

StorageIO Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification

By Greg Schulz

Founder and Senior Analyst, the StorageIO Group
Author “The Green and Virtual Data Center” (CRC)



July 3, 2009

This Industry Trends and Perspectives Paper is Compliments of:

The Green and Virtual Data Center

www.thegreenandvirtualdatacenter.com

Introduction

To say that there is no such thing as a data recession would be an understatement. Demand to store more data for longer periods of time is driving the need for more data storage which is second to servers in energy consumption and subsequent cooling demands. Building¹ on previous as well as recent works^{2,3,4} covering and addressing IT data infrastructure efficiency, optimization along with power, cooling, floor-space and other related topics, including closing the Green Gap⁵, StorageIO⁶ has prepared this Industry Trends and Perspective White Paper. This white paper provides comment and feedback to the U.S. EPA as an interested stake holder in their efforts to define an Energy Star for enterprise storage, or what might be more applicably referred to as Energy Star for data center storage.

Storage and related technologies also aid in sustaining business growth while building on infrastructure resource management functions, including data protection, business continuance and disaster recovery (BC/DR), storage allocation, data movement, and migration, along with server, storage, and networking virtualization topics. Although this chapter focuses on external direct attached and networked storage, either networked attached storage or a storage area network, the principles, techniques, and technologies also apply to internal dedicated storage. The importance of this chapter is to understand the need to support and store more data using various techniques and technologies to enable more cost effective and environmental as well as energy friendly data growth.

Background, Issues and Challenges

After facilities⁷ cooling for all IT equipment and server energy usage, external data storage has the next largest impact on power, cooling, floor space, and environmental (PCFE) considerations in most environments. In addition to being one of the large users of electrical power and floor space, with corresponding environmental impact, the amount of data being stored and the size of its the data footprint continue to expand.

Though more data can be stored in the same or smaller physical footprint than in the past, thus requiring less power and cooling, data growth rates necessary to sustain business growth, enhanced IT service delivery, and new applications are placing continued demands on available PCFE resources.

A key driver for the increase in demand for data storage is that more data is being generated and stored for longer periods of time as well as more copies of data in multiple

¹ Excerpt from “The Green and Virtual Data Center” (CRC) are included curtesy of CRC for perspective and background

² StorageIO Analysis on EPA Report to Congress (104-431) at <http://storageio.com/portfolio.html>

³ See StorageIO Blog: Shifting from enegy avoidance to energy efficiency: <http://storageio.com/blog/?p=562>

⁴ See StorageIO Blog: Storage Efficiency and Optimizaiton – Balnacing Time and Space: <http://storageio.com/blog/?p=510>

⁵ See and listen to closing the Green Gap Internet radio interview: <http://storageio.com/blog/?p=519>

⁶ See Eco-Tech Warrior: <http://storageio.com/blog/?p=515>

⁷ “The Green and Virtual Data Center” (CRC/Taylor and Francis) ©2009 at www.thegreenandvirtualdatacenter.com

locations. This trend toward increasing data storage will likely not slow anytime soon for organizations of all sizes.

The data footprint is the total data storage needed to support application and information needs and provides a perspective of where and how storage is used in a typical data center. Your data footprint may, in fact, be larger than the actual amount of data storage you have, or you may have more aggregated data storage capacity than actual data. A general approach to determine your data footprint is simply to add up all of your online, near-line, and offline data storage (disk and tape) capacity. For example, consider all the data being stored at home on personal computers and laptops, PDAs, digital cameras and video recorders, TiVo sets and DVRs, USB fixed and removable disk drive's among other media that support various data and information needs.

Suppose that a business has 20 TB of data storage space that is allocated and being used for databases, email, home directories, shared documents, engineering documents, financial, and other data in different formats, both structured and unstructured. For these 20 TB of data, the storage space is probably not 100% used; database tables may be sparsely allocated, and there is likely duplicate data in email and shared document folders. However, to keep the example straightforward, assume that of the 20 TB, two complete copies are required for BC/DR purposes, and 10 TB are duplicated to three different areas on a regular basis for application testing, training, and business analysis and reporting.

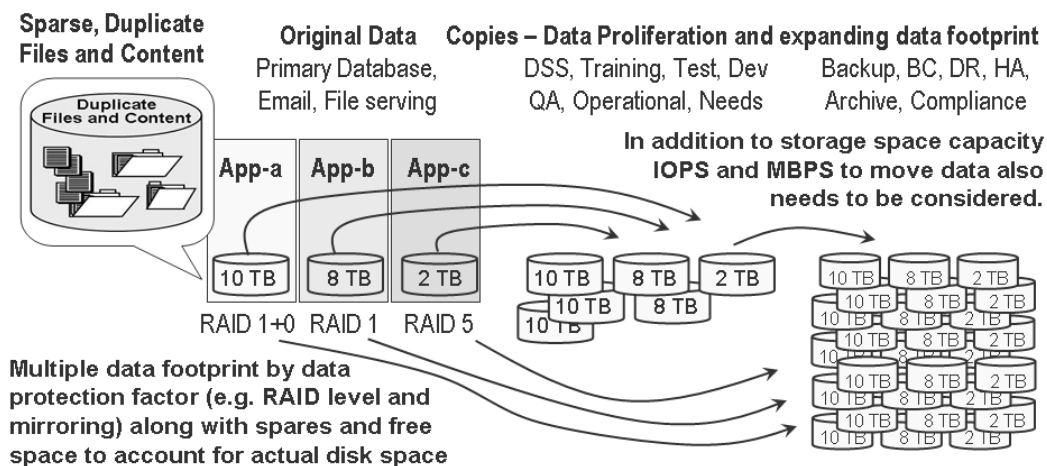


Figure 1 - Expanding Data Footprint due to Data Proliferation and Copies Being Retained
 Source and courtesy: "The Green and Virtual Data Center" (CRC)

The overall data footprint (figure 1) is the total amount of data, including all copies plus the additional storage required to support that data, such as extra disks for redundant array of independent disks (RAID) protection or remote mirroring. In this overly simplified example, the data footprint and subsequent storage requirement amount to several times the 20 TB of data. And the larger the data footprint, the more data storage

capacity and performance bandwidth are needed and that have to be powered, cooled, and housed in a rack or cabinet on a floor somewhere.

