
How Does Your Data Center Measure Up?
**Energy Efficiency Metrics and Benchmarks for Data
Center Infrastructure Systems**

Paul Mathew, Ph.D., Staff Scientist

Steve Greenberg, P.E., Energy Management Engineer

Srirupa Ganguly, Research Associate

Dale Sartor, P.E., Applications Team Leader

William Tschudi, P.E., Program manager

Environmental Energy Technologies Division

Lawrence Berkeley National Laboratory

April 2009

Disclaimer

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

Acknowledgements

This work was supported by the New York State Energy Research and Development Authority (NYSERDA) and the Assistant Secretary for Energy Efficiency and Renewable Energy, Building Technologies Program, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

How Does Your Data Center Measure Up? **Energy Efficiency Metrics and Benchmarks for Data Center Infrastructure Systems**

Paul Mathew, Ph.D., Staff Scientist

Steve Greenberg, P.E., Energy Management Engineer

Srirupa Ganguly, Research Associate

Dale Sartor, P.E., Applications Team Leader

William Tschudi, P.E., Program manager

Lawrence Berkeley National Laboratory

1 Introduction

Data centers are among the most energy intensive types of facilities, and they are growing dramatically in terms of size and intensity [EPA 2007]. As a result, in the last few years there has been increasing interest from stakeholders - ranging from data center managers to policy makers - to improve the energy efficiency of data centers, and there are several industry and government organizations that have developed tools, guidelines, and training programs.

There are many opportunities to reduce energy use in data centers and benchmarking studies reveal a wide range of efficiency practices. Data center operators may not be aware of how efficient their facility may be relative to their peers, even for the same levels of service. Benchmarking is an effective way to compare one facility to another, and also to track the performance of a given facility over time.

Toward that end, this article presents the key metrics that facility managers can use to assess, track, and manage the efficiency of the infrastructure systems in data centers, and thereby identify potential efficiency actions. Most of the benchmarking data presented in this article are drawn from the data center benchmarking database at Lawrence Berkeley National Laboratory (LBNL). The database was developed from studies commissioned by the California Energy Commission, Pacific Gas and Electric Co., the U.S. Department of Energy and the New York State Energy Research and Development Authority.

2 Data Center Infrastructure Efficiency (DCIE)

This metric is the ratio of the IT equipment energy use to the total data center energy use and can be calculated for annual site energy, annual source energy, or electrical power.

$$\text{DCIE}_{\text{site}} = \text{IT site energy use} / \text{Total site energy use}$$

$$\text{DCIE}_{\text{source}} = \text{IT source energy use} / \text{Total source energy use}$$

$$\text{DCIE}_{\text{elecpower}} = \text{IT electrical power} / \text{Total electrical power}$$

Note that the data center energy use is the sum of all energy used by data center, including campus chilled water and steam if present. The DC Pro tool [DOE 2008] can be used to assess DCIE for site and source energy.

DCIE provides an overall measure of the infrastructure efficiency i.e. lower values relative to the peer group suggest higher potential to improve the efficiency of the infrastructure

systems (HVAC, power distribution, lights) and vice versa. Note that it is not a measure of IT efficiency. Therefore a data center that has a high DCIE may still have major opportunities to reduce overall energy use through IT efficiency measures such as virtualization. It should also be noted that DCIE is influenced by climate and tier level, and therefore the potential to improve it may be limited by these factors.

Some data center professionals prefer to use the inverse of DCIE, known as Power Utilization Effectiveness (PUE), but both metrics serve the same purpose.

The benchmarking data in the LBNL database show a $DCIE_{elecpower}$ range from just over 0.3 to 0.75. Some data centers are capable of achieving 0.9 or higher [Greenberg et al. 2009].

