jobs running for that percent of the nodes are halted or killed, and new jobs are prevented from starting on these nodes. After the reservation block is executed, the percent of nodes can either be sent to idle mode or shut down completely. Figure 1 shows an example of the capacity reservation concept on both ends of the HPC cluster.

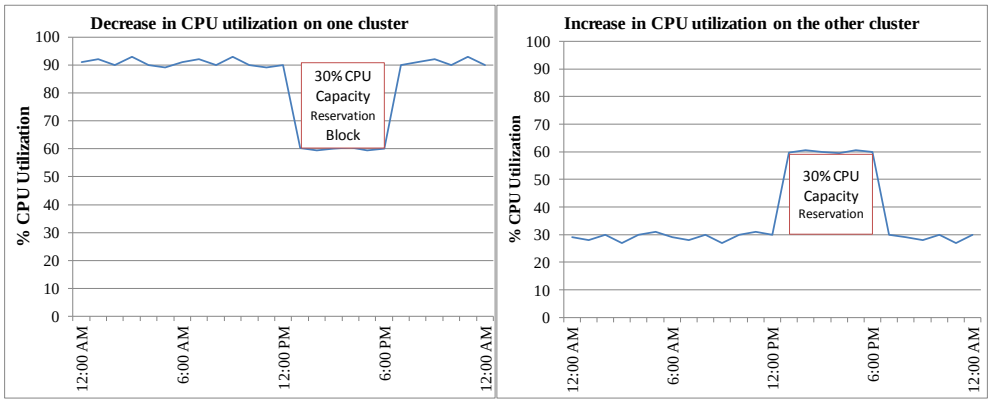

Figure 1: Example of Reservation Block Allocation for an HPC Cluster

7: Heterogeneous system (LR): For DR tests, new IT jobs were blocked from running on LR-1 clusters and migrated to run on LR-2 clusters. The old IT jobs on LR-1 decay or complete over a period of time (up to 72 hours, which is the maximum run-time limit for LR operation), resulting in DR shed. The DR savings come from potential increase in the energy use of LR-1 when new jobs are scheduled to run during the DR event window. However, measurement and verification (M&V) of these DR savings for such a strategy is challenging, as it may not be possible to accurately calculate the savings. Few options that would not have such M&V issues are: idling or shutting down the LR-1 equipment after job consolidation. This strategy was not implemented, as the data center operators did not have the resources and the infrastructure to manage the IT equipment remotely. Another strategy would be to implement a policy to allow the data center operator to kill running jobs at the start of the DR event and then restart them afterwards. Depending on the type of job, this policy may require that killed jobs have to be restarted from the beginning. Table 9 provides additional information about this strategy.

Table 9: Load Migration in High Performance Computing Clusters

Definition	Load migration in a High Performance Computing (HPC) Cluster
Applicability	IT infrastructure
End-Use type	Server, storage, and networking devices
Target loads	Computing nodes, processors, hard drives, routers, switches

Category	Load Migration
Sequence of Operations for Load Reduction	**6. Homogeneous System** ▪ Set a 30% (or appropriate %) capacity reservation on Thresher to prevent starting of new jobs during the DR event. ▪ Once the DR event starts, the capacity reservation will come into effect and push the HPC nodes into idle mode. ▪ Shutdown the IT equipment to achieve higher load shed if it meets the standard operating procedures of a participating data center. ▪ Hold the cluster in this state until the end of DR event.
	7. Heterogeneous System ▪ Monitor the status of CPU utilization and power consumption on both ends of the homogeneous cluster ▪ Set a 30% capacity reservation on LR-1 to prevent the start of new jobs. Migrate all news jobs to run on LR-2. ▪ Allow jobs currently running to finish processing and decay over a three-day period. After all the jobs are drained, the LR-1 nodes will be running in idle mode until the end of the DR event period. ▪ Nodes can be shut down to achieve higher load shed, provided the strategy meets standard operating procedures of the data center.
Sequence of Operations for Recovery	**6. Homogeneous System** ▪ After the DR event, the capacity reservation can be lifted to return the system to normal operation state. ▪ When nodes are idled, put them in active mode to accept new jobs. ▪ When nodes are shutdown, start and put them in active mode to accept new jobs.
	7. Heterogeneous System ▪ After the DR event, the capacity reservation can be lifted to return the system to normal operation state. ▪ The idle nodes should be sent to active mode to accept new jobs.
Notes	By staging the turn-on sequence, sudden spikes in demand after the DR event can be avoided. The nodes are stateless and boot from the network. As a best practice, implement a four-second-boot delay to prevent strain on the network.

The DR strategies for the IT equipment and site infrastructure for storage and HPC data centers can be used for similar data centers. Enabling technologies play a key role, as the manual DR strategies increase the response and recovery times and require a lot of human resources, and increase the cost of DR enablement and execution. A well-tuned building, also true of commercial buildings (Motegi et al. 2007), and a real-time energy monitoring is needed for aggressive cooling strategies. The study reaffirmed earlier findings that the cooling systems in data centers are becoming more efficient (Ghatikar et al. 2010). While this is a welcome development, in future this will decrease the cooling load availability for DR, thereby emphasizing the IT equipment where there is a larger potential of demand shed. Load migration strategy has significant potential, and this study conducted field tests for only a few strategies. The study results in the next section shows that reduction in IT loads reduce the supporting loads, leading to higher DR sheds.

5. Field Test Analysis and Results

This section presents results of the DR field tests for each data center. The results were applied to a common DR program framework in the U.S. This framework includes four key elements for DR programs: (1) Notification period, (2) response (or ramp) period, (3) active event period, and (4) recovery period (NIST 2010, NIST 2012, and OASIS 2012).[8] The framework can be used to map the results to commercial DR program requirements. Table 10 summarizes the results at whole building level, except for the notification period, which is strategy specific. An economic analysis for NetApp was conducted to determine potential savings by participating in the PG&E peak day pricing program.

Table 10: Summary of DR Test Results as per Common DR Framework

Data Center	DR Strategy	Event Date	Active Event Period	Response Period (min)		Recovery Period (min)
				5% shed	10% shed	
NetApp	Test 1 (Shift/Queue IT jobs – Storage)	19-Dec-11	2:30pm to 5pm	10	22	25
	Test 2 (Temperature set point adjustment)	21-Dec-11	12:00pm to 1:00pm	5	15	15
	Test 3 (Shift/Queue IT jobs – Storage w/ manual temperature adjustment)	13-Jan-12	2:00pm to 4:00pm	7	15	17
	Test 4 (Shift/Queue IT jobs – Storage)	11-Jan-12	1:00pm to 3.00pm	7	15	30
LBNL 50B	Test 1 (Server and CRAC units shutdown)	28-Oct-11	8:00am to 5:00pm	2	8	180
	Test 2 (Shift/Queue IT jobs – Server idling)	1-Nov-11	Midnight to 6:00am	n/a	n/a	n/a
	Test 3 (Temperature set point)	16-Nov-11	12:35 to Midnight	n/a	n/a	n/a
	Test 4 (Data Center Shutdown)	2-Dec-11 to 3-Dec-11	3:40pm 12:00pm	8	15	120
SDSC, UCB, and LBNL 50B	Test 1 (Load migration - Homogenous – Idling)	25-Apr-12	1:10pm to 2:.45pm	2	6	2
	Test 2 (Load migration - Homogenous – Shutdown)	25-Apr-12	2:46pm to 5:10pm	3	7	10
	Test 3 (Load migration - Heterogeneous – Decay)	3-Jul-12 to 5-Jul-12	10:45am to 11:00am	147	175	15

[8] *Notification period* is the time between event notification time and start time; *active event period* is the time between the start and end time of the event; *response (or ramp) period* is time required to reach the estimated shed; and *recovery period* is the time required for facility to be restored to normal operations.

A reliable baseline was established to calculate the demand savings. The study evaluated various baselines against load profiles (Coughlin et al. 2009). Results show that for flat load data centers, the demand savings are similar for all baselines (not weather sensitive). The least variance was observed for 10/10 baselines with morning adjustment (Adj. 10/10 BL) and for 20-day outside air temperature regression. The PG&E peak-day pricing program uses Adj. 10/10 BL. Appendix 9.3.1 describes the baselines.

The demand savings power metrics for each field test include percentage (%), amount (kW), and density in watts per square foot (W/ft^2) to understand the variations observed at the whole building level, and when data was available, at the data center level. Figure 2 summarizes the demand savings of all tests. The values include total kW and % demand savings, and where applicable, are broken by IT equipment and site infrastructure load.

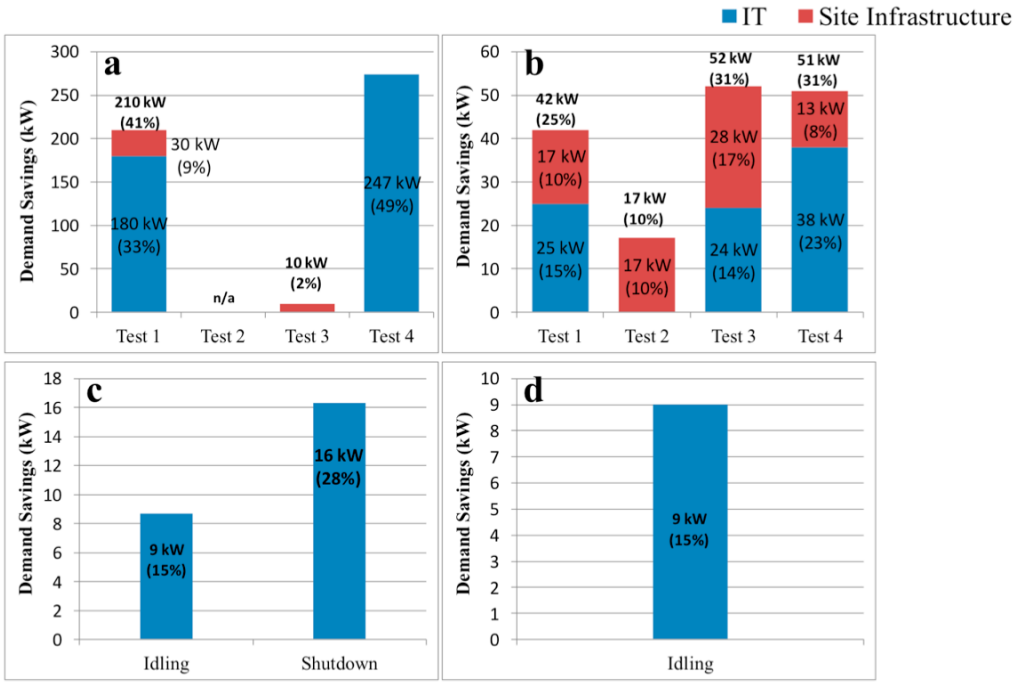

Figure 2: Demand savings summary (kW and %) of all field tests
(a) LBNL 50B; (b) NetApp; (c) SDSC and UCB; (d) SDSC and LBNL 50B

5.1. LBNL Building 50 Data Center

Four tests were conducted at LBNL's Building 50B Room 1275 (50B) data center to evaluate DR strategies for the IT equipment and cooling systems. The baseline and demand saving analyses for these tests are presented below. Detailed figures and tables for analysis conducted at LBNL 50B are shown in Appendix 9.3.2.

Field Tests: Baseline and Demand Savings Analysis

Test-1: Server and CRAC Unit Shutdown: Test 1 was conducted on October 28, 2011, from 8:00 a.m. to 5:00 p.m. For this test, the LR-2 servers were shut down as part of the maintenance activity. The LR-2 servers started shutting down as jobs were completed at 9:00 a.m., and the shutdown was complete by 12:00 p.m. During the maintenance period, the facility staff manually turned off the compressors in two of the CRAC units, as less

thermal conditioning was needed. Other end uses were not part of the tests. The resulting power profiles were monitored during this period to characterize demand savings. Figure 10 in Appendix 9.3 shows IT and HVAC power drop. The total power dropped from 489.5 kW to 280.5 kW, rebounding to 539.5 kW. Figure 3 shows the demand savings of 245 kW (40.8%) at the whole data center level against the 10/10 and the 10/10-MA baselines. The previous ten days of operation show higher energy consumption than the non-DR portion of the event day, thus shifting the 10/10 baselines upwards over the 10/10-MA baseline. Table 17, in the appendix, summarizes the hourly demand shed realized during the event, using the adjusted 10/10 baselines.