It is debatable how much energy in a typical data center is actually consumed by storage (internal to servers and external) as well as how much data is active or inactive. The major power draws for common storage systems are usually spinning hard disk drives (HDDs) and their enclosures, which account for, on average, 66–75%; controllers and related I/O connectivity components generally account for most of the balance of electrical power consumption. Consequently, data storage is an important area for energy optimization and efficiency improvements.

Putting it together

The following is the EPA Energy Star® Enterprise Storage Draft Specification Framework posted June 4, 2009 for review and initial comments. *StorageIO comments* are inserted in-line in *italics* followed at the end of the document with closing and general comments.

ENERGY STAR® Enterprise Storage Draft Specification Framework - June 4, 2009
Please send comments to Storage@energystar.gov no later than Friday, July 3, 2009

Overview

This document describes the key building blocks that form the basis for every ENERGY STAR specification; these items are intended to provide the framework around which the EPA can develop an effective energy efficiency program for Enterprise Storage. The principal objectives for this ENERGY STAR specification are threefold:

- (1) to encourage widespread adoption of appropriate hardware and software strategies to improve energy efficiency in enterprise storage systems,

StorageIO Commentary:

StorageIO is concerned with the usage and implication of “enterprise storage systems” as being too exclusive to that of only high-end, expensive price-band based storage solutions excluding mid-market and even entry level storage solutions and systems commonly deployed in IT data centers of all size. There needs to be some boundary on what is in and what is out of the scope of the intended Energy Star for Storage specification.

Such a specification to be effective for IT organization, the actual consumer and purchaser for this target set storage solutions needs to be applicable to their usage. Thus, encompassing various categories and price bands, different tiers and media types as well as multiple architectures across several market segments. The common denominator theme and naming aligned with other EPA initiatives including report to congress is that of the data center, exclusive of the term enterprise.

By simply using the term data center, or, data center storage, this would be encompassing enough to address storage that is being deployed by various market segments from Small medium business (SMB) to small medium enterprise (SME) to enterprise and even ultra large scale environments regardless of if they are an IT, managed hosting or service provider while excluding consumer or prosumer based storage solutions not commonly found in principal data center environments.

- (2) to provide purchasers with the means to identify the most energy efficient enterprise storage solutions for their specific end-use application, and

StorageIO Comment:

A key notion here will be to define what efficient means in terms of when the storage is being used for actual work, or, for storing data, or in-active and off-line. Some storage is designed to simply hold data while other storage is designed to meet performance and I/O acidity needs. Some storage is used for both.

- (3) to provide tools and information to designers and managers looking to improve the efficiency of data center operations.

StorageIO Comment:

A key notion here will be to define what efficient means in terms of when the storage is being used for actual work, or, for storing data, or in-active and off-line. Some storage is designed to simply hold data while other storage is designed to meet performance and I/O acidity needs. Some storage is used for both.

The purpose of each building block is provided under the subheadings below, along with EPA's preliminary thoughts on how each may ultimately be incorporated into the Version 1.0 Enterprise Storage specification. At the end of each section are a series of questions aimed at generating discussion about the proposed approach. Please note that this document is not intended to be a comprehensive review of the ENERGY STAR perspective on enterprise storage, rather it serves as a starting point for EPA's specification development efforts.

Stakeholders are encouraged to provide feedback on the specific concepts and definitions presented in this document, and are also welcome to submit comments of a more general nature. Communication between EPA and industry stakeholders is critical to the success of the ENERGY STAR program, especially in this early stage of the specification development process. Any and all creative suggestions for improvements to the basic ENERGY STAR approach outlined in this document will be considered for inclusion in subsequent draft and final specifications. ENERGY STAR representatives are available for additional technical discussions with interested parties at any time during the

specification development process. Please contact Steve Pantano, ICF International, at spantano@icfi.com to arrange a meeting.

Building Block #1: Definitions

a. **Purpose:** Establish a set of definitions to explicitly describe which products are covered by the specification and which are not, and to clearly differentiate between Enterprise Storage and other ENERGY STAR product categories. Provide definitions for operational modes, key components and sub-classes of products, etc. Note that a product may not be qualified as ENERGY STAR under more than one specification – manufacturers must select the product category that best describes the product they wish to qualify.

StorageIO Comment:

StorageIO would like to see some clarification in the case where multiple components may be combined to create a storage system or storage solution that is purchased by a data center. For example, assuming a storage system that combines software with off the shelf commercial servers covered under the Energy Star for server specification along with host bus adapters (HBAs) and network interface cards (NICs) be classified? Would the HBAs and NICs need to be covered under the server specification or a future server specification, or, would those components be covered under a storage specification?

Would the storage solution as a whole be covered under an Energy Star® for data center storage specification and simply include or allow the vendor (manufacture or var) to mention that Energy Star® server or other components are included? The assumption is that there is verbiage in existing Energy Star® material to prevent a misrepresentation of a Energy Star® for data center storage candidate solution as being Energy Star® if it only uses an Energy Star® compliant server, yet, the entire solution does not meet or fall under the Energy Star® for data center storage umbrella?