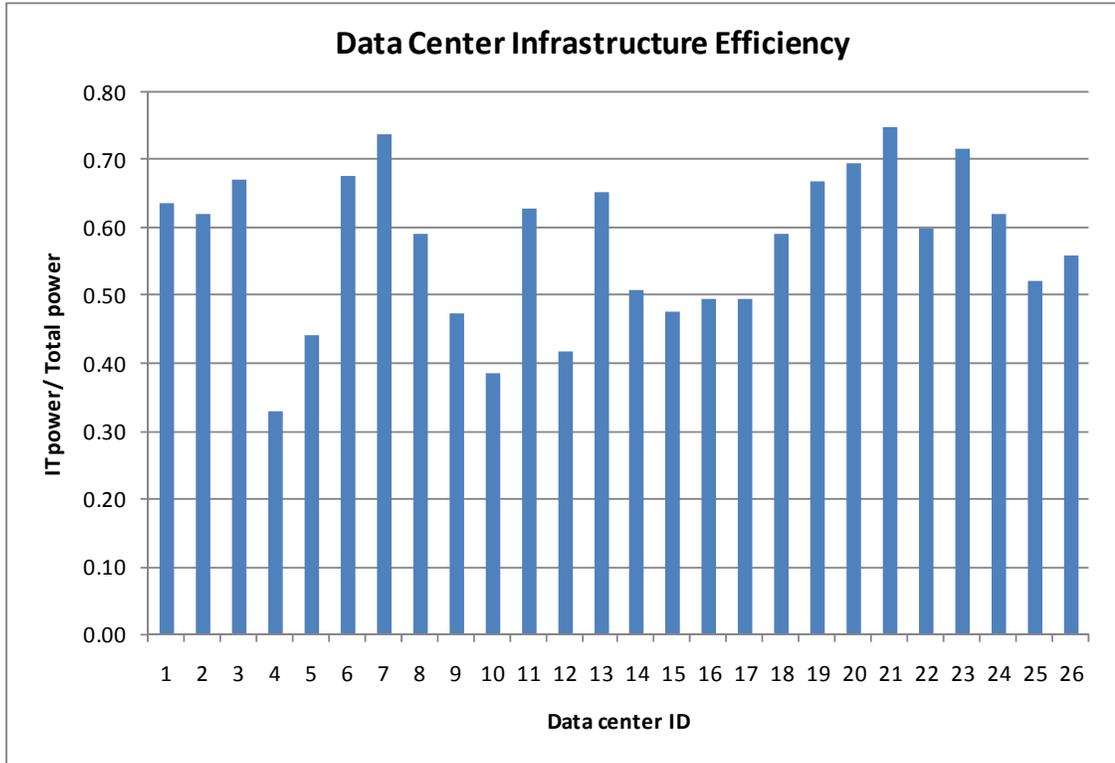


Figure 1. Data center infrastructure efficiency for data centers in the LBNL database. Note that these DCiE values are based on power, not energy.

3 Temperature and Humidity Ranges

The American Society of Heating, Refrigerating and Air-conditioning Engineers (ASHRAE) guidelines [ASHRAE 2008] provide a range of allowable and recommended supply temperatures and humidity at the inlet to the IT equipment.

The recommended temperature range is between 64F-80F, while the allowable is 59F-90F. A low supply air temperature and a small temperature differential between supply and return typically indicate the opportunity to improve air management, raise supply air temperature and thereby reduce energy use. Strategies to improve air management include better isolation between cold and hot aisles using blanking panels and strip curtains, optimizing configuration of supply diffusers and return grilles, better cable management, blocking gaps in floor tiles, etc.