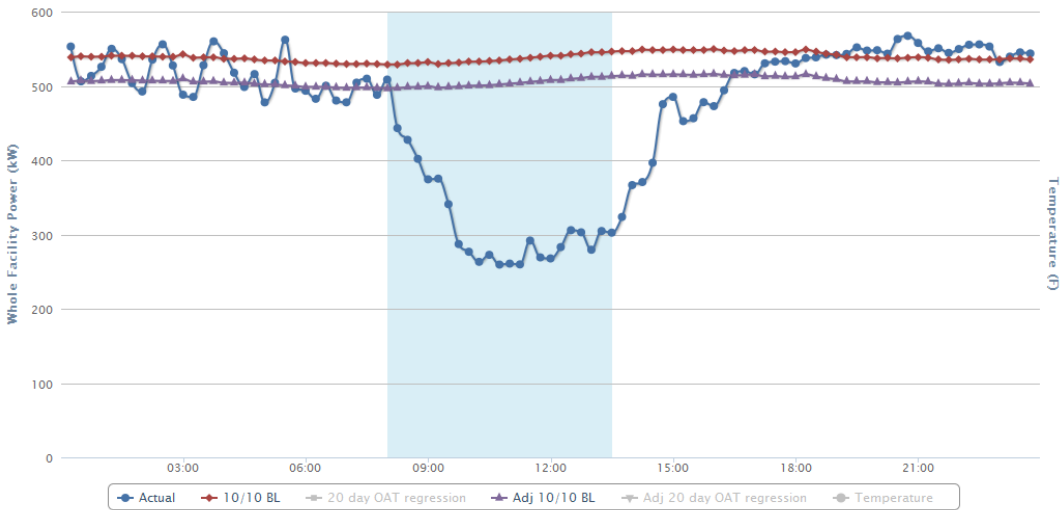

Figure 3: Baseline Analysis of Test 1 in the LBNL 50B Data Center

Test 2: Shift or Queue IT Jobs – Server Idling, no IT Load Migration*:* Test 2 was conducted on November 1, 2011, from midnight to 6:00 a.m. The LR-2 cluster computing loads are assigned using a scheduler program. To test whether a demand saving could be realized by setting the servers to idle, the scheduler program was used. A "reservation block" with zero computing loads was placed for 30% of the capability for a six-hour period on November 1. Power demand was observed before and after this period, and no appreciable change in IT power was seen. With further evaluation, it was determined that the reservation did not run as planned, and the test was repeated. Again, no appreciable change in demand was observed. More investigations were conducted and the IT staff measured the power versus the state of a representative server and found that there was a drop in power with a decrease in server load. It is clear that there is a potential for reduced power when this server is in idle mode. To explore further, additional sub-metering was put in place, and the scheduling algorithm was modified further so that additional tests could be conducted as part of the load migration tests.

Test 3: Temperature Set Point Adjustment*:* Test 3 was conducted on November 16, 2011, from 12:35 p.m. to midnight. This test was conducted after the data center operators had evaluated, adjusted, and tested the temperature settings, and resulting profiles and distributions within the data center to stabilize the thermal conditions. The test was conducted by raising the set points of the CRACs and CRAH by 2°F at a time.

Figure 4 shows the variation in total CRAC fan power. This test resulted in a small demand savings, averaging 1.6% of total data center facility demand.

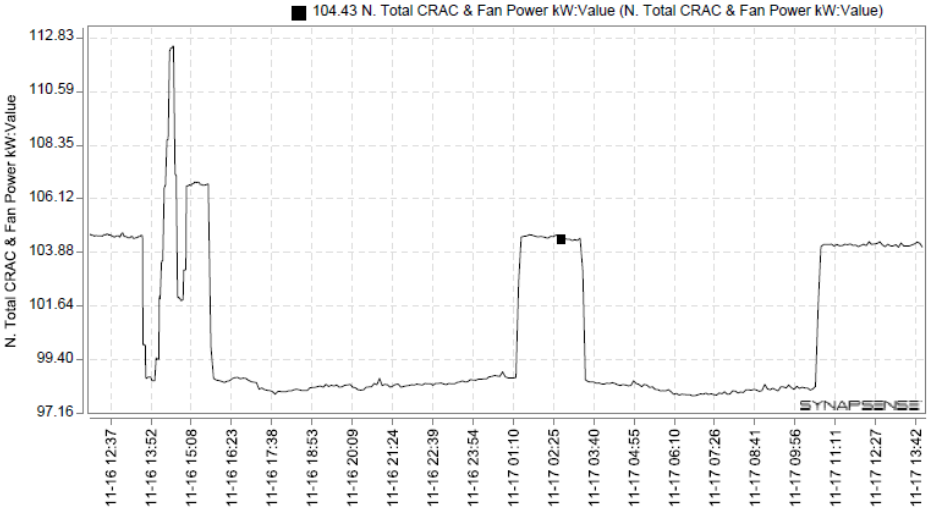

Figure 4: Recorded Demand Data by End Use: Test 3 (11/16/ 2011) in LBNL 50B

Test 4: Complete IT and CRAC Units Shutdown*:* Test 4 was conducted from December 2, 3:40 p.m. to December 5, 2011, until noon, to use the data center maintenance shutdown. While not a true DR test, it allowed the study to determine the demand savings and sequence of operations. It was not a complete shutdown, as the demand did not go to zero. More investigation would be needed to determine which portion of the data center was impacted by the shutdown. The total power dropped from 555 kW to 280 kW, resulting in a 50% load shed. The test resulted in average demand savings of 48.9% of the total data center facility demand. Appendix 9.3.2 show detailed analysis and results.

Summary*:* The DR field tests for LBNL's 50B cooling and IT systems were successful; some of them due to the maintenance. The results from server shutdown were most successful, with average demand savings ranging from 40.8% to 48.9% at the data center level. The demand savings from temperature set point adjustment strategies were low, averaging 1.6% at the data center level. Due to physical conditions (shallow under raised floor plenum and restrictions to airflow), this legacy data center may not be able to sustain a significant and long-duration temperature adjustment strategy. The results showed that well-tuned HVAC systems are necessary for the temperature adjustments.

5.2. NetApp Java-1 Data Center

Four tests were conducted at NetApp to evaluate DR strategies for the IT equipment and cooling systems. The baseline and demand saving analyses are presented below. Detailed baseline and demand savings analysis conducted at NetApp are shown in Appendix 9.3.3.

Field Tests: Baseline and Demand Savings Analysis

Test 1:Shutdown and idling of IT storage clusters: Test-1 was conducted on December 19, 2011, from 2:30 p.m. to 5:30 p.m. In this test, four filer heads and its storage hard drives were gracefully shutdown. The filers were first halted and then shutdown along with all their hard disk drives. The CRAC units could control the chilled water valve

position to limit the supply of chiller water in response to the outlet air temperature. As the outlet air temperature falls due to reduced heat from the IT equipment, the CRAC units responded by reducing the flow of the chilled water. Since the power draw of the chiller depends on the amount of chilled water supplied, significant demand savings were observed. The demand savings analysis showed a shed of 24.8 kW (14.5%) of IT load, and 16.5 kW (10.3%) of cooling and UPS loads, totaling 41.5 kW (24.8%) shed at the data center level with zero impact to data center reliability, redundancy, and operations.

Following the test, a restart sequence was initiated. During this period, the demand of the IT equipment increased by 4%. This was due to the increase in the processing power of the filer heads and rotation of few hard disks that were idle prior to the shutdown. The load was restored to pre-event condition in 1 hour 15 minutes. The storage IT equipment shed load in 3 minutes and was able to restore to normal operation in 5 minutes. Figure 5 shows the demand savings of 53 kW (12%) at whole building level against Adj. 10/10 BL and Adj. 20 day OAT regression baselines. These savings are greater than those at the data center level likely due to the decrease in the office load after 5:00 p.m. A small set shed of IT equipment (2 rows from a total of 8 rows) is detected at the whole building.

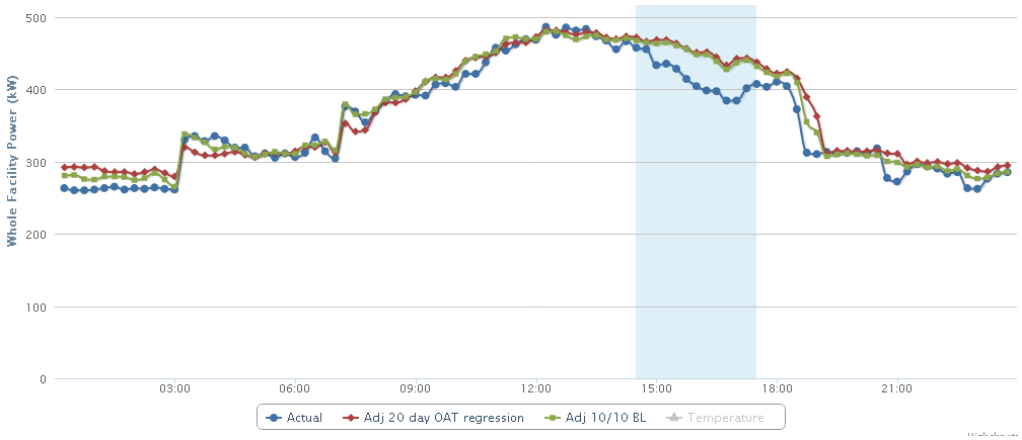

Figure 5: Baseline Analysis of NetApp Java-1 Building (12/19/2011)

Test 2: Temperature set point adjustment: Test 2 was conducted on December 21, 2011 by raising the zone temperature set points. The inlet air temperature was increased by 2°F (from 72°F). The 5-minute interval data showed that the zone 3 temperature increased by 8°F in 30 minutes, and the set points had to be restored to 72°F. The analysis at the whole building level did not show detectable shed due to a low operating PUE of 1.4. Since the CRAC units are of constant air volume type, no additional savings were realized from the fan power. The data center operator estimated a 17 kW demand savings from chiller, or 8% of the data center load. An aggressive set-point adjustment strategy is needed for or a detectable cooling load shed from low PUE data centers. A real-time (sub-minute) thermal monitoring system can build a confidence with the data center operator.

Test 3: Cooling relative to IT equipment load reduction: Test 3 was conducted on January 11, 2012, following the Test 1 and Test 2 analysis. This is similar to Test 1, however, the CRAC inlet air temperature was manually raised from 72°F to 78°F in response to storage equipment shutdown for rows 3 and 6. The backup jobs that run during the normal business hours were rescheduled to run after 6 p.m. Simultaneously,

the inlet temperature of the CRAC units was raised in increments of 2°F to a maximum of 6°F. The demand savings were 24 kW of IT load, 23 kW of cooling (CRAC units and chillers), and about 5 kW from UPS. The demand savings observed at the data center was 52.4 kW (32%). The manual rising of the CRAC inlet air temperatures resulted in faster shed response period. The relative changes in the response times for Test 1 and Test 3 are presented in Table 11. A data center with advanced cooling systems can pre-program the set-point adjustments to respond to IT equipment.

Table 11: Response and Recovery Period Comparison of Cooling Systems

	Test 1	**Test 3**
Control Strategy	Automatic*	Manual+
Response Period	60 min	15 min
Recovery Period	30 min	30 min

*Automatic = Automatic response by cooling systems in response to IT load reduction
+Manual = Manual adjustments made to the cooling systems in response to IT load reduction

Figure 6 shows the component-level demand savings. The savings from the CRAC units, IT equipment, and UPS were analyzed against Adj. 10/10 BL. The CRAC units showed a small shed, since the demand savings from cooling infrastructure are dominated by the central chiller plant, which is not represented in these data. The figure also shows the temperature changes in zones 2 and 3, which provided cooling to IT equipment that were part of tests. The baseline and demand savings analysis are shown in Appendix 9.3.

