

Chapter X: Data Center IT Efficiency Measures

The Uniform Methods Project: Methods for Determining Energy-Efficiency Savings for Specific Measures

Robert Huang, Cadmus

Eric Masanet, Ph.D., Northwestern University

Subcontract Report: NREL/SR-&A30-53827, XXXX 2013

Acknowledgements

The chapter author wishes to thank and acknowledge Dale Sartor of Lawrence Berkeley National Lab, Thomas Bolioli of Terra Novum, Dan Barbieri of DNV-GL, Naomi Cole of PEI, and Eric Winkler of ISO New England for their thoughtful contributions.

Acronyms

EM	Efficiency metric
EUL	Expected useful life
HDD	Hard disk drive
HVAC	Heating, ventilation, and air conditioning
IT	Information technology
M&V	Measurement and verification
MAID	Massive array of idle disks
PDU	Power distribution unit
PUE	Power usage effectiveness
RAID	Redundant array of independent disks
SAS	Serial attached small computer system interface
SATA	Serial advanced technology attachment
SERT	Server Efficiency Rating Tool
SSD	Solid-state drive
UPS	Uninterruptible power supply
VSD	Variable-speed drive

Chapter X – Table of Contents

- 1 Measure Description 1**
 - 1.1 Server Virtualization..... 1
 - 1.2 More Efficient Servers..... 1
 - 1.3 Data Storage Management..... 2
 - 1.4 More Efficient Data Storage Equipment..... 3
- 2 Application Conditions of the Protocol..... 4**
- 3 Savings Calculations 6**
 - 3.1 The Simple Algorithm 6
 - 3.2 Complicating Issues With the Simple Algorithm 6
 - 3.3 Calculating Data Center IT Savings 8
 - 3.4 Calculating Total Energy, Lifetime, and Peak Demand Savings..... 12
- 4 Measurement and Verification Plan..... 14**
 - 4.1 International Performance Measurement and Verification Protocol Option 14
 - 4.2 Verification Process..... 14
 - 4.3 Data Requirements/Collection Methods 15
- 5 Other Evaluation Issues..... 20**

List of Figures

- Figure 1. Challenges with determining gross savings of data center IT measures..... 7**
- Figure 2. Challenges with determining “burnout only” gross savings of data center IT measures .. 8**
- Figure 3. Watts per terabyte for various data storage types 19**
- Figure 4. Sample SERT data for Server A..... 24**
- Figure 5. Sample SERT data for Server B..... 25**

List of Tables

- Table 1. Examples of Data Center IT Incentives Across the Country as of October 2013 4**
- Table 2. Verification of Key Inputs Into Equations 15**
- Table 3. SERT Workload Types, Worklet Names, and Characteristics 23**

1 Measure Description

Data centers use about 2% of the electricity in the United States (Koomey 2011); a typical data center has 100 to 200 times the energy use intensity of a commercial building. Data centers present tremendous opportunities -- energy use can be reduced as much as 80% between inefficient and efficient data centers (DOE 2011). Data center efficiency measures generally fall into the following categories:

- Power infrastructure (e.g., more efficient uninterruptible power supplies [UPS], power distribution units [PDUs])
- Cooling (e.g., free cooling, variable-speed drives [VSDs], temperature and humidity set points)
- Airflow management (e.g., hot aisle/cold aisle, containment, grommets)
- Information technology (IT) efficiency (e.g., server virtualization, efficient servers, efficient data storage).

This chapter focuses on *IT measures in the data center* and examines the techniques and analysis methods used to verify savings that result from improving the efficiency of two specific pieces of IT equipment: servers and data storage. The discussion examines options in two categories:

- Using more efficient server and data storage equipment
- Managing servers and data storage equipment to work more efficiently.

Section 1.1 describes some common IT measures that save energy in data centers.

1.1 Server Virtualization

In the past, data center operators ran a single application on each server. This “one workload, one box” approach meant servers ran at a low “utilization rate”: the fraction of total computing resources engaged in useful work (EPA undated a). A 2012 *New York Times* article cited two sources that estimated average server utilization rate of 6% to 12% (EPA undated b). Another study stated that the “one workload, one box” approach resulted in 90% of all x86 servers running at less than 10% utilization, with a typical server running at less than 5% utilization (EPA undated b).

Administrators can use server virtualization to run multiple applications on one physical host server, thus consolidating server resources. In other words, multiple virtual servers can work simultaneously on a single physical host server. Therefore, instead of operating many servers at low utilization rates, virtualization combines the processing power onto fewer servers, operating at higher total utilization rates.

1.2 More Efficient Servers

ENERGY STAR[®]-certified servers have been available since 2009. The ENERGY STAR server specification covers four server form factors (blade, multi-node, rack-mounted, and pedestal) and

allows a maximum of four process sockets per server (or per blade or node). ENERGY STAR servers must have the following features:

- Efficient power supplies to limit power conversion losses
- Improved power quality
- Idle power draw limits for rack-mounted or pedestal servers with one or two processors;
- Results of the Server Efficiency Rating Tool (SERT) tests to accommodate comparisons of server efficiency under various usage scenarios
- Ability to measure real-time power use, processor utilization, and air inlet temperatures
- Advanced power management features and efficient components that save energy across various operating states (including idle)
- A Power and Performance data sheet for purchasers; this standardizes key information on energy performance, features, and other capabilities.

On average, ENERGY STAR servers operate about 30% more energy efficiently than standard servers. The servers operate particularly efficiently at low loads because processor power management requirements reduce power consumption when the servers are idle (EPA undated b).

1.3 Data Storage Management

Data storage resource management tools (Clark and Yoder 2008) help data storage administrators more efficiently and effectively provision and manage data storage. This entails using tools to create “maps” and “pools” of available storage across servers and disks, and using these disparate “chunks” of storage as if they operated as one system. These tools include:

- **Automated storage provisioning.** This improves storage efficiency through right-sizing, identifies and reallocates unused storage, and increases server capacity by improving existing storage use (Netapp 2014).
- **Deduplication software.** This condenses the data stored at many organizations by more than 95% by finding and eliminating unnecessary copies. Redundant copies consume more than half the total volume of a typical company’s data.
- **Thin provisioning.** This allocates just enough storage just in time by centrally controlling capacity and allocating space only as applications require it. Thus, administrators power only the storage currently in use.
- **Redundant array of independent disks (RAID).** This level is a storage technology that combines multiple disk drive components into a single logical unit. RAID 1 creates a duplicate copy of disk data and doubles the storage and power consumption. For storage that is not mission critical, RAID 5 guards against a single disk drive failure in a RAID set by reconstructing the failed disk information from distributed information on the remaining drives. Requiring only one extra, redundant disk, RAID 5 saves energy, although it sacrifices some reliability and performance. For a 10-disk array, increasing to

an 11-disk RAID 5 level (one extra disk) from a 20-disk RAID 1 level (duplicate copy) configuration would save 45% of data storage energy use.

- **Tiering storage.** This automatically stores low-priority data (rarely accessed information) on higher-latency equipment that uses less energy.

