

# Data analysis, Visualization and Knowledge Discovery in Sustainable Data Centers

Manish Marwah, Ratnesh Sharma, Rocky Shih, Chandrakant Patel  
Hewlett-Packard Laboratories  
Palo Alto, CA, USA

Vaibhav Bhatia, Mohandas Mekanapurath, Rajkumar Velumani, Sankaragopal Velayudhan  
Hewlett-Packard India Software Operations  
Bangalore, India  
{firstname.lastname}@hp.com

## ABSTRACT

A significant amount of energy consumption is now attributed to data centers due to their ever increasing numbers, size and power densities. Thus, there are efforts focused at making a data center more sustainable by reducing its energy consumption and carbon footprint. This requires an end-to-end management approach with requirements derived from service level agreements (SLAs) and a flexible infrastructure that can be closely monitored and finely controlled. The infrastructure can then be manipulated to satisfy the requirements while optimizing for sustainability metrics and total cost of operations. In this paper, we explore the role of data analysis, visualization and knowledge discovery techniques in improving the sustainability of a data center. We present use cases from a large, sensor-rich, state-of-the-art data center on the application of these techniques to the three main sub-systems of a data center, namely, power, cooling and compute. Furthermore, we provide recommendations for where these techniques can be used within these sub-systems for improving sustainability metrics of a data center.

## Categories and Subject Descriptors

H.5.m [Information Systems] Information Interfaces and Presentation-Miscellaneous; C.4 [Computer Systems Organization] Performance of Systems—reliability, performance; B.8 [Hardware] Performance and Reliability-Reliability, Testing, and Fault-Tolerance, Performance Analysis and Design Aids.

## General Terms

Management, Measurement, Design, Economics, Reliability, Verification.

## Keywords

**Data centers, sustainability, data analysis, visualization, knowledge discovery, power, cooling, compute**

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

*Compute 2009*, January 9–10, 2009, Bangalore, Karnataka, INDIA.  
Copyright 2009 ACM 1-58113-000-0/00/0004...\$5.00.


## 1. INTRODUCTION

In recent years, the computing paradigm has been undergoing a change. Software applications and services – such as email, word-processing and business applications – are migrating from local desktop machines to remote data centers. While the software and data reside at a data center, the application is accessible as a service via a web browser. This computing paradigm, referred to as software as a service (SaaS) [2] and hosted on cloud computing infrastructure [1], is attractive for a number of reasons: (1) it frees users from the issues and costs related to installing, maintaining and upgrading local software applications; (2) it allows easy access to applications and data from any location (using Internet connectivity); (3) it facilitates sharing and collaboration among multiple users who are geographically separated; and (4) it simplifies sending critical client software updates such as bug and security fixes to the clients (client scripts can be downloaded by a browser as needed).

This trend coupled with emerging web-based business, social networking and media applications and services have led to a tremendous growth in the number, size and power densities of data centers. Furthermore, with rising energy costs and global attention focused on carbon footprints, energy management and sustainability of a data center have become key issues. Data centers in the U.S. consumed 61 billion kilowatt-hour (kWh) of electricity in 2006 at the cost of \$4.5 billion [3]. This constitutes 1.5% of the total U.S. energy consumption for that year. Roughly half of this energy was consumed by the power delivery and cooling systems that provide the facility infrastructure for data centers, with the other half being consumed by servers, storage, and networking hardware. It is estimated that data center power consumption will increase 4% to 8% annually and is expected to reach 100 billion kWh by 2011. Furthermore, information technology (IT) as a whole is responsible for 2% of the global carbon emissions. Thus, even a small percentage savings in the energy consumption of data centers will have a huge economic and environmental impact.

Early recognition of energy management and sustainability as key issues in data centers led to research on dynamic smart cooling (DSC) [4]. In addition to static optimization of cooling resources through computational fluid dynamics (CFD) modeling, DSC provides additional energy savings by manipulation of flexible cooling infrastructure to dynamically provision cooling resources based on demand. Real-time feedback from a rack-level temperature sensor network distributed throughout a data center is

used to infer the cooling demand in different regions of the data center.