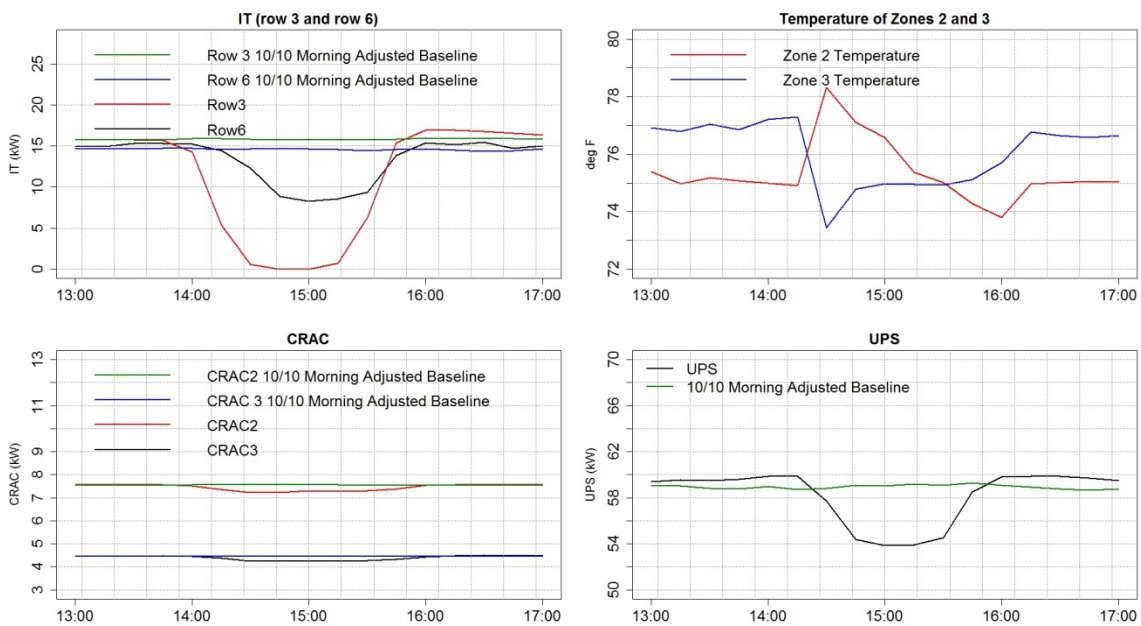

Figure 6: Demand Saving Analysis of Test 3 at NetApp Java-1

Test 4: Shutdown and idling of IT storage clusters: Test 4 was conducted on January 13, 2012, from 1:00 p.m. to 3:00 p.m. This test was similar to test 1 except that more filers and disk drives were shut down. A total of 6 filer heads and disk drives were gracefully shut down without any manual adjustments to the CRAC inlet air temperatures. The cooling systems responded automatically, with a total 12.1 kW of load shed. The baseline

and load shed analysis show that the total shed from IT equipment was 38 kW and shed from site infrastructure was 13 kW. The total data center load shed was 51kW, or 31%. Figure 22 in the appendix shows the demand savings at the data center level.

The demand savings observed at the utility meter is approximately 44kW (9%) at the whole building level. This is lower than 51 kW load reduction at the data center level, and is likely due to the increase in the non-data center or office load during the post-noon test window, which is peak period for office buildings.

***Summary**:* The DR field tests for NetApp Java-1 data center were successful and did not impact the data center's reliability and operations. The DR strategies for storage systems showed an average demand savings ranging from 23% to 31% at the data center level, and were detected at the whole building level. The temperature set point adjustment strategy did not show detectable shed at whole building level due to the low PUE of the data center. The analysis did not review the relationship between data storage and power levels due to lack of data. Since NetApp is a customer of PG&E, the study conducted economic analysis for this data center to identify cost savings from DR participation.

Economic Analysis

As a sample study, basic DR economics analysis was conducted for the NetApp Java-1 building, presuming its enrollment in the PG&E Peak Day Pricing (PDP) program. Figure 23 in Appendix 9.3 shows a relatively flat daily load shape, which is common in stand-alone data centers. The analysis using the PG&E's InterAct™ tool shows that by just enrolling in a PDP program, NetApp can save $7,500, or 1.4% of its annual energy bill.[9] This is largely due to the design of the PDP program, which is revenue neutral to class-average load shape. This means flat load data centers can realize benefits by just participating in the PDP program. By shedding load on event periods when the price of electricity is high, a data center can realize additional savings.

5.3. University of California Berkeley and San Diego Super Computer Center

One load migration test was conducted on April 25, 2012, from 2:30 p.m. to 5:00 p.m. for the ShaRCS homogenous system located in UCB (Mako) and SDSC (Thresher). The IT equipment for each cluster is housed in 10 racks at both data centers. Before the DR test, there were no jobs running on both Mako and Thresher.[10] To create a real-world HPC scenario, High Performance Linpack (HPL) benchmarking jobs were submitted to Mako and Thresher clusters to raise their CPU utilization rate to 95% and 40%, respectively. The HPL jobs loaded on the ShaRCS clusters were or random size and run time between 10 to 20 minutes. The data center operators allocated a 30% capacity reservation on Mako and migrated the jobs to Thresher, which had the capacity. At 2:30 p.m., the capacity reservation on Mako was initiated, which paused/halted its jobs and lowered the CPU utilization rate to 70%. Checkpoints were set on the paused jobs, and were migrated and run on Thresher. This raised the Thresher CPU utilization rate to 68%. After all the jobs were migrated to Thresher, 85 of the 268 Mako nodes (one-third) were idled to

[9] This analysis considered a total of 9 PDP events of 6-hour duration each in a year. However, up to 15 PDP events for 4 to 6 hours duration each can occur in a year.
[10] HPL is a global benchmarking package used to assess HPC system performance by loading test jobs.

achieve 30% capacity reservation. This reduced the Mako cluster demand by 8.7 kW (14%) and increased Thresher cluster demand by 6 kW (13.7%) in less than 4 minutes. The test was extended to study the demand savings of partial system shutdown for the 85 nodes, which were previously idled on Mako cluster. This resulted in a significant demand savings by 16.3 kW (27.6%) at the cluster level. These demands savings are small to be detected at the whole building level, which has a demand of 1 MW.

Based on scaling, for a 5% load shed at the whole building level, a load migration test would have to be conducted on 58 racks if idled, or 30 racks if they are shut down. It was noticed that the cluster management software was calculating the CPU utilization rate using the total number of active/online nodes, rather than the total number of nodes in the cluster. Thus, when the nodes were shut down at Mako, the CPU increased while the power decreased. Calculating the CPU utilization rate by the total nodes in the cluster normalized these data. Figure 7 presents the resulting important correlation between CPU utilization rate and power for both Mako and Thresher. The results using 15-minute interval data show a linear relationship between CPU and power for this equipment. The energy proportional computing started to occur as a generational improvement beginning circa 2006. The graphs show a decrease for Mako and an increase for Thresher clusters.

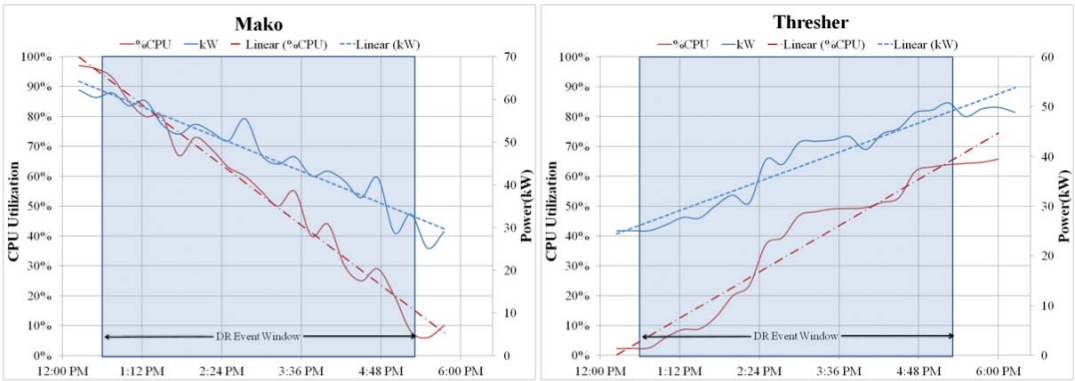

Figure 7: Correlation of CPU Utilization and Power

5.4. San Diego Super Computer Center and LBNL 50B

Two load migration tests were conducted for the LR heterogeneous system. The tests redirected all the new jobs, which were queued on LR-1 (SDSC) to LR-2 (LBNL 50B). This ascertained that the LR-2 system, with faster processing capabilities, can finish processing the IT jobs from LR-1 in similar or lesser amount of time. The study did not conduct DR tests where IT jobs were migrated to a data center with slower processing capabilities. The existing IT jobs on LR-1 were allowed to run and decay where the maximum allowed jobs duration is three days. The demand savings for LR-1 come from power savings when no new jobs are allowed to run, which gradually decline in time. This and the job decay will result in a lower CPU utilization rate and subsequent energy use. Compute nodes, which do not process any jobs, go to idle mode and lower the overall system power consumption. The response time for this load migration strategy depends on the size and duration of the job—The smaller the job, the faster the load reduction, and vice versa. The power use of LR-2 increases when the new jobs from LR-1 are transferred. Since LR is a production system, the data center operators decided to

allocate up to 30% capacity reservation block on LR-2 for DR tests. In our tests, this was not an issue, as the CPU utilization of LR-1 was less than 30% The first test on June 18, 2012, from 10:00 a.m. to 2:20 p.m. did not yield significant results. This was because on the DR test day there was just one job running on LR-1, with a CPU utilization rate of almost 0%. Also, no new jobs arrived at LR-1 to be migrated to LR-2.

The second DR test was scheduled from July 3, 2012, 10:45 a.m. to July 5, 2012, at 11:00 a.m. At the start of the test, there were eight jobs running on LR-1, at 45% CPU utilization rate. All these jobs were processed by 4:45 p.m. the same day, which dropped the of CPU utilization rate to 0%. Hence in this scenario, the job decay time was 6 hours, which is technically the end of DR test. However, the test was allowed to run until July 5, as the data center operators wanted to understand the impact of migrating the jobs to LR-2. A demand saving of 9 kW (17%) was observed at the LR-1 clusters. Since LR-1 is a portion of SDSC IT equipment and total data center load, demand savings would not be noticeable at the whole building level, which has a total demand of 2.3 MW. During the DR event period, two jobs on LR-1 were migrated to LR-2. This did not result in a noticeable load increase on LR-2. Figure 8 shows a side-by-side analysis of power and CPU utilization rate of LR-1 and LR-2 clusters for the duration of the DR event.

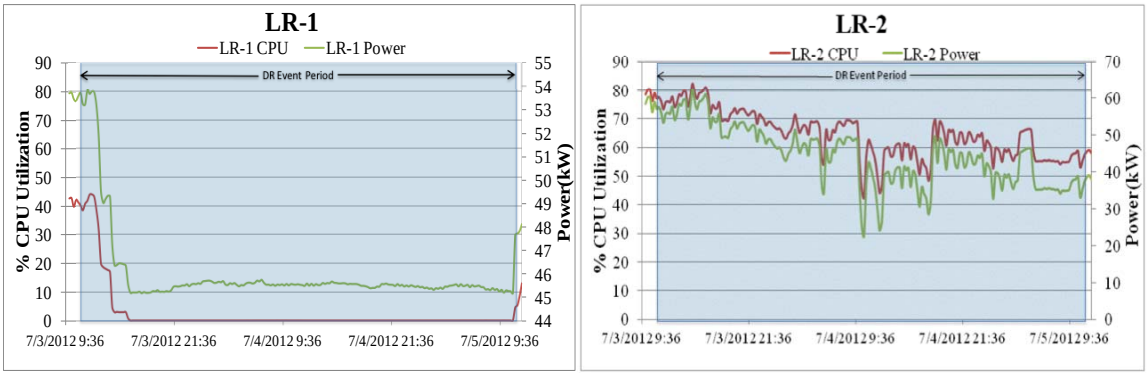

Figure 8: CPU and Power of LR-1 and LR-2 Clusters

The IT equipment CPU utilization rate depends on the IT job characteristics (size, duration, etc.), which influences its energy use. A higher number of jobs on a DR event day will result in underestimating the baseline based on utility baseline models described earlier. To accurately calculate baselines at the IT equipment level, a "job-adjusted baseline" may be needed to provision for the increase in energy use during a DR event day. This concept is similar to the "temperature adjusted" baseline used in the weather-sensitive buildings (Coughlin et al. 2009). This issue is true of most of the industrial facilities, which have high variability (Coughlin et al. 2008). However, in data centers, at the whole building or the data center level, the variability is minimal to none.