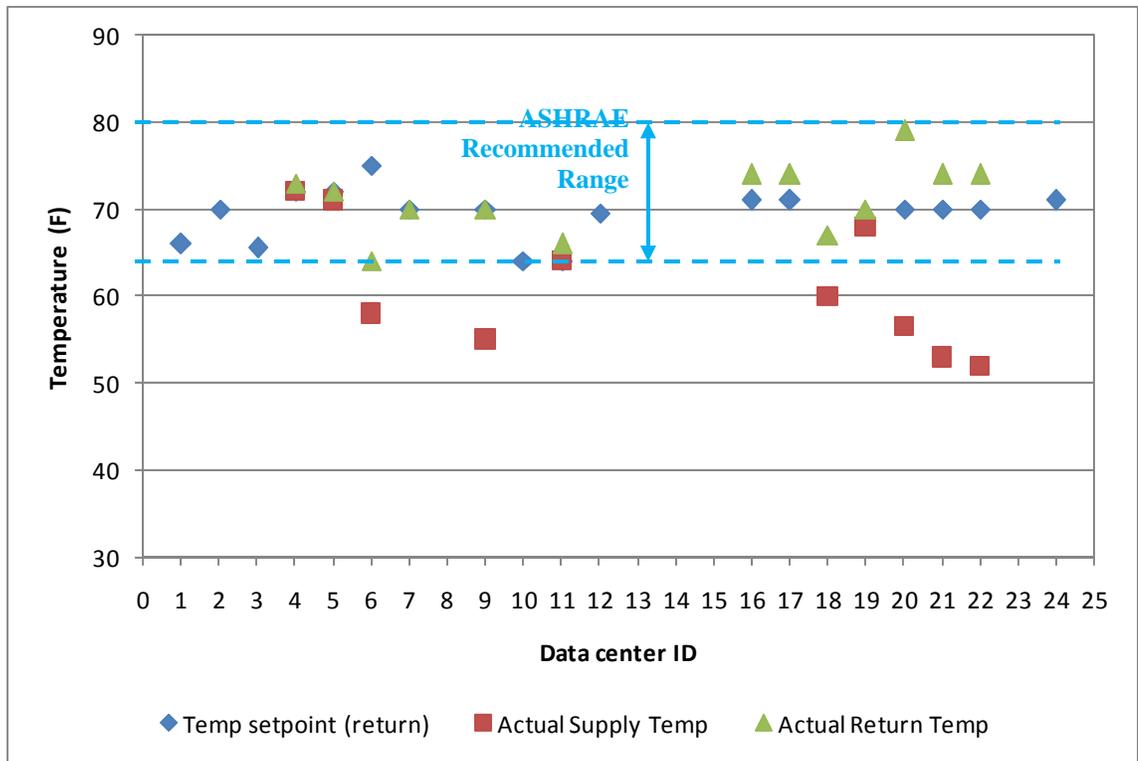


Figure 2. Return air temperature setpoints, measured supply and return temperature for data centers in the LBNL database show that many data centers are operated at lower temperatures than required.

The recommended humidity range is between a lower end defined as a minimum dew point of 42F and the upper end set at 60% relative humidity and 59F dewpoint. The allowable relative humidity range is between 20%-80% and 63F maximum dewpoint. A small, tightly controlled, relative humidity range suggests opportunities to reduce energy use, especially if there is active humidification and dehumidification. Centralized active control of the humidification units reduces conflicting operations between individual units, thereby improving the energy efficiency and capacity.

