Data Center Demand Response via Server Power Management

DR13.08.00

Prepared by:

Emerging Markets & Technologies
Customer Programs & Service
Southern California Edison

June 2017

SOUTHERN CALIFORNIA
EDISON®
An EDISON INTERNATIONAL® Company

Acknowledgements

Southern California Edison's Emerging Products (EP) group is responsible for this project. It was developed as part of Southern California Edison's Emerging Markets & Technologies Program under internal project number DR13.08.00. Jay Madden conducted this technology evaluation, with overall guidance and management from Paul Delaney. Contact Jay.Madden@sce.com for more information on this project.

The project team, led by Electric Power Research Institute (EPRI – Micah Sweeney, Mukesh Khattar), conducted the technology evaluation and prepared this report. The project team included the California Institute for Telecommunications and Information Technology (Calit2) at the University of California, Irvine (Dr. G.P. Li, Chris Battista); Intel (Derek Collier, Enrique Castro-Leon); Schneider Electric (Phil Davis); and IPKeys (Jim Boch, Philip Waldvogel, Aaron Jeskie).

EXECUTIVE SUMMARY

Data centers, spaces that house computers, have been considered by utilities to be a good candidate for energy efficiency improvements, due to their high energy intensity. Yet the Demand Response (DR) load impact capabilities and participation in utility programs of data centers have not been fully explored. In recent years, some efforts have sought to investigate the technical potential for DR in data centers, but these are still in the early exploration stages.

The goal of this study is to investigate and test the capability of utilizing recently-developed server management software to limit server power in response to external signals, for potential data center DR. One of the largest chip manufacturers has recently developed a technology for managing server nodes. This technology takes advantage of the built-in functionality of limiting power supply in modern Information Technology (IT) equipment hardware, such as servers, storage, and network equipment. When total power delivery is curtailed to a server, the software adjusts power to its various components, such as the Central Processing Unit (CPU), memory, and storage, to stay below the software's power limit. For example, it would adjust the supply CPU's voltage and operating frequency to stay within the power limit. In restricting the power consumed by the servers and other IT equipment in data centers, it is expected that the power of supporting infrastructure—including both cooling systems and losses in the power delivery and conditioning equipment—will respond correspondingly, for a cascaded effect.

In a project funded by SCE, and in collaboration with several industry partners, the Electric Power Research Institute (EPRI) deployed a small proof-of-concept demonstration of this technology in its Innovations in Datacenter Efficiency Advances Laboratory (IDEA Lab, as well as a field site at Calit2 (California Institute for Telecommunications and Information Technology) at the University of California, Irvine (UCI). The technology selected for this demonstration is the power-capping feature in servers using the latest line of Intel® processors, enabled by Intel's Node Manager technology. This study seeks to better understand the capabilities of this feature in terms of power reduction, time to respond, impact to operations, and post-event rebound, as well as to show the ability for this feature to be initiated as an automated response to an external DR signal.

Laboratory testing was conducted by applying a power limit to servers running a set CPU-intense benchmarks (specifically using LINPACK, an industry-accepted test protocol used to rank and measure peak performance of supercomputers on the TOP500 list in terms of Floating Point Operations per Second, or FLOPS). Results from laboratory tests of power capping demonstrate this feature can successfully limit the instantaneous power draw of a server to a defined level. Various-sized benchmarks were run, to determine the impact of power capping on increasingly CPU-intensive workloads. These workloads were run with decreasing power caps (lower power levels) from a baseline of 152 W (100%) down to a minimum power level of 72 watts (W) [~50%]. These results indicate power caps successfully limit the power draw of the server under workload, and show more pronounced impact on all workloads at the lower power levels.

When the power supply to a server was constrained, the total time to execute the compute-intensive LINPACK workload was increased. Though the power demand was reduced successfully, the power capping was found to increase the energy required to complete the LINPACK workload, due to the increased time required to complete the benchmark workload. This suggests power capping can work successfully in reducing power demand in data centers, but may increase energy use to complete finite workloads, particularly for

CPU-intensive workloads such as the LINPACK benchmark. However, the primary purpose of DR is to reduce demand due to limited availability of power during periods of constrained power supply, which was demonstrated successfully. It is possible that less-computationally-intense workloads, such as transaction-based or web-services workloads, may exhibit differently in execution under power caps. It is recommended that future work examine power capping under a variety of real-world applications, to better examine the impact on different types of workloads.

The quick response of servers to a power cap command indicate this technique may be capable of responding to fast dispatch DR signals, or providing ancillary services to the grid, such as frequency or voltage regulation. In addition, the inherent intelligence and communications abilities of servers offer the potential for quick, flexible, and complex DR behavior, with limited additional hardware costs. As power grids integrate more renewable resources, which are intermittent by their nature, increasingly quick and fast-acting DR capabilities will be needed. Data centers could be good resources for fast acting DR. Although data center workloads centers are often considered mission critical—which is one reason data centers are reluctant to adopt DR—not all data center workloads are the same. They have different criticality levels, and those that are less time sensitive may be well suited for regulation, fast and flexible DR, and the integrated grid.

ABBREVIATIONS AND ACRONYMS

BMC	Baseboard Management Controller
Calit2	California Institute for Telecommunications and Information Technology
CPU	Central Processing Unit
DCIM	Data Center Infrastructure Management
DCM	Data Center Manager (Intel product)
DOE	Department of Energy
DR	Demand Response
EISS®	Energy Interop Server & System
EPA	Environmental Protection Agency
EPRI	Electric Power Research Institute
FLOPS	Floating Point Operations per Second
HPC	High-Performance Computing
HVAC	Heating, Ventilation, and Air Conditioning
IDEA Lab	Innovations in Datacenter Efficiency Advances Laboratory
iDRAC	Integrated Dell Remote Access Controller
IoT	Internet of Things
IPMI	Intelligent Platform Management Interface
IT	Information Technology
LAN	Local Area Network
LBNL	Lawrence Berkeley National Laboratory
PMBus®	Power Management Bus
PSU	Power Supply Unit
PUE	Power Usage Effectiveness
PV	Photovoltaic
RAM	Random Access Memory
SLA	Service Level Agreement
SUT	Servers Under Test
UCI	University of California, Irvine
UPS	Uninterruptible Power Supply
VEN	Virtual End Node
VTN	Virtual Top Node

CONTENTS

FIGURES

TABLES

INTRODUCTION

Data centers have become the backbone of the modern digital society, where countless computers and terabytes of data are housed. One of the most energy-intense of commercial/industrial building types, data centers, can consume more than 40 times the power density of conventional office space. Industry experts estimate that data centers are responsible for roughly 2% of the electricity consumed in the United States [1] and it is growing at 8% annually. An Environmental Protection Agency (EPA) report estimates that 20% of U.S. data center energy use is in the Pacific region alone [2]. However, the industry has made great strides in improving infrastructure energy efficiency, such as cooling and power delivery in recent years. A recent report from the Department of Energy's (DOE) Lawrence Berkeley Laboratory suggests that in recent years, the electricity growth rate has slowed down to less than 1% a year [3]. With projected growth in internet connected devices, commonly known as the Internet of Things (IoT), electricity demand for networking, data storage, computing, and infrastructure should continue to increase.

Data centers have typically been considered business-sensitive electric loads that are not suitable for participation in demand response (DR) programs. DR is defined as a voluntary dynamic change in electricity usage, coordinated with power supplier or market needs [4]. Specifically, DR is strategy electric utilities can use to send a signal for the end user to momentarily curtail the load of their end-use devices and processes on the power system during periods of constrained power availability or high prices. In return, utilities provide incentives to program participants. Utility DR programs provide valuable tools to grid operators in providing reliable and affordable electricity in the face of increasing demand peaks that strain grid capacity. To date, a variety of technologies have been utilized to achieve DR, including lighting, water heaters, and Heating, Ventilation, and Air Conditioning (HVAC) equipment, as well as some industrial processes that are not time sensitive and can be rescheduled after DR events.

Despite the size of data center loads on the grid, they have not yet been considered suitable for DR, perhaps for reasons such as: the flat nature of their load profiles on the grid, the mission-critical nature of the workload, and a reluctance to introduce any risk of interruption to business operations. Previous studies have identified several opportunities to reduce demand in data centers [5] [6], which have focused on the infrastructure, e.g. cooling equipment thermal set points, turning off idle servers and IT equipment (which also cascaded into additional reduction in cooling demand), and shifting work load in High-Performance Computing (HPC) environments to different times. In the last 10 years, technologies have been shown to reduce data center load and sustain operations over a specified period-of-time, offering a significant opportunity for balancing demand on the grid [5] [6]. In addition, most servers are equipped with built-in sensors and controllers that enable power adjustments; yet this telemetry has not yet been exploited for DR in practice today.

This project investigates a novel method of DR of the IT equipment itself, by directly restricting the power supply to servers. Such an approach is expected to yield compound demand reductions by reducing the load on the overhead systems required to support IT equipment in the data center, including power conditioning and cooling systems.

The goal of this study is to gain insights into how data centers and their large electrical loads can be utilized as a grid resource by participating in DR. First, an emerging technology solution (Intel's Node Manager) was identified to curtail server demand at the hardware level. To respond to external DR signals, code was developed to enable the server's power management software to communicate with industry-standard DR protocols. This was developed by project partners Schneider and IPKeys. Finally, a proof-of-concept demonstration was tested in the laboratory and at a data center field site (Calit2 at UCI) to validate the performance of this technology.

BACKGROUND

ENERGY CONSUMPTION IN DATA CENTERS

Data centers support a variety of digital services that can be critical to modern business operations, including email, e-commerce, website hosting, financial transaction, database, and file server support, among many others. Given the importance of these services, data centers are designed to provide a very high level of service availability. Typically, data center operators have Service Level Agreements (SLAs) with their users (clients, in the case of third-party facilities, or IT staff within companies that operate their own data centers) to ensure IT availability matches their business objectives.