6. Conclusions and Future Research

Conclusions

The results from the implementation of manual DR strategies for both IT equipment (server, storage) and site infrastructure (cooling) identified that similar data centers can participate in DR programs with no impact to operations or service-level agreements,

which were set by the data center operators. The data centers, which provide hosted services for other customers, need to consider the service level agreements for DR participation. As data center site infrastructure becomes more efficient with lower power usage effectiveness (PUE), a large DR opportunity is within the IT equipment, with larger savings when combined with DR automation and integration with the cooling systems. With the growth in cloud computing, the load migration strategies that are unique to data centers can enable participation in DR and price-responsive programs such as ancillary services and address intermittency issues with renewables and other issues with minimal or no impact to data center's day-today operation.

However, some key barriers remain for data centers to participate in DR programs, and future studies must consider them. More field-tests of component-level performances, data center types, and automation is necessary for wide scale DR adoption. The enabling technologies need to link data center operational requirements with the supply side systems, and provide aggregation to visualize metered DR information. The evolving DR market design must describe the value proposition, the measurement and verification models, and how it can benefit a data center's social, economic, and efficiency goals. Few specific examples that lead to these conclusions are as follows:

- The DR strategies vary by data center function, IT equipment, cooling system, value to each customer, data center management, and operator's comfort.
- The IT equipment DR tests reduced other support loads. This, plus the need for a well-tuned data center for cooling strategies and reducing data center PUE, presents the largest DR opportunities within the IT equipment.
- The NetApp tests show that: (a) higher temperature set-point adjustments are possible with real-time monitoring to alleviate any equipment risks; (b) the recovery of storage equipment has to consider staged restart to avoid any rebound.
- The manual load migration tests in LBNL 50B, SDSC, and UCB required four staffs to plan and execute DR tests. This process can be made less resource intensive with technologies for automation, remote monitoring, and management.
- Enabling technologies can motivate data center operators to participate in DR programs and reduce the response and recovery times, although they are more cost effective when used to meet other data center energy-efficiency goals.
- The DR strategies for IT jobs in LBNL 50B, SDSC, and UCB—and parameters that characterize the jobs—are important, since they influence the IT equipment shed. Such jobs identify specific load shed and migration strategies.

Future Research

Although the results are encouraging, it represents a small dataset. Further studies must:

- assess the results for a diverse group of data centers to validate these findings,
- consider the integration of data centers with cost-effective automation technologies for grid integration,
- combine IT equipment and load migration strategies with the enabling technologies for different DR programs and electricity-price markets,
- evaluate end-uses and its requirements for grid responses such as price-response, ancillary services, renewable integration, distributed resources, and others.

7. Glossary of Key Terms

ASHRAE	American Society of Heating, Refrigerating, and Air-Conditioning Engineers
Auto-DR	Automated Demand Response
CAV	Constant Air Volume
CEC	California Energy Commission
CPU	Central Processing Unit
CRAC	Computer Room Air Conditioning
CRAH	Computer Room Air Handler
DDN	Data Direct Networks
DOE	United States Department of Energy
DR	Demand Response
DRAS	Demand Response Automation Server
DRRC	Demand Response Research Center
EM/4	Power Assure's Energy Management Platform
EMCS	Energy Management Control System
EPA	United States Environmental Protection Agency
Gbps	Gigabits Per Second
HPC	High Performance Computing
HPL	High Performance Linpack
HVAC	Heating, Ventilation, and Air Conditioning
IP	Internet Protocol
IT	Information Technology
kWh	Kilowatt-Hour
LBNL	Lawrence Berkeley National Laboratory
LR	Lawrencium
M&V	Measurement and Verification
MW	Megawatt
OAT	Outside Air Temperature
OEM	Original Equipment Manufacturer
OpenADR	Open Automated Demand Response Standard
PG&E	Pacific Gas and Electric Company
PIER	Public Interest Energy Research
PUE	Power Usage Effectiveness
RAM	Random Access Memory
RPM	Rotations Per Minute
SDG&E	San Diego Gas and Electric
SDSC	San Diego Supercomputer Center
ShaRCS	Shared Research Computing Service
SLA	Service Level Agreement
SVP	Silicon Valley Power
TAG	Technical Advisory Group
TCP	Transmission Control Protocol

UCB	University of California, Berkeley
UCOP	The University of California Office of the President
UCSD	University of California, San Diego
UPS	Uninterruptible Power Supply
VAV	Variable Air Volume
VFD	Variable Frequency Drive
WAPA	Western Area Power Administration
ΔT (Delta T)	Temperature Difference

8. References

ASHRAE. 2008. 2008 ASHRAE Environmental Guidelines for Datacom Equipment - Expanding the Recommended Environmental Envelope. Available at http://tc99.ashraetcs.org/documents/ASHRAE_Extended_Environmental_Envelope_Final_Aug_1_2008.pdf.

Coughlin K., M. A. Piette, C. Goldman, and S. Kiliccote. 2009. *Statistical analysis of baseline load models for non-residential buildings.* LBNL-4984E

Coughlin K., M. A. Piette, C. Goldman, and S. Kiliccote. 2008. *Estimating Demand Response Load Impacts: Evaluation of Baseline Load Models for Non-Residential Building in California.* Lawrence Berkeley National Laboratory Report LBNL-63728.

EPA. 2007. *EPA Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431.* Available at http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf. Both accessed on July 11, 2012.

Ghatikar, G., M. A. Piette, S. Fujita, A. T. McKane, J. Q. Han, A. Radspieler, K. C. Mares, and D. Shroyer. 2010. *Demand Response and Open Automated Demand Response Opportunities for Data Centers.* LBNL-3047E.

Koomey J. G. 2011. Growth in data center electricity use 2005–2010. Report by Analytics Press. Available at: http://www.mediafire.com/file/zzqna34282frr2f/koomeydatacenterelectuse2011finalversion.pdf.

Motegi, N., M. A. Piette, D. S. Watson, S. Kiliccote, and P. Xu. 2007. I*ntroduction to Commercial Building Control Strategies and Techniques for Demand Response.* California Energy Commission. LBNL-59975.

National Institute of Standards and Technology (NIST). 2010. NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 1.0. January. Available at http://www.nist.gov/public_affairs/releases/upload/smartgrid_interoperability_final.pdf.

National Institute of Standards and Technology (NIST). 2012. NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 2.0, NIST special publication 1108R2. February.

Organization for Advancement of Structured Information Standards (OASIS). 2012. Energy Interoperation Version 1.0 Committee Specification. February. Available at http://docs.oasis-open.org/energyinterop/ei/v1.0/energyinterop-v1.0.html

9. Appendices

These appendices in this section correspond to those mentioned in the different sections of the main report.

9.1. Research Methodology

TAG Members and Roles:

The following are the members of the TAG (ordered alphabetically by organization):

Organization	**Member Name**
California Energy Commission	Anish Gautam, Paul Roggensack
Critical Facilities Roundtable	Bruce Myatt
Electric Power Research Institute	Brian Fortenbery
The Green Grid	Henry Wong (from Intel)
Individual	Christian Belady (from Microsoft)
Individual	Mark Bramfitt (Consultant)
Lawrence Berkeley National Laboratory	Dale Sartor, William Tschudi
Lawrence Berkeley National Laboratory 50B	Ed Ritenour, Greg Kurtzer (Participant)
NetApp, Inc.	Stan Cox (Participant)
Pacific Gas and Electric Company	Albert Chiu, Jonathan Burrows
San Diego Supercomputer Center	Matt Campbell (Participant)
San Diego Gas and Electric	Eric Martinez
Quantum Energy Services & Technologies, Inc.	Irina Krishpinovich

During the study period, five formal TAG meetings (virtual) between August 2011 to May 2012 to review and revise the materials presented in this report.
In particular, the TAG members provided guidance and answers to the following areas to answer few key specific questions.

Review and guidance on site recruitment:

- Any additional key sites that might be appropriate for this project?
- Any additional site or technology selection considerations?

Review and guidance on documents:

- Any additional site selection criteria appropriate for the project?
- Any additional technology selection criteria appropriate for the project?
- What specific questions for the data center DR survey should be included?
- What are the opportunities and barriers for data center DR adoption?

Review and guidance on DR strategies:

- Any additional strategies you think should be considered?

The TAG members' expertise was solicited to review the field test results and further strengthen the strategies, conclusions, and recommending next steps. Before finalizing the study, TAG member meetings are scheduled to review the final field test results and the draft report.

Site Selection Criteria and Framework

The site selection criteria and framework emphasized on:

- Willingness and volunteering support for DR tests for identified strategies.
- Meet the requirements for enabling technologies.
- Evaluation of existing infrastructure for sub-metering.
- Data center location in California and at least one within each of the IOU territories.
- Data Center load, energy use, and access to data
- Different data centers types (e.g., research, non-mission critical).

In certain cases, we had to relax some of the requirements as the recruitment of data center for field tests was challenging and we had to go with the primary criteria of willingness of a data center to participate in tests and provide support. The detailed site selection criteria and the framework used for the data center recruitment are listed below:

Technology Selection Criteria and Framework

The technology selection criteria and the framework were based on the Phase one study findings and closely aligned with the project objectives.

The study of enabling technologies and to identify the ability of data centers to reduce electric power demand for pre-defined periods (typically hours) in response to a DR event is an important objective of this project. The technologies considered may offer other demand-side management services such as energy efficiency and daily peak load management. The technology selection criteria included considerations such as:

- Providing enabling technologies for IT equipment and site infrastructure.
- Relation of technologies and their integration with different equipment types.
- The scalability of enabling technologies for different data center types.
- The maturity of technologies, adoption by the market, and meet the project needs.

The detailed technology selection criteria and the framework used for the data center DR tests and identification of enabling technologies is listed below:

Assessment, Outreach, and Site Recruitment

LBNL led the recruitment for field tests with support from other organizations (e.g., Data center member associations such as the Silicon Valley Leadership Group).[11]
With a good understanding that data center recruitment will be a challenge, LBNL worked with a subcontractor, Megawatt Consulting, and other organizations. The recruitment emphasized (not mandatory) participating sites within the IOU territories and that can participate in DR. For example, we obtained the data center list from the Pacific Gas and Electric Company (PG&E) energy efficiency assessment program and conducted evaluation and outreach for recruitment. Although the original project scope included considering up to two data centers, four data centers were recruited in total for field-tests. The list of participants were reviewed and confirmed with the sponsors and the TAG members. The three participating sites are as follows:

- LBNL Building 50B 1275 (50B) data center, Berkeley, California
- NetApp Java 01 data center, Sunnyvale, California

[11] These organizations can be potentially considered for outreach for data center recruitments in commercial programs by presenting the study findings at conferences organized by such organizations.

- San Diego Supercomputer Center (SDSC), La Jolla, California

Outreach was conducted for many data centers to participate in the study. Between August and November 2011, several data centers were approached Data center operators approached included large end users, supercomputing centers and co-location data centers. The table below shows a list of data centers that engaged in discussions. Data center operators have expressed the following challenges.

- Lack of financial incentives to cover the direct cost of their resource and staffing.
- Lack of real benefits to the data center operator other than minor energy savings during the DR tests.
- Rebound impacts, which can put the IT equipment at risk of failure to return to normal state following a DR test. Possible downtown to equipment, data center or customer operations
- Lack of software that can turn off computer equipment, whether automatically or with minimal manual control.
- Their geographic presence is outside California.

Table 12 below summarizes the data centers pursued for recruitments for field tests.

Table 12: Summary of data centers engaged in outreach and recruitment

Site	Status	Location	Utility	Strategies	Notes
LBNL	Confirmed	Berkeley	WAPA	IT equipment, migration, and cooling.	R&D and super computing data center
NetApp	Confirmed	Sunnyvale	PG&E	IT equipment and cooling.	Non-mission critical data center
SDSC	Confirmed	La Jolla	SDG&E	IT equipment, migration, and cooling.	R&D and supercomputing data center
4	Declined	Silicon Valley	SVP		Cannot disclose data center information.
5	Declined	San Jose	PG&E	IT equipment power capping	Project timeline and support not aligned.
6	Declined	San Diego	SDG&E		
7, 8, 9	Declined	Silicon Valley	PG&E	Cooling	Energy Efficiency sites for PG&E. DR not feasible with timeline.

Furthermore, the methods of reducing load in a DR test for data centers were often risky or difficult to enact. Some of these methods include: temporarily increasing supply temperatures, relocating computer load to another location, turning off unnecessary equipment (e.g. redundant UPS unit(s), and other methods. Of these, most of the data centers that were approached were already operating at their maximum comfortable temperature set points (comfortable meaning for equipment, people, or to meet service level agreements or other customer needs).

