

Assessing and Improving Data Center Storage-Related Energy Efficiency

Technology Concepts and Business Considerations

Abstract

Taking a holistic view of data center operations provides an objective assessment of overall efficiency and value, as well as enables guidance for improving processes. This white paper examines the storage aspects of data center operations and demonstrates how many energy-related indications and assessments can be achieved with existing tools and information.

December 2008

Copyright © 2008 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number H6021

Table of Contents

Executive summary	4
Introduction	4
About the authors	5
Audience	5
Radial Graph	5
Potential indicators	7
Capacity Allocation and Usage category	8
Storage Power category	8
Data Classification category	8
Protection category	8
Data Center Performance category	8
Proposed initial indicators	9
Usage	10
Conclusions and moving forward	11
References	11
Acknowledgements	11

Executive summary

The IT industry is experiencing an increasing number of issues pertaining to data center energy consumption and efficiency. With the rising cost of electricity, energy usage has become a significant percentage of overall operating expense. In addition, many data centers are facing energy, cooling, and functional capacity limitations. Inefficiencies result not only in excessive energy requirements but also in other unnecessary costs and overheads.

Before data center managers consider solutions to these growing problems, they must first understand how effectively the current IT infrastructure is being utilized and how well its activities meet business requirements. This can be achieved through the use of clear and easy-to-use measurement tools.

A number of products have recently entered the market offering various measurement solutions. Many of these products oversimplify the problem, make it too complicated, and/or fail to take business requirements into account. Furthermore, most of the focus to date has been on server power dissipation and utilization. Storage, on the other hand, not only consumes a significant share of data center energy but has a capacity growth expectancy of nearly 60% per year¹. As a result, measurement methods dealing with the storage component of a data center are a must.

A number of Key Performance Indicators (KPIs) exist that can be used to assess current levels of data center storage-oriented operating power, efficiency, and performance, as well as provide guidance toward improvement. This white paper defines those KPIs and describes how to use them.

Introduction

Data center power and energy usage efficiency issues have been well documented². Several initiatives, practices, standards, and products aimed at continued improvement are in development or have been released by government agencies, independent technology groups, and device vendors³. These efforts are not only focused on individual device efficiency but, perhaps more importantly, on actual power and energy measurement at the device and system level. Simple and clear measurement tools, together with appropriate policies, arguably provide the basis for information processing power and energy usage efficacy. However, the definition and methods associated with such measurements have proven difficult to establish and, in many cases, require significant lead time for implementation into products.

While definition and standardization efforts continue, IT managers require some kind of energy usage assessment methods now. They expect their device and system management vendors to provide such capability in a framework usable across a heterogeneous data center environment inclusive of processing, storage, and interconnect infrastructures. Such methods must allow for indication and assessment of present operational characteristics against target goals using existing information, but with a vision toward tomorrow's measurement methods.

The Green Grid Productivity Indicator (a.k.a. Radial Graph) has already demonstrated great potential for assessing efficiencies in the data center⁴. IT managers have shown enthusiasm for its clear approach to both visualizing existing system indicators against goals and modeling the impact of proposed changes. It also enables clear results communication with management.

¹ *The Diverse and Expanding Digital Universe*, IDC white paper, sponsored by EMC, March 2008

² *Report to Congress on Server and Data Center Energy Efficiency*, U.S. Environmental Protection Agency, ENERGY STAR Program

³ More information can be found on the [Storage Networking Industry Association](#) and [ENERGY STAR](#) websites

⁴ *The Green Grid Productivity Indicator*, The Green Grid

This white paper intends to demonstrate how the Radial Graph method can be applied immediately by IT managers to plot indications and trends of storage energy usage within the data center. The methods described utilize existing and easily obtainable information while providing an extensible model allowing use of new (standardized) measurement indicators as they become available.

About the authors

Michael Dunbar is a Consulting Engineer in the Backup Platforms Group at EMC Corporation. He has been developing computer and storage systems for over 20 years. Most recently he led the effort to integrate disk drive spin down into the EMC® Disk Library. Michael has also architected several remote backup solutions with the EMC Disk Library for offsite data protection. Currently he is developing data de-duplication systems and integrating power monitoring and management into EMC storage systems. Prior to joining EMC, Michael was part of Data General Corporation's Central Processor Development organization focusing on the I/O side of Data General's computer systems. As part of CPD he architected and developed numerous I/O systems for microcomputers based on proprietary, Motorola 88K, and Intel CPUs. This included the first integration of a SCSI controller into a DG computer system. Outside of his professional career, Michael became a certified ENERGY STAR homebuilder and acted as a general contractor to build his ENERGY STAR home, which features a geothermal heating system. He holds BSEE and BSCS degrees from the University of Connecticut.