1.4 More Efficient Data Storage Equipment

A number of data storage equipment types use less energy (Yoder 2012), including the following:

- **Lower speed drives.** Higher-spin speeds on high-performance hard disk drives (HDDs) (e.g., 15 K rpm serial attached small computer system interface [SAS]¹ drives) mean faster read/write speeds. All things being equal, power use is proportional to the cube of the disk spin speed. To reduce storage energy use, storage administrators should look for slower drives (e.g., 7.5 K rpm serial advanced technology attachment [SATA]² drives) that are available to accommodate specific tasks at hand.
- **Massive array of idle disks (MAID).** MAID operates more energy efficiently than older systems and often offers an effective solution for Tier 3 storage (data accessed infrequently). MAID saves power by shutting down idle disks, then powering the disks back up only when an application must access the data.
- **Solid-state drives (SSDs).** Energy-saving, solid-state storage increasingly offers an energy-efficient option. Without powering spinning disks, SSDs provide “read” speeds 10 times faster than hard disks. For example, compared to a 7.2 K rpm SATA disk, an SSD consumes one ninth the power per byte stored (Pflueger 2010). SSDs are, however, more expensive than conventional hard disk options.
- **ENERGY STAR-certified data storage (EPA undated b).** EPA’s ENERGY STAR program certifies energy-efficient online data storage that meets the following criteria:
 - Employs efficient power supplies that limit power conversion losses.
 - Relies on internal variable-speed fans for cooling.
 - Provides features to help better manage data, leading to reduced storage and energy consumption.

¹ SAS is a faster and historically more expensive interface that moves data to and from storage devices.

² SATA is the next-generation computer bus interface that moves data to and from storage devices.

2 Application Conditions of the Protocol

Unlike other efficiency measures in the Uniform Methods Project, data center IT measures present a new target for utility programs.³ As shown in [Table 1](#)~~Table 4~~, most utilities offer custom incentives for data center IT measures, where applicants must calculate and demonstrate savings from data center IT equipment. Utilities pay incentives based on actual verified savings. [Table 1](#)~~Table 4~~ shows a range of \$0.06 to \$0.16/kWh saved. In general, standard custom programs work in the following manner:

- A customer submits a project application that includes energy use of existing equipment, equipment required by code or standard, and the efficiency measure (PG&E 2013). In addition, customers must specify whether they install the efficiency measure as an early replacement (where an existing unit has remaining useful life) or at burnout (where the existing unit no longer operates).
- The utility inspects and approves the project before removing the existing equipment/systems and installing the new equipment/systems.
- Upon completion of the project, the utility inspects and approves installation of the measures and finalizes the incentive amounts.

Sometimes utilities offer prescriptive incentives for server virtualization. For example, Seattle City Lights and the Energy Trust of Oregon offer prescriptive incentives based on the number of servers retired. A company in the Seattle City Light territory could receive \$900 for retiring six servers through a virtualization effort. In developing the prescriptive incentive, utilities calculated predefined fixed average energy savings, or deemed values, for existing and efficient IT equipment.

Server virtualization also improves scalability, reduces downtimes, enables faster deployments, reduces IT footprints, and has become commonplace, especially in large data centers. A 2011 survey of more than 500 large enterprise data centers found that 92% use virtualization to some degree (Veeam 2011). Free-ridership concerns have caused some utilities to remove server virtualization from their data center efficiency programs. Silicon Valley Power’s Data Center Program (limited to larger data centers) does not provide incentives for server virtualization. (The program also does not allow IT equipment incentives, unless specifically approved.) PG&E and BC Hydro also stopped offering server virtualization incentives. This trend may continue as organizations redesign data center programs to adjust to market conditions.

Table 1. Examples of Data Center IT Incentives Across the Country as of October 2013

Utility	Measure	Incentive Amount	Notes
Seattle City Light (2013)	Custom IT Equipment—Plug Loads	\$0.06/kWh saved	Energy savings from custom projects where software or

³ As discussed in *Considering Resource Constraints* in the Introduction of this UMP report, small utilities (as defined under the Small Business Administration regulations) may face additional constraints in undertaking this protocol. Therefore, alternative methodologies should be considered for such utilities.

Utility	Measure	Incentive Amount	Notes
			hardware deployments save energy in IT equipment.
	Server Virtualization	\$150/server removed	Maximum of 200 servers removed.
NYSERDA (2014)	Examples listed: <ul style="list-style-type: none"> • Energy-efficient servers, storage, and switches • Server virtualization • Server refresh • Storage consolidation and optimization • High-performance computing systems 	\$0.12/kWh saved upstate \$0.16/kWh downstate	Capped at \$5 million per facility.
ComEd (2014)	Examples listed: <ul style="list-style-type: none"> • Virtualization • Consolidation • Thin-provisioning • Solid state storage 	\$0.07/kWh saved	Up to 100% of the incremental cost and 50% of the total cost of the project.
Energy Trust of Oregon (2014a, 2014b)	Virtualization	\$350 per server decommissioned	10 server minimum
Arizona Public Service (2014)	Example listed: server virtualization	\$0.09/kWh	Virtualization listed as “typical custom project,” up to 75% of incremental costs.
Southern California Edison (2012)	Reduced process load	\$0.08/kWh	Also \$100/kW.
Silicon Valley Power (2014a, 2014b)	Virtualization and consolidation of servers, IT equipment	Not Allowed	Large data centers (greater than 350 kW IT load or greater than 100 tons cooling) denied server virtualization/consolidation incentives. General, IT measure savings are not allowed unless specifically approved by SVP.

3 Savings Calculations

3.1 The Simple Algorithm

Unique challenges arise in calculating savings for data center IT measures. On one hand, savings estimates can appear straightforward. For custom incentives, calculations can use data center IT equipment power and energy readings taken from UPS, PDUs, or rack power strips. Estimated energy savings can use power draw readings (in kW) taken before and after measure implementation. Annual savings can be estimated using Equation 1 below:

$$\text{Annual Energy Savings} = 8760 * (\text{Power Draw}_{\text{Pre-Existing}} - \text{Power Draw}_{\text{Efficient Measure}}) \quad (1)$$

3.2 Complicating Issues With the Simple Algorithm

A number of challenges can, however, arise when calculating typical energy savings for a data center IT efficiency measure using Equation 1. [Figure 1](#) shows the typical factors involved in calculating early replacement and burnout energy savings for efficiency measures, including power draws (of efficient, standard/code, and preexisting measures) and the useful life (of existing measures and efficiency measures). The challenges include:

- The first challenge (represented by the red circles) arises from the difficulty in determining useful life. IT equipment generally does not stop working: rather, customers replace it for a variety of other reasons. For example, organizations often purchase new servers at the end of the old servers' service agreement or if new server features and capabilities require upgrades. Various International Data Corporation studies indicate organizations replace their servers once every 3 to 5 years (IDC 2010, 2012a, 2012b).