**Figure 1 The three sub-systems of a data center and the architectural policies and techniques used for its sustainable management**

A sustainable data center solution extends this research to focus on “end-to-end” holistic management of compute, power and cooling resources based on the workload demand and service level agreements (SLAs) with the goal of optimizing sustainability criteria such as carbon emissions, energy consumed, exergy loss [5]. This not only entails workload management (demand side), but also cross-layer supply side management of power, cooling and compute resources. In order to make efficient decisions, trade-offs need to be quantifiably evaluated in relation to the optimizing criteria. For example, what is the cost of temporarily allocating additional resources to a particular workload in order to meet its SLA goals compared with the penalty of letting the SLA be violated?

The three main sub-systems of a data center, shown in Figure 1, are compute, power and cooling [10]. The management architectural framework consists of principles and techniques that span the three sub-systems. These principles and techniques include: flexible and configurable building blocks, which allow data center infrastructure to be “sized” based on the demand; sensing infrastructure, which collects important information related to a data center’s health and operational state; and, policy-based integrated management that optimizes for sustainable operation of a data center.

Data center facilities and systems produce huge amounts of data related to their physical and operational state for the purposes of monitoring and troubleshooting. This includes environmental sensor data (e.g. temperature), operational state of systems and devices, workload information (e.g. user requests), etc. Since the sheer volume of such data precludes manual inspection, automated data mining and knowledge discovery techniques are used to glean vital information. The knowledge thus gathered in

the form of models, trends and patterns can then be exploited for furthering the goals of sustainable operation in a number of ways:

- Anomalous behavior diagnosis techniques can be used to detect, localize, and perform root cause analysis of an anomaly. Timely and specific detection of such anomalies may allow corrective actions that save energy. For example, early detection of a computer room air-conditioning (CRAC) unit failure in a data center allows redirection of optimal (minimum power consumption) alternate cooling resources to the affected racks. Similarly, prompt detection of a fan unit failure in a blade server enclosure can allow the enclosure management system to compensate with the most energy-efficient combination of functioning fans units.
- Summarization and visualization of sustainability metrics and raw data in a data center to obtain a high-level view and gain insights into its operation.
- Optimization of user-defined criteria, such as, sustainability metrics like CO<sub>2</sub> emissions or exergy loss.
- Building models for efficient control of devices and processes.
- Prediction of anomalous behavior or significant events facilitating preemptive resource reallocation to optimize user-selected sustainability metrics.

We envision that data analysis, visualization and knowledge discovery will play a key role in the efficient management of a sustainable data center. In Figure 1, these techniques are shown as a layer between the sensing infrastructure layer from where it receives data, and, management layer to which it provides knowledge. Specifically, this paper makes the following contributions:

- Presents use cases of application of data analysis, visualization and knowledge discovery to a large production data center for its sustainable management.
- Describes lessons learnt from real-life use cases and provides recommendations for use of the three techniques.
- Discusses how greater insight and efficiency is possible by considering all the data center sub-systems – cooling, power and compute – together.

## 2. OUR TEST BED – A SUSTAINABLE DATA CENTER


Figure 2 displays a typical state-of-the-art data center infrastructure. Our test bed consists of 70,000 sq. ft. of data center space as well as office space for around 3000 employees. The data center houses nearly 2000 racks of IT equipment with an average power consumption of 3-5kW per rack. Typically the racks are laid out in rows separated with hot and cold aisles. The cold aisles supply cold air to the systems and the hot aisles remove hot air from the systems. Computer room air conditioning (CRAC) units cool the exhaust hot air from the computer racks. Energy consumption in data center cooling comprises work done to distribute the cool air and to extract heat from the hot exhaust air. A refrigerated or chilled water cooling coil in the CRAC unit

extracts the heat from the air and cools it within a range of 10°C to 18°C. The site in question is powered by utility and onsite diesel generators. The cooling and power sub-systems of the data center are shown in Figure 3 and Figure 4, respectively.



**Figure 2** A data center showing racks of servers arranged in rows, and, some attached office space

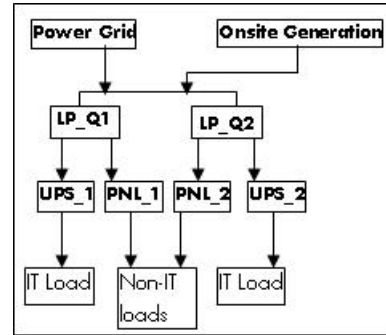
Figure 3 shows the data center cooling infrastructure. Key elements of this infrastructure include cooling tower, chillers and CRAC units and chilled water distribution. Heat dissipated from IT systems equipment is extracted by CRAC units and transferred to the chilled water distribution systems. Chillers extract heat from the chilled water system and reject it to the environment through cooling towers. Apart from IT equipment, the data center cooling infrastructure accounts for 50% of the total power demand.