Upon recommendation by the TAG members and the subcontractor, a two-page factsheet was prepared to provide an overview of the study and detail the benefits/value to the data center for participating in the tests and in DR programs. The detailed factsheet for outreach and recruitment is attached with this report. This factsheet expands from an earlier factsheet that LBNL has prepared before the Phase 1 study.

Memorandum of Understanding

The objective of the MOU was to set forth the roles and intentions of each data center participant in implementation of the upcoming field tests for the project. The MOU is not a legally binding document and establishes expectations from a participant. The MOU introduced the project, described the project, the intentions for conducting tests and collaboration, data and testing requirements, and information sharing, including general provisions such as the rules of engagement for participating in the study.

Site Survey Questionnaire

The main purpose of the survey was to assess the data center characteristics to prepare test and monitoring plans before conducting the tests. Although the preparation of a site survey questionnaire was not part of the original scope of the project, it was a useful document to understand additional information on data centers and access DR potential. The survey is useful for application to other data centers for more in-depth evaluation of both DR and enabling technologies. The survey included information such as:

- Site name and contacts
- Energy usage
- IT and facility management
- Data Center services
- Data access/availability
- Back-up and/or onsite generation capabilities

The responses from the site survey questionnaire during the data center recruitment were used to understand the data center characteristics and other information relevant to the field tests. LBNL recruited a subcontractor, SCG, for this activity and provided assistance in the process. Few sites that were part of the recruitment discussion expressed concern in disclosing any information on the data centers although the site information was kept confidential and that the results would be aggregated.

Develop Custom Field Tests and Monitoring Plans

The field test and monitoring plans were prepared based on availability of IT equipment and site infrastructure loads (e.g., cooling, servers) for DR and potential integration with California utilities' Open Auto-DR programs, including the willingness of the participating sites in the tests.

The detailed test plans were useful to outline the monitoring plans, which could be applied to similar data centers and other potential industries. These test plans and their results will provide insights on how data center types and DR strategies are applicable for participation in different DR programs and markets.

Existing metering and sub-metering systems were considered as much as possible and custom infrastructure will be installed as needed. Successful efforts were made to recruit sites that are already planning or have sub-metering infrastructure (e.g., for energy efficiency assessment). For the LBNL 50B, the enabling technology vendor, SynapSense Corporation provided the technology and sub-metering infrastructure (along with the existing 50B infrastructure). For NetApp, the enabling technology vendor, Power Assure, provided the technology and sub-metering infrastructure (along with the existing NetApp infrastructure and Automated Logic Corporation for site infrastructure data). In case of SDSC, although there is a lot of existing infrastructure, for the purposes of this study and load migration strategy, some additional sub-metering will be necessary. LBNL is worked with SDSC to identify and set up the monitoring infrastructure for the study.

DR Strategies for Field-Tests

The DR strategies for consideration included those outlined within the objectives of the project. This is another important aspect of the project study. These include considering a set of potential DR strategies for site infrastructure (HVAC, lighting, etc.) and IT infrastructure (servers, storage, etc.) loads for data centers with in-depth analysis of sequence of operations and engineering analysis to understand the opportunities and any issues for data center participation in DR programs.

Other innovative strategies potentially included virtualization and control technologies, methods and strategies to deploy standard-based Open Automated DR technologies (e.g., OpenADR) for Automated DR within the utilities' commercial programs. Load migration strategies subject to the availability of technology infrastructure were also considered within the participating data center. While the focus of this project and field-tests were to answer the basic question of if data centers are good candidates for DR, the strategies and their enabling technologies in the project looked at the potential of automation. During the recruitment of data center for field tests, we managed to get a significant support from the participating data centers to try a multitude of DR strategies for tests. These strategies and other relevant details are explained in Section 4 for each site.

9.2. Data Center Characterization

The following section provides details of data center characterization, primarily the IT equipment description at each of the participating sites.

Equipment description

This section provides technical description of the row wise configuration of the Storage equipment tested in this study. Details of the Lawrencium and ShaRCS clusters are also provided for a deeper understanding of the IT equipment used for field tests in this study. The information presented below relates to the section 3.3- IT equipment and cooling systems in the main report.

NetApp

Table 13 shows the configuration of the two rows, which were part of the DR tests conducted at NetApp. It can be seen that Row 3 has only two filers while Row 6 has 6 filers. Another key difference is the storage capacity of these two rows.

Table 13: NetApp Java-1 data center IT equipment and row configuration

NetApp Java-1	Row 3		Row 6			
Filer Name	Filer 1	Filer 2	Filer 3	Filer 4	Filer 16	Filer 17
Total Power (kW)	6.9	10.1	2.3	2.3	2.3	2.3
Filer Manufacturer	NetApp	NetApp	NetApp	NetApp	NetApp	NetApp
Filer Model	FAS 3070	FAS 3070	FAS 6080	FAS 6080	FAS 6080	FAS 6080
Storage Shelf Model	DS14MK2AT	DS14MK2AT	DS4243	DS4243	DS4243	DS4243
Number of Disk Shelves	24	35	10	10	10	10
Drives per Shelf	14	14	24	24	24	24
Drive Capacity	500 GB	500 GB	2 TB	2TB	2 TB	2TB
Power (W/shelf)	288	288	230	230	230	230
Total Capacity (TB)	164	239	480	480	480	480

Lawrencium Cluster at SDSC and LBNL 50B

Each node is a Dell PowerEdge 1950 equipped with two 2.66Ghz Xeon Intel quad-core 64-bit Harpertown processors (8 cores in all) and 16GB RAM Each node contains 16GB of memory. The Lawrencium cluster is coupled with a BlueArc high performance NFS storage server that provides a total of 48TB home and a Data Direct Networks 300TB Lustre parallel file system that provides high performance scratch space to users. Scratch space is much like temporary cache that is intended for transitory data generated during a calculation and is usually deleted shortly after computation. Additional hardware specifications are provided in Table 14 below:

Table 14: Specifications of Lawrencium cluster

Component	Description
System Name	Lawrencium
Operating System	CentOS 5 x86_64 (based on Redhat Linux)
Number of nodes (cores)	198 (1584)
Total Aggregate Memory	3.1 TB
Performance	16.8 TFlops (Peak); 12.6 TFlops (Linpack)
Processor	Intel 2.66Ghz Xeon Quad core 64-bit Harpertown
Network interconnect	Double Data Rate Infiniband (20 gb/s)
Resource Manager	PBS Torque
Job Scheduler	Moab
Disk (Storage)	Bluearc NFS (9TB Scratch & 48TB Home) Data Direct Networks Lustre (300TB)
Serial Compilers	Intel ver 10.1: Fortran 90/95, C, C++ GNU ver 4.1.2: Fortran 90/95, C, C++
Parallel Compilers	OpenMPI ver 1.2.7 (built with above Intel compilers)

ShaRCS Cluster at SDSC and UC Berkeley

The Shared Research Computing Services (SRCS) Pilot Project built by The University of California Office of the President (UCOP). The SRCS Pilot consists of two Linux clusters: the SRCS North cluster, known as Mako, located at the University of California, Berkeley (UCB), and the SRCS South cluster, known as Thresher, located at the San Diego Supercomputer Center (SDSC) in the University of California, San Diego (UCSD). Each SRCS cluster is a 272 node, 2176 processor IBM system based on the 2.4Ghz Intel Xeon (Nehalem) processors with 24GB RAM and 8MB of L3 cache.. Each compute cluster is a equipped with a high performance, low latency Infiniband interconnect ideal for a stable, high performing resource to run a wide diversity of scientific applications. Each cluster is equipped with BlueArc NFS storage and Data Direct Networks (DDN)-based Lustre file system. Additional hardware specifications are provided in Table 15 below:

Table 15: Specifications of ShaRCS cluster

Component	Description
System Name	SRCS (Shared Research Computing Services Pilot Project).
Login Host	(1) http://mako.berkeley.edu (North Cluster) and (2) http://thresher.sdsc.edu (South Cluster)
Operating System	CentOS 5.4 x86_64 (based on RedHat Linux 5.4)
Number of nodes (cores)	272 (2176) each cluster
Total Aggregate Memory	6.43 TB each cluster
Processor	Intel Xeon Quad core 64-bit Nehalem
Network interconnect	Quad Data Rate Infiniband
Resource Manager	PBS TORQUE
Job Scheduler	Moab
Disk (Storage)	Bluearc NFS Data Direct Networks Lustre
User Environment Management	Modules
Serial Compilers	Intel ver 11.1: Fortran 90/95, C, C++ GNU ver 4.1.2: Fortran 90/95, C, C++
Parallel Compilers	OpenMPI ver 1.4 (built with above Intel compilers)

Figure 9 shows the network topology of the ShaRCS cluster. The homogenous nature of this cluster can be seen clearly along with all the components, which make up the cluster.

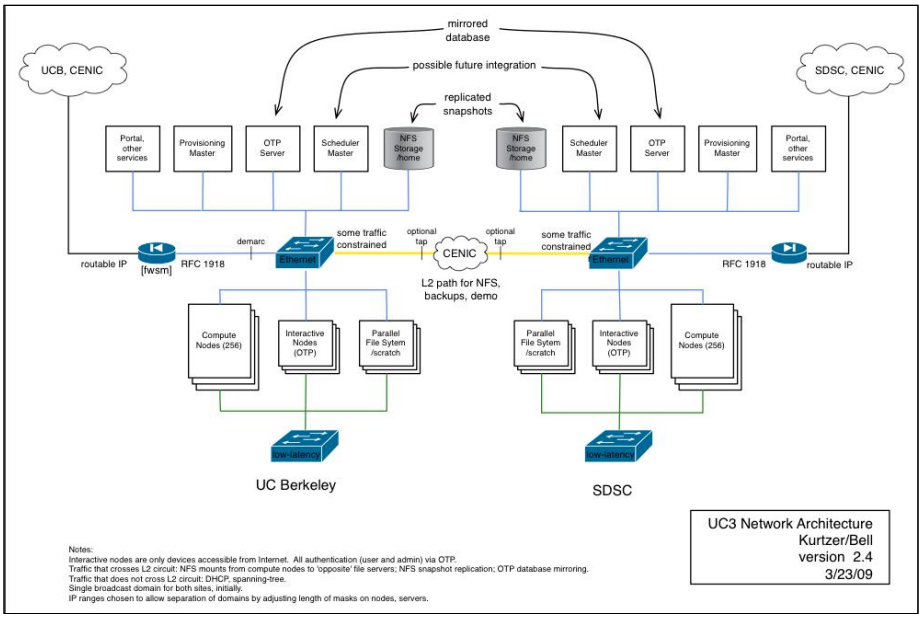

Figure 9: Network Topology of ShaRCS cluster

9.3. Results

This section presents details on different baselines used for this study along with demand savings analysis for each test conducted at the LBNL 50B. NetApp, UC Berkeley and SDSC. This information relates to the section 5 and its subsections in the main report.

Baseline Analysis

In this study we have conducted baseline analysis on all four participating data center sites. Taking the interval data from utility meters, we have carried out analysis using 3/10, 10/10 and outside air weather regression models along with morning adjustments. The models are defined as follows.

- **3/10 Baseline with morning adjustment:** Simple average over the highest 3 out of 10 most recent admissible days, with morning adjustment.
- **10/10 Baseline with morning adjustment**: 10-Day simple average with morning adjustment. The average of the hourly load over the 10 most recent admissible days before the event is used to predict the load on the event day
- **20 day Weather Regression model**– A model developed by considering 20 working days worth of demand and outside air temperature prior to the event day

LBNL Building 50B 1275 Data Center

This section presents a summary of the demand savings analysis for LBNL 50B data center as mentioned in section 5.1 of the main report.

Test-1 Server and CRAC Compressor Shut-Down Strategy – October 28, 2011

Figure 10 and table 16 shows IT power dropped from 310 kW to 130 KW, rebounding to 350 kW when the servers were turned back on and IT jobs were started. HVAC power dropped from 160 kW to 130 kW, returning to 160 kW when the compressors were

turned back on. Lights and plugs losses stayed at 24 and 5.5 kW, respectively as they were not part of tests. Table 17 shows the detailed demand savings analysis.