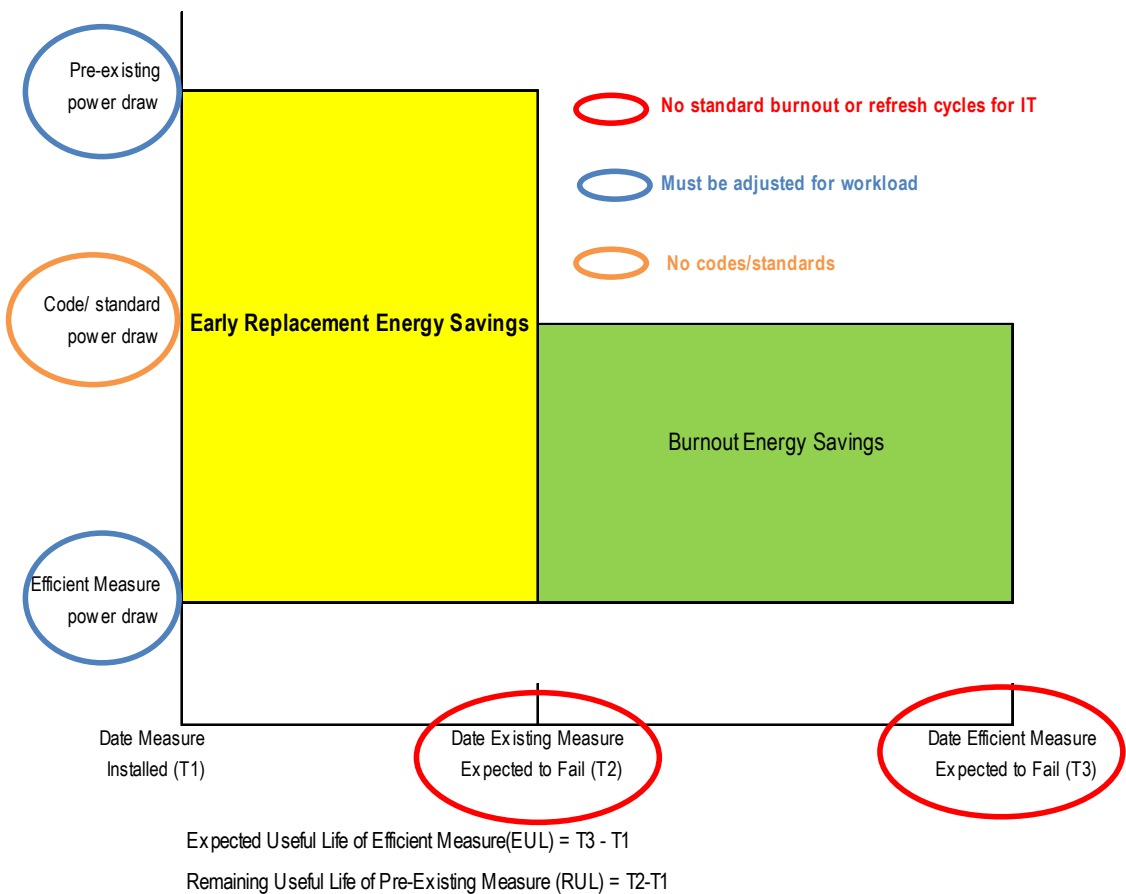


Figure 1. Challenges with determining gross savings of data center IT measures

- The second challenge (represented by the blue circles) arises from the varying power draws of IT equipment over time and per business demands, due to changes in the useful work output required of a device (e.g., email server workloads after large-scale layoffs). One would thus ideally normalize energy use for the data center workload to ensure accurate savings estimates. For example, if the data center workload increases just before ENERGY STAR servers are installed, the resulting power draw of the ENERGY STAR servers will be higher, producing underestimated savings. Conversely, if the data center workload decreases before new servers are installed, savings will be overestimated. Many ways to define workload-per-Watt have been proposed and used for data centers (e.g., CPU utilization/Watt, kB transmitted/Watt, GB storage/Watt, various benchmark workloads) (The Green Grid 2009; Pflueger 2010). There is, however, no single metric or industry standard for consistent measurement.
- The third challenge (represented by the orange circles) arises because—unlike many other efficiency measures in other sectors—energy codes or U.S. Department of Energy standards do not define “typical” or “standard” efficiencies for IT equipment. For such savings estimates, data center operators typically have information about the efficiency

measure and preexisting measure, but rarely have information about the “standard” unit, making calculation of burnout savings difficult.

3.3 Calculating Data Center IT Savings

As stated earlier, although in perfect working condition, data center IT equipment often undergoes upgrades when no longer useful (remaining useful life = 0) for reasons other than breaking down (e.g., expired service level agreements, antiquated feature sets, unsatisfactory workload performance issues, incompatibility with hardware-based management systems) (Search Data Center 2012). In other words, “early replacement” savings do not typically apply to data center IT equipment.

Therefore, the following sections present only savings calculations that focus on estimating burnout savings: the energy use difference between the hypothetical “standard” or “typical” equipment available on the market (not the existing equipment) and the efficient equipment to be installed. [Figure 2](#) shows the challenges that remain for calculating the burnout savings of IT equipment.




Figure 2. Challenges with determining “burnout only” gross savings of data center IT measures

3.3.1 Calculating Savings When Upgrading to More Efficient Servers

As stated, manufacturers have just started offering server efficiency metrics (EMs) that allow comparisons of server efficiencies.⁴ Server EMs soon will allow for simple comparisons between an efficient server and a “baseline” server, which will be established by examining the EMs of servers with similar configurations (e.g., chip sets, memory, and hard drives), computational outputs, and manufacturer years. Equation 2 shows the savings equation when server EMs increase when units becomes more efficient (e.g., operations/Watt), as with the new “efficiency score” generated by SERT. See the Appendix for an example of how to use SERT “efficiency scores” and Equation 2 to determine the savings from purchasing an energy-efficient server.

$$\text{Annual Energy Savings}_{\text{Efficient Servers}} = kW_{EE} * (EM_{EE}/EM_{\text{baseline}} - 1) * 8760 \quad (2)$$

Where,

kW_{EE} = power draw in kilowatts of new efficient server equipment

EM_{EE} = efficiency metric for efficient server

EM_{baseline} = efficiency metric for baseline server

8760 = number of hours in a year as servers run 24/7 in a data center

Another way to calculate savings for servers is to consider ENERGY STAR-certified servers as “efficient servers.” Using EPA estimates of percentage savings compared to standard or typical servers, savings can be calculated as shown in Equation 3.

$$\begin{aligned} \text{Annual Energy Savings}_{\text{ES Servers}} &= (kW_{\text{baseline}} - kW_{\text{ENERGY STAR}}) * 8760 \\ kW_{\text{ENERGY STAR}} &= \sum_{ES=1}^n (kW_{ES,\text{idle}} + U_{ES} * (kW_{ES,\text{full load}} - kW_{ES,\text{idle}})) \\ kW_{\text{baseline}} &= kW_{\text{ENERGY STAR}} / (1 - a) \end{aligned} \quad (3)$$

This approach leads to the following simplified expression shown in Equation 4.

$$\text{Annual Energy Savings}_{\text{ES Servers}} = \left(\frac{1}{(1-a)} - 1 \right) kW_{\text{ENERGY STAR}} * 8760 \quad (4)$$

⁴ EPA requires reporting of the results of SERT, developed by the Standard Performance Evaluation Corporation,

Where,

$kW_{ENERGY STAR}$	= power draw in kilowatts of ENERGY STAR server
ES	= ENERGY STAR servers, numbered 1 to n
$kW_{ES, idle}$	= power draw in kilowatts of ENERGY STAR server at idle
$kW_{ES, full load}$	= power draw in kilowatts of ENERGY STAR server at full load
U_{ES}	= utilization of ENERGY STAR server
$kW_{baseline}$	= power draw of baseline servers
a	= percentage ENERGY STAR server is more efficient than baseline “standard” or “typical” unit
8760	= number of hours in a year (servers run 24/7 in a data center)