**Figure 3** Data center cooling infrastructure

The CRAC units provide two actuators that can be controlled. The variable frequency drive (VFD) controls the blower speed and can be varied between 60% and 100%. The chill water regulates the amount of chilled water flowing into the unit (between 0% and 100%). These built-in flexibilities allow the units to be adjusted according to the workload demand in the data center. The demand is detected via temperature sensors installed on the racks across the data center.

Figure 4 shows a logical representation of power distribution in a typical data center. Power from utility and on-site generators are used to power the IT and non-IT infrastructure. Demand centers in the facility can be broadly sub-divided into IT and non-IT load. IT load includes servers, storage and network equipment. Non-IT load includes chillers, air handling units, pumps and lighting. All IT load is powered through uninterruptible power supplies (UPS) that not only provide emergency power but ensure power quality at the equipment. Load centers marked PNL\_\* represent non-IT loads including pumps, cooling tower, chillers and air handling units. Panels marked LP\_Q\* contain transfer switches that allow the power source to be switched between utility and onsite generation.



**Figure 4** Data center power infrastructure

An intensive sensor network monitors the status of the data center infrastructure. Alerting mechanisms detect and report the following failure conditions: (1) faulty or disconnected sensor, (2) CRAC Failure, (3) hotspot alert (when a sensor crosses its maximum allowed threshold), (4) chiller failure, (5) diesel generator tripping, and, (6) chilled water flow disruption.

Various data sources, consisting primarily of the rack sensor network and equipment belonging to data center subsystems, generate data at different rates. The sensor network, with thousands of sensors, produces temperature data every 10 seconds. Similarly, temperature, utilization and other operating parameters are generated by CRACs, chiller units, diesel generators, blade servers, and other data center components.

### 3. POWER

Sub-systems within a data center have diverse power requirements. The compute components (servers, storage and network equipment) need high-grade, uninterrupted power, while the quality of power supplied to lighting and other non-essential entities in a data center is usually not a concern. Similarly, on the supply side, there are diverse sources. Power can be drawn from the utility grid or generated onsite with technologies such as diesel powered generation, fuel cells, and photo-voltaic solar panels. Each of these sources has distinct characteristics pertaining to reliability, cost and impact on sustainability as shown in Table 1.

The challenge is to formulate power allocation and usage policies based on optimizations that match the power demands of a data center with the available power supply such that: (1) power requirements are quantitatively and qualitatively met, and (2) sustainability metrics and/or total cost of operation are optimized depending on goals set by the administrator of the data center. For

example, the policies can exploit demand (usage) patterns based on time of day, season, or location.

A rich sensing infrastructure at the power generation (for onsite generation), delivery and distribution sites is essential for measuring power consumption. The data thus collected can then be used for analysis.

Power Source	Cost		Sustainability Impact	Reliability
	Initial	Operational		
Wind	High	Low	Low	Low
Solar	High	Low	Low	Low
Fuel Cell	High	Low	Low	High
DG	Medium	High	High	High
Utility	Low	Medium	Medium	Low/High

**Table 1 Comparison between different power sources**

### 3.1 Reliable Operation


Since power is a critical resource, it is essential that component failures in the power infrastructure be efficiently and seamlessly handled. Furthermore, predictions of failures will allow preemptive actions to be taken. To this end, data analysis techniques can be applied to model behavior and further sustainability goals.

#### 3.1.1 Utility Power

Similar power sources exhibit vastly different reliability behavior in different regions. While utility power is considered reliable in the U.S. and other western countries, it is considered unreliable in developing nations such as India, where rolling blackouts (also referred to as load shedding) are commonplace. In the U.S., grid power has an availability of about 99.95% [6] (around 5 minutes of outage per week), whereas in India it is estimated to be about 88% (around 21 hours of outage per week). Furthermore, the quality of power available from the grid also varies. Thus, the strategy to handle utility power outages depends on the geographical region. While a backup solution comprising only UPS units may be sufficient for regions with highly available power (short-lived, rare outages), onsite generation, in addition, is necessary for areas where power availability is low (frequent, long lasting outages).