James W. Espy is a Senior Consulting Engineer and staff member of the Hardware Platform Architecture Group at EMC Corporation. James has a long history of innovation in storage and systems design, and has been involved in many strategic projects during his present tenure at EMC and previously at CLARiiON® when it was a division of Data General. He architected the most recent EMC Centera® Content Addressable Storage (CAS) cluster platform and developed the electromechanical and interconnect platform for the world's first full back-end and front-end Fibre Channel storage array while at CLARiiON. Current interests include utilization of flash technology and energy efficiency in EMC platforms. Prior to EMC, Jim held staff positions at Cisco Systems, Kendall Square Research, and Data General where he began his career designing high-performance disk drives. James holds a MSEE degree from the University of Illinois – Urbana and a BSEE from the University of Wisconsin – Madison, and holds several patents related to storage systems.

Audience

The intended audience of this paper includes IT professionals and facilities personnel responsible for data center operations, as well as CIOs and those in other high-level management roles.

Radial Graph

The Radial Graph provides visualization of multiple performance indicators against desired targets. It can be used to show existing state, change progress, and the impact of “what if” scenarios. Its graphical nature allows for easy and clear interpretation. The following figures demonstrate its construction and usage.




Figure 1. Pentagonal Radial Graph and axis points (KPIs)

Each axis point of the example pentagon in Figure 1 represents a performance indicator defined to have a value range of 0%-100%. Performance indicators for a single graph may or may not be related, but selecting indicators with interdependencies and relationships can be very effective for holistic impact analysis.

In this paper, we will refer to an axis point as a Key Performance Indicator (KPI) and define three associated values (others are possible):

- Theoretical Maximum – 100%
- Target – A value less than 100% considered a desirable/achievable goal
- Actual – A value provided by actual system data

Note that the Radial Graph does not have to be pentagonal. Any number of KPIs could be plotted up to a level allowing for reasonable viewing. One could also have subgraphs that may share a common KPI, for example, one indicating overall system performance, as another way to show the impact and relationship among many indicators. For communication purposes, it's generally best to keep the number of indicators at a minimum (typically 4-6).




Figure 2. Hexagonal Radial Graph with defined axis, target, and actual points

In the example graph shown in Figure 2, the dark red-colored polygon represents a target boundary with each line segment connecting two adjacent KPI target points. The lighter green-colored polygon similarly represents actual measured value points. Improvements are seen as the actual points and lines move outward and closer to the target points and lines. Negative impacts are seen as the actual points and lines move inward and farther from the targets. A particular system change may result in some actual points becoming worse while others improve. Regardless, the tool’s graphical nature allows for the “immediate” and clear visual indication of changes against multiple KPIs.

The Radial Graph not only makes for a good performance, change impact, and trends assessment tool, it is also a good communication method to management and related lines of business, aiding in the determination of capital equipment and expense budget planning. Target KPIs can be interpreted as business value parameters.

Potential indicators

In this paper we propose using the Radial Graph to plot data center-wide, storage-oriented KPIs instead of those focused on individual devices. However, we believe such a set of indicators should include at least one representing performance as seen by data center customers, for example, Service Level Agreements (SLAs).

A good storage KPI is defined as having the following attributes:

- Clearly shows, is related to, or otherwise impacts key business objectives (for example, opex, SLAs)
- Can be designated as a percentage (dimensionless) with an ideal of 100%
- Can be easily derived from data accessible and available from all storage vendor devices
- Supports the identification of key improvement options
- Addresses the 80th percentile (or more) of data center use models

The third bullet represents one of the immediate challenges associated with using the Radial Graph tool, that being the availability and easy accessibility of actual data. For instance, while it would be good to have access to storage device power measurements, such capabilities are not, at the time of publication, widely

available. Therefore, other less direct power and energy assessment methods and related KPIs are necessary to provide the indicators data center managers need today. However, direct power measurement-oriented KPIs can be added in the future as such data becomes available.

In the following sections we categorize potential storage-oriented indicators and then follow up in a later section with those specific KPIs that presently have a high likelihood of being implemented. Note that these categories are in no particular order of importance.

Capacity Allocation and Usage category

Various “storage utilizations,” or selected ratios of actual storage capacity usage versus raw and allocated capacities, provide a good indication of efficiency. A large amount of unused capacity implies power wasted while operating storage devices in idle or low duty cycle modes. While this is not exact and may not consider certain power saving techniques (for example, spin-down), it is a reasonable indicator of storage system efficiency.

Data to calculate related KPIs is both available and easily accessible from most storage system management tools. Hence KPIs under this category have a high potential for immediate use. However, while such data is readily available for file and object storage systems, it is not as easily obtainable for block storage systems.