3.3.2 Calculating Savings for Server Virtualization

Server virtualization savings compare baseline energy use of a large set of single application servers that would have been purchased normally during a server upgrade, without virtualization to a smaller set of virtual host servers, as shown in Equation 5. See the Appendix for an example of how to use SERT data to determine savings from server virtualization.

$$\begin{aligned}
 kW_{baseline} &= \sum_1^n (kW_{sa, idle} + U_{sa} * (kW_{sa, full load} - kW_{sa, idle})) \\
 kW_{w virt} &= \sum_1^m (kW_{vh, idle} + U_{vh} * (kW_{vh, full load} - kW_{vh, idle})) \\
 Annual Energy Savings_{virt} &= (kW_{baseline} - kW_{w virt}) * 8760
 \end{aligned} \tag{5}$$

Where,

$kW_{baseline}$	= total power draw in kilowatts of all single-application servers without virtualization during server refresh
sa	= single application servers, numbered 1 to n
$kW_{sa, idle}$	= power draw in kilowatts of a single-application server at idle
$kW_{sa, full load}$	= power draw in kilowatts of a single-application server at full load
U_{sa}	= average utilization of a single-application server over the year
$kW_{w virt}$	= total power draw in kilowatts of all virtual hosts
vh	= virtual host servers, numbered 1 to m
$kW_{vh, idle}$	= power draw in kilowatts of a virtual host server at idle

$kW_{vh, full\ load}$ = power draw in kilowatts of a virtual host server at full load
 U_{vh} = average virtual host server utilization over the year

3.3.3 Calculating Savings for Using More Efficient Storage

Savings from upgrading to more efficient storage equipment (Section 1.4) can be calculated using Equations 6 and 7. Equation 6 uses efficiency metrics of the efficient and baseline unit to estimate savings. Equation 7, similar to Equation 4 (in Section 3.3.1), uses the percentage savings for an ENERGY STAR-certified data storage to estimate savings. To calculate savings from software management tools (Section 1.3), Equation 8 relies on measuring power draws before and after storage management tools are implemented. These power measurements pre- and post-storage management tool should be taken AFTER the efficient storage equipment is installed (if that was also part of the measure) to avoid double counting with savings estimated in Equations 6 and 7.

$$EM_{baseSE} = \left(\sum_{j=1}^m f_{baseSE(j)} EM_{baseSE(j)} \right) \quad (6)$$

$$EM_{EERE} = \left(\sum_{i=1}^n f_{EERE(i)} EM_{EERE(i)} \right)$$

$$Annual\ Energy\ Savings_{Efficient\ Storage} = kW_{EERE} * (EM_{baseSE} / EM_{EERE} - 1) * 8760$$

Where,

kW_{EERE} = power draw of new energy-efficient storage equipment
 EM_{EERE} = efficiency metric for energy-efficient storage equipment
 EM_{baseSE} = efficiency metric for baseline storage equipment
 $EM_{EERE(i)}$ = Watts per terabyte (TB) of energy-efficient storage device/array i (this value can come from product specifications for devices and/or arrays)
 $EM_{baseSE(j)}$ = Watts per TB of baseline device/array j (this value can come from product specifications for devices and/or arrays)
 $f_{EERE(i)}$ = fraction of total TB stored on energy-efficient device/array i
 j = baseline devices/arrays, numbered 1 to m
 $f_{baseSE(j)}$ = fraction of total TB stored on a baseline device/array j
 i = energy-efficient devices/arrays, numbered 1 to n
8760 = number of hours in a year as servers run 24/7 in a data center

$$\text{Annual Energy Savings}_{ES\text{ Storage}} = \left(\frac{1}{(1-b)} - 1\right) kW_{ES\text{ STOR}} * 8760 \quad (7)$$

Where,

- $kW_{ES\text{ STOR}}$ = power draw in kilowatts of ENERGY STAR storage
- b = percentage of ENERGY STAR storage more efficient than typical or standard storage
- 8760 = number of hours in a year (servers run 24/7 in a data center)

$$\text{Annual Energy Savings}_{DS\text{ Man}} = 8760 * (kW_{Pre\ DS\ Man} - kW_{Post\ DS\ Man}) \quad (8)$$

Where,

- $kW_{Pre\ DS\ Man}$ = total power draw in kW of data storage *before* data storage management tool measures implemented (or with tool turned off) and after efficient data storage equipment is installed, if that was part of the measure (the savings from the efficient storage equipment can be calculated using either Equation 7 or 8)
- $kW_{Post\ DS\ Man}$ = total power draw in kW of data storage *after* data storage management tools are implemented and after efficient data storage equipment is installed, if that was part of the measure (the savings from the efficient storage equipment can be calculated using either Equation 7 or 8)
- 8760 = number of hours in a year (servers run 24/7 in a data center)

3.4 Calculating Total Energy, Lifetime, and Peak Demand Savings

Total energy savings, which include additional cooling and power infrastructure savings, can be calculated by multiplying energy savings from an IT upgrade by the data center’s power usage effectiveness (PUE):⁵ the total data center energy use (e.g., lights; heating, ventilation, and air conditioning [HVAC]; UPS losses; IT) divided by the IT energy use. As a data center becomes more efficient, PUE moves toward 1.

⁵ This savings calculation assumes the data center’s cooling system will be controlled to maintain a given interior temperature set point. When reducing IT power use, less heat must be rejected from the data center. Thus, to maintain a constant temperature set point, cooling system power consumption will be reduced proportional to IT power use reductions. Energy losses at the UPS and transformers also will be reduced proportionally to IT energy-use reductions. Lighting loads may remain constant, but represent only a small fraction of a data center’s non-IT energy use. Therefore, PUE remains nearly constant with reduced IT power use. Consequently, total annual energy savings (IT equipment savings plus energy savings in cooling, UPS, and transformer systems) can be reasonably estimated by multiplying PUE by annual IT energy savings.

Equation 9 calculates total energy and demand savings.

$$Annual\ Energy\ Savings_{Total} = PUE * Annual\ Energy\ Savings_{IT} \quad (9)$$

Where,

$$PUE = \text{average PUE determined over the entire year}$$

Equation 10 calculates IT lifetime savings for server virtualization, efficient server upgrades, or efficient storage.

$$Lifetime\ Energy\ Savings_{IT} = Annual\ Energy\ Savings_{Total} * EUL \quad (10)$$

Where,

$$EUL = \text{expected useful life based on IT upgrade cycle of data center}$$

Equation 11 calculates seasonal peak demand savings, based on server and storage 24/7 operations.

$$Peak\ Demand\ Savings_{Winter} = PUE_{Winter} * Annual\ Energy\ Savings_{IT} / 8760 \quad (11)$$

$$Peak\ Demand\ Savings_{Summer} = PUE_{Summer} * Annual\ Energy\ Savings_{IT} / 8760$$

Where,

$$PUE_{Winter} = \text{average PUE over the winter peak demand period,}^6 \text{ which can be tracked over an entire year. } PUE_{Winter} \text{ may be smaller in winter due to free cooling)}^7$$

$$PUE_{Summer} = \text{average PUE over the summer peak demand period. } PUE_{Summer} \text{ may be much higher during the summer as free cooling options may not be available as often.}$$

⁶ Summer and winter peak demand periods usually vary by state. In Massachusetts, for example, the summer on-peak period is 1:00 pm–5:00 pm on non-holiday weekdays in June, July, and August; the winter on-peak period is 5:00 pm–7:00 pm on non-holiday weekdays in December and January.