Additional clarification for what a category is would also be applicable given that some storage solutions can be deployed in a data center via customer or vendor configurable settings for multiple usage cases, for example on-line vs. near-line, high-performance vs. high-capacity, active vs. archive, database vs. email or file sharing among many other permutations. Consequently some clarification is needed or will need to be considered as to how different storage systems or solutions can be classified and covered in the proposed specification.

b. Initial Approach: EPA prefers to make use of existing definitions that are generally accepted by industry. In cases where industry accepted definitions are not available or appropriate, EPA will work with stakeholders to develop acceptable definitions.

c. Preliminary List of Definitions^{8 9}:

StorageIO Comment:

StorageIO provides the following as is in addition to available industry trade group, vendor or other applicable standard, or defecto standard terms, definitions, abbreviations, acronyms along with links and URLs.

	<i>Binary Number of Bytes</i>	<i>Decimal Number of Bytes</i>	<i>Abbreviation</i>
<i>Kilo</i>	<i>1,024</i>	<i>1,000</i>	<i>K, ki, kibi</i>
<i>Mega</i>	<i>1,048,576</i>	<i>1,000,000</i>	<i>M, Mi, bebi</i>
<i>Giga</i>	<i>1,073,741,824</i>	<i>1,000,000,000</i>	<i>G, Gi, gibi</i>
<i>Tera</i>	<i>1,099,511,627,776</i>	<i>1,000,000,000,000</i>	<i>T, Ti, tebi</i>
<i>Peta</i>	<i>1,125,899,906,842,620</i>	<i>1,000,000,000,000,000</i>	<i>P, Pi, pebi</i>
<i>Exa</i>	<i>1,152,921,504,606,850,000</i>	<i>1,000,000,000,000,000,000</i>	<i>E, Ei, exbi</i>
<i>Zetta</i>	<i>1,180,591,620,717,410,000,000</i>	<i>1,000,000,000,000,000,000,000</i>	<i>Z, Zi, zebi</i>
<i>Yotta</i>	<i>1,208,925,819,614,630,000,000,000</i>	<i>1,000,000,000,000,000,000,000,000</i>	<i>Y, Yi, yobi</i>

Source and Courtesy: "The Green and Virtual Data Center" (CRC)

www.80plus.org

www.afcom.com

www.aiim.com

www.ashare.org

www.cdproject.net

www.climatesaverscomputing.com

www.cmg.org

www.cmp.com

www.computerworld.com

www.dsireusa.org

www.drj.com

www.echannelline.com

www.eere.energy.gov

www.eia.doe.gov/fuelelectric.html

www.energyshop.com

Energy efficient power supply trade group

Data center industry user group

Archiving and records management trade group

HVAC Engineers Association

Carbon Disclosure Project

Green computing industry trade group

Computer Measurement Group and capacity planners

CMP Media Group

Publication covering servers, storage and networking

Database of State Incentives and Renewable Energy

Disaster Recovery Journal

Channel focused publication venue


U.S. Department of Energy (DoE) website

Portal for electrical power generation, use and costs

Portal for energy pricing and options

⁸ EPA is interested in information from stakeholders on industry standard definitions for those terms marked as TBD. Stakeholders are also encouraged to provide additional suggestions or clarifications regarding the proposed definitions.

⁹ The proposed definitions were compiled from a variety of reference sources, including the ENERGY STAR specification for Computer Servers and the Storage Networking Industry Association (SNIA) Web site.



Information Technology Industry
StorageIO Group Analysts and Consultancy
Industry Trends and Perspective White Paper

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

www.energystar.gov	United States EPA Energy Star® website
www.enterprisestorageforum.com	Publication coverage data and storage management
www.epa.gov/stateply/	United States EPA Climate Leaders' Initiative
www.epeat.net	Site for comparing desktop and related products
www.fcoe.com	Web site pertaining to Fibre Channel over Ethernet
www.fibrechannel.org	Fibre Channel Industry Association website
www.fueleconomy.gov	US Government site for energy efficiency
www.greendatastorage.com	Site pertaining to green data storage and related topics
www.greenpeace.org	Greenpeace website
www.greenwashing.net	Information about green washing
www.ieee.org	Institute of Electrical and Electronics Engineers
www.ietf.org	Internet Engineering Task Force
www.iso.org	International Standards Organizations
www.jupitermedia.com	Jupiter Media Group including Internet News
www.naspa.org	System administrator's user group
www.pcisig.com	Peripheral Component Interconnect (PCI) trade group
www.scsita.org	SCSI trade association
www.snia.org	Storage Networking Industry Association
www.spec.org	Server and storage benchmarking site
www.storageio.com	Website for the StorageIO Group
www.storageioblog.com	Author's blog site
www.storageperformance.org	Storage performance benchmarking site
www.svlg.net	Silicon Valley Leadership Group
www.t11.org	Fibre Channel and related standards
www.techtarget.com	IT data center publications, conferences and events
www.thegreengrid.org	Industry trade group
www.tiaonline.org	Telecommunications Industry Association
www.top500.org	List of top 500 supercomputing sites
www.tpc.org	Transaction performance council benchmark site
www.trustedcomputinggroup.org	Security related website
www.uptimeinstitute.org	Uptime Institute
www.usenix.org	LISA and data center forums
www.usgbc.org	United States Green Building Council
http://communities.vmware.com	VMware technical community website
www.wwf.org	World Wildlife Fund
www.zjournal.com	Enterprise focused servers, storage and networking

Source and Courtesy: "The Green and Virtual Data Center" (CRC)

Additional glossaries and dictionaries can be found on various industry trade groups including SNIA among others websites including those listed above. In addition to the above links for more information, industry and technology solution providers, trade groups, publications and related information can be found at:

www.storageio.com/interestinglinks.htmwww.greendatastorage.com/interestinglinks.htm

AES	Form of encryption
AFR	Annual failure rate measured or estimated failures per year
AFR	Annual failure rate measured or estimated failures per year