**Efficient failover.** In our test-bed data center in Bangalore, diesel generators (DGs) augment the grid supply as well as provide power backup on grid power failures. The failover time can be high since starting up a DG takes a few minutes. Although the IT load is backed up by UPSs, prompt take over by DGs will extend the life of UPSs. While keeping the backup generators running solves this problem, it is at the expense of using more energy. Figure 5 shows the power demand and onsite generation profiles of a data center over a week. Observe the surges in onsite generation caused by the interruptions in utility power supply. Collecting data and constructing models that predict utility power outages can allow DGs to be turned on only when their need is anticipated. To construct such a model, the important parameters are utility power voltage, utility power frequency and the current power load. Usually voltage or frequency droop indicate an insufficient supply. A simple failure detection model could monitor these parameters and signal failure based on threshold

values. To estimate these values historical logs of these parameters together with information on outages could be used. More complex modeling techniques include regression and machine learning techniques.



**Figure 5 Data center power demand profile and the contribution of onsite generation. The peaks in onsite generation power indicate utility outages.**

#### 3.1.2 UPS

As described earlier, UPSs condition power supplied to the computing equipment (servers, network switches, storage arrays, etc.) as well as provide seamless battery backup in the event of a power outage. Typically, they can provide power for up to 20 minutes. Data analysis and knowledge discovery techniques can be used to construct models for (1) estimating the remaining lifespan of a UPS, (2) failure prediction. The input parameters to these models are properties that have a significant impact on UPS operation. For example, the following parameters adversely affect the battery life of a UPS.

- Number of times it is charged and discharged.
- High ambient temperatures. In fact, every 10° C increase in temperature reduces battery life by half.

The models can be used for maximizing battery life. Furthermore, models predicting the health of a UPS can be used to preemptively replace a unit before it fails. In our test-bed data center, two UPS units supply power to each floor, which ensures that if one fails, there is still some power available on that floor.

#### 3.1.3 Onsite power generation

There are several reasons for deploying onsite power generation in a data center.

- Utility power may not have enough capacity
- Utility power may be too erratic or ‘dirty’
- Onsite generation may have better TCO, sustainability metrics
- Onsite generation may have higher reliability
- To shave peak load (see Section 3.2).

**Diesel powered generation.** Using diesel generators is a popular onsite power generation mechanism. In our test-bed data center, utility power is insufficient and, thus, DGs are used on a regular basis to provide power during normal operation. In fact, they provide two-thirds of the total power. Furthermore, to provide business continuity during utility failures, as described earlier, DGs also provide reliable backup power. As discussed in Section 3.1.1, being able to determine when a failure occurs allows efficient management of backup DGs. In this case, a failure prediction model for DGs is needed. Data collected to build such a model should include a DG's efficiency, which, if too low, can indicate a problem with the DG.

Furthermore, sensors within a DG, in components, such as, a fuel pump or turbo charger, can provide important information about the nature of a failure.

### 3.2 Peak Power Shaving and Optimal Usage

Typically a differential pricing structure is used for utility power, making power costs significantly higher if the power consumption exceeds a base load. This provides incentive for using onsite generation technologies to *shave peak power demand* and improve TCO, and, in many cases, sustainability metrics. Onsite power generation using solar panels is attractive as it can provide power during the day when power consumption is likely to peak. For example, use of solar power can shave peak power usage in an office building, where typically the power consumption peaks during the day when user area cooling, lighting and other amenities are used.

For a data center, workload demand models can be used for predicting peak demand levels and their time of occurrence. This information can then be used for planning and sizing of onsite power generation technologies.

### 3.3 Equipment Usage Policies

Different equipment in a data center works efficiently under different operating policies and conditions. For example, while the lifespan of a UPS depends more on the number of times it is power cycled, that of a DG depends on the cumulative number of hours of operation (the number of times it is started/stopped is not as significant). Knowledge of such behavior can be used to increase equipment lifespan. DGs can be rotated so that they clock the same number of hours and wear evenly.

In our test-bed data center, the total hours of operation of each DG is logged. The generators are rotated every 20 hours of continuous operation to ensure equal wear. Furthermore, such data can also be used to estimate the lifespan of a DG, allowing a facilities administrator take preemptive actions such as pre-ordering a new unit to substitute a DG close to its end of life.