Due to the critical services data centers support, downtime can be costly. A 2013 study on outages in data centers (commissioned by Emerson Network Power, a leading provider of critical power products) found that unplanned downtime costs large data centers an average of $7,900 per minute [7]. For this reason, data centers typically employ several types of power protection equipment to prevent digital service interruption. To provide power during extended utility power interruptions, many data centers rely on backup generators, typically fueled by diesel. To prevent sensitive IT equipment from disruption due to small dips in utility voltage (called "voltage sags") data centers employ Uninterruptible Power Supplies (UPS). To provide "bridge power" as backup generators initialize, some form of energy storage is incorporated, typically as batteries, ultracapacitors, or flywheels that may be integrated into the UPS. In many cases, data centers install redundant equipment to protect against failure of any single device. The number of redundant devices (or power paths) is usually specified as N, N+1, 2N, etc., determined by the required level of availability. Each of these devices incurs electrical losses, which directly reduce the efficiency of the data center and can increase the load on cooling systems when located in a conditioned space.

The energy used within data centers can be attributed to three primary loads: IT equipment, cooling system, and power delivery (power conditioning, conversion, and backup sources, such as batteries and generators). One of the primary metrics used by industry to measure the efficiency of data centers is Power Usage Effectiveness (PUE), which is calculated as the ratio of total data center energy use divided by the energy used by the IT equipment. Such a metric gives insight into the energy used by "overhead" (infrastructure) equipment relative to the amount of IT equipment, with an ideal value of 1.0 (all energy used by IT equipment, and zero consumed by infrastructure). Yet PUE is often criticized for ignoring the efficiency of the IT equipment itself. Several industry studies in recent years have reported industry-average PUE to fall between 1.8 and 2.0, indicating that for every Watt consumed by IT equipment, 0.8 to 1.0 Watts are consumed to support this equipment.

Unlike most electrical loads, data centers exhibit near-constant load on the grid, with relatively flat demand profiles in terms of daily and seasonal changes. This is primarily attributed to servers being continuously-powered—without entering sleep or hibernation mode—to provide uninterrupted services. In addition, data center HVAC load exhibits much less dependence on outdoor conditions than occupied space, because the primary source of cooling load is the IT equipment, which operates continuously. Yet data center cooling systems can exhibit changing efficiency curves due to improved cooling system efficiency at low outdoor temperatures, introducing some variation in cooling system demand with daily and seasonal changes. These changes are especially pronounced in data centers that utilize economizers.

Background on Demand Response

DR is defined as a reduction in end-use electricity consumption, often to balance supply and demand on the power system. This approach is also typically used to relieve the strain placed on grid assets that are utilized near capacity on days when demand on the grid is high, avoiding the need to expand grid capacity or utilize expensive "peaker" plants. In some regions, peak system demand is experienced during late afternoon and early evening hours, on the hottest days of the summer when nearly every air conditioning unit is running. Yet in other regions, system peak can occur on the coldest winter mornings, as heaters and water heaters are activated to provide warmth and hot showers. In either case, DR can be a valuable resource to mitigate the cost of peak demand to utilities—and by extension, to their customers. Note that DR requires the average power of electrical load be reduced during peak periods, meaning DR must reduce energy use kilowatt-hours (kWh) over the interval.

In addition to relieving strain on the grid caused by peak demand, DR can also be used to react to unplanned events that threaten grid stability. For example, when a generator or transmission line trips offline due to an unforeseen circumstance or disruption of planned operation, DR can be used to quickly match supply and demand on the grid, to avoid wide-scale disruption. In addition, fast-reacting DR is increasingly viewed as a tool to respond to the rapid variation of renewable energy sources, such as wind and solar photovoltaics (PV) on the grid. To provide the quick response required to react to these sources, markets have formed for ancillary grid services, such as frequency regulation, voltage support (providing reactive power), and reserve operating capacity (historically known as "spinning reserves").

The earliest examples of DR involved grid operators calling upon large customers (via telephone) to shift or curtail load during peak periods. Such requests were typically made one day in advance, according to agreed-upon quantity in kilowatts (kW) or megawatts (MW) and duration of load shed at a specified price. Ideally, this request would have a minimal impact on productivity or occupant comfort. Some examples of traditional DR include industrial processes that are not time sensitive, thermostat set points for HVAC, and water heaters with storage capacity. Recent efforts have sought to automate DR, developing communications technologies to give insight to grid conditions and enhance load resource flexibility. As these technologies mature, they seek to enhance the value of DR to the grid such that it can be relied upon in much the same way as dispatchable "peaker" plants or large-scale energy storage on the grid.

To enable large-scale load response automation for DR, the Demand Response Research Center (DRRC) at Lawrence Berkeley National Laboratory (LBNL) was foundational in the development of the Open Automated Demand Response (OpenADR) communications standard. This communications protocol defines a data model by which a utility or Independent System Operator (ISO) can send DR signals to DR program participants on the power grid. The objective of the OpenADR standard is to enable large electrical loads, such as building and industrial control systems, to make predefined actions in response to power system conditions, without manual intervention. This standard in addition to facilitating DR for today's needs is expected to enable the development of an efficient transactive retail market in the future whereby energy consumers can react to peak conditions without human interaction.

Data Center Demand Response

As previously noted, data centers exhibit a relatively flat load profile to the grid, due to the continuous operation of IT equipment, and the dependence of cooling load and electrical losses on IT load. For this reason, data centers are typically considered to be attractive loads to power providers, and have not been considered significant opportunities for DR. Yet data centers may present suitable platforms for automated DR, due to the highly-automated nature of their operations.

Previous research efforts have investigated the potential of DR in data centers, and have identified several opportunities [5] [6]. Findings from these studies have identified DR opportunities in adjusting cooling system set points, rescheduling batch processes, and activating on-site generation. In particular, data centers that support "non-mission-critical" systems—such as those performing research and supporting laboratories—are excellent candidates for DR. Yet there are several barriers to widespread adoption of these techniques. First is the perceived risk to IT equipment lifetime if operated—however briefly—beyond recommended temperature ranges. Second, batch processes represent a small minority of the IT workload in data centers. Finally, use of on-site diesel generators for DR is not permitted in many regions, due to emissions requirements.

This study seeks to evaluate the DR potential of the IT equipment itself by managing server hardware power usage. Such an approach is expected to have a compounded effect at the data center level, due to dependence on electrical losses and cooling load on the IT equipment. By reducing the demand of IT equipment, it is expected that demand savings will be realized in power delivery and cooling systems, cascading to facility-level savings while avoiding the barriers associated with previously-identified DR measures.

EMERGING TECHNOLOGY/PRODUCT

This study seeks to evaluate the capability of Intel Node Manager technology to manage server power consumption from the hardware level. Compatible with Intel's latest generations of server processors (Xeon® E3, E5, and E7), the server motherboard reports power and temperature over standard communications protocols for monitoring with Data Center Infrastructure Management (DCIM) tools offered by numerous vendors to manage data center assets and operations. In addition to monitoring capabilities, these technologies enable energy consumption management by placing specific power limits on each server. The server hardware responds to these commands by adjusting its operating state in terms of both performance (P state) and throttle (T state), manipulating the voltage and frequency of the CPU (P state) and introducing a small-time delay between each computation cycle (T state). In this way, the technology is able to place a cap (or limit) on the power draw of a server or group of servers under management.

It should be noted that although Intel dominates the market for server processors, its primary competitor—AMD—offers a similar feature. On its latest CPUs, AMD has integrated a feature called TDP Power Cap (TDP refers to Thermal Design Power, which specifies to the system designer the maximum amount of heat that must be removed from the CPU under real-world workloads). However, this study chose to focus on Intel processors, due to their predominance in the server market in recent years.

This technology employs the Baseboard Management Controller (BMC) on the motherboard, to serve as the interface between the server hardware and management software. The BMC communicates using an Intelligent Platform Management Interface (IPMI), an industry-standard communications protocol for managing data center equipment. Each server manufacturer offers its own BMC technology: integrated Dell Remote Access Controller (iDRAC), HP integrated Lights Out (iLO), and others from Cisco, IBM, etc.

Power capping is enabled by Intel's Data Center Manager (DCM), a DCIM tool for managing IT assets, provisioning power, and monitoring IT equipment power and temperature. DCM's power-capping feature is designed to maintain continuity of server operations under critical power shortage, without loss of data or hardware availability. This feature allows the servers to overcome temporary power shortages. In full data centers, DCM can prioritize and respond by keeping important tasks running, and rescheduling or postponing less important workloads.

In its current configuration, DCM's simplest method of power management is provided by an "Emergency Power Reduction" mode, which can be configured to apply specified power cap levels to different servers (or groups of servers) in data centers. The default setting for this feature is to set a power cap at the minimum power level of all non-critical machines. Yet DCM allows more precise server power control by manual input of power limits.

Despite the relatively flat load profile of data centers, the power draw of each individual server modulates with the IT workload it supports. Depending on the volume and type of workload provided, server power may rarely reach its peak rating. In fact, many servers exhibit power levels close to idle, with only sporadic bursts of activity. Depending on server hardware, idle power may represent 50% or less of peak power. In emergency power mode, the DCM manager caps server power draw to little more than idle power.

COMMERCIAL MATURITY AND BARRIERS TO ADOPTION

Data Center Manager is Intel's offering for DCIM tools, and was released in 2009. Intel DCM is distinct from other DCIM products—such as those by Emerson Network Power (Trellis), Schneider Electric (StruxureWare), and CA Technologies—in that it can act as middleware between the DCIM software and the server hardware. By using Intel Node Manager technology, Intel DCM can place limits on the power usage of the servers being monitored. To take advantage of this feature in their own DCIM tools, several vendors—including Nlyte, Dell, Schneider, and Rackwise—have integrated Intel DCM into their own DCIM solutions. To simplify the evaluation of the core technology, EPRI has chosen to evaluate Node Manager using Intel's own DCM tool in this study.

A DCIM tool can monitor IT equipment across geographical boundaries as long as it is networked. A central DCIM can manage many servers across an enterprise. The policy of power limit can be applied to a group of servers that are located in a designated utility service area.

Despite its growing hardware and software compatibility, this technology has not been widely adopted. The barriers to adoption of this technology are largely the same as those common to all DCIM tools. Namely, the initial cost and complexity limit the adoption of DCIM across the data center industry. The DCM also faces the added barriers of limited hardware compatibility.

ASSESSMENT OBJECTIVES

The primary objective of this study is to provide a proof-of-concept demonstration of the DR capability of servers using power-capping technology through testing in a laboratory and the field. By demonstrating the functionality and assessing this technology's maturity, the results of this study will demonstrate the technology's potential, and provide insight into its viability as a demand response solution.