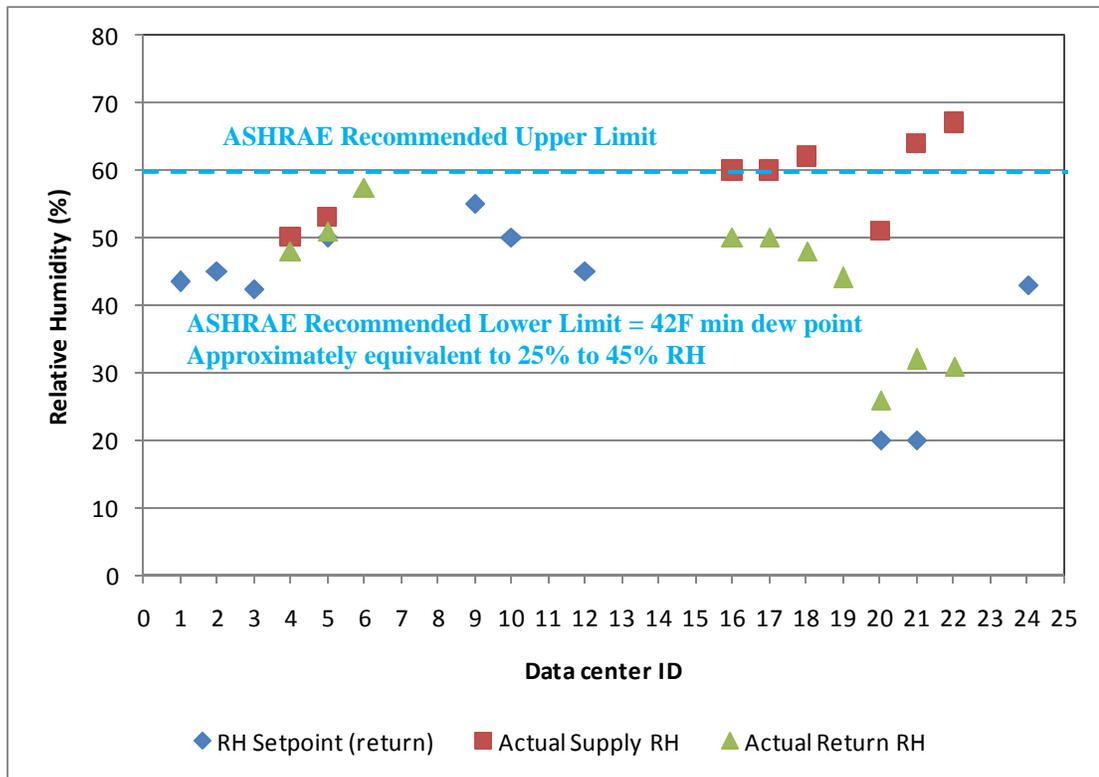


Figure 3. Return air relative humidity setpoints, measured supply and return relative humidity for data centers in the LBNL database

Since temperature and humidity affect the reliability and life of IT equipment, any changes to the air management and temperature and humidity settings should be evaluated with metrics such as the Rack Cooling Index (RCI) [Herrlin 2005], which can be used to assess the thermal health of the IT equipment. Many data centers operate well without active humidity control. While humidity control is important for physical media such as tape storage, it is generally not critical for the rest of the data center equipment. Studies by LBNL and the Electrostatic Discharge Association suggest that humidity may not need to be as tightly controlled.

4 Return Temperature Index

This metric is a measure of the energy performance resulting from air management [Herrlin 2007]. The primary purpose of improving air management is to isolate hot and cold airstreams. This allows elevating both the supply and return temperatures and maximizes the difference between them while keeping the inlet temperatures within ASHRAE recommendations. It also allows reduction of the system air flow rate. This strategy allows the HVAC equipment to operate more efficiently. The return temperature index (RTI) is ideal at 100% wherein the return air temperature is the same as the temperature leaving the IT equipment.

RTI is also a measure of the excess or deficit of supply air to the server equipment. An RTI value of 100% is ideal. An RTI value of less than 100% indicates that the some of the supply air is by-passing the racks, and a value greater than 100% indicates that there is recirculation of air from the hot aisle. The RTI value can be close to ideal (100%) by improving air management.

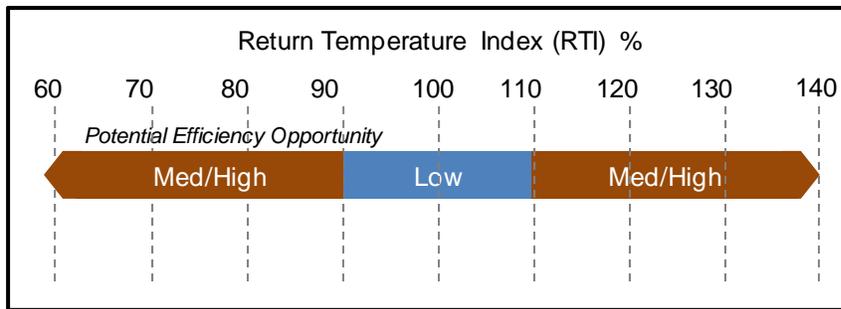


Figure 4. Benchmarks for Return Temperature Index

5 UPS load factor

This metric is the ratio of the load of the uninterruptible power supply (UPS) to the design value of its capacity. This provides a measure of the UPS system over-sizing and redundancy. A higher value of this metric indicates a more efficient system. UPS load factors below 0.5 may indicate an opportunity for efficiency improvements, although the extent of the opportunity is highly dependent on the required redundancy level. The load factor can be improved by several means, including the following:

- Shutdown some UPS modules when Redundancy Level exceeds N+1 or 2N
- Install a scalable/modular UPS
- Install a smaller UPS size to fit present load capacity
- Transfer loads between UPS modules to maximize load factor % per active UPS