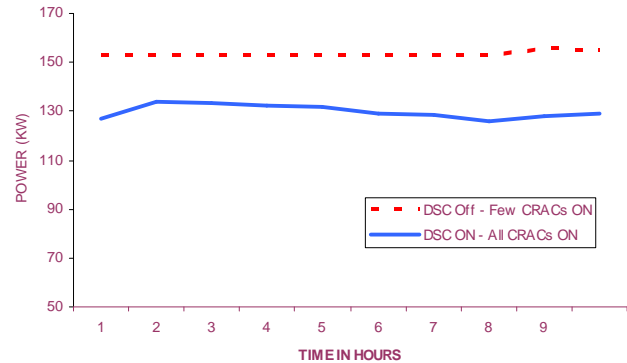
### 3.4 Power Budgeting

If power needs to be capped or budgeted, what is the most effective way to distribute the limited power? Using ad-hoc mechanisms or intuition to decide power allocation to the data center entities may not be most efficient. For example, during a period of power shortage at the test-bed data center, the facility administrators decided to turn servers on only in one part of the data center. Based on an ad-hoc policy, 18 (out of a total of 55) CRAC units that were closest to the operational servers were powered on. However, this arrangement turned out to be non-optimal in the total CRAC power consumption. Quantifying the

cooling requirement and knowledge of the influence of a CRAC unit at a rack together with its operational efficiency curves will allow optimization of the total CRAC cooling power.

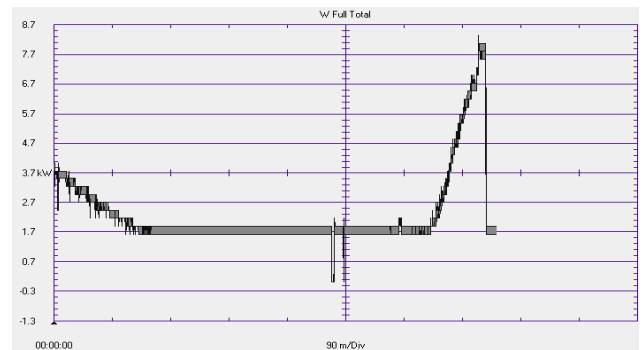
Using dynamic smart cooling (DSC) [4], which deploys a sensor network to obtain real-time rack inlet temperatures, resulted in more efficient operation. Although DSC used all 55 CRAC units compared to 18 used earlier, less power was consumed.

Figure 6 shows that on an average 25 KW of power were saved through use of DSC.



**Figure 6: Comparison of CRAC power consumption using ad-hoc policy and DSC.**

The reason for power savings despite using more CRAC units lies in the operational power curve of a CRAC unit which shows a cubic behavior with respect to fan speed. As shown in Figure 7, as the fan speed increases by a factor of 1.7 (from 60% to 100%), the power consumption increases by a factor of about 5 (from 1.7 KW to 8.4 KW). Thus, in this case, running more units at lower speed is more power efficient than running a smaller number of units at high fan speed.



**Figure 7: This graph shows the power consumption of a CRAC unit with time as its fan speed is varied from 60% to 100%.**

### 3.5 Equipment Placement

Before new equipment can be installed at a particular location inside a data center, it must be ascertained that adequate power is available at that location. Although an estimate based on power rating of equipment can be made, it may not be accurate and likely to be an overestimate. Power measurements at power distribution units (PDUs) can be collected and analyzed to determine the actual usage in a particular segment of the data center. In case



there are multiple potential placement locations, these usage statistics can be used to compare and pick the location with maximum spare capacity.

#### 4. COOLING


Power consumption of cooling equipment can be a substantial part of a data center's total power usage. In fact, prior studies [4] [8] show that it can be up to 50% of the total power consumption. Cooling resources consist of those internal to a data center, such as, CRAC units and enclosure fans, and external, such as, chillers that typically supply chilled water to the CRAC units. A sensing infrastructure monitors the cooling devices, for example, CRAC supply and return temperatures, and, blower speed are measured and logged. External to the data center, chiller supply and return temperatures, and, utilization values are also monitored and logged.