This study seeks to assess three separate functions of the technology to demonstrate its DR feasibility. First, the technology should be demonstrated to have the ability to effectively limit server power. Second, this technology should be shown to reduce the average server power use under typical workload over a useful period of time, from a few minutes to several hours. Finally, the technology should be able to initiate a power restriction in response to an external signal, showing the feasibility of utilizing this technology to automatically respond without human intervention to a DR signal from a utility or other power system operator – for example, an Independent System Operator (ISO).

To determine the technology's ability to accomplish the functions described above, the authors implemented mirrored testing in a laboratory environment, and in an operational data center at Calit2, University of California, Irvine (UCI), using five Dell servers equipped with Intel processors compatible with its Node Manager technology. Power capping is managed by Intel's DCM, a basic DCIM tool to monitor and manage power within the data center. Finally, a custom version of EISS Client software was delivered by IPKeys to provide end-node response to OpenADR signals from a utility or power system operator.

TECHNICAL APPROACH

To verify the power-capping software's performance, a proof-of-concept demonstration was conducted in both laboratory and field site settings. Laboratory testing was used to verify proper operation of the technology before field deployment in order to mitigate disruption of business-critical operations. Field demonstration allows the software's performance to be evaluated under real-world workloads, which exhibit inherent variability and unpredictability. To allow comparing laboratory results with field results, identical server hardware was acquired and installed in EPRI's IDEA Lab, located in Knoxville, TN.

For this testing, new server hardware was specified to be compatible with Intel's Node Manager technology, allowing demonstration of its power capping feature as a DR instrument. Server hardware was selected to meet the selected field site's IT needs, and near-identical hardware was installed at both the laboratory and field test sites. To better understand the basic impact of power capping on server performance, testing was begun in the laboratory, with various benchmarks run unconstrained and under power capping limits. After benchmark tests, power capping was tested using more realistic workloads to evaluate its real-world impact.

FIELD SITE SELECTION

To provide a site for power-cap technology field testing, the California Institute for Telecommunications and Information Technology (Calit2) at University of California, Irvine (UCI) was selected as a partner for this study. This partner selected a small server room in its facility on the UCI campus to serve as an appropriate IT environment for this demonstration. This server room provides support to much of the research and administrative staff of Calit2, providing file server, web hosting, active directory, and network services to the staff. In this way, the site serves as a good example of typical small-business IT needs, with some limited influence from academic and research use.

SERVER HARDWARE UNDER TEST

To meet the institute's IT needs, five new servers were selected. Table 1 shows the five server models chosen for this testing. All five servers are 12th-generation Dell PowerEdge models, supporting Intel Xeon E5-2400 series processors with power management features. The R520 models are 2U wide, meaning they occupy two standard rack unit slots. Three identical R320 models were installed, each 1U in wide. Each of these servers is equipped with a single CPU from the Intel E5-24xx series. The models were selected to supply the field site's needs in terms of application, Operating System (OS), memory, and storage, as listed in Table 1.

TABLE 1: SERVER HARDWARE UNDER TEST

Model	CPU	RAM	Drives	Application(s)
Dell PowerEdge R520	Intel Xeon E5-2430 (2.5 GHz, 6-core, 95W)	2x 8GB DIMM (1600 MT/s, low-voltage)	3x 2TB (7.2k SATA HDD), 2x 300GB (15k SAS HDD)	(Virtualized) web server, data center manager
Dell PowerEdge R520	Intel Xeon E5-2450 (2.5 GHz, 8-core, 95W)	2x 16GB DIMM (1600 MT/s, low-voltage)	2x 500GB (7.2k SATA HDD), 3x 900GB (10k SAS HDD)	(Windows) file server
Dell PowerEdge R320	Intel Xeon E5-2407 (2.2 GHz, quad-core, 80W)	2x 8GB DIMM (1600 MT/s, low-voltage)	2x 500GB (7.2k SATA HDD)	(Windows) active directory, domain name service

DATA MONITORING PLAN

To assess the behavior of servers in response to DR via power capping, the researchers installed monitoring equipment capable of capturing the data points listed below. These data points were collected with one-minute resolution with a data acquisition system for subsequent archival in EPRI's data servers for analysis.

ELECTRICAL CHARACTERISTICS

- Voltage [volts]
- Current [amps]
- Power [W]
- Power factor
- Energy consumption [kWh]
- Total Harmonic Distortion (THD) [% of amps]

THERMAL CHARACTERISTICS

- Temperature of air exhausted from servers [°F]; inlet air temperature was available from the server internal monitoring sensors

DATA MONITORING EQUIPMENT

The equipment used for metering and data collection at the field site is listed below:

- Power meter: Elkor WattsOn (revenue grade)
- Current transformers (CTs): Continental Controls Accu-CT (15 A, revenue grade)
- Sealed temperature sensors: Omega HSTH-44031

- Communications — Obvius products:
  - Data acquisition unit: AcquiSuite A8810
  - Input/output module: Flex IO remote

Figure 1 shows the basic layout of the data monitoring equipment used at the field site for this study.

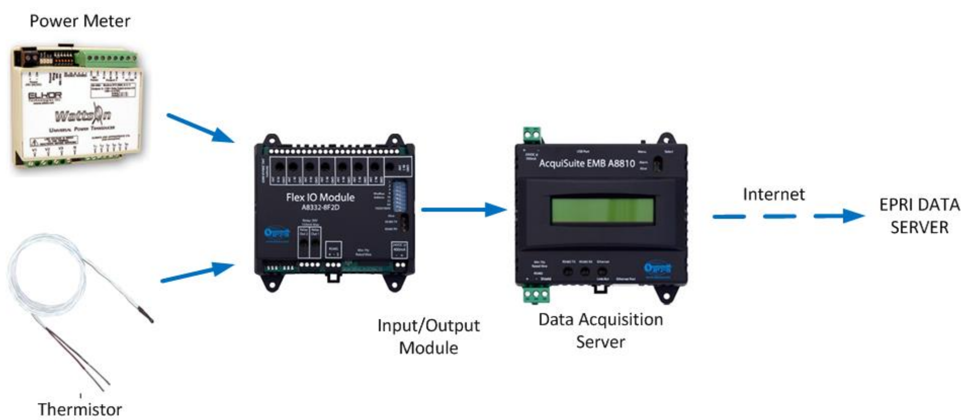

FIGURE 1: DIAGRAM OF INSTRUMENTATION CONFIGURATION

Table 2 lists the accuracy of the sensors used in this study.

TABLE 2: ACCURACY OF SENSORS USED

Instrument	Accuracy
Elkor WattsOn	<0.2% @ 25°C
Accu-CT	±0.75%
Omega HSTH-44031	±0.1°C @ 0° to 70°C

Figure 2 shows the layout of the sensors used to monitor behavior in this study.

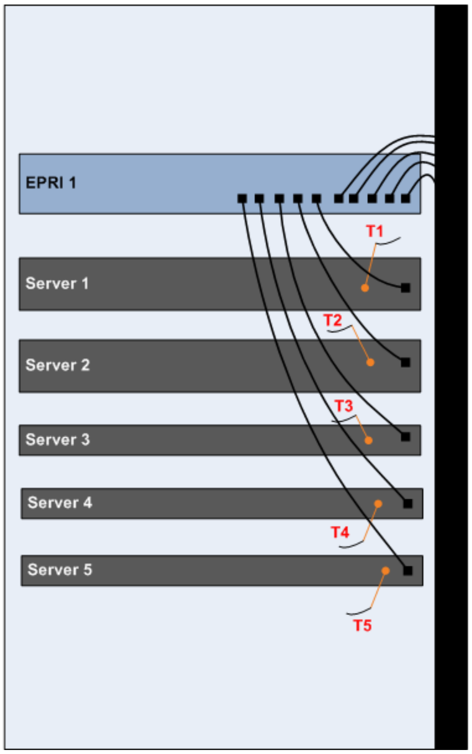

FIGURE 2: LAYOUT OF MONITORING POINTS

LABORATORY TESTING

To evaluate server power capping in a controlled environment, nearly identical specifications to those at the field site servers were acquired in the lab. The lab servers had identical CPU and Random Access Memory (RAM) specifications, and the same number and type of disc drives. (Although the number, type, and speed of disk drives were matched, the storage capacity was not duplicated, due to its limited impact on drive energy use.) Intel's DCM was installed on a separate (sixth) server that was not operated under the power cap policy. Figure 3 shows the Servers Under Test (SUT) installed in EPRI's IDEA Lab.

FIGURE 3: SERVERS UNDER TEST IN EPRI'S IDEA LAB

The Intel DCM also monitors server power and temperature by communicating with the Baseboard Management Controller (BMC) on each server. Dell's BMC features are named iDRAC.

To enable remote management on Dell servers, each was equipped with a Power Supply Unit (PSU) that supports Power Management Bus® (PMBus) communications protocol. In addition, the BMC was configured to enable Intelligent Platform Management Interface (IPMI) over a Local Area Network (LAN). The remote power management feature on Dell servers requires an upgraded enterprise iDRAC license, which added a license fee of $300 per server.

Industry standard LINPACK benchmarks were used to load the SUT to verify basic power-capping functionality. This testing utilized a version of the LINPACK benchmark that was optimized and compiled by Intel for use with its processors. Table 3 shows the specifications of the server used for testing power capping with the LINCPACK benchmark.

TABLE 3: SPECIFICATIONS OF SERVER USED FOR BENCHMARKING AND POWER CAP TESTING IN LAB

Model	Dell PowerEdge R520
Processor (single)	Intel Xeon E5-2430v2 (2.5 GHz, 6 cores, 80 W)
Memory	Two (2) 8GB (1600 MT/s, low-voltage)
Storage	Three (3) 500GB 7.2k SATA HDD Two (2) 300GB 10k SAS HDD

At its core, the benchmark tasks a computer with solving a large N by N system of linear equations, requiring a vast number of floating point addition and multiplication calculations be performed. LINPACK is expected to illustrate the level of curtailment attainable at the high end of the virtualized cloud workloads above 80% CPU utilization. As such, the benchmark is an extreme, worst-case processing requirement, and offers the most challenging workload a server would ever experience. Traditional enterprise workloads load the CPU to less than 20% utilization, and in many cases, less than 10%. It is expected that the amount of power curtailment available with LINPACK is greater than that available from more conventional workloads.