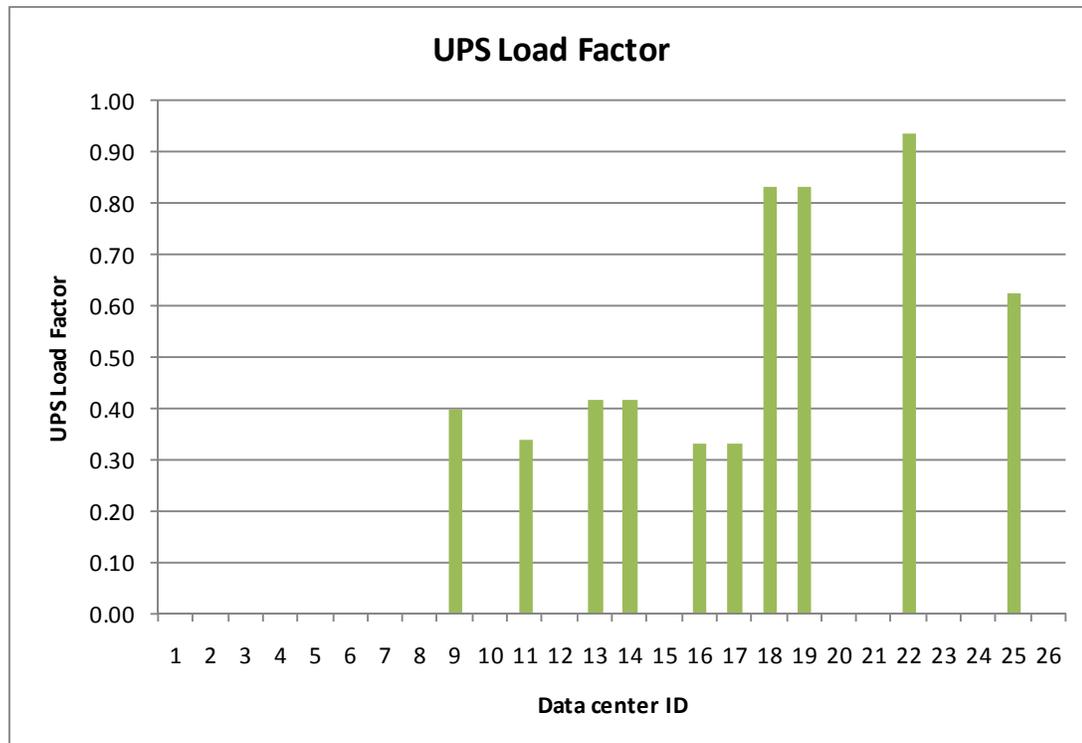


Figure 5. UPS load factor for data centers in the LBNL database

6 UPS System Efficiency

This metric is the ratio of the UPS input power to the UPS output power. The UPS efficiency varies depending on its load factor and therefore the benchmark for this metric depends on the load factor of the UPS system. At UPS load factors below 40% the system can be highly inefficient due to no load losses. Figure 6 shows the range of UPS efficiencies from factory measurements of different topologies. Figure 7 shows the UPS efficiencies for data centers in the LBNL database. These measurements taken several years ago illustrate that efficiencies vary considerably. Manufacturers claim that improved efficiencies are available today. When selecting UPS systems, it is important to evaluate performance over the expected loading range.

Selection of more efficient UPS systems, especially the ones that perform well at expected load factors (e.g. below 40%) improves energy savings. For non-critical IT work bypassing the UPS system using factory supplied hardware and controls may be an option. Reducing the level of redundancy by using modular UPS systems also improves the efficiency.