⁷ Free cooling can include water-side and air-side economization, drastically reducing or eliminating the need for mechanical cooling loads. This is used more often in winter.

4 Measurement and Verification Plan

The following two major savings components must be examined for measures in a data center:

- The power draw of the efficient data center IT equipment.
- The efficiency standards for the measure and for the available IT equipment. (This information allows for development of savings estimates.)

On the surface, the requirements of a typical measurement and verification (M&V) plan for data center IT appear very similar to other energy efficiency measures (e.g., HVAC, lighting). However, given the limited data for EMs in IT spaces and the varied access to data center power draw data, an M&V plan must be flexible and accommodate a wide range of available data.

4.1 International Performance Measurement and Verification Protocol Option

International Performance Measurement and Verification Protocol Option A (Partially Measured Retrofit Isolation) offers the best and only approach for measuring data center IT measures, given its flexibility. Option A relies on field measurements of key performance parameters and estimates of key parameters not selected for field measurements. Data center IT measure energy-use estimates rely on estimates drawn from historical data, manufacturers' specifications, or engineering judgment. Other International Performance Measurement and Verification Protocol options do not provide this flexibility:

- Option B (Retrofit Isolation/Metered Equipment) requires measurement of all energy quantities to compute savings. It does not offer a viable approach because:
 - Data center IT equipment "burnout" savings calculations require using current codes or standards as baseline equipment. As this baseline equipment is not installed, it cannot be metered, and hence cannot fit into an Option B methodology (which requires metering).
 - Generally, a risk-averse manager will not allow metering of IT equipment in a data center. The manager may, however, be able to share data gathered from metering equipment installed at the UPS, PDUs, or in-rack smart power strips.
- Option C uses pre- and post-billing analysis. It also does not present a viable approach. As with Option B, the baseline used in the "burnout" savings calculation draws on current codes or standards, which are not represented in preimplementation electricity bills.

4.2 Verification Process

The verification process involves examining the core assumptions used in developing the savings estimate; this should include the following steps:

- Desk reviews of information pertaining to:
 - Energy-efficient IT equipment

- Baseline standard or typical IT equipment
- EMs
- Efficiency of ENERGY STAR server and storage
- Power draws
- EUL
- PUE
- On-site audits to confirm:
 - Installation of efficient IT equipment
 - Power draws of efficient IT equipment, based on spot readings of UPS, PDU, power strips, and server power
 - Utilization of servers
 - PUE.

4.3 Data Requirements/Collection Methods

Table 2 provides details on the types of data needed to verify key inputs for an energy-saving calculation of data center IT equipment, along with methods used for collecting the data:

Table 2. Verification of Key Inputs Into Equations

Key Inputs Into Equations	Verification of Data
Number of Energy-Efficient IT Equipment Units Installed	Reviewers should examine work orders and invoices, and conduct site visits to confirm purchases of efficient units and their installation.
“Baseline” unit	As savings estimates are limited to burnout savings estimates, reviewers should carefully examine how applicants determined baseline standards or typical IT equipment. Baseline IT equipment should: (1) provide the same performance as the energy-efficient it unit (i.e., the same storage capacity in data storage units, same chip set, memory, storage in servers, same computational capacity); and (2) be manufactured in the same year as the energy-efficient IT unit.
Efficiency Metrics for Servers EM_{EE} = efficiency metric for efficient server $EM_{baseline}$ = efficiency metric for baseline server	Reviewers of these metrics should examine SERT. Manufacturers of ENERGY STAR-certified servers must include SERT. Please see the Appendix for an example of how to interpret and use SERT “efficiency score” data to calculate savings for efficient servers and server virtualization.

Key Inputs Into Equations	Verification of Data
<p>Efficiency Metrics for Storage $EM_{EESB(j)}$ = Watts/TB of energy-efficient storage device/array j (this value can come from product specifications for devices and/or arrays) $EM_{baseSB(j)}$ = Watts/TB of baseline device/array j (this value can come from product specifications for devices and/or arrays)</p>	<p>As shown in Figure 3, the energy use required for data storage varies by technology and disk speed. Energy use can decrease by an order of magnitude with equipment upgrades, if an organization replaces faster spinning (15 K rpm) fiber channel hard disks to energy-efficient SSDs. In addition, the Storage Networking Industry Association Emerald Power Efficiency effort (http://snia.org/emerald/view) is gathering data on storage device efficiency. However, as of May 2014, very limited data are available on EMs for data storage equipment.</p>
<p>Percent savings for ENERGY STAR IT Equipment <i>a</i> = percentage ENERGY STAR server is more efficient than baseline “standard” or “typical” unit <i>b</i> = percentage ENERGY STAR storage is more efficient than baseline “standard” or “typical” unit</p>	<p>Reviewers should confirm the estimates for servers and data storage, as provided at the ENERGY STAR website www.energystar.gov/products.</p>
<p>Power Draws of Servers and Data Storage Based Off Measurements kW_{EE} = power draw of new efficient server equipment $kW_{ENERGY STAR}$ = power draw of ENERGY STAR server $kW_{w virt}$ = total power draw in kilowatts of all virtual hosts kW_{EESB} = power draw of new energy efficient storage equipment $kW_{ES STOR}$ = power draw in kilowatts of ENERGY STAR storage $kW_{Pre DS Man}$ = total power draw in kW of data storage before data storage management measures implemented $kW_{Post DS Man}$ = total power draw in kW of data storage after data storage management tools are implemented</p>	<p>Power draw measurements can be taken from: data center energy management systems, storage management tools, UPS, PDUs, power strip with metering capability, or even the actual server or data storage units directly. For example, ENERGY STAR-certified servers must “provide data on input power consumption (W), inlet air temperature (°C), and average utilization of all logical CPUs.” (EPA 2013). When examining measured power draw data, reviewers should look to: (1) review data averaged over a month to account for differences in server loads on weekends and nights or in differing storage levels used due to data storage resource management tools; and (2) account for PDU or UPS power losses when measuring IT equipment at the PDU or UPS. Although the data center manager probably will not allow confirmatory metering of power draws of IT equipment, options may be available to meter at electrical panels feeding specific data center loads.</p>
<p>Full Load and Idle Load Power Draws of Servers Based Off Manufacturer’s Data $kW_{sa, idle}$ = power draw in kilowatts of a single-application server at idle $kW_{sa, full load}$ = power draw in kilowatts of a single-application server at full load $kW_{vh, idle}$ = power draw in kilowatts of a virtual host server at idle $kW_{vh, full load}$ = power draw in kilowatts of a virtual host server at full load</p>	<p>Reviewers of these metrics should examine SERT. Manufacturers of ENERGY STAR-certified servers must include SERT data that will include full-load and idle load data. Please see the Appendix for an example of how to interpret and use SERT idle and full-load power draw data to calculate savings for efficient servers and server virtualization.</p>

Key Inputs Into Equations	Verification of Data
<p>Utilization of Servers U_{vh} = average virtual host server utilization over the year U_{ES} = utilization of ENERGY STAR server U_{sa} = average utilization of a single-application server over the year</p>	<p>For the installed virtual host server or installed ENERGY STAR server, utilization of servers should be derived from a data center's server performance software. Utilization of a baseline single application server may be estimated based on past implementations before server virtualization was implemented.</p>
<p>EUL</p>	<p>Reviewers should recognize that IT upgrades generally occur every 3 to 5 years, but can vary by organization. IT managers should base EULs on historical data from past hardware purchases and refresh cycles, as those EULs will be much more accurate for a given organization. When such information is not available, an IT manager might use 5 years for smaller data centers and 3 years for larger data centers, based on national average refresh cycles. The reviewer should also ask to compare a recommended EUL to:</p> <ul style="list-style-type: none"> • Length of data center service-level agreements. • Time period since last IT upgrade.