##### 4.1 Optimal Usage

Data center cooling infrastructure comprises of various devices like chillers, CRAC units, air handlers, vent tiles. Each of these devices has a different region of influence, operating characteristics and useful life. Downtime and cost of operation is a function of these operational characteristics. Sustainable datacenters also need to maintain certain sustainability targets (energy efficiency, resource efficiency) during the lifetime. In this scenario, offline and online data analytics are crucial. While online analytics provides a mechanism to optimize runtime operation, offline analytics ensures that the dynamic operational model is current and suggests redesign, if needed.

##### 4.2 Equipment Placement

Computational fluid dynamic models enable proper placement of racks and equipment at design phase. Figure 8 shows the output of one such model indicating the temperature distribution in a datacenter. Observe the variations in temperature across the racks in the datacenter denoted by varying colors. Purple indicates low temperature regions while shades of orange/red indicate hot spots. Advanced techniques like thermal zone mapping can identify regions of influence of CRAC units and determine the availability of cooling resources at different locations within a data center.




**Figure 8: Temperature distribution in a cross-section of a datacenter**

In operational datacenters, temperature sensors within the data center provide valuable information on the current thermal state in the data center. The sensors, placed in the front and rear of server racks, provide data at regular intervals [4]. Analysis of this data is critical in deciding where new equipment should be placed inside

the data center. Without such data, one may go by intuition and common knowledge of the data center which may be misleading and simplistic. Placing new equipment in areas around hot spots would lead to higher probability of the systems powering off due to high temperature.

For example, a data center administrator may decide to place 20 new servers in a particular zone inside the data center, because it is relatively cool in that space at the time when the equipment was being placed. But this may be a temporary phenomenon, the temperatures at that point may be low because other machines are switched off and may come on later. Collecting data from temperature sensors would have shown that historically this particular area is loaded heavily and this is not the best place for new IT equipment. Figure 7 illustrates a visual representation of diurnal data from five sensors that belong to one rack. Observe the changes in color shades denoting temperature variations during the day and those among the sensors within the rack. Ultimately, a multi-dimensional context that includes time, space, and rack configuration is important [9].



**Figure 9: Visual representation of rack temperature data**

##### 4.3 Eliminating Inefficiencies

Visual analytics also help in determining which zones of the data center are more critical in terms of cooling and administrators can have stricter monitoring in place for those zones only. These techniques also help to isolate recirculation zones inside the data centers where hot and cold air mixes. Elimination of such zones is key to improving energy efficiency. Areas where administrators see high temperatures could be because of high density or recirculation which can be fixed by placing blanking panels on the racks. Such knowledge also helps in deciding redeployment of the IT equipment from a zone which is ineffectively cooled to one where adequate cooling is available. These decisions can be made by analyzing the thermal trends inside the data center. It makes sense to move equipment within a data center only if a cooler zone has remained cool for a considerable duration. Making decisions based on instant thermal situations can be disastrous. As an example, the temperature variation among sensors shown in Figure 9 may be indicative of thermal inefficiencies within the rack. There may be gaps between servers leading to mixing of hot and cold air.

Advanced data analytics and knowledge models are needed to identify inefficiencies in infrastructure. Apart from visualization, monitoring the operations of the chiller unit can also provide valuable information about infrastructure problems to involve the knowledge experts for problem solving at the right time.

##### 4.4 Anomaly detection

Several techniques have been explored for anomaly detection among which visualization and principal component analysis are notable. Our visualization tool displays the thermal map of a data center with physical locations of racks, sensors, and CRAC units.

Visual inspection of the map can help detect hotspots, especially useful after reorganizations of racks, computers, or CRAC units. It can also be used to detect abnormal thermal operating conditions, for example, where temperature variations are unexpected and could indicate air recirculation. The visualization tool also aids in analyzing the dynamics of the data center when any major or minor event occurs.

Correlation engine provides additional information around the cause and impact of such events. Figure 10 shows the anomaly detection process of one such event. Observe the trace of the short term perturbation. Further, advanced queries can be made to identify similar variations and possible causes. Perturbation and long term transients in air flow and temperature can radically change the temperature distributions at the inlet of racks.