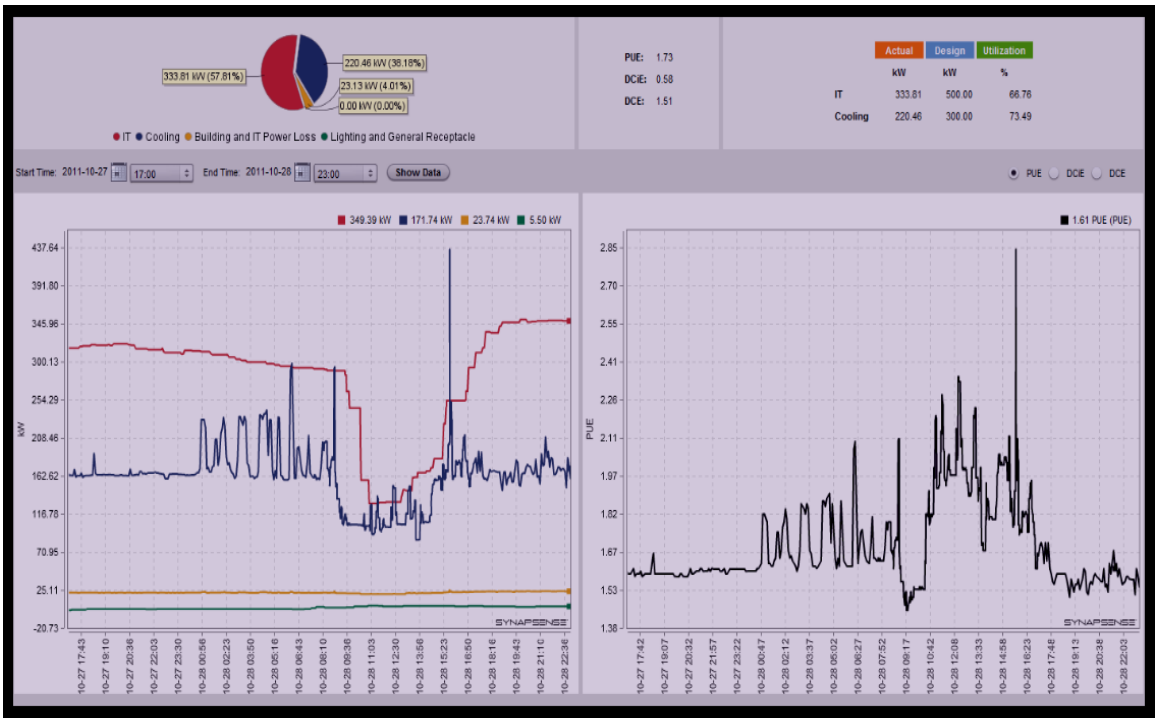

Figure 10 Demand Data by End Use for Test 1 (10/28/2011) in LBNL 50B
(IT Power shut down Started at 9 a.m., Stopped at 12 p.m.)[12]

Table 16: Analysis of Demand Savings Analysis by End Uses in LBNL 50B (Total Reduction= 272kW, IT =212 kW, Cooling =62 kW, (328 & 118))

Demand Response Test Data							
		LOAD in kW					
Date	**Time**	**IT**	**Cooling**	**Power loss**	**Lighting**	**Total**	**PUE**
11-Oct	10:00	429	177	23	5	634	1.48
26-Oct	11:00	343	166	22	2	533	1.55
28-Oct	8:00	290	163	22	3.5	478.5	1.65
28-Oct	9:30	289	117	21	4	431	1.49
28-Oct	10:00	244	104	21	4.5	373.5	1.53
28-Oct	10:35	157	103	21	5.5	286.5	1.82
28-Oct	11:00	129	97	20	6.5	252.5	1.96
28-Oct	11:30	129	136	21	6.5	292.5	2.27
28-Oct	11:45	131	104	20	6	261	1.99
28-Oct	12:00	130	101	20	5.5	256.5	1.97
28-Oct	12:30	131	150	20	5.5	306.5	2.34
28-Oct	13:00	146	104	20	6	276	1.89

[12] **Note**: The single data point spikes in power demand are a meter issue and not actual demand.

28-Oct	14:00	166	86	21	6	279	1.68
28-Oct	15:00	183	157	22	6	368	2.01
28-Oct	15:30	225	158	22	6	411	1.83
28-Oct	15:50	253	253	25	6	537	2.12

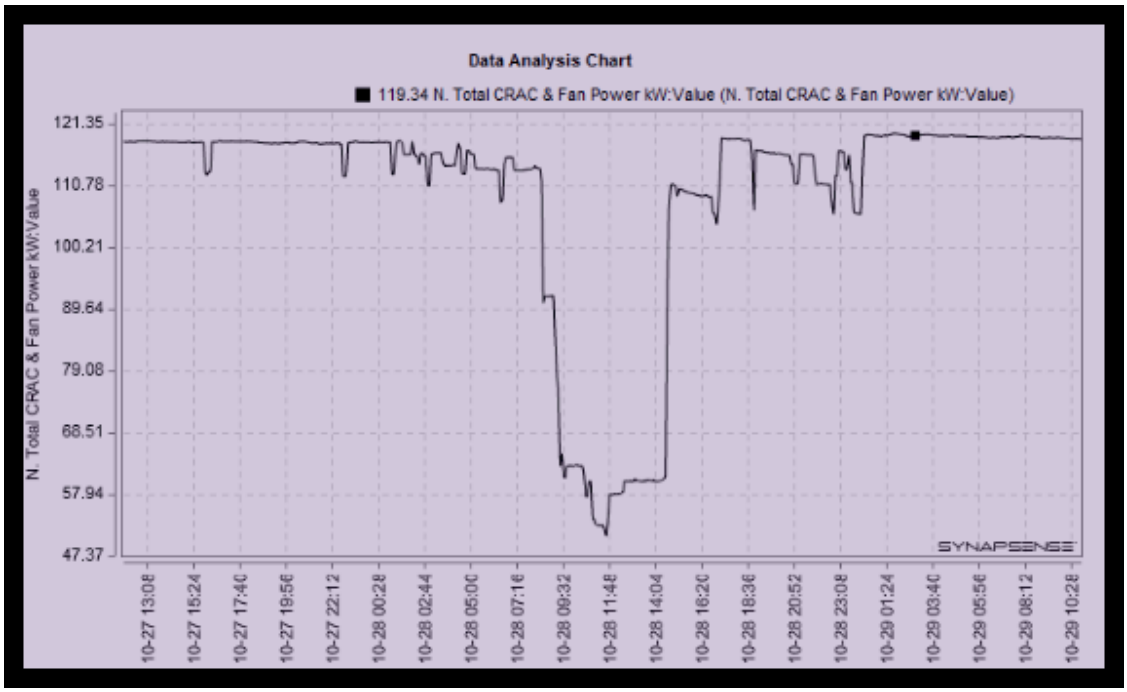

Figure 11: CRAC and Fan Power

Table 17: Whole Facility Power Demand Savings for Test 1 in LBNL 50B

Period	kW			W/ft^2			WFP%		
	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
08:30:00 - 09:30:00	98	159	126	24.43	39.63	31.62	19.6	31.8	25.3
09:30:00 - 10:30:00	213	238	226	53.25	59.55	56.56	42.6	47.5	45.1
10:30:00 - 11:30:00	215	245	237	53.66	61.32	59.16	42.4	48.6	46.9
11:30:00 - 12:30:00	205	241	228	51.18	60.29	56.88	40.1	47.4	44.7
12:30:00 - 13:30:00	208	234	216	52.11	58.43	53.9	40.6	45.5	42
08:30:00 - 13:30:00	98	245	206	24.43	61.32	51.62	19.6	48.6	40.8

Figures 12 and 13, respectively, show the whole data center level and the sub-level of the total CRAC and fan power against the corresponding adjusted 10/10 baseline and the 10/10-MA baseline. Table 18 shows a summary of hourly demand savings from these results.

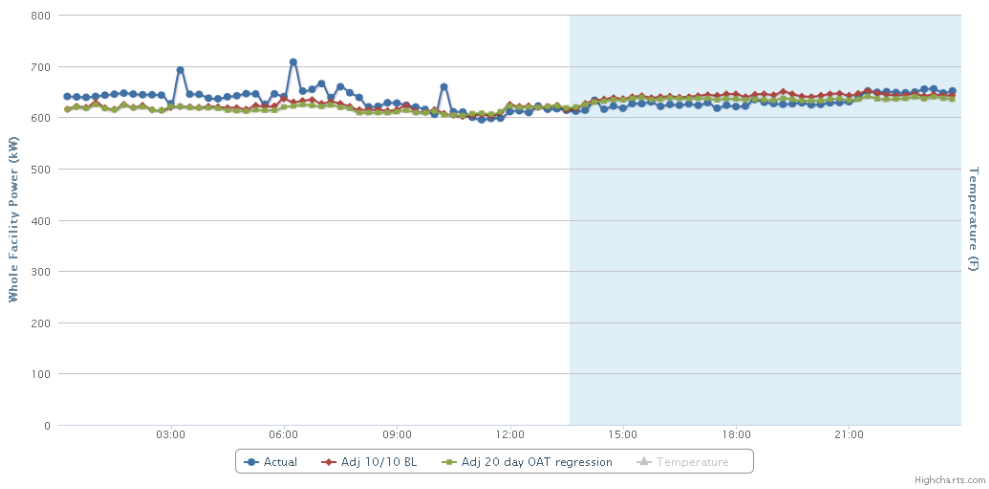

Figure 12 - Baseline analysis of LBNL Building For Test 3 in LBNL 50B

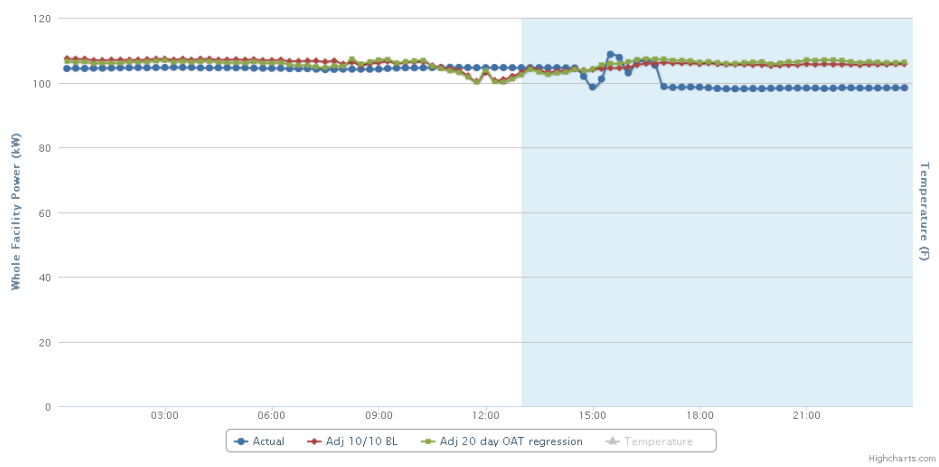

Figure 13: Baseline and demand of CRAC & Fans for Test 3 in LBNL 50B

Table 18: Whole Facility Power Demand Savings for Test 3 in LBNL 50B

	kW			W/ft2			WFP%		
Period	Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
12:35:00 - 13:35:00	6	6	6	1.42	1.42	1.42	0.9	0.9	0.9
13:35:00 - 14:35:00	-1	19	12	-0.14	4.82	2.98	-0.1	3	1.9
14:35:00 - 15:35:00	8	18	14	2.09	4.6	3.41	1.3	2.9	2.1
15:35:00 - 16:35:00	14	18	16	3.52	4.56	3.98	2.2	2.9	2.5
16:35:00 - 17:35:00	15	25	20	3.78	6.14	5	2.3	3.8	3.1
17:35:00 - 18:35:00	10	25	17	2.42	6.13	4.26	1.5	3.8	2.6
18:35:00 - 19:35:00	13	25	18	3.26	6.23	4.59	2	3.8	2.8
19:35:00 - 20:35:00	16	18	17	3.98	4.47	4.26	2.5	2.8	2.6
20:35:00 - 21:35:00	-1	12	5	-0.36	2.99	1.16	-0.2	1.9	0.7
21:35:00 - 22:35:00	-5	-3	-5	-1.37	-0.84	-1.15	-0.9	-0.5	-0.7
22:35:00 - 23:35:00	-13	-5	-10	-3.37	-1.22	-2.39	-2.1	-0.8	-1.5
13:30:00 - 23:45:00	-13	25	10	-3.37	6.23	2.58	-2.1	3.8	1.6

Test 4 - Complete IT and CRAC Units Shutdown:

This partial shutdown resulted in the IT power demand dropping from 330 kW to 55 kW while the facility power demand stayed about the same at 23 kW. Plug loads and losses stayed constant at 22 kW and 0.5 kW, respectively.