This server was used to test the basic functionality and impact of power capping and its impact with the LINPACK benchmark. Two benchmark parameters were chosen to represent a light workload (1,000 by 1,000 system of equations) and a heavy workload (30,000 by 30,000). Each benchmark was run without power restrictions as a reference, followed by a test under a minimum 'emergency level' power cap. Under this scenario, the CPU power is held to its lowest operating level.

RESULTS

The results of preliminary LINPACK testing under a maximum power cap (minimum power level) are shown in Table 4, indicating the time required to complete each benchmark, measured as FLOPS score, and average power of the server over the duration of the benchmark. These results show that power capping successfully limits the server's power use to 72 W, yet impacts LINPACK performance metrics by an order of magnitude for both workload levels. For reference, this server exhibits idle power of about 60 to 62 W.

Though the average power consumption reduced by 35% to 47% when power was capped with light and heavy loads respectively, the time to complete each of these workloads increased by a factor of 22x for light loads and 26x for heavy loads. This clearly indicates that this level of power capping has an enormous impact on the performance of these workloads. Benchmark performance is reported as computations per second, computed from the time required to complete the benchmark. The computing performance—measured in floating point operations per second (FLOPS)—was found to be cut by over 95% for both workloads under the maximum 'emergency' power cap.

TABLE 4: RESULTS OF LINPACK BENCHMARK UNDER MAXIMUM POWER CAP

		BASELINE	POWER CAP	CHANGE
Light Load	Time (s)	0.016	0.341	22x
	GFlops	42.8	2.0	
	Power (W)	110	72	35%
Heavy Load	Time (s)	157	4151	26x
	GFlops	114.5	4.3	
	Power (W)	135	72	47%

These benchmarks were repeated under several additional settings to determine the impact of varying power cap levels on the system performance. Preliminary testing of this server revealed peak power of 152 W under heavy workload, which was limited to 72 W under maximum power cap. With these characteristics in mind, the two benchmarks were repeated with decreasing power limits, beginning with a limit of 152 W and decreasing the cap by 10 W for each successive test. Figure 4 shows the results of this test, with benchmark score plotted as a function of average server power during the benchmark. Performance is seen to exhibit a non-linear dependence on power, falling off dramatically at lower power limits. Interestingly, benchmark performance was seen to exceed baseline performance for high power limits. For the heavy workload, superior performance was measured under 132 W, 142 W, and 152 W limits than without a power cap. In addition, the 152 W limit was shown to increase average power over the test period.

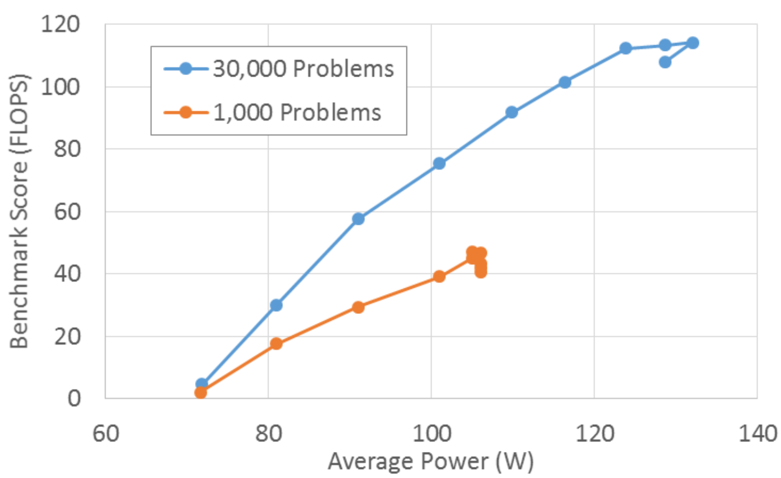

FIGURE 4: PERFORMANCE SCORE (FLOPS) OF TWO LINPACK BENCHMARKS UNDER INCREASING POWER CAP LEVELS

To examine the dependence of performance on power cap for benchmarks of different sizes, a series of six benchmarks was run under five power cap levels, and compared with unconstrained performance. The five power cap levels chosen were 72 W (minimum power), 92 W, 112 W, and 132 W. Under each of these power limits, the following benchmark sizes were run: 1000, 5000, 10000, 20000, 30000, and 35000. Figure 5 shows benchmark performance against workload, demonstrating similar performance curves under each power cap level, with the lowest limits having the greatest negative impact on performance. These results indicate that a restriction in power has the least impact on the smallest benchmark workload, but serves to “flatten” the performance of the larger workloads.

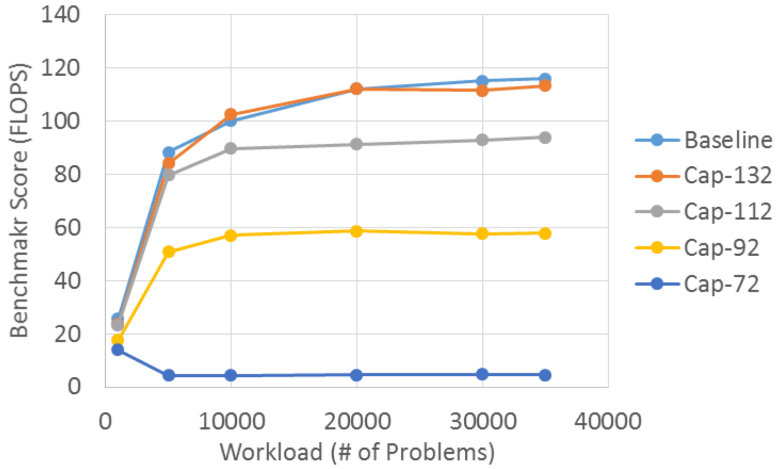

Figure 5: Performance score (FLOPS) of various LINPACK benchmarks under select power cap levels

Figure 1 shows the power level of each benchmark over the five power cap conditions, showing that each cap restrained server power below the specified limit. These results demonstrate that power capping is able to limit the average demand over the period required to complete each benchmark. Compared to the baseline power of 127 W, the power cap was shown to reduce server demand by as much as 55 W, signifying a 43% reduction. Yet given the significant impact to server performance, such a limit may only be feasible in extreme situations (e.g. when grid stability is threatened or when power is transitioning from utility to UPS). Less-stringent limits yielded lower savings: 35 W (28%) for the 92-W cap, 17 W (13%) for the 112-W cap, and virtually no savings for the 132-W cap. However, due to the high level of load the LINPACK benchmark places on the CPU, it is expected that these results would only be observed for real-world applications with very high, near-continuous CPU utilization.

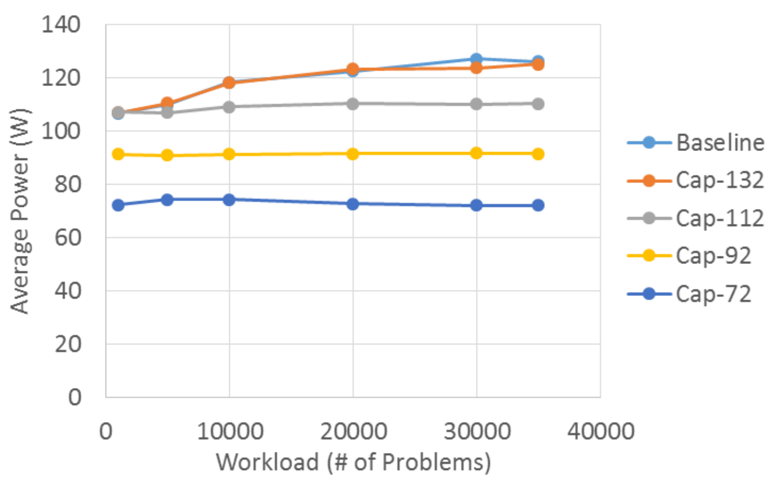

Figure 6: Server power (W) for LINPACK benchmarks under several power cap levels

Although the power cap was shown to maintain server power below specified limits, the impact on performance was found to dramatically increase the time required to complete each benchmark. Figure 7 shows the amount of time to complete each workload under the five power limits, with a logarithmic scale marking the duration of benchmark tests. These results show that the maximum power limit increased the time to complete each benchmark by an order of magnitude over the baseline and 132-W cap benchmarks.

Figure 7 shows that the effects of capping on performance are nonlinear. A small impact on performance is observed until 92 W. One explanation is that Node Manager removes cycles "off the top". A CPU running at less than 100-percent utilization has headroom left to soak up variations in demand. Below 92 W, there is no more headroom left, and we start seeing resource contention.

LINPACK presents a constant workload to the CPU. This underestimates a side effect of capping that does not show up as reduced performance. As CPU cycles are removed off the top, the CPU's ability to respond to demand upticks is impaired. As an example, assume a virtualized workload is at 80 percent utilization. If the server is capped so it is not allowed to go beyond the current power level, a 10-percent uptick in workload will put the system into resource contention, with significant performance degradation. Without the cap, utilization might go from 80 to 90 percent, with a very small performance degradation.

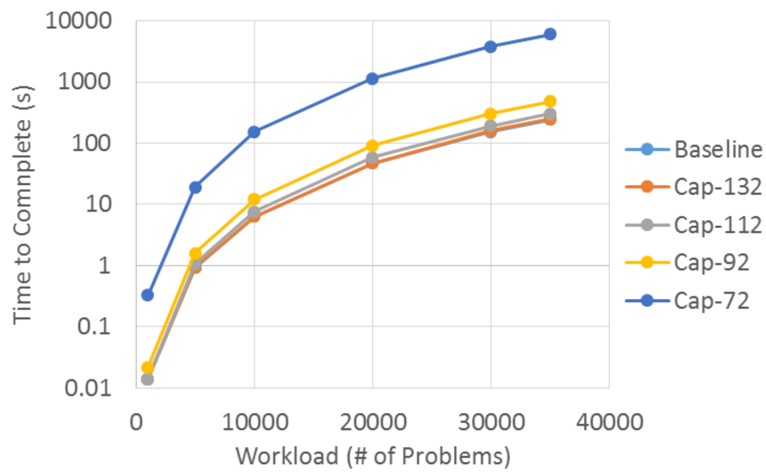

FIGURE 7: TIME TO COMPLETE LINPACK BENCHMARKS UNDER SEVERAL POWER CAP LEVELS

Figure 7 may indicate that throughput-related workloads (e.g. transactions) have increased opportunities for power reduction, while minimizing impact to the transaction throughput. The case would be especially applicable for constant workloads. Also, although the baseline performance might not change, capping will affect the system's ability to respond to a demand uptick.