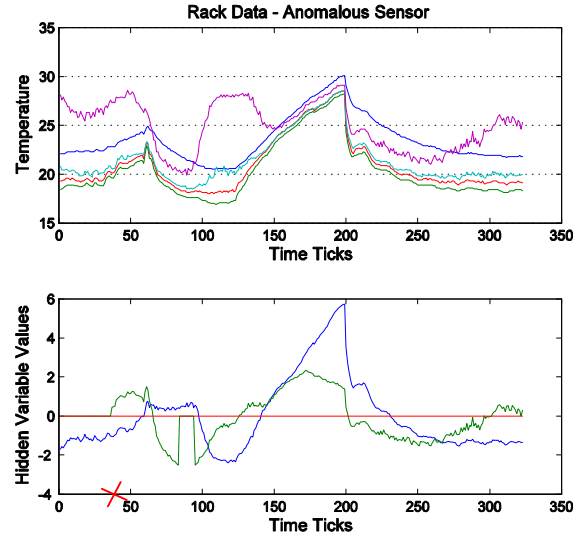


**Figure 10: Identification and Analysis of thermal events**

Another technique to discover anomalous behavior in a data center is to analyze sensor data using principal component analysis (PCA) [11]. PCA is a generic technique to reduce the dimensionality of correlated variables by introducing a new orthogonal basis. These are called the principal components (PCs). The original variables in the transformed space are called hidden variables. Change in the number of principal components indicates change in the correlations between variables. This fact can be used to detect anomalous behavior that manifests as broken correlations between variables. Furthermore, PCA can be performed in an incremental fashion [12] on sensor data streams to expeditiously detect anomalous behavior. Figure 11 shows PCA analysis for five temperature sensors. From the raw temperature plot, it can be seen that one temperature shows erratic behavior. This fact is autonomously discovered (without visual inspection) during the PCA analysis by the emergence of a second hidden variable around time tick 40 (marked with a cross on the x-axis). During normal operation one hidden variable is sufficient to summarize all five sensors. PCA is useful in detecting changes in correlation among variables and thus is able to identify hard to detect anomalies where thresholds may not have been violated.

## 5. COMPUTE

Given the importance of power efficiency, energy proportional computing [13], implying close to zero power consumption at zero utilization and linear increase in power consumption with increase in utilization, has been proposed for compute devices and components. However, until energy proportional computing is realized, techniques such as workload consolidation, turning off idle machines, and moving workload to more cooling efficient locations can be used to make a data center more sustainable. For



these techniques to be effective, data analysis and models are essential.

**Figure 11: The top figure shows raw temperature values at rack inlet; the bottom one shows the resulting hidden variables.**

## 5.1 Power Management


Dynamic voltage and frequency scaling (DVFS) [14] allows machines to be put in multiple power states (called p-states) depending on their utilization levels. Low power states provide less capacity while consuming lower power. Furthermore, machines can be turned off when they are idle. Server usage patterns based on collected data can be used to turn idle machines off with least impact on the users. For example, daily, weekly, or seasonal usage patterns can be modeled in order to predict the resource requirements at any time. This would allow resources/machines to be turned on or off in anticipation of an increase or decrease in workload, respectively. In our test bed data center, we have deployed a server provisioning tool which can be used to classify servers within the data center into different categories, such as, test machines, production machines, business critical machines and storage machines. Also associated with each category is a usage pattern. For instance, using this information, it can decide which machines are not used during nights and weekends, and thus can be turned off during those times. These optimizations have resulted in significant power savings for the facility. Figure 12 shows the server power consumption over a week.

## 5.2 Workload Placement

Given the flexibility provided by techniques, such as, server virtualization [15] and load balancing [16], workload can be distributed in a data center based on some criterion. Server virtualization facilitates migration of workload between machines while load balancing can be used to direct incoming workload to a specific machine. The criteria for workload placement can be derived from optimization of desired goals such as sustainability. An example of such a policy is to place workload in areas of a data center that are easier to cool resulting in cooling power savings [7]. A metric, called local workload placement index

(LWPI), is defined that quantifies the efficiency of cooling at a particular time and location in a data center. In order to calculate LWPI values at different locations, it is assumed that a temperature sensor network exists to obtain temperature measurements of those locations. In addition to cooling efficiency, reliability of cooling at a particular location – measured by the degree of overlap of cooling resources (e.g. CRAC units) – is another metric to be considered by a workload placement algorithm. Given multiple locations with the same cooling efficiency, more critical loads could be placed at locations or zones with higher cooling reliability. Depending on user preferences, utility functions could be used to capture the trade-offs between these two and other potentially competing criteria.