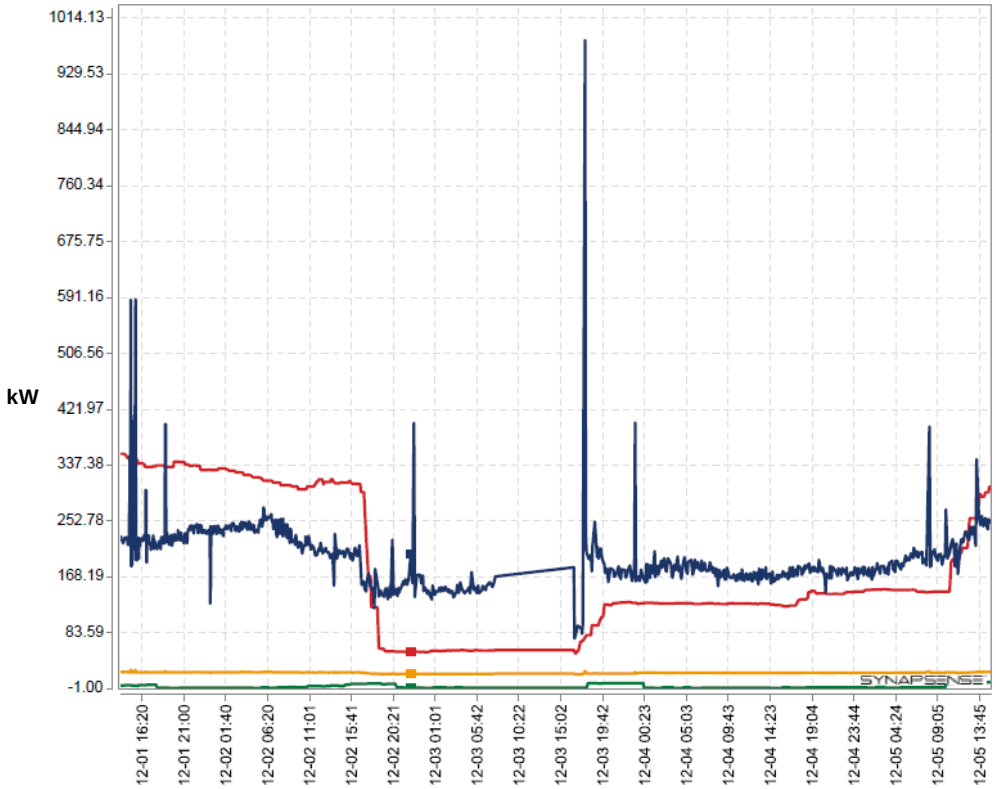

***Legend**: Red: IT Load, Blue: Facility Load, Yellow: Plug Load, Green: Losses*

Figure 14 Recorded Demand Data by End Use – LBNL Building 50 Data Center 1275 – Test4 - Complete IT and CRAC Units Shutdown – Dec 2, 2011[13]

Figures 15 shows the data center level against the corresponding adjusted 10/10 baseline and the 10/10-MA baseline for December 2nd, complete shutdown, and December 5th, post recovery back to normal state. Note that since the 10/10-MA baseline can do "per day" analysis for the DR-event day, the 10/10-MA baseline is shown just for December 2nd. As shown in Figures 15 and 16 as well as in Table 19, hourly demand savings, this test resulted in demand savings averaging 48.9% of the total data center facility demand.

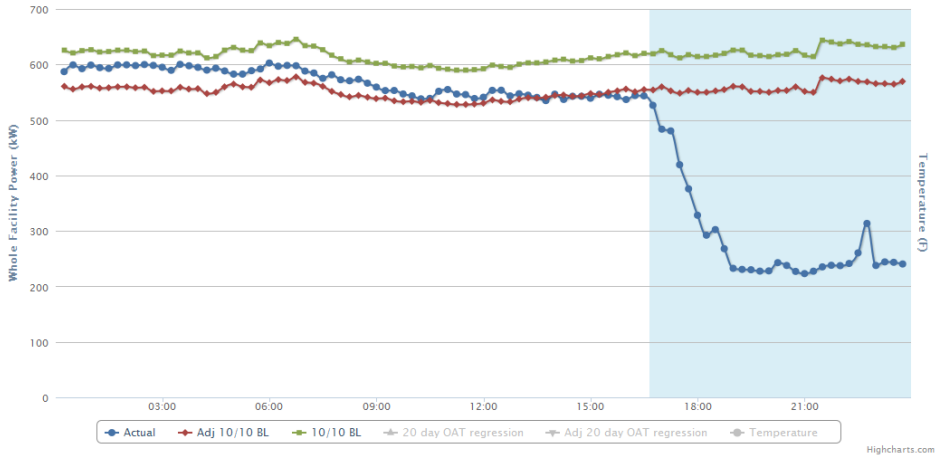

Figure 15 - Baseline analysis of LBNL Test 4 (12/2/2011)

[13] **Note**: The single data point spikes in power demand are a meter issue and not actual demand.

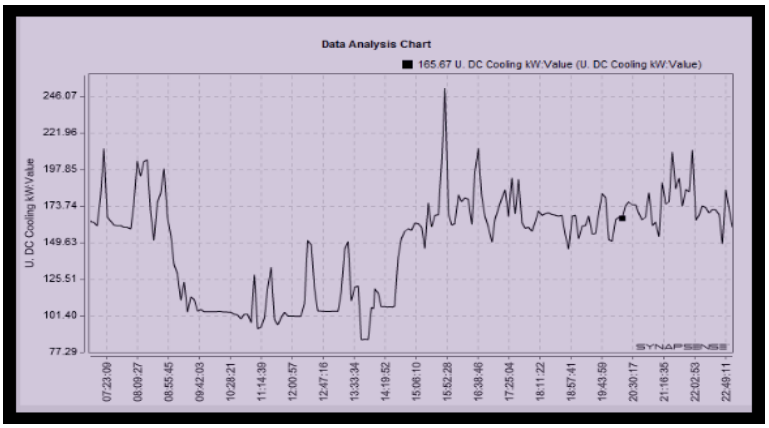

Figure 16- Total data center cooling power (12/2/2011)

Table 19: Whole Facility Power Demand Savings - Test 4

Baseline	Period	kW			W/ft^2			WFP%		
		Min	Max	Avg	Min	Max	Avg	Min	Max	Avg
Adj 10/10 baseline	15:40:00 - 16:40:00	28	28	28	6.99	6.99	6.99	5	5	5
	16:40:00 - 17:40:00	72	177	114	18.11	44.22	28.38	13.1	32	20.5
	17:40:00 - 18:40:00	221	287	254	55.31	71.69	63.47	40.3	51.7	46
	18:40:00 - 19:40:00	322	329	326	80.43	82.34	81.43	58.3	58.8	58.6
	19:40:00 - 20:40:00	310	332	320	77.55	83.11	80	56.1	59.4	57.8
	20:40:00 - 21:40:00	323	341	332	80.63	85.24	82.96	58.5	59.6	59
	21:40:00 - 22:40:00	255	333	308	63.84	83.3	76.88	44.9	58.4	53.9
	22:40:00 - 23:40:00	321	329	325	80.3	82.34	81.24	56.8	58	57.4
	16:30:00 - 23:45:00	28	341	274	6.99	85.24	68.43	5	59.6	48.9

Thermal Mapping Technology: The thermal maps by SynapSense® provided real-time visualization of the heat generated by the IT equipment and thermal conditions of the LBNL-B50 data center during the DR test conducted on Oct 28, 2011 from 8:00 a.m. to 5:00 p.m. The technologies is useful for temperature set point adjustment DR strategies as it provides continuous monitoring and real-time visualization of the temperature and any occurrence of hot spots that may lead to equipment failures or to manage SLAs in hosted data centers. For DR, application faster refresh rates of temperature maps enable customers make larger set point adjustments to achieve higher demand savings. Figure 17 to 20 show the top and bottom rack temperatures at the start and end of the test respectively.

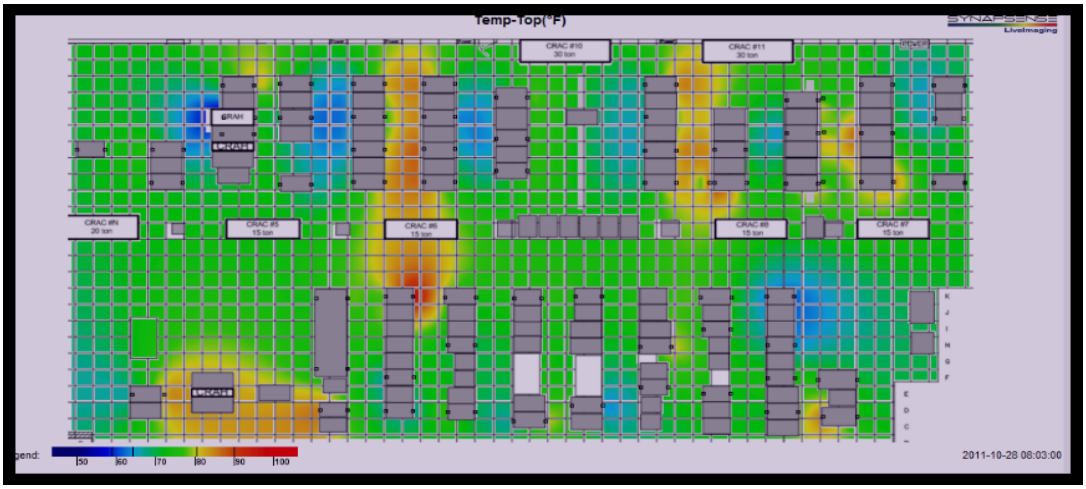

Figure 17- Rack Top Air Intake Temperature Map at 8:00am, Oct 28, 2011
Legend: Blue = 60°F, Green = 70°F, Orange= 80°F, Red = 90°F and greater

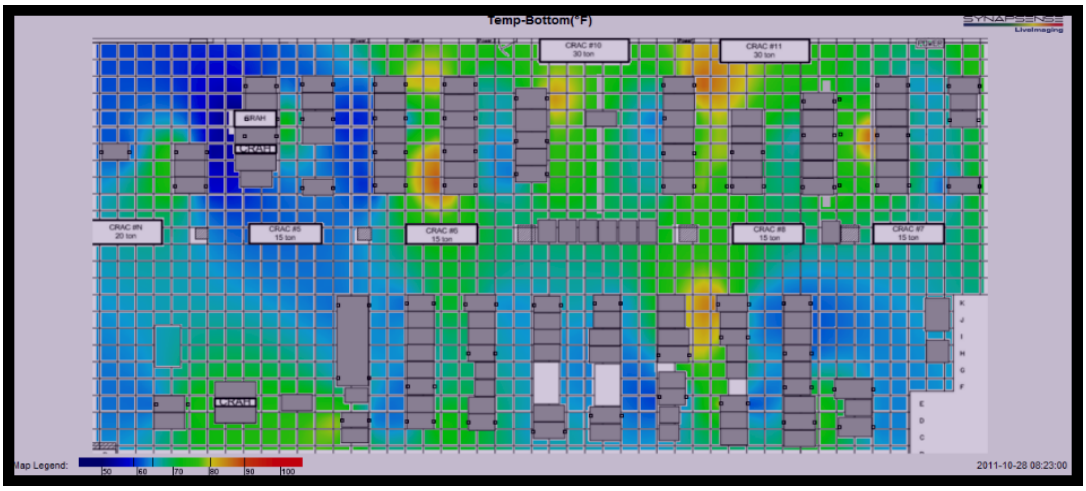

Figure 18- Rack Bottom Air Intake Temperature Map at 8:00am, Oct 28, 2011
Legend: Blue = 60°F, Green = 70°F, Orange= 80°F, Red = 90°F and greater

Figure 19- Rack Top Air Intake Temperature Map at 5:00pm, Oct 28, 2011
Legend: Blue = 60°F, Green = 70°F, Orange= 80°F, Red = 90°F and greater

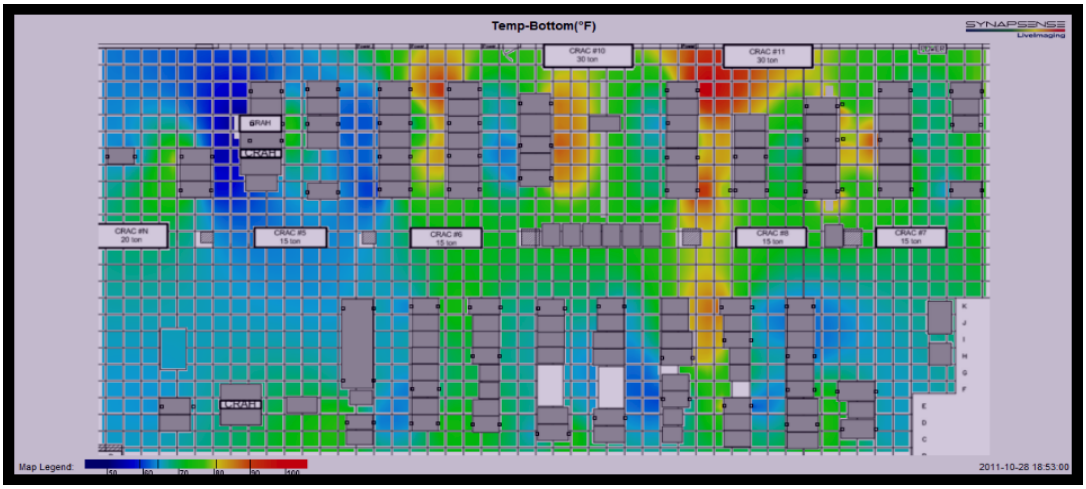

Figure 20- Rack Bottom Air Intake Temperature Map at 5:00pm, Oct 28, 2011
Legend: Blue = 60°F, Green = 70°F, Orange= 80°F, Red = 90°F and greater

NetApp

This section presents a summary of the demand savings and economic analysis for NetApp Java-1 building as mentioned in section 5.2 of the main report.