In terms of shifting load over a sustained period, another metric of interest is the total energy required to perform each benchmark routine. Table 5 shows the energy used to complete each benchmark (J, equal to one watt-second), calculated as the average power over the benchmark (W) times the average period of completion (s). These results echo earlier findings that lower limits have greater impact on overall energy use, with the maximum power limit increasing execution energy by an order of magnitude, in all cases. However, it should be noted that the primary purpose of power capping is to maintain a reliable power supply when grid power supply is under stress, and not necessarily to reduce energy use. A longer execution time and more energy consumption is expected when the power supply is capped. It should be noted that the 132-W limit had minimal impact on energy use, since this is close to the chip's peak power demand.

TABLE 5: ENERGY USED TO COMPLETE LINPACK BENCHMARK (J) UNDER VARIOUS POWER CAP LEVELS

Workload	1000	5000	10,000	20,000	30,000	35,000
Baseline	1.5	103.0	751.9	5,720	19,672	30,725
Cap-132	1.5	104.1	750.7	5,795	19,710	31,341
Cap-112	1.5	113.5	836.9	6,476	20,924	33,728
Cap-92	2.0	144.3	1097	8,410	28,028	44,292
Cap-72	23.6	1435	11363	83,415	276,103	435,718

OPENADR SIGNAL COMMUNICATIONS

One of the objectives of this study was to verify the ability of a server to respond to an external DR signal. Within OpenADR, this signal is communicated between a Virtual Top Node (VTN) at a utility or third-party aggregator and a Virtual End Node (VEN) at the load. The plan of communicating the OpenADR signal, its receipt and confirmation with the VEN at site, and subsequent command by the VEN to the server load for power capping required software and communications coordination among three vendors—IPKeys, Schneider Electric, and Intel. For the laboratory tests, project partner IPKeys installed a custom version of its EISS client on an EPRI server, which also hosted the Intel DCM program to act as a VEN and confirm DR signal receipt. In turn, the EISS client commanded the DCM to cap server power. IPKeys provided access to its VTN developmental dashboard to set up tests for sending signals to the test server VEN. Figure 9 depicts the communications architecture used in this testing. The DR signal communication receipts and its DCM triggering were successfully verified in the lab.

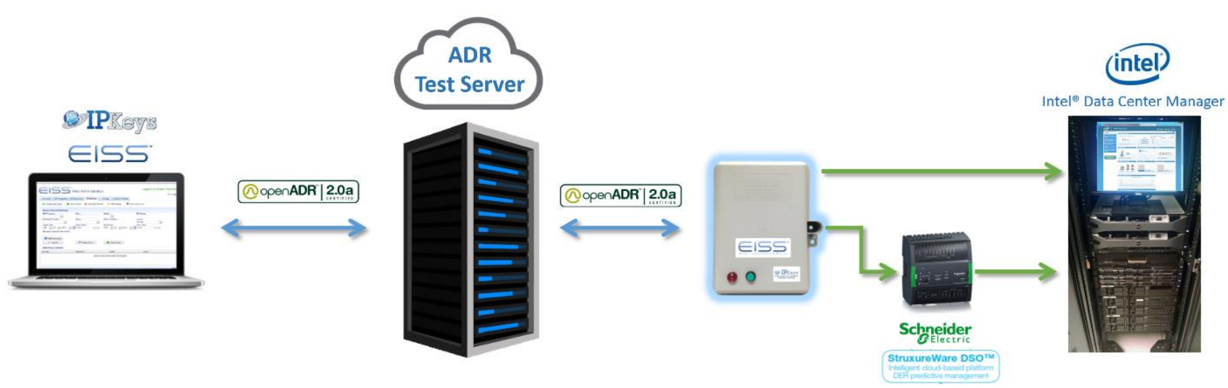

FIGURE 8: DIAGRAM OF COMMUNICATIONS PATHWAY

A series of OpenADR signals were scheduled through the dashboard, and at the defined times, a curtailment signal was sent to the server in the lab. The results of this testing demonstrate that a power cap can be set from an external control. The SUT responded very quickly (in a manner of seconds) representing the latency between the server hosting the OpenADR dashboard in Virginia and EPRI's IDEA lab in Tennessee. It is expected that real-world adoption of such a technique would also be quite fast, depending upon communication latencies between the VTN server sending the OpenADR signal (utility or third-party) and the VEN for load to be curtailed. This quick response indicates that power caps could potentially be useful for fast DR programs (for example, to provide frequency support to the grid).

The VEN sent a command to DCM for power capping after receiving the signal from the VTN. The Intel DCM was programmed to apply an "emergency power reduction" policy that limited each server to minimum power level for these tests, which was set at 70% of design power. There is potential for more sophisticated functionality to be programmed in the OpenADR signal and the DCM. For example, the OpenADR signal may request a specific power reduction (say 10%) and the DCM would, in turn, only power cap the requested amount.

For application in a large data center with multiple servers, however, IPKeys EISS Client is embedded in Schneider Electric's StructureWare product that acts as a VEN, which in turn sends commands to the Intel DCM. In such applications, StructureWare can also be used for DR from other data center equipment, such as cooling units.

FIELD TESTING

The field tests were performed at Calit2 for this study. Figure 9 shows the five servers installed in the Calit2 server rooms, along with power and thermal monitoring provided by EPRI for this study. The five servers in this study provide support to the research staff at Calit2, including a file server, web hosting, active directory, and network services. These applications represent some of the most common services provided by small server rooms in commercial buildings of all types.

FIGURE 9: SERVERS UNDER TEST, WITH METERING INSTALLED

Table 6 lists the specifications and workloads on the five servers installed at Calit2. Note that Servers 3, 4, and 5 have identical hardware specifications, and that Server 4 functions as a redundant "mirror" of Server 3. Server 1 is unique in that it is configured in a virtualized environment, with several Windows and Linux Virtual Machines (VMs) operating on it, including the research group's website (http://iot.calit2.uci.edu/) and the Intel DCM instance used in this study.

TABLE 6: SERVERS INSTALLED AT CALIT2

Server	Model	CPU TDW	RAM	Drives	Environment	Application(s)
1	R520 (2U)	E5-2430 (95W)	2x 8GB	3x SATA, 2x SAS	VMWare	Web server, data center manager, energy apps
2	R520 (2U)	E5-2450 (95W)	2x 16GB	2x SATA, 3x SAS	Windows	File server
3	R320 (1U)	E5-2407 (80W)	2x 8GB	2x SATA	Windows	Active directory, domain name service
4	R320 (1U)	E5-2407 (80W)	2x 8GB	2x SATA	Windows	Active directory, domain name service
5	R320 (1U)	E5-2407 (80W)	2x 8GB	2x SATA	Windows	Misc. applications

Figure 10 shows the power (hourly average) of the five field servers over the month of June 2016, with colors representing Server 1 (blue), Server 2 (orange), Server 3 (green), Server 4 (red), and Server 5 (purple). In general, all of these servers exhibit very little variation. In fact, except for Server 1, these servers exhibit near-constant load. With its numerous applications and external-facing website, it is expected that Server 1 would exhibit the most variation in workload. Yet the other four servers exhibit very little workload, suggesting that they experience very little workload.

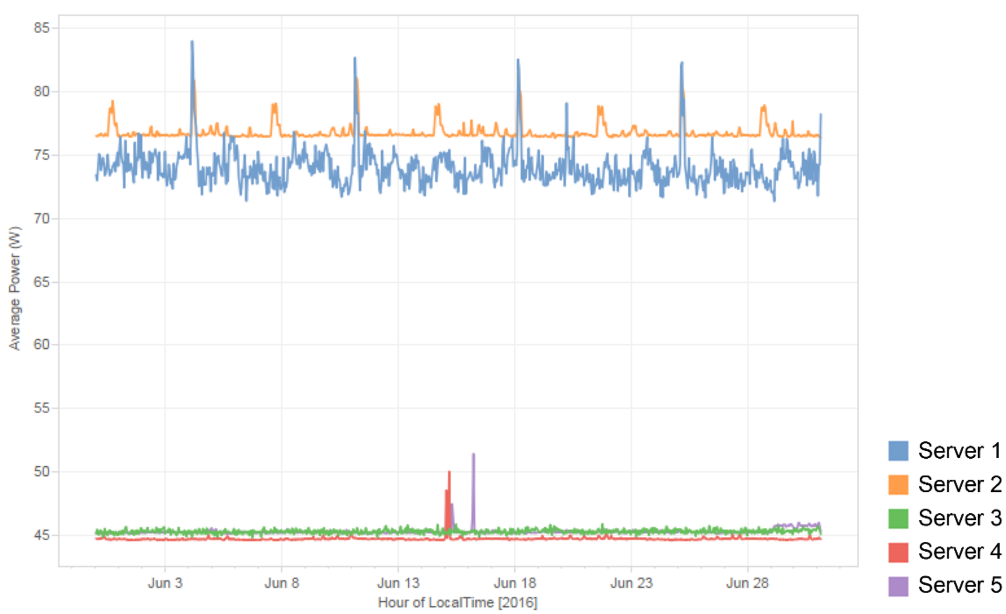

FIGURE 10: HOURLY AVERAGE POWER OF SERVERS IN FIELD

Figure 11 shows the weekly profile of average power for each server. Once again, Server 1 is the only server that exhibits regular activity. Both Server 1 and Server 2 demonstrate increased activity on Saturday morning, which is evidently when regular "batch" processes, such as software updates and data backups, are scheduled. In addition, Server 2 shows a regular increase in activity on Tuesday afternoons.