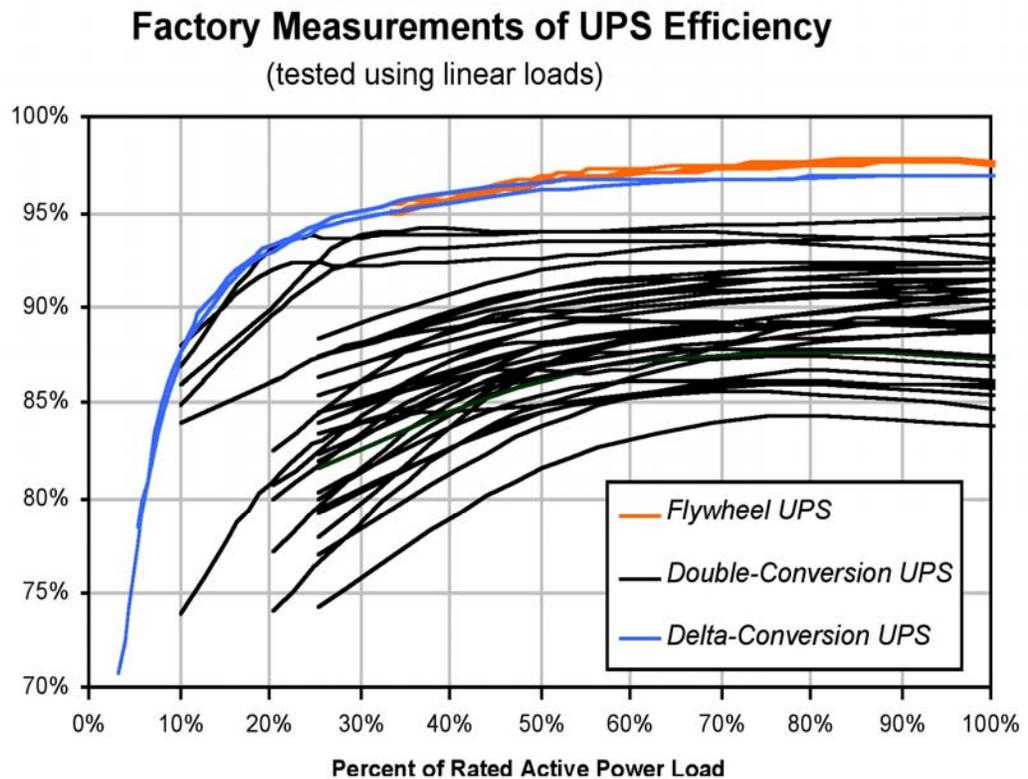


Figure 6. Range of UPS system efficiencies for factory measurements of different topologies

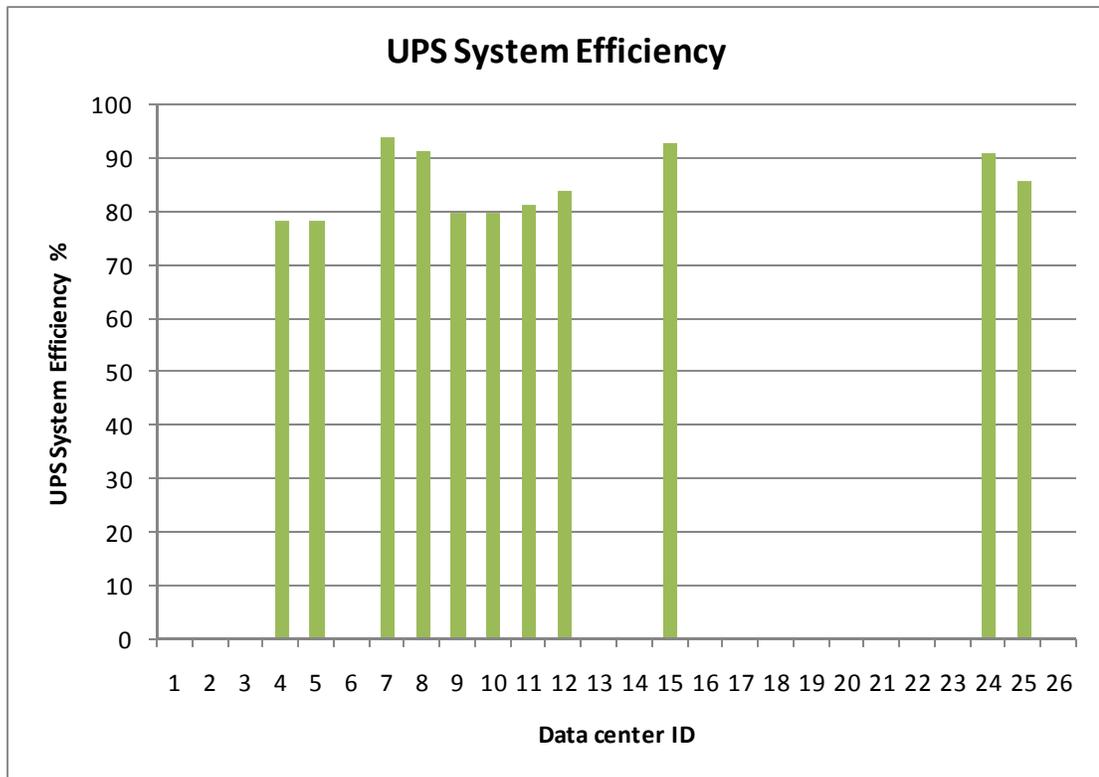


Figure 7. UPS efficiency for data centers in the LBNL database

7 Cooling system efficiency

The key metrics and benchmarks to evaluate the efficiency of cooling systems in data centers are no different than those typically used in other commercial buildings. These include chiller plant efficiency (kW/ton), pumping efficiency (hp/gpm), etc. Since these are well-documented elsewhere, they are not further discussed here. Figure 8 shows the cooling plant efficiency for LBNL benchmarked data centers. Based on data from the LBNL database, 0.8 kW/ton could be considered as a good practice benchmark and 0.6 kW/ton as a better practice benchmark.