Key Inputs Into Equations	Verification of Data
<p>PUE PUE is the total data center energy use (e.g., lights, HVAC, UPS losses, IT) divided by the IT energy use.</p>	<p>Reviewers must recognize that the quality of the PUE estimate varies a great deal across data centers. Larger data centers have an in-house PUE estimates, tracked over time. According to a 2013 recent Uptime Institute industry survey of large data centers, PUE averages roughly 1.65 and 66% of large data centers measure PUE. Google’s large chiller-less data centers have achieved a PUE of 1.1 (Miller 2011). Many small data center spaces (e.g., server closets and rooms, localized data centers smaller than 1,000 ft²) may have never measured PUE.</p> <p>Evidence suggests that in some cases, poorly managed data center cooling, lack of variable-speed fans, load reduction that leads to reduced UPS efficiency, and other issues cause PUE to worsen (rise) after reducing IT load in data centers and not stay constant. Therefore, reviewers are encouraged to use PUE estimates after IT load is reduced.</p> <p>Numerous online models (developed by the Green Grid, APC, and others) exist, some are simple and some relatively complex, to estimate PUE (Karthi 2008). We recommend two guides for measuring PUE:</p> <ul style="list-style-type: none"> • A multiparty task force (composed of 7x24 Exchange, ASHRAE, The Green Grid, Silicon Valley Leadership Group, U.S. Department of Energy Save Energy Now Program, U.S. Environmental Protection Agency’s ENERGY STAR Program, United States Green Building Council, and Uptime Institute) developed a 12-page guidance for measuring and reporting PUE (EPA 2011). • The Green Grid developed an 80-page document titled “PUE: A Comprehensive Examination of the Metric” in 2012. This document supersedes previous white papers and consolidates all information that The Green Grid has developed and published relating to PUE (The Green Grid 2012).




Figure 3. Watts per terabyte for various data storage types
(Pflueger 2010)

5 Other Evaluation Issues

Two issues can complicate evaluation of data center IT equipment savings (EPA 2012):

- **Long lead times.** Data center deployments often take longer to complete than other types of energy efficiency engagements. All projects, whether related to IT equipment or to its supporting infrastructure, require careful planning and execution. These long lead times may complicate evaluating savings, as the project simply may not be completed by the time evaluation takes place. Evaluating savings before completion of an IT upgrade may result in significantly smaller savings than originally estimated.
- **Short production cycles.** Servers and many other types of IT equipment have annual production cycles due to frequent technological upgrades. These production cycles differ from product categories such as HVAC equipment, food service equipment, and residential appliances, which generally advance over multiyear timeframes. Technological advances can cause data center equipment to become antiquated with relative frequency. Thus, savings calculations for IT equipment should be based on a “burnout” scenario, comparing the efficiency measure to the baseline standard or typical equipment available at the time of installation. During the evaluation, reviewers must carefully examine the baseline equipment available at the time of the IT efficiency measure’s installation. If the baseline equipment does not represent the equipment available at the time of the efficiency measure’s installation, savings could be significantly underestimated, given the short production cycles and how quickly IT equipment efficiency increases over time.

References

- Arizona Public Service (2014). "Retrofit & New Construction Rebates." <http://www.aps.com/en/business/savemoney/rebates/Pages/custom-rebates.aspx>
- Clark, T.; Yoder, A. (2008). "Best Practices for Energy Efficient Storage Operations Version 1.0." Storage Networking Industry Association's Green Storage Initiative, www.snia.org/sites/default/files/GSI_Best_Practices_V1.0_FINAL.pdf.
- ComEd (2014). "Stay Cool While Saving Energy in Your Data Center." <https://www.comed.com/business-savings/commercial-industrial/Pages/data-centers.aspx>.
- DOE (2011). "Energy 101: Energy Efficient Data Centers." U.S. Department of Energy www.youtube.com/watch?v=xGSdf2uLtlo.
- Energy Trust of Oregon (2014b). "IT Power: Rack Up the Savings." <http://energytrust.org/commercial/incentives/equipment-upgrades-remodels/Software/it-power/>;
- Energy Trust of Oregon (2014a). "Incentives: Data Centers." http://energytrust.org/library/forms/be_pi0195d.pdf
- EPA (undated a). ENERGY STAR "Consolidation of Lightly Utilized Servers." U.S. Environmental Protection Agency, www.energystar.gov/index.cfm?c=power_mgt.datacenter_efficiency_consolidation
- EPA (undated b). "ENERGY STAR Server Specifications." U.S. Environmental Protection Agency www.energystar.gov/products
- EPA (2011). "Recommendations for Measuring and Reporting Overall Data Center Efficiency." U.S. Environmental Protection Agency, www.energystar.gov/ia/partners/prod_development/downloads/Data_Center_Metrics_Task_Force_Recommendations_V2.pdf?5311-22ed.
- EPA (2012). "ENERGY STAR Understanding and Designing Energy-Efficiency Programs for Data Centers." U.S. Environmental Protection Agency, www.energystar.gov/ia/products/power_mgt/ES_Data_Center_Utility_Guide.pdf?ff29-42fa.
- EPA (2013). "ENERGY STAR Program Requirements for Computer Servers." Section 5.1. U.S. Environmental Protection Agency, www.energystar.gov/products.
- IDC (2012b). "The Cost of Retaining Aging IT Infrastructure." Sponsored by HP. http://mjf.ie/wp-content/uploads/white-papers/IDC-White-Paper_the-cost-of-retaining-aging-IT-infrastructure.pdf
- IDC (2010). "Strategies for Server Refresh." Sponsored by Dell. <http://i.dell.com/sites/content/business/smb/sb360/en/Documents/server-refresh-strategies.pdf>
- IDC (2012a). "Analyst Connection: Server Refresh Cycles: The Costs of Extending Life Cycles." Sponsored by HP/Intel. Online. <http://resources.itworld.com/ccd/assets/31122/detail>
- Karthi (2011). Site 24x7. "Five Online Calculators Available to Calculate Energy Efficiency of a Data Center." <https://blogs.site24x7.com/2011/08/26/five-online-calculators-available-to-calculate-energy-efficiency-of-a-data-center/>.

Koomey, J. (2011). "Growth in Data Center Electricity Use in 2005 to 2010." www.analyticspress.com/datacenters.html.

Miller, R. (2011). "Chillerless Data Center is Google's Top Performer." www.datacenterknowledge.com/archives/2011/05/26/chiller-less-data-center-is-googles-top-performer/.

Netapp (2014). "Simple Provisioning of Your Unified Storage Environment." www.netapp.com/us/technology/unified-storage/simple-provisioning.aspx

NYSERDA (2014). "Data Center Incentives Through the Industrial Process and Efficiency Program." <https://www.nysersda.ny.gov/Energy-Efficiency-and-Renewable-Programs/Commercial-and-Industrial/Sectors/Data-Centers.aspx>.