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>Agent</i>	<i>Software for performing backup or other IRM functions on a server</i>
<i>AHU</i>	<i>Air handling unit for cooling and HVAC systems</i>
<i>Air scrubber</i>	<i>Device for removing contaminants or emissions from air</i>
<i>AMD</i>	<i>Manufacturer of processing chips</i>
<i>ANSI</i>	<i>American National Standards Institute</i>
<i>Antimony</i>	<i>Hazardous substance found in some computer related technologies</i>
<i>API</i>	<i>Application Program Interface</i>
<i>APM</i>	<i>Adaptive power management; varies energy use to service delivered</i>
<i>Application blades</i>	<i>Server, storage and I/O networking blades for different applications</i>
<i>Applications</i>	<i>Programs or software that performance business or IRM services</i>
<i>Archiving</i>	<i>Identifying and moving inactive data to alternate media for future use</i>
<i>ASP</i>	<i>Application service provider delivering functionality via the Internet</i>
<i>Asynchronous</i>	<i>Time delayed data transmission used for low cost, long distances</i>
<i>ATM</i>	<i>Asynchronous Transfer Mode networking technology</i>
<i>Availability</i>	<i>The amount or percent of time a system is able and ready to work</i>
<i>Availability</i>	<i>The amount or percentage of time a system is able and ready to work</i>
<i>AVS</i>	<i>Adaptive Voltage Scaling; varies energy used to work performed</i>
<i>AVSO</i>	<i>Adaptive voltage scaling optimized or enhanced energy savings</i>
<i>Bandwidth</i>	<i>Measure of how much data is moved in a given amount of time</i>
<i>BC</i>	<i>Business continuance</i>
<i>Blade center</i>	<i>Packaging combining blade servers, I/O and networking blades</i>
<i>Blade server</i>	<i>Server blade packaged as a blade for use in a blade center</i>
<i>CapEx</i>	<i>Capital expenses</i>
<i>Carbon emissions</i>	<i>CO2 emissions as a byproduct of electricity generation</i>
<i>Carbon offset credits</i>	<i>Means of offsetting emissions by paying or buying credits</i>
<i>Carbon tax</i>	<i>Emissions tax placed on carbon as a resulted of energy used</i>
<i>CAS</i>	<i>Content addressable storage</i>
<i>CD</i>	<i>Compact disc</i>
<i>CDP</i>	<i>Continuous Data Protection or Complete Data Protection</i>
<i>Chiller</i>	<i>Cooling device to remove heat from coolant</i>
<i>Chlorine</i>	<i>Substance used in some electronic equipment</i>
<i>CIFS</i>	<i>Common Internet File system (NAS) for file and data sharing</i>
<i>CIM</i>	<i>Common information model for accessing information</i>
<i>Citix</i>	<i>Virtualization solutions provider</i>
<i>Cloud computing</i>	<i>Internet or web remote based application or IRM related services</i>
<i>Cluster</i>	<i>Collection of servers or storage working together, also known as a grid</i>
<i>Clustered file system</i>	<i>Distributed file system across multiple servers or storage nodes</i>
<i>CMDB</i>	<i>Configuration Management Database or repository</i>
<i>CMG</i>	<i>Computer Measurement Group; capacity and performance planners</i>
<i>CNA</i>	<i>Converged Network Architecture</i>
<i>CNIC</i>	<i>Converged Network Interface Card or Chip</i>
<i>CO2 emissions</i>	<i>Carbon dioxide emissions</i>
<i>Cold Aisles</i>	<i>Aisles in between equipment racks or cabinets with cold air</i>
<i>Consoles</i>	<i>Management interfaces for configuration or control of IT devices</i>
<i>Cooked storage</i>	<i>Formatted, usable storage with a file system or non-raw storage</i>
<i>Cooling ton</i>	<i>12,000 BTUs to cool one ton of air</i>
<i>COS</i>	<i>Console operating system also known as a console or boot system</i>
<i>CP</i>	<i>Capacity planning</i>
<i>CPU</i>	<i>Central Processing Unit</i>
<i>CRAC</i>	<i>Computer room air conditioning</i>
<i>Cross technology domain</i>	<i>Solution or tools that address multiple technologies and disciplines</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