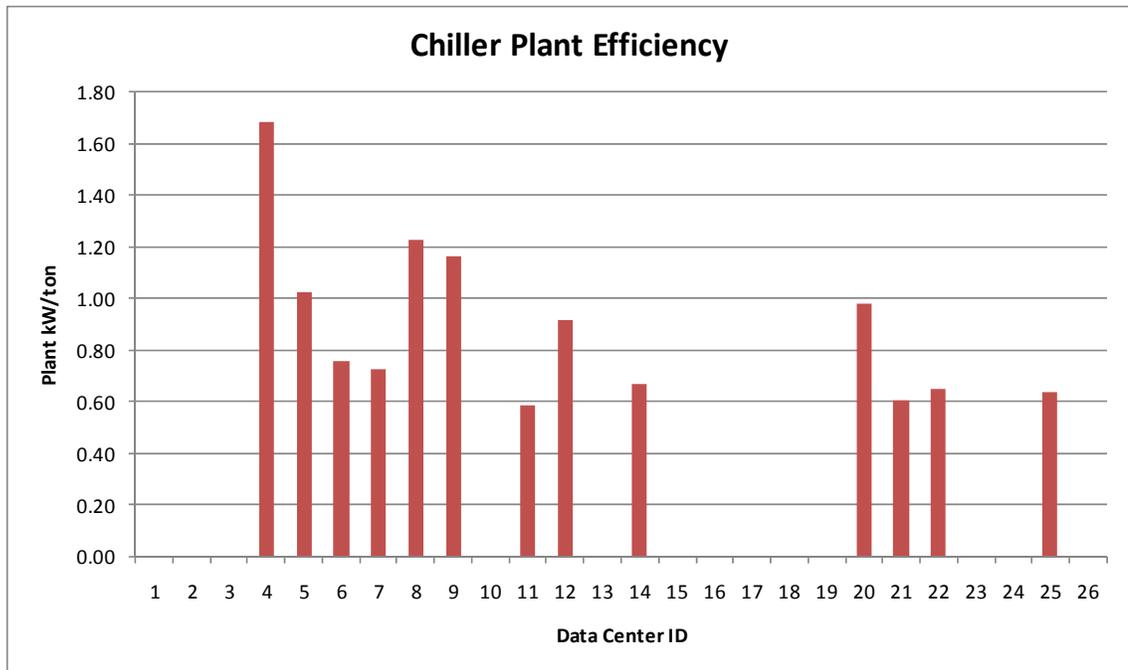


Figure 8. Cooling plant efficiency for LBNL benchmarked data centers

8 Air economizer utilization

This metric characterizes the extent to which air-side economizer system is being used to provide “free” cooling. It is defined as the percentage of hours in a year that the economizer system can be in full or complete operation (i.e. without any cooling being provided by the chiller plant). The number of hours that the air economizer is being utilized could be compared to the maximum possible for the climate in which the data center is located. This can be determined from simulation analysis. As a point of reference, Figure 9 shows results from simulation analysis for four different climate conditions. The Green Grid has developed a tool to estimate savings from air- and water-side free cooling [TGG 2009], though the assumptions are different than those used for the results presented in Figure 9.