Pflueger, J. (2010). "Understanding Data Center Energy Intensity." A Dell Technical White Paper.

Veeam (2011). "Veeam Launches V-Index To Measure Virtualization Penetration Rate." www.veeam.com/news/veeam-launches-v-index-to-measure-virtualization-penetration-rate.html.

Search Data Center (December 7, 2012). "The server lifecycle: Is it time for that aging server to go?" <http://searchdatacenter.techtarget.com/tip/The-server-lifecycle-Is-it-time-for-that-aging-server-to-go>.

Seattle City Light (2013) "2013 Energy Conservation Incentives: Medium and Large Commercial." Online. <http://www.seattle.gov/light/Conserve/Business/2013Incentives.pdf>

Silicon Valley Power (2014b). "2013–2014 Data Center Program Rebate Application." <http://siliconvalleypower.com/Modules/ShowDocument.aspx?documentid=5561>.

Silicon Valley Power (2014a). "2013–2014 Customer Directed Rebate Application." <http://siliconvalleypower.com/Modules/ShowDocument.aspx?documentid=5102>.

SPEC (2014). "Server Efficiency Rating Tool Design Document 1.0.2." Standard Performance Evaluation Corporation, www.spec.org/sert/.

Southern California Edison (2012). "Data Center Energy Efficiency Program." <http://www.willdan.com/energy/DCEEP.aspx>.

The Green Grid (2009). "Proxy Proposals for Measuring Data Center Efficiency." Online. www.thegreengrid.org/Global/Content/white-papers/Proxy-Proposals-for-Measuring-Data-Center-Efficiency

Yoder, A. (2012) "Technologies for Green Storage." Storage Networking Industry Association presentation," <http://snia.org/emerald>

The Green Grid (2012). "PUE: A Comprehensive Examination of the Metric." www.thegreengrid.org/en/Global/Content/white-papers/WP49-PUEAComprehensiveExaminationoftheMetric

Appendix: Hypothetical Calculations of Savings From an Efficient Server or Server Virtualization Using SERT Data

SERT was created by Standard Performance Evaluation Corporation to measure server energy efficiency by using a set of synthetic⁸ workloads or worklets as they are called, to test discrete system components such as processors, memory and storage, providing detailed power consumption data at different load levels. SERT, required by the ENERGY STAR program, allows a data center manager to compare server “efficiency score” and power draw across various workload types and at various loads between two or more candidate systems. As SERT becomes more widely adopted, it will allow for more standardized energy savings determinations. We provide some background on SERT and the manner in which savings could be determined. The worklets are grouped into four workload types shown in Table 3:

Table 3. SERT Workload Types, Worklet Names, and Characteristics
(SPEC 2014)

Workload Type	Worklet Names	Characteristics
CPU	Compress, CryptoAES, LU, SHA256, SOR, SORT, XML Validate	<ul style="list-style-type: none"> • The worklet requires consistent processor characteristics per simulated “user” regardless of number of processors, cores, enabled threads, etc. • At the 100% load level, the performance bottleneck is the processor subsystem. • The worklet’s performance should increase with more processor resources, including the number of processors, the number of cores, possibly the number of logical processors, increased frequency, larger available cache, lower latency, and faster interconnect between CPU sockets. • Readings at loads of 25%, 50%, 75%, and 100%.
Memory	Flood, Capacity	<ul style="list-style-type: none"> • The worklet contains consistent memory access characteristics per simulated “user” regardless of size and number of dynamic inline memory modules. • At the 100% load level, the performance bottleneck is the memory subsystem. • The worklet’s performance should measure a higher (better) performance score with improved memory characteristics (e.g., higher bandwidth, lower latency, total memory size). • The worklets as a group should reflect a combination of random and sequential reads and writes, and small and large memory accesses. • Readings at loads of 50% and 100% for Flood. • Readings at 4, 8, 16, 128, 256, 512, 1024 GB for Capacity.

⁸ Synthetic workloads or worklets are discrete operations of a specific type that are repeated over and over again. They represent theoretical capabilities of the system that are rarely exercised in such a repetitious and discrete manner in the real world. Application benchmarks more represent typical activities but don’t allow specific performance capabilities to be isolated.

Workload Type	Worklet Names	Characteristics
Storage I/O	Random, Sequential	<ul style="list-style-type: none"> The worklets reflect consistent input/output characteristics per simulated “user” regardless of system size and number of disks or the installed memory. The worklets consist of a combination of random and sequential accesses, reads and writes, and small and large inputs and outputs. At the 100% load level, the performance bottleneck is the storage subsystem. The worklets should score a higher (better) performance result for higher bandwidth and lower latency. The worklets are limited to testing individual internal storage devices only. RAID arrays and external storage devices are not supported. Readings at loads of 50% and 100%.
Hybrid	SSJ	<ul style="list-style-type: none"> The worklet reflects a combination of a wide variety of processor and memory-intensive tasks. At the 100% load level, the performance bottleneck is due to multiple subsystems. The combined worklets should measure a higher (better) performance score for improved processor and memory characteristics. Readings at loads of 12.5%, 25%, 37.5%, 50%, 62.5%, 75%, 87.5%, 100%.
Idle	Idle	No transactions occur during this measurement although the server is in a state in which it is capable of completing a transaction

Figure 4 and Figure 5 show SERT sample outputs on two hypothetical servers – Server A and Server B. (These data do not represent any particular server models.) The “efficiency score” column is the “normalized performance” divided by “Watts” across different loads. Idle power draws are indicated by the vertical blue line in the “Watts” column. Hash marks along each horizontal line represent data at the load levels specified in Table 3 for each worklet.