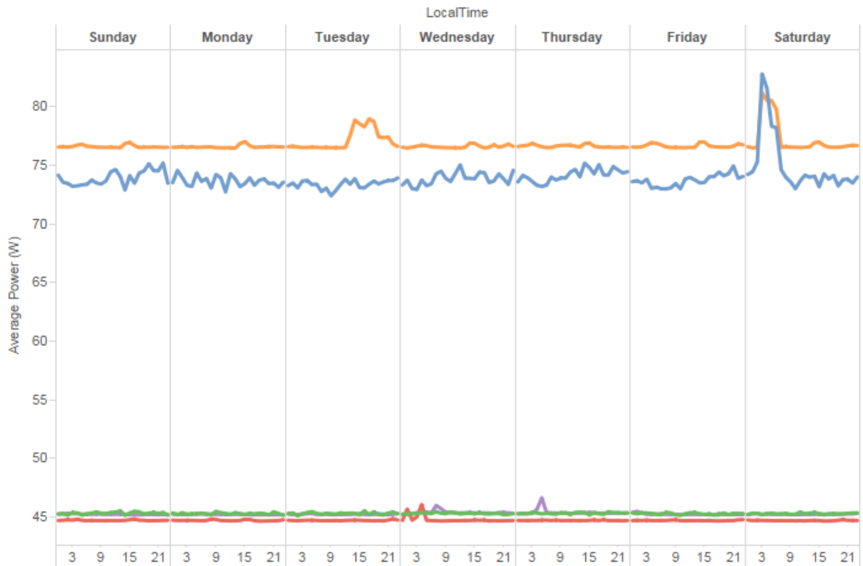

FIGURE 11: WEEKLY PROFILE OF AVERAGE SERVER POWER

Figure 12, showing the weekly profile of the average server exhaust air temperature, is included here incidentally as an observation of the cooling system operation, and not directly related to the server DR. All five of the servers demonstrate daily and day-to-day variations in exhaust temperature, with elevated temperatures measured outside regular business hours (Monday through Friday, from roughly 8:00 a.m. to 6:00 p.m.). Such behavior may indicate that the HVAC is programmed to follow a typical schedule for an occupied business space. In addition, Servers 1 and 2 exhibit a slight increase in temperature on Saturday morning, when their batch processes run.

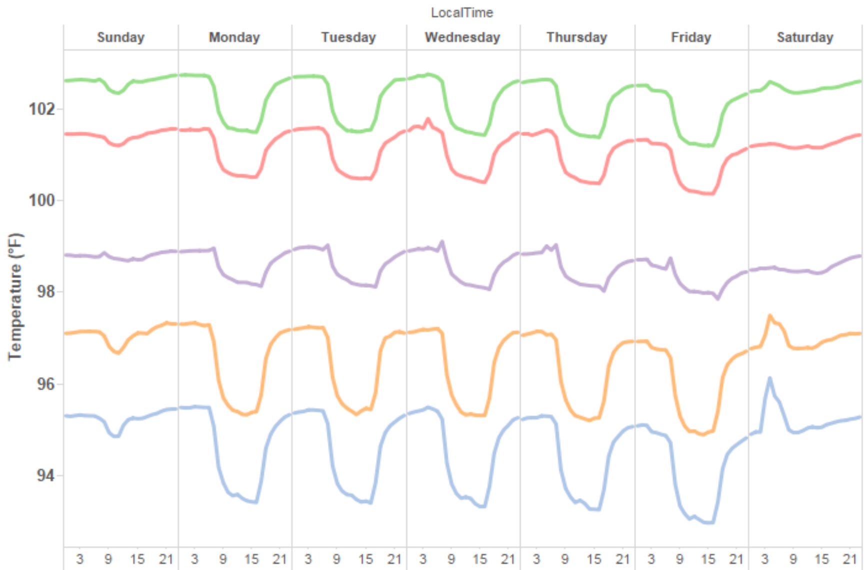

FIGURE 12: WEEKLY PROFILE OF AVERAGE SERVER EXHAUST TEMPERATURE

RESULTS

Power cap tests were applied to Server 1, which supported several real-world workloads in a virtualized environment including web servers, DCM, and an energy management application. This server exhibits a moderately dynamic load profile as shown in Figure 10, which is expected to provide more informative findings from power cap testing. First, a series of power limits (85W, 80W, 77W, and minimum power) were scheduled on the server for 30-minute periods, with 30 minutes period in between the tests to return the server to normal operation, to evaluate its ability to limit server power. These limits were selected based on an average server power use of 75W. Testing was conducted during periods of low user load—very early morning hours—so as not to impact business operations. Figure 13 shows the results of this testing. For power limits of 80W and below, power capping was found to increase average server power use by about 4W.

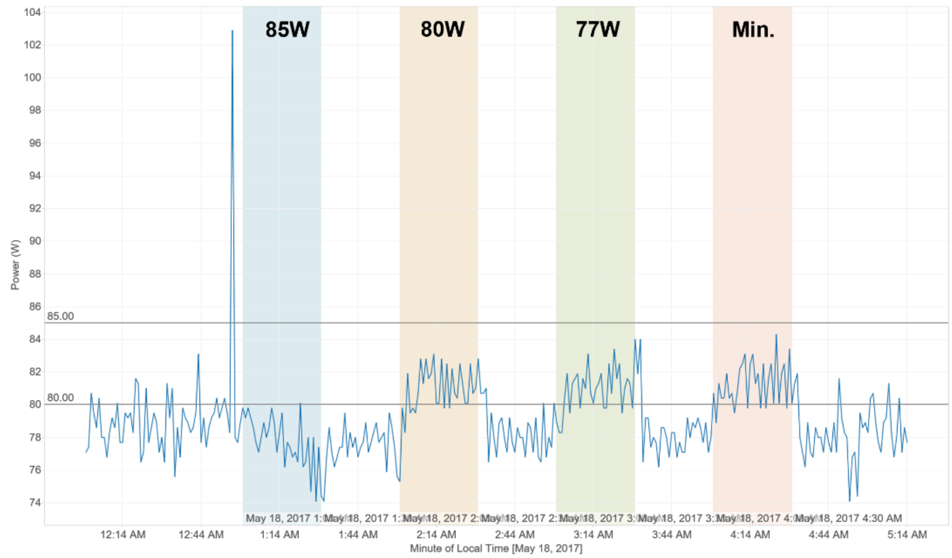

FIGURE 13: POWER CAP TESTS ON SERVER IN FIELD WITH LOW LOAD

The measured power use of this server indicates that the average load is relatively low, however. For this reason, additional load was injected on the server in short intervals to test the impact of power capping on moderately loaded servers. Researchers used Apache JMeter™—a software designed to measure the performance of servers under several different workloads—to add load and measure the performance of the web server with and without power capping policies in place. The HTTP request function of JMeter allows stress testing and performance benchmarking of a web server by making HTTP requests from a customizable group of simulated users and measuring the time for the web server to respond.

Two different web server tests were run using JMeter to simulate loads of increasing size. Test 1 made website requests from 5 simulated users, and Test 2 from 50 users. Each test was looped continuously for 1 hour during early morning hours (to minimize impact to normal business operations), meaning that the simulated users would continue making website requests for the duration of the test period. A series of power caps were applied in 5-minute intervals which were separated by 5-minutes of unconstrained operation. Table 7 lists the power caps applied to a test operated from 4:00 to 5:00 AM local time.

TABLE 7:POWER CAPS APPLIED TO WEB SERVER IN FIELD

Start	End	Cap (W)
4:05	4:10	110
4:15	4:20	100
4:25	4:30	90
4:35	4:40	80
4:45	4:50	Minimum power

Figure 14 shows server power (maximum, average, and minimum power) when the series of power limits noted in Table 7 were applied to the web server test with 5 users (test 1).

Results indicate that power capping is able to limit the power use of a web server under load.

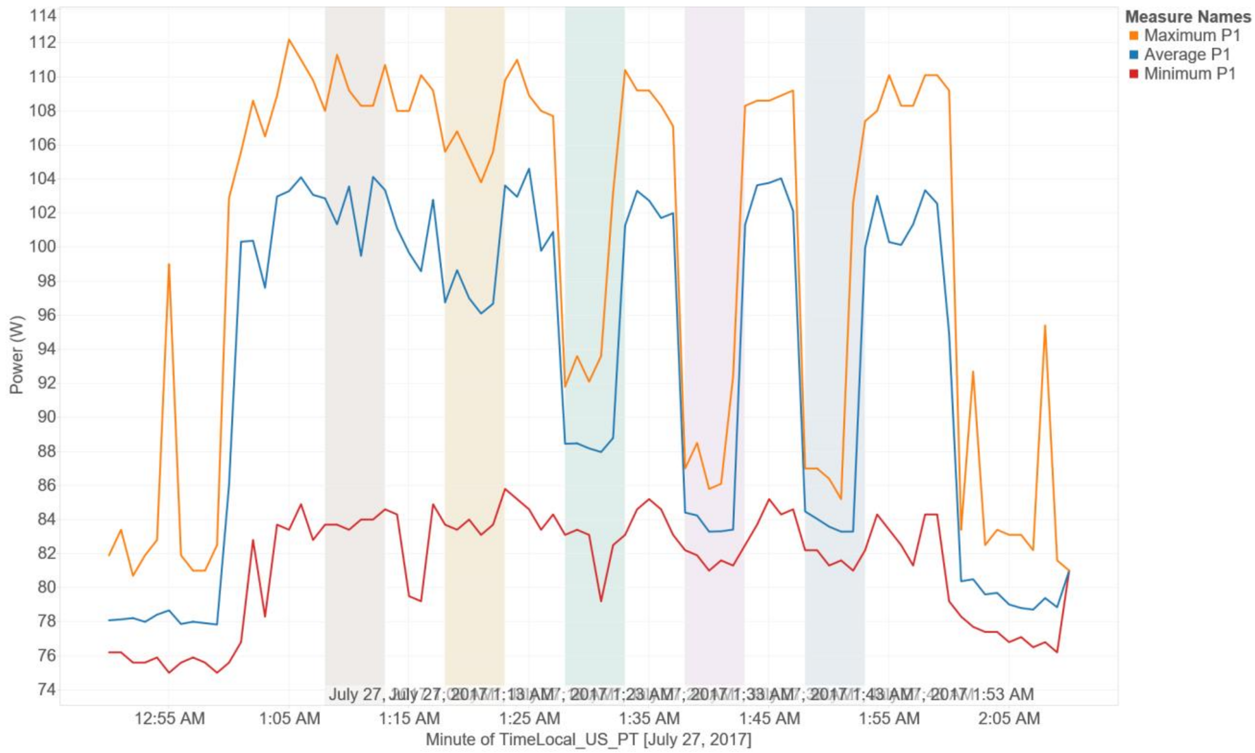

FIGURE 14: SERVER POWER WHEN LIMITS APPLIED TO WEB SERVER TEST 1 (5 USERS)

Table 8 shows the impact that power capping had on average server power for the web server testing with 5 simulated users, along with percent savings compared to baseline. These results indicate that power capping successfully reduced the average power of the web server under load by as much as 18.2W (17.9%) from baseline. Note that the highest power cap level (110W) was found to slightly increase the average power use of the web server by 0.4W (0.4%), similar to results from LINPACK testing in the lab.