CSV	<i>Comma separated variable format used for spread sheet data</i>
D2D	<i>Disk to disk snapshot, backup, copy or replication</i>
D2D2D	<i>Disk to disk to disk snapshot, backup, copy or replication</i>
D2D2T	<i>Disk to disk to tape snapshot, backup, copy, replication or archive</i>
D2T	<i>Disk to tape backup, copy or archive</i>
DAS	<i>Direct Attached Storage, either internal or external to a server</i>
Data barn	<i>Large repository for holding large amounts of on-line or off-line data</i>
Data in-flight	<i>Data being moved between locations, between servers or storage</i>
Database	<i>Structured means of organizing and storing data</i>
DB2/UDB	<i>IBM database software</i>
DBA	<i>Database administrator</i>
DBS	<i>Dynamic bandwidth switching varies energy use to performance</i>
DC	<i>Data center</i>
DC	<i>Direct current electricity</i>
DCE	<i>Data center Ethernet for converged I/O and networking</i>
DCE	<i>Data communications equipment</i>
DCE	<i>Distributed computing environment</i>
DCiE	<i>Data center infrastructure efficiency</i>
DCPE	<i>Data center performance efficiency</i>
DDR/RAM	<i>Double Data Rate Random Access Memory</i>
De-dupe	<i>De-duplication or elimination of duplicate data</i>
DEFRA	<i>UK Department for Environment, Food, Rural Affairs</i>
Desktop	<i>Workstation or laptop computer, also known as a PC</i>
DFS	<i>Distributed File systems for distributed and shared data access</i>
DHCP	<i>Dynamic host configuration protocol for network management</i>
DIO	<i>Direct IO operations addressing specific storage addresses</i>
Director	<i>I/O and networking large scale, multi-protocol resilient switch</i>
DL	<i>Disk library used for storing backup and other data, alternative to tape</i>
DLM	<i>Data lifecycle management</i>
DMA	<i>Direct memory access</i>
DMTF	<i>Distributed Management Task Force</i>
DNS	<i>Domain name system for managing internet domain names</i>
DOE FEMP	<i>US DoE Federal Energy Management Program</i>
DoE	<i>Department of Energy</i>
DPM	<i>Data protection management</i>
DR	<i>Disaster recovery</i>
DRO	<i>Data replication optimization</i>
DRP	<i>Disaster recovery planning</i>
DSM	<i>Demand side management</i>
DVD	<i>Digital Video Disc</i>
DVR	<i>Digital Video Recorder</i>
DWDM	<i>Dense wave division multiplexing</i>
ECC	<i>Error correcting code</i>
ECCJ	<i>Energy Conservation Center Japan</i>
Economizer	<i>Cooling device that utilizes outside cool air for cooling</i>
eDiscovery	<i>Electronic search and data discovery</i>
EHS	<i>Environmental Health and Safety</i>
ELV	<i>End of Life Vehicle Directive - European Union</i>
Emissions tax	<i>Tax for emissions such as CO2 from energy generation or use</i>
EMS	<i>Environmental management system</i>
Energy Star	<i>US EPA Energy Star® program</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>EPA</i>	<i>Environmental Protection Agency</i>
<i>EPEAT</i>	<i>Electronic Product Environmental Assessment Tool</i>
<i>ESRP</i>	<i>Microsoft Exchange Solution Reviewed Program benchmark</i>
<i>Ethernet</i>	<i>Network interface</i>
<i>ETS</i>	<i>Emissions trading scheme</i>
<i>EU</i>	<i>European Union</i>
<i>E-waste</i>	<i>Electronic waste associated with disposal of IT equipment</i>
<i>FAN</i>	<i>File area network or file based storage management</i>
<i>FC</i>	<i>Fibre Channel</i>
<i>FCIA</i>	<i>Fibre Channel Industry Association</i>
<i>FCIP</i>	<i>Fibre Channel on IP for long distance data movement and mirroring</i>
<i>FCoE</i>	<i>Fibre Channel over Ethernet</i>
<i>FCP</i>	<i>Fibre Channel SCSI_Protocol</i>
<i>FC-SB2</i>	<i>FICON Upper Level Protocol (ULP)</i>
<i>FEMP</i>	<i>US DoE Federal Energy Management Program</i>
<i>File data access</i>	<i>Accessing data via a file system locally or remotely</i>
<i>Firewall</i>	<i>Security device or software to block un-authorized network access</i>
<i>FLASH memory</i>	<i>Non-volatile memory</i>
<i>FTP</i>	<i>File transfer protocol</i>
<i>Fuel Cell</i>	<i>Alternative source for producing energy including electricity</i>
<i>G&T</i>	<i>Generating and transmission network for electrical power</i>
<i>GbE</i>	<i>Gigabit Ethernet</i>
<i>Generator</i>	<i>Device for producing electrical power for standby or co-generation</i>
<i>GHG</i>	<i>Green house gas</i>
<i>Ghz</i>	<i>Gigahertz frequency measure of speed</i>
<i>Global namespace</i>	<i>Directory name space to ease access across multiple file systems</i>
<i>GPS</i>	<i>Global Positioning System (or initials of the author of this book)</i>
<i>Green supply chain</i>	<i>Suppliers and partners adoption and implementation of being green</i>
<i>Green technology</i>	<i>Technology that addresses power, cooling, floor space or EHS</i>
<i>Green washing</i>	<i>Using green to market or being seen as green instead of being green</i>
<i>Grid</i>	<i>Local or wide area cluster of servers or storage working together</i>
<i>Guest OS</i>	<i>Guest operating system in a VM or LPAR, also known as an image</i>
<i>GUI</i>	<i>Graphical User Interface</i>
<i>HA</i>	<i>High availability</i>
<i>HBA</i>	<i>Host Bus Adapter for attaching peripherals to servers or storage</i>
<i>HCA</i>	<i>Host channel adapter for InfiniBand</i>
<i>HD</i>	<i>High definition broadcast or video</i>
<i>HDD</i>	<i>Hard disk drive such as Fibre Channel, SAS, SATA or USB</i>
<i>HDTV</i>	<i>High definition TV</i>
<i>HFC</i>	<i>Hydro fluorocarbons found in computer related technologies</i>
<i>HHDD</i>	<i>Hybrid HDD with RAM, FLASH and/or magnetic media</i>
<i>Hosting site</i>	<i>Facility or service provider that hosts IT components and services</i>
<i>Hot aisles</i>	<i>Aisles in between equipment cabinets where warm air exhausts</i>
<i>HPC</i>	<i>High performance compute</i>
<i>HSM</i>	<i>Hierarchical Storage Management</i>
<i>HTTP</i>	<i>Hypertext transfer protocol for serving and accessing web pages</i>
<i>HVAC</i>	<i>Heating, ventilation and air conditioning</i>
<i>HyperV</i>	<i>Microsoft virtualization infrastructure software</i>
<i>Hypervisor</i>	<i>Virtualization framework emulates and partitions physical resources</i>
<i>I/O rate</i>	<i>How many I/O operations (IOPS) read or write in a given timeframe</i>
<i>I/O size</i>	<i>How big the I/O operations are</i>