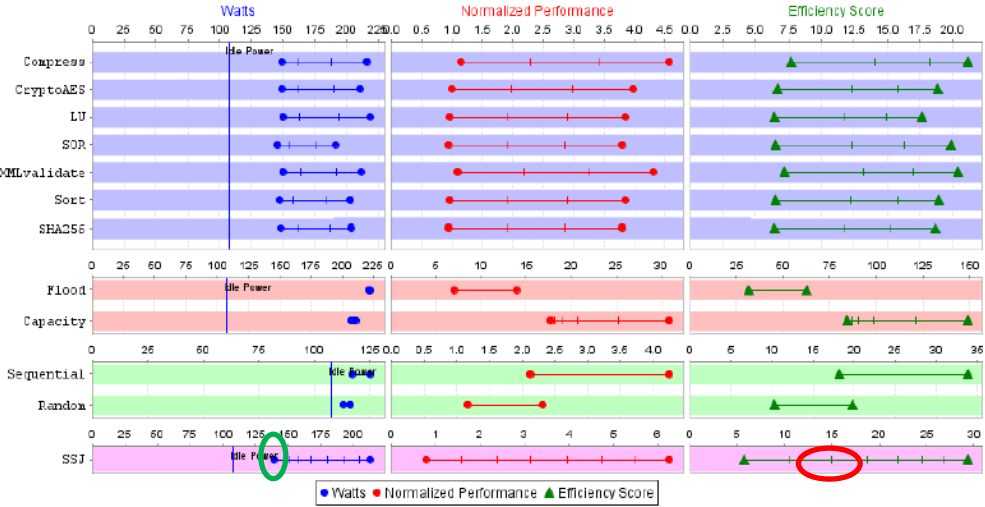


Figure 4. Sample SERT data for Server A

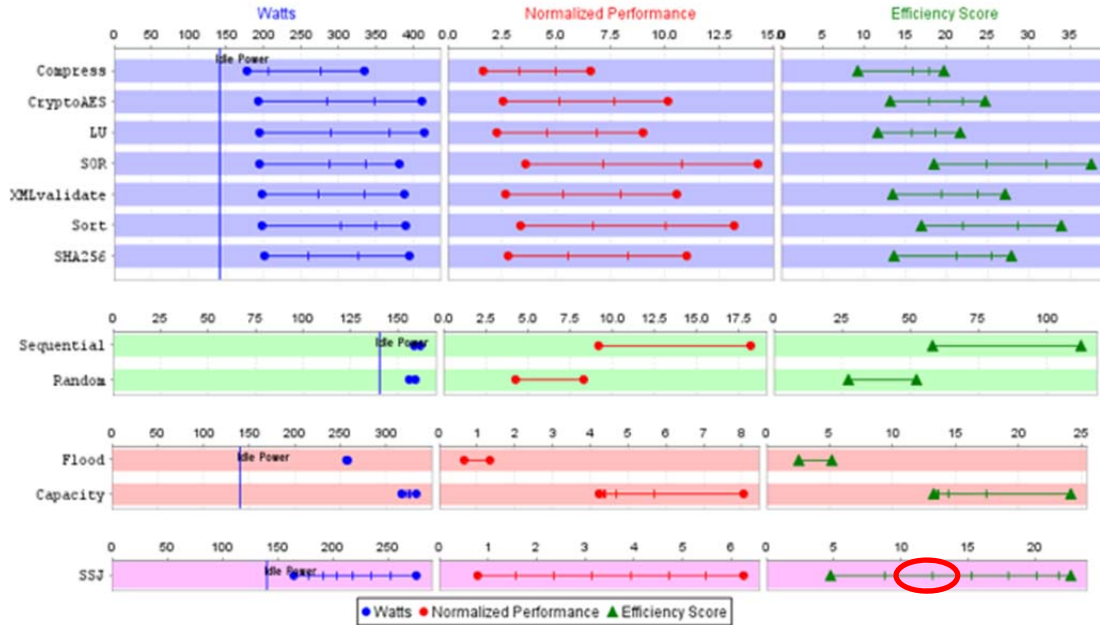


Figure 5. Sample SERT data for Server B

Savings From an Efficient Server

The example below demonstrates using SERT data to estimate savings from purchasing Server A instead of Server B. In order to calculate savings from purchasing the Server A, Equation 2 (from Section 3.3.1) would be used to take advantage of the availability SERT efficiency metrics:

$$\text{Annual Energy Savings}_{\text{Efficient Servers}} = kW_{EE} * (EM_{EE}/EM_{\text{baseline}} - 1) * 8760$$

Where,

kW_{EE} = power draw of new efficient Server A (kilowatts)

EM_{EE} = efficiency metric (SERT efficiency score) for efficient Server A

EM_{baseline} = efficiency metric (SERT efficiency score) for baseline Server B

8760 = number of hours in a year as servers run 24/7 in a data center

The steps to be taken include:

- **Determine which workload is most appropriate.** SERT has different workload types for CPU, memory, and storage-intensive loads. If you are unclear what type of load is present, you can use as an exemplar, the Hybrid SSJ worklet. This workload reflects a combination of synthetic loads to a wide variety of processor and memory-intensive tasks.
- **Determine the appropriate baseline model.** During the purchase of Server A, other alternative servers, such as Server B, were examined with similar CPU, memory and storage capacity. Server B was selected as the baseline model.

- **Measure the wattage of the efficient Server A.** Using the data center infrastructure management system data, the average power draw of the efficient Server A, kW_{EE} , is **160 Watts**.
- **Estimate the percentage load on the efficient server.** Using the data center's server performance software, utilization is estimated at **25%**.
- **Determine the efficiency scores at the appropriate worklet and load level.** In our example, the SERT efficiency scores for the SSJ worklet at 25% load were determined to be **15** for the efficient Server A (EM_{EE}) and **12.5** for the baseline Server B ($EM_{baseline}$). (See red circles on Figures 4 and 5.)

Using the values determined above, annual energy savings for an efficient server are estimated as:

$$Annual\ Energy\ Savings_{Efficient\ Servers} = 280\ kWh = 0.16\ kW * \left(\frac{15}{12.5} - 1\right) * 8760hr$$

Savings From Server Virtualization

The example below demonstrates using SERT data to estimate savings from server virtualization. In order to calculate savings from server virtualization, Equation 5 (in Section 3.3.2) would normally be used:

$$kW_{baseline} = \sum_1^n (kW_{sa, idle} + U_{sa} * (kW_{sa, full\ load} - kW_{sa, idle}))$$

$$kW_{w\ virt} = \sum_1^m (kW_{vh, idle} + U_{vh} * (kW_{vh, full\ load} - kW_{vh, idle}))$$

$$Annual\ Energy\ Savings_{virt} = (kW_{baseline} - kW_{w\ virt}) * 8760$$

Where,

$kW_{baseline}$ = total power draw in kilowatts of all single-application servers (assumed Server A) without virtualization during server refresh

sa = single application servers, numbered 1 to n

$kW_{sa, idle}$ = power draw in kilowatts of a single-application server (assumed Server A) at idle

$kW_{sa, full\ load}$ = power draw in kilowatts of a single-application server (assumed Server A) at full load

U_{sa} = average utilization of a single-application server over the year

$kW_{w\ virt}$ = total power draw in kilowatts of all virtual hosts

vh = virtual host servers, numbered 1 to m

$kW_{vh, idle}$ = power draw in kilowatts of a virtual host server at idle

$kW_{vh, full\ load}$ = power draw in kilowatts of a virtual host server at full load
 U_{vh} = average virtual host server utilization over the year

Server A, depicted in Figure 4, was assumed to represent the baseline single application server. In addition, it was assumed the baseline scenario is 20 single application server, $n = 20$, and the virtualization scenario uses two virtual host servers, $m = 2$. Because of the available metering data and SERT data of power draws at different loads, the equations relying on idle and full-load power draws to estimate savings are not necessary. Instead, the steps to be taken include:

- **Determine which workload is most appropriate.** Since it is unclear what type of load will be present, a Hybrid SSJ worklet, which reflects a combination of loads, is selected.
- **Estimate the wattage of the installed virtual hosts.** Using the DCIM system data, the average power draw of the two virtual host servers, $kW_{w\ virt}$, is **400 Watts**.
- **Estimate the wattage of the single application servers that would have been purchased.** For the baseline estimate, the alternative scenario is a conventional server upgrade where 20 single old application servers were replaced with 20 new single application servers (assumed to be Server A). The average utilization for the single application servers, U_{sa} , was assumed to be **12.5%**, based on IT manager estimates of the load on single application servers run in the past. As shown in the green circle in Figure 4, using the Hybrid SSJ worklet, the wattage at 12.5% load for Server A is **140 Watts**.

Using the values determined above, annual energy savings for a virtualization effort are estimated as:

$$kW_{baseline} = 2.8\ kW = 20 * 0.14\ kW$$

$$kW_{w\ virt} = 0.40\ kW$$

$$Annual\ Energy\ Savings_{virt} = 21,240\ kWh = (2.8\ kW - 0.40\ kW) * 8760$$