TABLE 8: AVERAGE SERVER POWER UNDER POWER CAPPING FOR WEB SERVER TEST 1 (5 USERS)

Power Cap	**Baseline**	**110W**	**100W**	**90W**	**80W**	**Minimum**
Average Power	101.9 W	102.3 W	97.0 W	88.4 W	83.7 W	83.7 W
Savings	N/A	-0.4%	4.8%	13.2%	17.9%	17.9%
HTTP Request Time	960	957	1041	1459	43074	43047

Figure 15 shows average server power and the average time needed to complete HTTP requests (including latency and length of response) under each power cap. (Note that the scale for sample time is logarithmic.) These results indicate that the lowest power caps (80W and “minimum power”) increased the time to deliver HTTP requests by an order of magnitude, with only an 18% reduction in average server power. These results suggest that power capping to such significant levels may not be feasible except in extreme situations. Testing with 50 users yielded very similar results, which are included in the Appendix for reference.

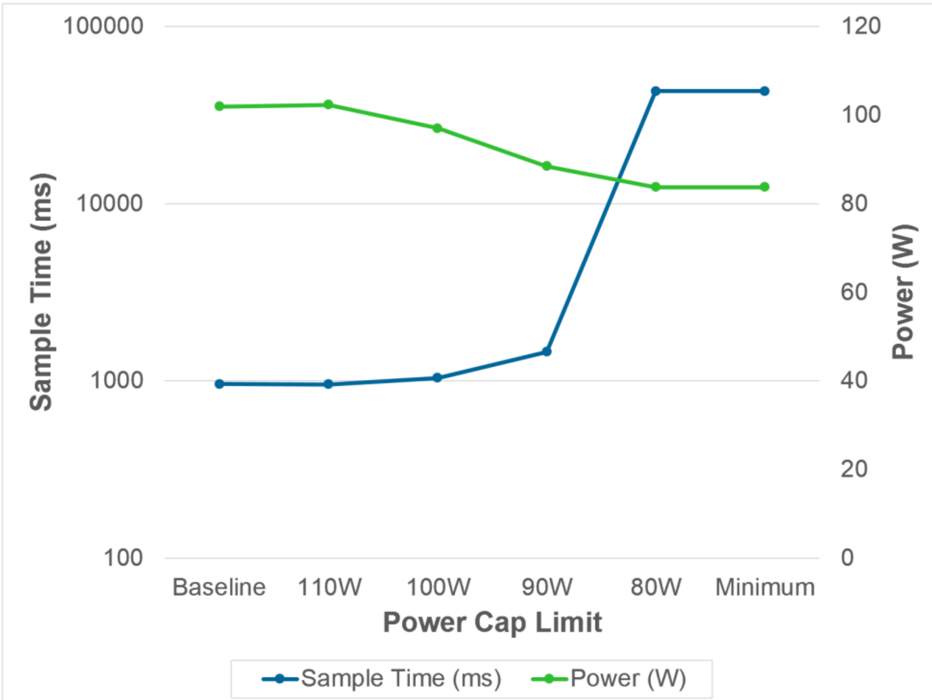

FIGURE 15: WEB SERVER PERFORMANCE UNDER POWER CAP FOR TEST 1 (5 USERS)

On the other hand, a power limit of 90W was found to reduce server power by 13% while increasing sample response time by 52%. Such a performance impact is not likely acceptable in data center applications. However, the sensitivity of performance impact from 90W to 80W indicates that this technique is very sensitive to the load on the server. Thus, it is possible that additional development can enable power capping that adjusts to a reduced power level without severely impacting performance. It is recommended that automatically adjusting algorithms be developed for power capping before this technique is evaluated for DR at a wider scale.

In the test scenario, the web request from the JMeter were continuous. In real life, the web requests are not continuous but random based on user's requests. There are times in between requests for the server to recover to unconstrained operation. Considering the test data, and random nature of web requests, smaller level of capping of 5 to 10% are likely to be practical.

CONCLUSION

This report describes the efforts the project team led by EPRI has undertaken in a project funded by SCE to demonstrate the feasibility of server power capping for electricity demand response in data center. Preliminary testing using LINPACK benchmarks demonstrates that power capping can successfully limit a server's instantaneous power. When applied to a loaded web server in the field, power capping was able to reduce average server power by up to 19% (19.6 W) under more heavily loaded testing. Yet these power reduction levels increased the time required to deliver HTTP requests by more than one order of magnitude, but increase in time was smaller with smaller level of power capping. Also, the tests were conducted with continuous web requests whereas in practical applications, the requests are random in nature. Thus, power capping may be practical for smaller level of capping, say 5 to 10%, for short periods of time.

Though power capping may provide longer-term DR in some scenarios where workloads are flexible, it is likely that it will find most practical use for shorter-term durations, even in seconds and minutes. Power capping was demonstrated to successfully limit the instantaneous power of multiple servers at a moment's notice. This capability may provide value to the grid (for example, to respond to grid emergencies or to provide ancillary services, such as frequency or voltage regulation) with the very quick response required to compensate for variable generation sources on the grid (such as wind and solar).

With increased penetration of renewables on the grid, and the very intermittent nature of power generation from such renewables, which can vary almost instantaneously (for example, when a cloud passes over a large PV array) there is a need for an electrical load that can provide fast and flexible response. Using power capping for IT equipment may provide that resource. Power consumption from silicon-based IT equipment can respond nearly as quickly as the output of PV modules can change. What's more, not all IT equipment needs to be at a single physical location, as long as they are connected via internet to the DCIM from which it can receive DR signals. With continued growth in data center power demand, and data centers becoming modern-day factories, this technology could provide a mechanism to allow for broad and coordinated control between utilities and end users.

The power capping feature of the DCM can be greatly helpful during power transition from the utility to the backup generation. Although backup generators can fire up within a few cycles of transition and can ramp up and synchronize to take the full data center electrical load within a very short time, these transition times can sometimes result in failure to transfer and loss of continuous power supply. The use of power capping at such transition times can greatly reduce the IT loads for brief periods until power transition is complete, and therefore reduce the risk of power loss during transition. The DCIM can greatly improve the data center reliability, an added benefit to the customer besides demand response.

HOW TO USE THIS INFORMATION

The information generated from the R&D activities can be used in the following ways:

- A 5-10% power demand capping is likely practical for short durations in less critical operations. The small delay in response, or increased time to execute, may not be very perceptible to end users. In a large 10 MW data center, a 250-500kW DR is possible.
- All the IT servers do not need to be in the same physical location as long as they have internet connectivity with the DCM server. A single DCM instance can manage many servers.
- Although DCIM tools are expensive, these provide many other business functions and features like asset management, etc. The power capping is a one of features of the software. Once DCIM tools are used in a data center, adding the power capping features may not be very expensive.
- The server DR has a very fast response. It can likely provide the most value in grid balancing (e.g., frequency regulation) where significant penetration of renewable energy exists. For example, it can ramp power consumption down and up as fast as power generation from a solar PV system when a patch of cloud passes by.
- The power capping feature can provide great customer benefit in improving reliability during transition of power to backup generators.

FUTURE RECOMMENDATIONS

There are a few gaps in the state-of-the-art for managing server power for DR at present. First, a power system operator cannot be expected to decide how to limit each of its servers. Such information needs to be supplied by the data center operator, including the amount of power curtailment possible, and performance indications, such as response time, duration, and pricing signals. What's more, this cannot be expected using the tools that are currently available for local server management. The technology evaluated in this study elevates these tools to commands that can be called by software, but the operational knowledge to set power limits is still absent. Certain curtailment requests may provide more incentives than others, which needs to be reflected through pricing signals.

The DR field tests were conducted with real-world applications for a file server, web hosting, active directory, and network services. It is recommended that additional testing be done to evaluate power capping with additional real-world applications, such as e-mail server, database, etc., so the impact to workloads with different needs may be evaluated (comparing processor-intensive workloads to memory or data-limited applications). Such a study could be performed in both laboratory and field conditions, so that basic functionality, response to stochastic load, and user impacts can be quantified.

Finally, it is commonly seen in hierarchical control systems that a control policy exposed at one level can become a "control knob" (i.e. a controllable parameter) for the level above. It is recommended that future efforts pursue the development of more robust DR signal communication pathways that allow an intelligent load to communicate its current state and make an informed decision about how to respond.

REFERENCES

[1] National Resources Defense Council (NRDC), "Data Center Efficiency Assessment: Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers," New York, NY, 2014.

[2] EPA, "EPA Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431.," 2007.

[3] Arman Shehabi, Sarah Smith, Dale Sartor, Richard Brown, Magnus Herrlin, "United States Date Center Energy Usage Report," Lawrence Berkeley National Laboratory, Berkeley, CA, 2016.

[4] EPRI, "Evaluation Framework for Sustainable Demand Response Implementations," EPRI, Palo Alto, CA, 2011.

[5] G. Ghatikar, M.A. Piette, S. Fujita, A. McKane, J.H. Dudley, A. Radspieler, "Demand Response and Open Automated Demand Response Opportunities for Data Centers," California Energy Commission, PIER Program and Pacific Gas and Electric Company (PG&E), 2010.

[6] Girish Ghatikar, Venkata Ganti, Nance Matson, Mary Ann Piette, "Demand Response Opportunities and Enabling Technologies for Data Centers: Findings from Field Studies," Lawrence Berkeley National Laboratory (LBNL), 2012.