Storage Power category

This category represents a more direct approach based on power measurements. Simple notions of Watts/GB do not fit the need for a dimensionless indicator; a different denominator is required. Note that power-per-performance indicators are purposely left out of this discussion as they are difficult to normalize across differing vendor systems.

At present, data to calculate related KPIs is less available. Therefore, most KPIs under this category have a low potential for immediate use. However, there are efforts under way in the Storage Networking Industry Association (SNIA) to create a normalized value for storage efficiency. Once this activity is complete, storage power-related KPIs may become more feasible.

Data Classification category

Policy-based data classification and placement across storage tiers result in effective storage resource usage and minimized power. Indicators may include those appraising the appropriate location of data across multiple tiers or whether data should be deleted or archived (and active storage capacity reclaimed).

Availability of indicator data may vary from system to system and can be difficult to summarize. However, there are ways to ascertain data activity, such as the size of incremental backups or levels of data inactivity.

Protection category

This category assesses the power needed to meet low data loss and high data availability customer requirements. For instance, what percentage of system power is dedicated to redundant hardware in order to meet a particular level of 9s availability? How much data is purposely duplicated for protection reasons? Such information can be used in power efficiency assessments but has to be viewed as the price to pay for the realities of system reliability and the methods necessary to meet business and mission-critical goals.

While valuable in providing an indication of power required for protection, such parameters are presently difficult to measure or derive.

Data Center Performance category

There are many who believe performance is binary; it either meets SLAs, or it does not. Most data centers have multiple customers (different departments or business units). Meeting all customer SLAs all the time

is no doubt difficult. Regardless, an SLA-oriented KPI on the same graph as energy-oriented KPIs enables the viewing of power saving method impacts on system performance as seen by customers. Note that it is not uncommon for an SLA to be overserved to meet overall targets. Such behavior can result in energy usage inefficiencies.

Data on how well customer-level performance parameters are being met are available to the data center manager.

Proposed initial indicators

From the previous discussion it is clear that several different KPIs are possible. However, information for many of them is not readily available, making their usage infeasible at present. We next present several KPIs that can be used today to provide a holistic view of data center operational efficiency.

- **Capacity Allocation and Usage** - This KPI indicates the percentage of raw capacity used for data storage. Two KPIs are defined, one for file or object-based storage systems and one for block-based storage systems. Actual usage capacity information is readily available via file- and object-based storage system management interfaces and hence its KPI is easily calculated as actual usage capacity in GB divided by raw capacity in GB (expressed as a percentage). However, block storage systems may only provide access to raw and allocated capacity information. Actual usage data would have to be collected via attached host servers. Efforts to collect such information may be acceptable in smaller configurations, but could prove tedious in larger installations, even if only sampled. Therefore, the second KPI is suggested to be block allocated capacity in GB divided by raw capacity in GB (again expressed as a percentage). This is a rough estimate but is reasonable in that it provides indications of state and trends as well as being better than nothing in most cases. In general, these two indicators represent a clear example of ones that should not have a 100% target. At the other end of the spectrum, anything significantly below 50% is not a good target either. For many situations, targets in the area of 50%-80% are probably appropriate, the key being to first find out where one stands. For organizations with rapidly growing amounts of information, numbers in the 50%-70% range may be acceptable, while more static growth rates could live with less headroom, so 80%-90% capacity utilizations would be reasonable.
- **Average Spindle Size** - This KPI represents a percentage of implemented spindle size against today's current largest capacity disk, for instance, 1 TB. It makes sense in two ways. First, spindle size is often indicative of spindle age with newer disks being of larger capacity. Second, larger capacity disks tend to be more efficient than smaller ones simply because of their higher storage densities; they require less power/GB. One argument against this indicator is that it makes no sense to use large capacity disks for all applications, as the larger capacity disks also tend to be less performant. For this reason the target value for such an indicator should be less than 100%.
- **Optimizations** - This KPI represents the utilization of particular efficiency features. Many storage vendors offer a variety of features providing value to the business as well as energy efficiency. Since quantifying these features is difficult, we provide a framework that can be customized for individual usage. Conceptually, each of these features provides some efficiency, so each one is assigned a point value. The total number of potential points represents 100%. Table 1 defines this framework with example entries.