<i>I/O type</i>	<i>Reads, writes, random or sequential</i>
<i>I/O</i>	<i>Input output operation, read or write</i>
<i>IBA</i>	<i>InfiniBand Architecture</i>
<i>IC</i>	<i>Integrated Circuit</i>
<i>IEEE</i>	<i>Institute of Electrical and Electronic Engineers</i>
<i>IETF</i>	<i>Internet Engineering Task Force</i>
<i>ILM</i>	<i>Information lifecycle management</i>
<i>IM</i>	<i>Instant messaging</i>
<i>Image</i>	<i>Guest operating system or workload residing in a VM or LPAR</i>
<i>IMPI</i>	<i>Intelligent Platform Management Interface,</i>
<i>Intel</i>	<i>Large processor and chip manufacturer</i>
<i>Iometer</i>	<i>Load generation and simulation tool for benchmark comparisons</i>
<i>IOPS</i>	<i>I/O operations per second for reads and writes of various sizes</i>
<i>lostat</i>	<i>I/O monitoring tool</i>
<i>IOV</i>	<i>I/O virtualization including converged networks and PCI switching</i>
<i>IP</i>	<i>Intellectual property</i>
<i>IP</i>	<i>Internet protocol part of TCP/IP</i>
<i>IPM</i>	<i>Intelligent power management; varies energy used to service delivered</i>
<i>IPSec</i>	<i>IP based security and encryption</i>
<i>IPTV</i>	<i>IP based TV</i>
<i>IRM</i>	<i>Infrastructure resource management</i>
<i>iSCSI</i>	<i>SCSI command set mapped to IP</i>
<i>ISO 14001</i>	<i>Environmental management standards</i>
<i>ISO</i>	<i>International Standards Organization</i>
<i>ISV</i>	<i>Independent software vendor</i>
<i>IT</i>	<i>Character (Cousin IT) from the Addams family</i>
<i>IT</i>	<i>Information technology</i>
<i>JEDEC</i>	<i>Joint Electronic Device Engineering Council</i>
<i>JEITA</i>	<i>Japan Electronics Information Technology Industry Association</i>
<i>J-MOSS</i>	<i>Japan version of RoHS</i>
<i>Joule</i>	<i>Measure of energy usage</i>
<i>JRE</i>	<i>Java Runtime Environment</i>
<i>JVM</i>	<i>Java Virtual Machine</i>
<i>Key management</i>	<i>Managing encryption keys</i>
<i>Kill a Watt by P3</i>	<i>Device for monitoring watts, volts and amperage</i>
<i>KPI</i>	<i>Key performance indicator</i>
<i>KVM</i>	<i>Keyboard video monitor</i>
<i>Kyoto Protocol</i>	<i>Multi-national protocol and treaty to reduce emissions and GHGs</i>
<i>LAN</i>	<i>Local area network</i>
<i>Laptop</i>	<i>Portable computer</i>
<i>LEED</i>	<i>Leadership in Energy Efficiency Design</i>
<i>Linux</i>	<i>Open source operating system</i>
<i>LiveMigration</i>	<i>Virtual Iron function similar to VMware VMotion</i>
<i>LPAR</i>	<i>Logical partition or virtual machine</i>
<i>LPG</i>	<i>Liquid propane gas</i>
<i>LUNS</i>	<i>Logical unit numbers addressing for storage targets or devices</i>
<i>MAC</i>	<i>Media access control layer for networking interfaces</i>
<i>Magnetic tape</i>	<i>Low cost, energy efficient, removable media for storing data</i>
<i>MAID 2.0</i>	<i>Second generation MAID with intelligent power management</i>
<i>MAID</i>	<i>Massive array of idle disks that avoids power when not in use</i>
<i>Mainframe</i>	<i>IBM legacy large server; generic name for a large frame based server</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>MAN</i>	<i>Metropolitan area network</i>
<i>Metadata</i>	<i>Data describing other data including how and when it was used</i>
<i>MHz</i>	<i>Megahertz frequency or indicator or speed</i>
<i>MIB</i>	<i>Management information block for SNMP</i>
<i>MO</i>	<i>Magneto Optical storage medium</i>
<i>MPLS</i>	<i>Multi-protocol labeling switching WAN networking protocol</i>
<i>MR-IOV</i>	<i>PCI multi-root IOV capability</i>
<i>MSDS</i>	<i>Material Safety Data Sheet for products</i>
<i>MSP</i>	<i>Managed service provider or airport code for Minneapolis/St. Paul</i>
<i>MTBF</i>	<i>Mean time between failures; measured or estimated reliability</i>
<i>MTBF</i>	<i>Mean time between failures—measured or estimated reliability</i>
<i>MTTR</i>	<i>Mean time to repair or replace a failed item</i>
<i>MTTR</i>	<i>Mean time to repair or replace failed item back into service</i>
<i>Mw</i>	<i>Megawatts</i>
<i>NAND</i>	<i>Non-volatile computer memory such as FLASH</i>
<i>NAS</i>	<i>Network attached storage such as NFS and CIFS windows file sharing</i>
<i>Near-line</i>	<i>Non primary active data storage that does not need fast access</i>
<i>NEMA</i>	<i>National Electronic Manufactures Association</i>
<i>NERC</i>	<i>North American Electric Reliability Association</i>
<i>NFS</i>	<i>Network File System (NAS) for file and data sharing</i>
<i>Nfstat</i>	<i>Operating system utility for monitoring NFS activity</i>
<i>NIC</i>	<i>Network interface card or chip</i>
<i>NIST</i>	<i>National Institute of Standards and Technology</i>
<i>NO₂</i>	<i>Nitrate oxide</i>
<i>NOCC</i>	<i>Network operations control center</i>
<i>NPIDV</i>	<i>N_Port ID Virtualization for Fibre Channel I/O networking</i>
<i>NPVID</i>	<i>N_Port Virtual ID, similar to NPIDV</i>
<i>NRDC</i>	<i>Natural Resources Defense Council</i>
<i>NVRAM</i>	<i>Non-volatile RAM</i>
<i>Object data access</i>	<i>Data access via application specific API or descriptors</i>
<i>OC</i>	<i>Optical carrier network</i>
<i>Off-line</i>	<i>Data or IT resources that are not on-line and ready for use</i>
<i>OLTP</i>	<i>On-line transaction processing</i>
<i>On-line</i>	<i>Data and IT resources that are on-line, active and ready for use</i>
<i>OpEx</i>	<i>Operational expenses</i>
<i>Optical</i>	<i>Optical based networking or optical based storage medium</i>
<i>Orphaned storage</i>	<i>Lost, misplaced, or forgotten about storage or storage space</i>
<i>OS</i>	<i>Computer operating system such as Linux, UNIX, Microsoft Windows</i>
<i>Outage</i>	<i>Systems or subsystems are not available for use or to perform work</i>
<i>Over-supporation</i>	<i>Allocating common shared service to multiple users to reduce costs</i>
<i>P2V</i>	<i>Physical-to-virtual migration or conversion of a server</i>
<i>PACE</i>	<i>Performance Availability Capacity Energy</i>
<i>PACE</i>	<i>Popular dipping sauce for snack food</i>
<i>Para-virtualization</i>	<i>Optimized virtualization requiring custom software change.</i>
<i>Parity</i>	<i>Technique using extra memory or storage to ensure that data is intact</i>
<i>PATA</i>	<i>Parallel ATA I/O interface</i>
<i>PC</i>	<i>Payment card industry for credit card security</i>
<i>PC</i>	<i>Personal computer or program counter</i>
<i>PCFE</i>	<i>Power, Cooling, Floor-space, EHS (PCFE)</i>
<i>PCI IOV</i>	<i>PCI Sig I/O virtualization implementation</i>
<i>PCI</i>	<i>Peripheral computer interconnect for attaching devices to servers</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>PCIe</i>	<i>PCI express is the latest implementation of the PCI standard</i>
<i>PDA</i>	<i>Personal Digital Assistant (e.g. Blackberry, Apple, Windows based)</i>
<i>PDU</i>	<i>Power distribution unit</i>