[7] Ponemon Institute, "2013 Cost of Data Center Outages," Traverse City, MI, 2013.

BIBLIOGRAPHY

Data Center Demand Response: Avoiding the Coincident Peak via Workload Shifting and Local Generation, Zhenhua Liua, Adam Wiermana, Yuan Chenb, Benjamin Razona, Niangjun Chena; Performance Evaluation 70 (2013) 770–791, 2013, Elsevier B.V.

Opportunities and Challenges for Data Center Demand Response, Green Computing Conference (IGCC), 2014 International, 10.1109/IGCC.2014.7039172, Adam Wierman
Computing and Mathematical Sciences, California Institute of Technology
Zhenhua Liu, Iris Liu, Hamed Mohsenian-Rad

2011 Smart Energy Demand Coalition (SEDC)

Grid Integration for Data Centers, Girish Ghatikar, Vish Ganti, Nance Matson, and Mary Ann Piette; SVLG Data Center Summit October 24, 2012; http://svlg.org/wp-content/uploads/2012/10/2.15-Girish-Ghatikar-Opening-Slides-DR-WEBSITE.pdf

A Data Center Perspective on Demand Response, Feb 2013, http://www.datacenterdynamics.com/content-tracks/power-cooling/a-data-center-perspective-on-demand-response/73892.fullarticle

A Hierarchical Demand Response Framework for Data Center Power Cost Optimization Under Real-World Electricity Pricing; Cheng Wang, Bhuvan Urgaonkar, Qian Wang†, George Kesidis; http://www.cse.psu.edu/~buu1/papers/ps/hieropt.pdf

Microsoft Says 'Computational Demand Response' Could Lower Data Center Emissions 99%; http://www.greentechmedia.com/articles/read/Microsoft-Says-Computational-Demand-Response-Could-Lower-Data-Center-Emis

Demand Response in Data Centers; https://sites.google.com/a/lbl.gov/try-this-one/

Demand Response and Open Automated Demand Response Opportunities for Data Centers; G. Ghatikar, M.A. Piette, S. Fujita, A. McKane, J.H., Dudley, A. Radspieler; Lawrence Berkeley National Laboratory; K.C. Mares, Megawatt Consulting; D. Shroyer, SCG; January 2010. https://sites.google.com/a/lbl.gov/try-this-one/documents

Data Center Demand Response-Coordinating IT and the Smart Grid, Zhenhua Liu, December 18, 2013. Data Center Demand Response - dimacs

Pricing Data Center Demand Response, Zhenhua Liu, Iris Liu, Steven Low, Adam Wierman; http://smart.caltech.edu/papers/DCDRpricing.pdf

Towards Continuous Policy-driven Demand Response in Data Centers; David Irwin, Navin Sharma, and Prashant Shenoy; http://www.ecs.umass.edu/~irwin/greennet.pdf

Data Center Energy Systems: Current Technology and Future Direction Santosh Chalise, Amir Golshani, Shekhar Raj Awasthi, Shanshan Ma, Bijen Raj Shrestha, Labi Bajracharya, Wei Sun and Reinaldo Tonkoski Department of Electrical Engineering and Computer Science, South Dakota State University, Brooking SD santosh.chalise@sdstate.edu; http://www.eecs.ucf.edu/~weisun/papers/conferences/date-center-review.pdf

Reducing Electricity Demand Charge for Data Centers with Partial Execution; Hong Xu Department of Computer Science City University of Hong Kong Kowloon, Hong Kong, henry.xu@cityu.edu.hk

Baochun Li, Department of Electrical and Computer Engineering University of Toronto Toronto, ON, Canada, bli@eecg.toronto.edu; http://iqua.ece.toronto.edu/papers/hxu-eenergy14.pdf

Greening Multi-Tenant Data Center Demand Response; Niangjun Chena, Xiaoqi Rena, Shaolei Renb, Adam Wiermana, http://www.ece.ucr.edu/~sren/doc/paper/performance_2015.pdf

National Resources Defense Council (NRDC), "Data Center Efficiency Assessment: Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers," New York, NY, 2014.

APPENDIX

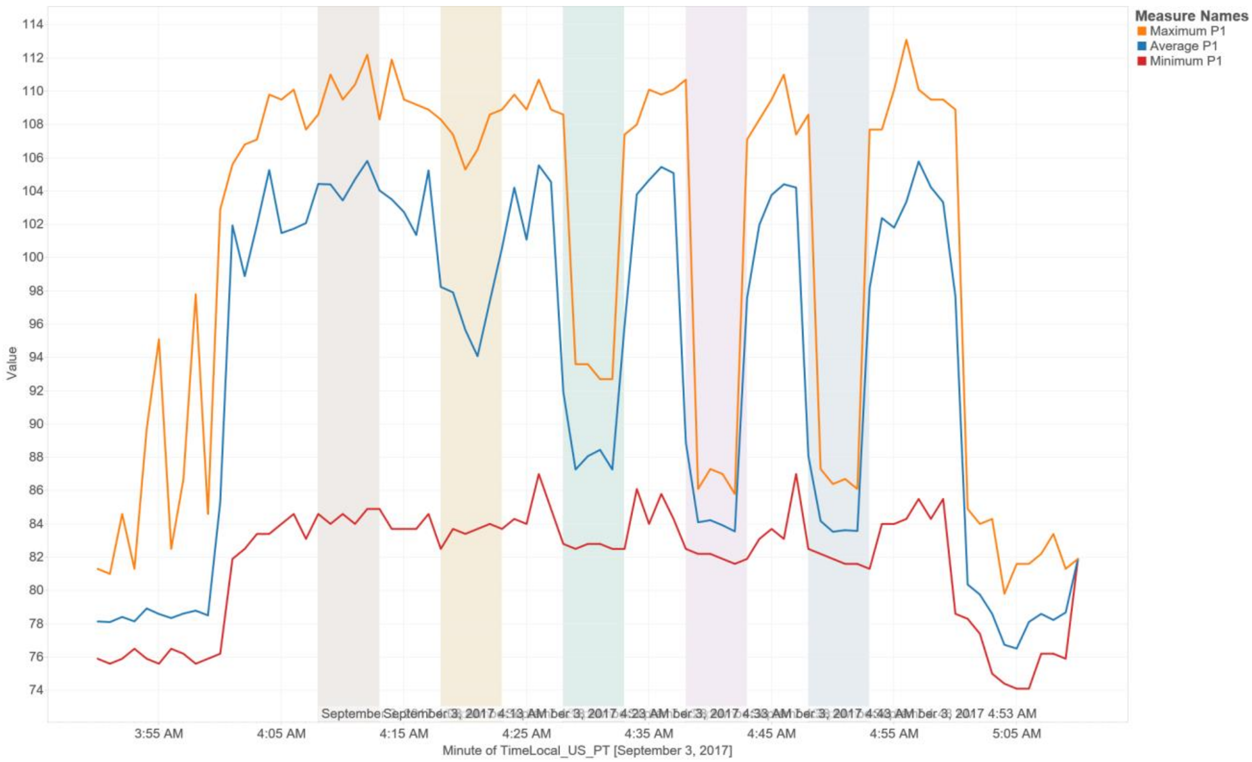

FIGURE 16: SERVER POWER WHEN LIMITS APPLIED TO WEB SERVER TEST 2 (50 USERS)

TABLE 9: AVERAGE SERVER POWER UNDER POWER CAPPING FOR WEB SERVER TEST 2 (50 USERS)

Power Cap	**Baseline**	**110W**	**100W**	**90W**	**80W**	**Minimum**
Average Power	103.3	104.6	96.3	87.8	84.0	83.7
Savings	N/A	-1.2%	6.8%	15.1%	18.7%	19.0%
HTTP Request Time	10598	9240	10924	15589	176631	170131

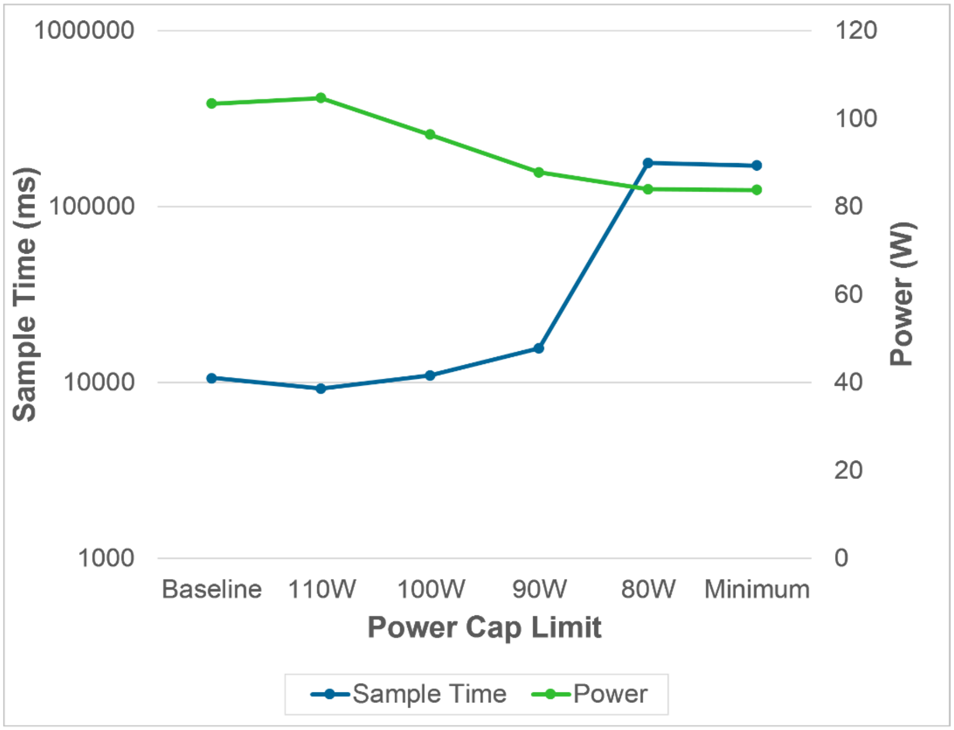

FIGURE 17: WEB SERVER PERFORMANCE UNDER POWER CAPPING FOR TEST 2 (50 USERS)