Field Tests: Baseline and Demand Savings Analysis

Table 20: NetApp DR Test Demand Savings Summary
Load Shed (kW) and (% savings at data center level)

DR Test No.	IT load	Site infrastructure	Total
1	25 kW (15%)	17 kW (10%)	41kW (25%)
2	No IT Shed	17 kW (10%)	17 kW (10%)
3	24kW (14%)	28 kW (17%)	52 kW (31%)
4	38 kW (23%)	13 kW (8%)	51 kW (31%)

The DR savings analysis required studying the actual load shed in comparisons to the baselines for IT load, site infrastructure and demand at data center and at the whole building level. The figures below show the charts derived from such an analysis. In case of the first test, we did not have access to enough historical data for IT equipment to develop baselines. However, historical data for site infrastructure and demand at whole building was available which was helpful in the analysis. The charts presented in this section are plotted with demand on the y-axis and time on x-axis. The shaded blue portion of the chart indicates the DR event period. These charts include DR savings of IT and site infrastructure at the data center level. The figures below also show the total demand savings at the data center level and the whole building level. Figure 21 shows these charts for DR test-1 at NetApp. Similar figures are presented below for Test 2 (Figure 22), Test 3 (Figure 23) and Test 4 (Figure 24).

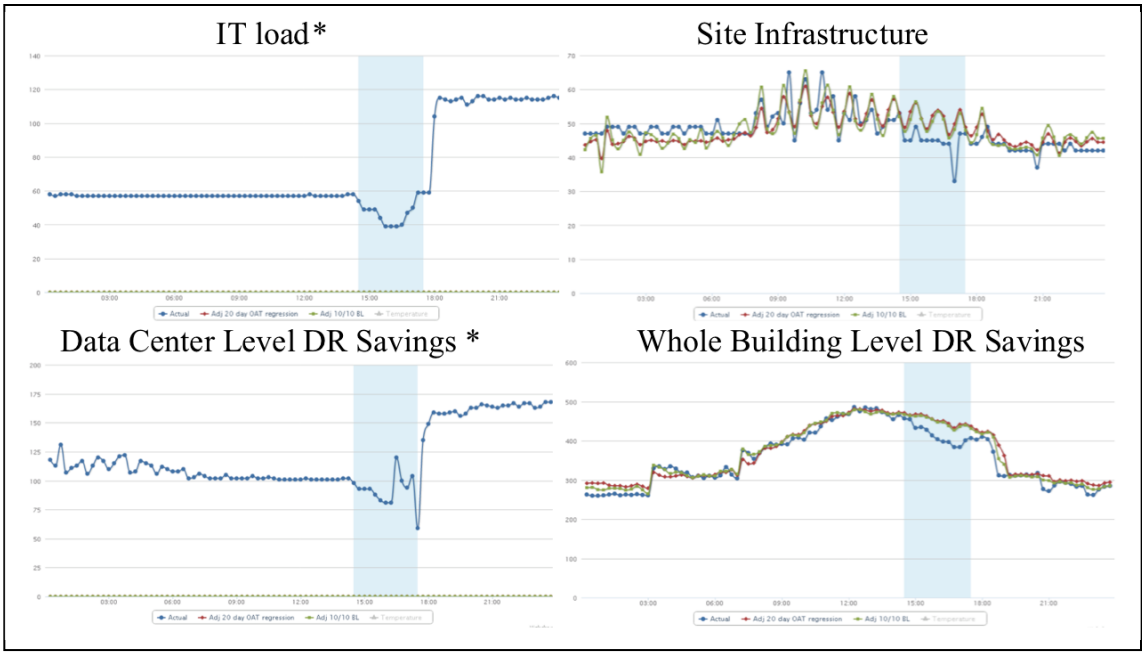

Figure 21: NetApp DR Savings Analysis for Test 1

Figure 22: NetApp DR Savings Analysis for Test 2

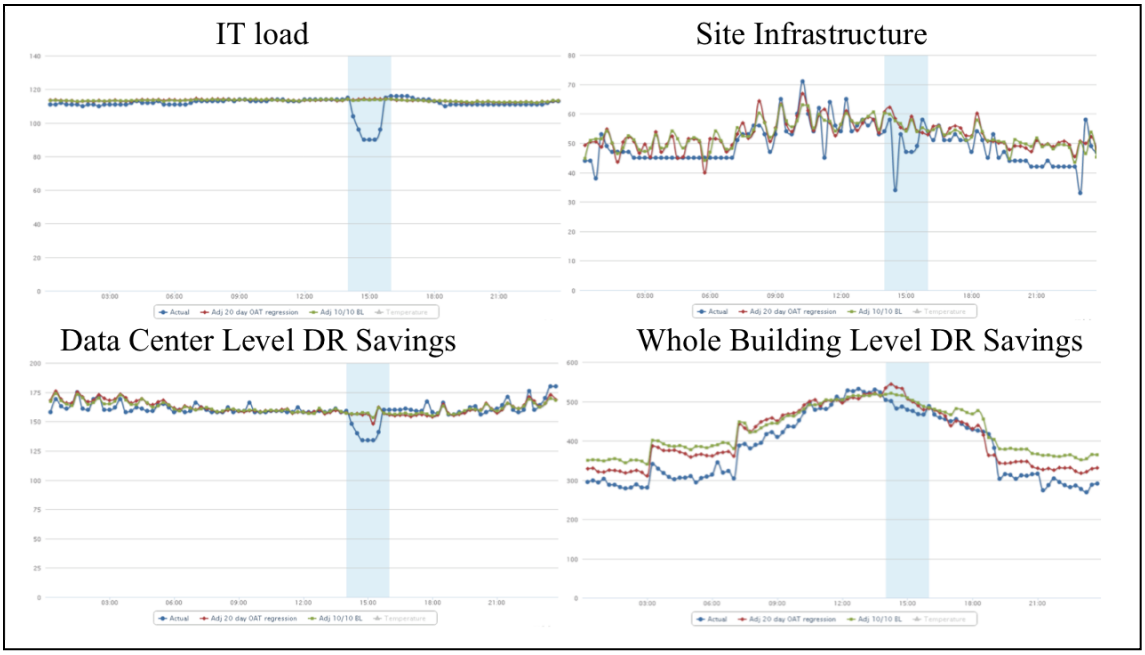

Figure 23 NetApp DR Savings Analysis for Test 3

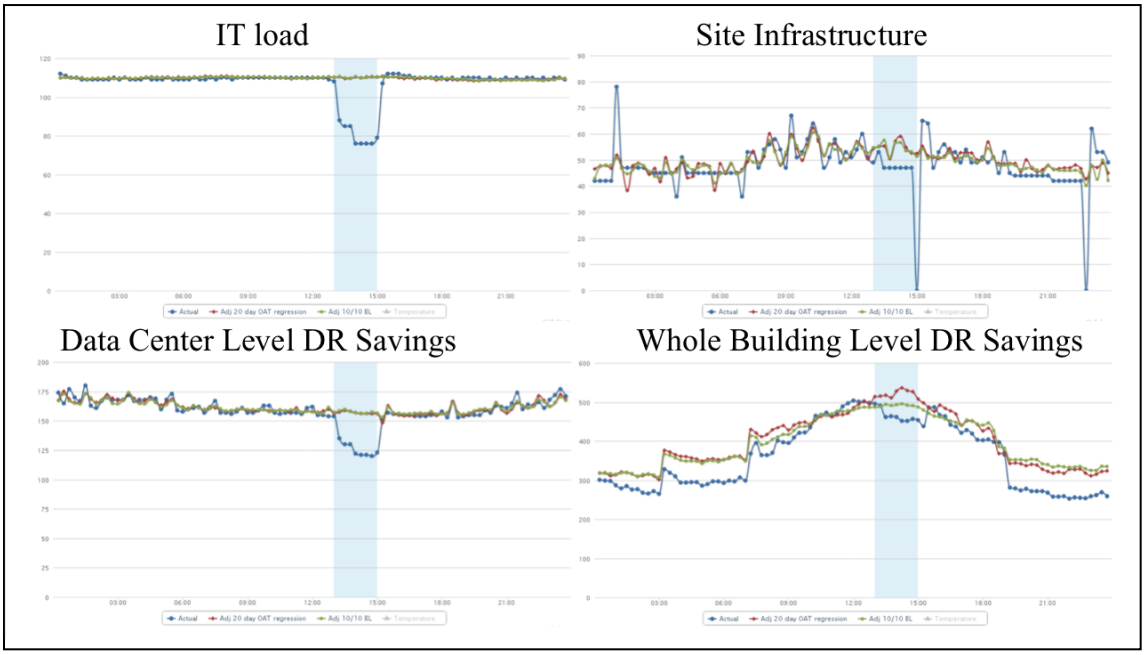

Figure 24: NetApp DR Savings Analysis for Test 4

Economic Analysis

To understand and highlight the economic value of data centers to participate in DR programs, this study has looked at a simple economic analysis. This section relates to section 5.2.2-economic analysis conducted at NetApp. The figure 25 shows the load profile characterization of NetApp data center to indicate it is relatively flat load. The slight increase in the energy during afternoon period is from the office spaces.

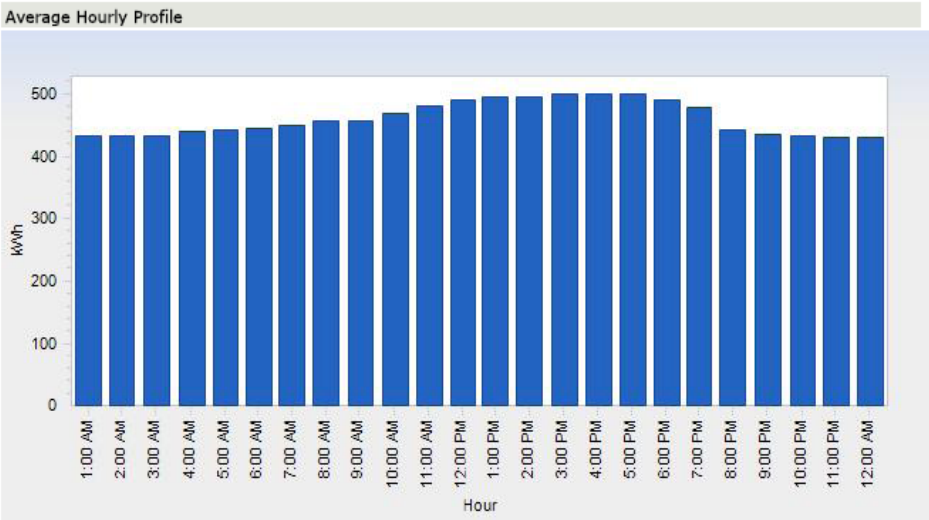

Figure 25: Flat Load Profile of NetApp Java-1 Building