<i>Pf</i>	<i>Power factor of how efficiently a power supply utilizes power</i>
<i>PFC</i>	<i>Per fluorocarbons</i>
<i>PG&E</i>	<i>Pacific Gas & Electric utility in Northern California</i>
<i>Physical volume</i>	<i>A disk drive or group of disk drives presented by a storage system</i>
<i>PIT</i>	<i>Point in time</i>
<i>PM</i>	<i>Physical machine or a real physical server or computer</i>
<i>PMDB</i>	<i>Performance Management Database</i>
<i>pNFS</i>	<i>Parallel NFS (NAS) for parallel high performance file access</i>
<i>POTS</i>	<i>Plain old telephone system</i>
<i>Power plant</i>	<i>Electrical power generation facility</i>
<i>Precision cooling</i>	<i>Pin-pointing cooling closest to heat sources for efficiency</i>
<i>Primary</i>	<i>Storage, server or networks used day to day for service delivery</i>
<i>Provisioning</i>	<i>Allocating and assigning servers, storage and networking resources</i>
<i>Proxy</i>	<i>Server configured to off-load servers for certain IRM functions</i>
<i>PS</i>	<i>Power supply</i>
<i>PST</i>	<i>Microsoft Exchange email personal storage folder file</i>
<i>PUE</i>	<i>Power usage effectiveness measurement</i>
<i>PVC</i>	<i>Polyvinyl chloride substance used in various products</i>
<i>QA</i>	<i>Quality assurance</i>
<i>QoS</i>	<i>Quality of service</i>
<i>Quad Core</i>	<i>Processor chip with four core CPUs</i>
<i>RAID</i>	<i>Redundant array of independent disks</i>
<i>RAID</i>	<i>Substance for stopping bugs and other pests</i>
<i>Raised floor</i>	<i>Elevated floor with cabling and cooling located under the floor</i>
<i>RAM</i>	<i>Random access memory</i>
<i>RASM</i>	<i>Reliability availability serviceability management</i>
<i>Raw storage</i>	<i>Storage un-configured or un-formatted with file system or RAID</i>
<i>RDM</i>	<i>Raw device mapped storage as opposed to file mapped storage</i>
<i>REACH</i>	<i>Registration, Evaluation, Authorization and Restriction of Chemicals</i>
<i>REC</i>	<i>Renewable energy credit</i>
<i>Reliability</i>	<i>Systems function as expected when expected with confidence</i>
<i>Reliability</i>	<i>Systems function as expected, when expected, with confidence</i>
<i>Remote mirroring</i>	<i>Replicating or mirroring data to a remote location for BC and DR</i>
<i>Renewable energy</i>	<i>Energy sources such as wind, solar or hydro</i>
<i>Replication</i>	<i>Mirroring of data to a second system or alliterative location</i>
<i>RFID</i>	<i>Radio frequency ID tag and reader</i>
<i>RHDD</i>	<i>Removal HDD</i>
<i>ROBO</i>	<i>Remote office branch office</i>
<i>RoHS</i>	<i>Restriction of hazardous substances</i>
<i>ROI</i>	<i>Return on investment</i>
<i>Router</i>	<i>Networking or storage device for protocol conversion and routing</i>
<i>RPC</i>	<i>Remote procedure call for program to program communications</i>
<i>RPO</i>	<i>Recovery point objective from which data can be recovered</i>
<i>RPO</i>	<i>Recovery-point objective—the point to which data can be restored</i>
<i>RTO</i>	<i>Recovery time objective of when data and applications are usable</i>
<i>RTO</i>	<i>Recovery-time objective—when recovered data will be usable again</i>
<i>RTSP</i>	<i>Real-time streaming protocol for streaming data</i>
<i>RU or U</i>	<i>Rack unit</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>RUT</i>	<i>Rules of thumb</i>
<i>S/390</i>	<i>IBM mainframe architecture now referred to a "Z" or "Zed" series</i>
<i>SaaS</i>	<i>Software as a service</i>
<i>SAN</i>	<i>Storage area network</i>
<i>SAR</i>	<i>System analysis and reporting tool</i>
<i>SAS</i>	<i>Serial attached SCSI I/O interface and type of disk drive</i>
<i>SAS</i>	<i>Statistical analysis software</i>
<i>SATA</i>	<i>Serial ATA I/O interface and type of disk drive</i>
<i>SB20/50</i>	<i>California Electronics Waste Recycling Act of 2003</i>
<i>SCADA</i>	<i>Supervisory control and data acquisition</i>
<i>Scheduled downtime</i>	<i>Planned downtime for maintenance, replacement, and upgrades</i>
<i>Scheduled downtime</i>	<i>Planned downtime for maintenance, replacement, and upgrades</i>
<i>SCSI</i>	<i>Small Computer Storage Interconnect I/O interface and protocol</i>
<i>SCSI_FCP</i>	<i>SCSI command set mapped to Fibre Channel also known as FCP</i>
<i>SDK</i>	<i>Software development kit</i>
<i>Semi-structured data</i>	<i>Email data that has structured or index and unstructured attachments</i>
<i>SFF</i>	<i>Small form factor disk drives, servers, I/O and networking blades</i>
<i>SFP</i>	<i>Small form factor optical transceiver</i>
<i>SharePoint</i>	<i>Microsoft software for managing documents utilizing SQLserver</i>
<i>SIS</i>	<i>Single instance storage also known as de-duplicated or normalized</i>
<i>SLA</i>	<i>Service level agreement defines service availability and performance</i>
<i>SLO</i>	<i>Service level objective to manage service delivery towards</i>
<i>SMB</i>	<i>Small medium business</i>
<i>SMIS</i>	<i>Storage Management Interface Specification</i>
<i>Snapshots</i>	<i>A picture or image of the data as of a point in time</i>
<i>SNIA</i>	<i>Storage Networking Industry Association</i>
<i>SNMP</i>	<i>Simple network management protocol for device management</i>
<i>SoA</i>	<i>Service oriented architecture</i>
<i>SOHO</i>	<i>Small office home office</i>
<i>SONET/SDH</i>	<i>Synchronous optical networking/synchronous digital hierarchy</i>
<i>SPC</i>	<i>Storage Performance Council benchmarks</i>
<i>SPEC</i>	<i>Performance benchmarks</i>
<i>SQL database</i>	<i>Microsoft based SQL database product that supports SharePoint</i>
<i>SQL database</i>	<i>Structure query language based database</i>
<i>SR-IOV</i>	<i>Single root PCIe IOV</i>
<i>SRM</i>	<i>Server, storage or system resource management</i>
<i>SRM</i>	<i>System or storage resource management tools</i>
<i>SRM</i>	<i>VMware site recovery manager for data protection management</i>
<i>SSD</i>	<i>Solid state disk device using FLASH, RAM or a combination</i>
<i>SSP</i>	<i>Storage solution provider also known as MSP or cloud storage</i>
<i>Structured data</i>	<i>Data stored in databases or other well-defined repositories</i>
<i>Super computer</i>	<i>Very fast and large performance oriented server</i>
<i>SUT</i>	<i>System under test</i>
<i>SUV</i>	<i>Sport utility vehicle</i>
<i>SUV</i>	<i>System under validation</i>
<i>Switch</i>	<i>I/O and networking connectivity for attaching multiple devices</i>
<i>Synchronous</i>	<i>Real time data movement based communications</i>
<i>T11</i>	<i>ANSI standards group for Fibre Channel</i>
<i>Tape</i>	<i>Magnetic tape used for storing data off-line</i>
<i>TCP/IP</i>	<i>Transmission control protocol/internet protocol; networking protocols</i>
<i>Thin provisioning</i>	<i>Virtually allocate or overbook physical storage to multiple servers</i>