Zephyr [17] uses thermal models for minimizing the total power consumption – both server power and fan cooling power – of a server blade enclosure (with sixteen blades cooled by ten fans). Consolidation of servers is performed through use of server virtualization and turning off idle machines. For a given workload distribution, the cooling (fan) power is minimized by formulating a convex optimization problem [17] and solving it to determine fan speeds of the ten fans that minimize the fan power consumption. Zephyr also uses server power states, described in the previous Section 5.1, to further reduce the power consumption of each powered on blade.



Average Savings due to Daily Shutdown: 60.5 KW  
Savings due to Weekend Shutdown: 64 KW per day

Figure 12: Server power load in the test bed data center. Power dips are seen due to machines being turned off.

## 6. CONCLUSIONS

As data centers increase in number, size and power densities, their sustainable operation and management become critical. In this paper, we presented use cases and recommendations for application of data analysis, visualization and knowledge discovery to make a data center more sustainable. These techniques can play a key role in improving energy efficiency and increasing the lifetime of cooling, power and compute infrastructure. Furthermore, the lessons learnt during operation and management can be applied to synthesis of a sustainable data center.

## 7. REFERENCES

[1] Cloud Computing, [http://wikipedia.org/Cloud\\_computing](http://wikipedia.org/Cloud_computing)

[2] Software as a Service (SaaS), [http://wikipedia.org/Software\\_as\\_a\\_service](http://wikipedia.org/Software_as_a_service)

[3] U.S. Environmental Protection Agency, Report to Congress on Server and Data Center Energy Efficiency Public Law 109-431, Aug 2007

[4] Patel, C.D., Bash, C.E., Sharma, R.K, Beitelmal, A, Friedrich, R.J., “Smart Cooling of Datacenters”, Proceedings of ASME International Electronics Packaging Technical Conference and Exhibition, Maui, HI, IPACK-03, July 2003

[5] A.J. Shah, V.P. Carey, C.E. Bash, C.D. Patel, “Exergy Analysis of Data Center Thermal Management Systems,” *ASME Journal of Heat Transfer*, vol. 130, no. 2, pp. 021401-10, February 2008.

[6] C. Marnay, “Microgrids and Heterogeneous Security, Quality, Reliability, and Availability,” *Proc. 4th Power Conversion Conf.*, IEEE Press, 2007, pp. 629–634.

[7] Bash, C.E., Forman, George, “Cool Job Allocation: Measuring the power savings of placing jobs at cooling efficient locations in the data center.”, Proceedings USENIX Annual Technical Conference, 2007

[8] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt. Best practices for data centers: Results from benchmarking 22 data centers. In Proc. of the 2006 ACEEE Summer Study on Energy Efficiency in Buildings., Pacific Grove, CA, Aug. 2006.

[9] Hao et. al., Interactive Poster: Visual Monitoring of Temperature Data In a Smart Data Center, IEEE VisWeek 2008.

[10] Sharma R.K., Shih R., Bash, C.E., Patel, C.D., Varghese, P., Mekanapurath, M., Velayudhan, S., Kumar M.V., “On building Next Generation Data Centers”, proceedings of Compute 2008, Jan 18-20, Bangalore, Karnataka, India

[11] I. T. Jolliffe. Principal Component Analysis. Springer, 2002.

[12] S. Papadimitriou, S. Jimeng, C. Faloutsos, “Streaming Pattern Discovery in Multiple Time-Series”, Proceedings of the 31st VLDB Conference, 2005

[13] L. A. Barroso , U. Holzle. The case for energy-proportional computing. *IEEE Computer*, 40(12):33–37, 2007.

[14] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam. Managing server energy and operational costs in hosting centers. In Proc. of SIGMETRICS, pages 303–314, Banff, Canada, June 2005.

[15] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live migration of virtual machines. In Proc. of the 2nd NSDI, pages 273–286, Boston, MA, May 2005.

[16] James Aweya, Michel Ouellette, Delfin Y. Montuno, Bernard Doray, and Kent Felske. An adaptive load balancing scheme for web servers. *Int. Journal of Network Management*, 12(1):3–39, 2002.

[17] N. Tolia, Z. Wang, M. Marwah, C. Bash, P. Ranganathan, and X. Zhu. Zephyr: a unified predictive approach to improve server cooling and power efficiency. Technical Report HPL-2008-107, HP Laboratories, Palo Alto, CA, 2008.