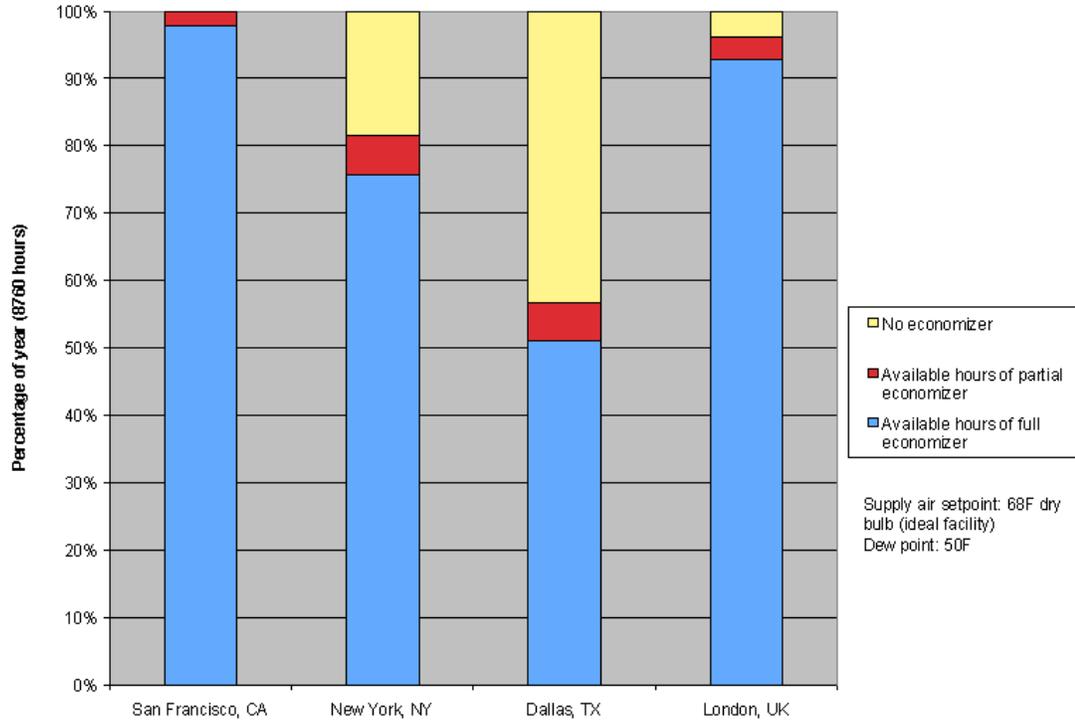


Figure 9. Simulated air-side economizer utilization potential for four different locations.
Data source: Syska Hennessy 2007

9 Summary and Outlook for Productivity Metrics

This article presented key metrics that data center operators can use to track the efficiency of their infrastructure systems. The system level metrics in particular allow operators to help identify potential efficiency actions. The article also presented data from the LBNL benchmarking database which shows that there is a wide range of efficiency across the surveyed datacenters.

DCIE is gaining increasing acceptance as a metric for overall infrastructure efficiency, and can be computed in terms of site energy, source energy or electrical load. However, DCIE does not address the efficiency of IT equipment. Organizations such as the Green Grid are working to develop productivity metrics (e.g. Data Center energy Productivity (DCeP)) which will characterize the work done per unit of energy. The challenge is to categorize the different kinds of work done in a data center and identify appropriate ways to measure them. As they become available, they will complement the infrastructure metrics described in this article.

10 References

- ASHRAE 2008. 2008 ASHRAE Environmental Guidelines for Datacom Equipment - Expanding the Recommended Environmental Envelope. ASHRAE Datacom series. American Society of Heating, Refrigerating and Air-Conditioning Engineers, Inc.

- DOE 2008. Data Center Profiling tool “DC Pro” <http://dcpro.ppc.com/>
- EPA 2007. Report to Congress on Server and Data Center Energy Efficiency Public Law 109-43. U.S. Environmental Protection Agency ENERGY STAR Program.
- Greenberg, S., Khanna, A., and Tschudi, W. 2009. "High Performance Computing with High Efficiency" ASHRAE Transactions TRNS-00232-2008. In press. To be presented at the ASHRAE Annual Conference, Louisville KY, June 2009.
- Herrlin, M.K. 2005. “Rack cooling effectiveness in data centers and telecom central offices: The rack cooling index (RCI).” ASHRAE Transactions 111(2):725-731
- Herrlin, M.K. 2007. Improved Data Center Energy Efficiency and Thermal Performance by Advanced Airflow Analysis. Digital Power Forum, 2007. San Francisco, CA, September 10-12, 2007. <http://www.ancis.us/publications.html>
- Syska Hennessy 2007. “The Use of Outside Air Economizers In Data Center Environments.” White paper 7. Syska Hennessy Group.
http://www.syska.com/critical/knowledge/wp/wp_outsideair.asp
- TGG 2009. Free cooling estimated savings calculation tool. The Green Grid.
http://cooling.thegreengrid.org/calc_index.html .