Table 1. Features and point values

Feature	Points	Converted to %
Thin provisioning	1	7
Disk spin-down (tier 3+)	1	7
Solid state storage (tier 0)	1	7
File and/or e-mail tiering	2	13
Snapshots	2	13
Storage virtualization	2	13
De-duplication	2	13
Compression	1	7
Archiving	2	13
Other	1	7

- **Right-Sized Tiering** - This KPI is the percentage of storage located below tier 1 (tier 0, the highest priority tier, is defined as solid state storage and tier 1 is high performance disk drives such as FC, 15k rpm). Although many organizations have huge amounts of data, in practice only a small percentage of that data is appropriately termed mission-critical, or even business-critical. Much, if not the vast majority, of an organization's data is most likely not mission- or business-critical and as such does not warrant the absolute fastest, most reliable, and most feature-rich storage. So, a large percentage of storage should, in fact, be located in lower tiers. The actual KPI is the amount of storage, in GB, classified as tier 2 and lower priority (tiers 2, 3, 4, etc.) divided by the total amount of all storage tiers (in GB).
- **Storage Consolidation** - This KPI is the percentage of data consolidated on external storage. Storage inside individual servers is typically accessible only to the server within which it resides, and hence is often underutilized. Such storage is difficult to reallocate and manage. Networked storage, be it NAS, FC SAN, or iSCSI, is typically much more flexible, easily managed, and highly available.
- **SLA Compliance** – This KPI indicates the percentage of time data center customers' SLAs are met. To make use of this KPI, the data center operations and IT departments must have an SLA tracking process in place. For each SLA, one must know how many minutes of each day the SLA was met. The KPI is the ratio of the aggregate number of minutes in a day that all the SLAs were met divided by the number of SLA minutes in a day (the number of SLAs multiplied by 1440, or the number of minutes in one day). As an example, suppose there were six SLAs defined for data availability, one for each of six departments and all but one of the SLAs were met. The KPI for that day would be $(5*1440 + 1*1430)/(6*1440)$, or 99.88%.

Usage

Once a set of KPIs is defined and agreed upon, the next step is to collect relevant information to construct a Radial Graph. Existing management tools can be used to access raw data from which KPIs are derived. Spreadsheets programs allow one to easily construct the Radial Graph.

The next step is to define target values for each KPI. These would be a new set of data shown as another series on the chart. With that activity completed, one now has an initial and fairly clear picture of overall data center efficiency, including a vision of where it could go.

As target values are defined, note that most target KPIs probably should not be set at 100%. For example, it's unrealistic to have 100% storage utilization as that would allow absolutely no room for growth. Also, the defined targets are likely to be highly specific to a particular data center, taking into consideration factors such as rate of growth, procurement cycles, and tolerance for risk.

Collecting KPI information and updating the Radial Graph should be a periodic activity, perhaps quarterly or monthly. Having a flow of assessment data allows one to observe trends and changes in data center

operations over time. Such trends can be compared with changes in data center equipment, policies, process, and overall data center operations. For example, utilizing virtualization can reduce power requirements but may create or reveal a performance bottleneck. On the other hand, one may find that performance is still meeting SLAs while overall operational costs have improved. Being able to correlate cause and effect is a powerful tool that will help drive good business decisions.

These techniques can be used as an effective means to justify changes and/or expenditures in the data center and as a communications tool with business units and management. Additionally, they may also serve as an excellent basis for comparison with other data centers, not because they should be the same, but because they invite examination of underlying assumptions, exposing opportunities for improvement.

Perhaps one of the more interesting applications of this tool is the visualization of experimental or modeled test results. For instance, if one were to migrate a certain class of data to a different tier, what would be the impact on system power as well as performance levels seen by customers? Or, one may discover an inordinate amount of tier 3 class data residing on high performance, highly available storage. This data could probably be moved to a much lower-cost storage tier. The interdependency aspect of the graph will clearly show the composite results of such a move. However, the method itself requires validation with actual data before modeled results can be accepted.

Conclusions and moving forward

The use of KPIs and Radial Graphs provides a means to observe several key data center attributes. These graphs provide insight into how changes in one area may affect others. They can be used to help make decisions regarding changes in equipment and policies as well as provide tools to track, chart, and direct improvements. And, most importantly, they can be used to graphically and intuitively show the impact of changes. Of no less importance, minimal investment is required. Existing and familiar tools are used to extract information.

One can envision a time when all KPI data is automatically gathered and real-time Radial Graphs displayed. Such a future management tool would allow the user to model potential changes and observe their effects on all KPIs.

References

- *The Diverse and Expanding Digital Universe*, IDC white paper, sponsored by EMC
- *Report to Congress on Server and Data Center Energy Efficiency*, U.S. Environmental Protection Agency, ENERGY STAR Program
- [Storage Networking Industry Association website](#)
- [ENERGY STAR website](#)
- *The Green Grid Productivity Indicator*, The Green Grid

Acknowledgements

The authors wish to thank Kathrin Winkler, Philip Tsihlis, Chris Mancini, Gregory Robidoux, and Ken Cowan, all of EMC Corporation, for their contributions to this paper.