**Comments and Feedback for EPA Energy Star® for Enterprise Storage Specification
 July 3, 2009**

<i>Thumb drive</i>	<i>FLASH memory based devices with USB interface for moving data</i>
<i>Tiered access</i>	<i>Different I/O and network interfaces aligned to various service needs</i>
<i>Tiered protection</i>	<i>Different data protection techniques and RTO/RPO for service needs</i>
<i>Tiered servers</i>	<i>Different types of servers aligned to various cost and service needs</i>
<i>Tiered storage</i>	<i>Different types of storage aligned to various cost and service needs</i>
<i>TPC</i>	<i>Transaction Processing Council benchmarks</i>
<i>Transaction integrity</i>	<i>Ensuring write order consistency of time based transactions or events</i>
<i>ULP</i>	<i>Upper level protocol</i>
<i>UltraSCSI</i>	<i>Most recent form of parallel SCSI cabling and SCSI command set</i>
<i>UNIX</i>	<i>Open systems operating system</i>
<i>Un-scheduled</i>	<i>Un-planned downtime for emergency repair or maintenance</i>
<i>Unscheduled</i>	<i>Unplanned downtime for emergency repair or maintenance</i>
<i>Unstructured data</i>	<i>Data including files, videos, photos, slides stored outside of databases</i>
<i>Usable storage</i>	<i>Amount of storage that can actually be used when formatted</i>
<i>USB</i>	<i>Universal Serial Bus for attaching peripheral to workstations</i>
<i>VCB</i>	<i>VMware consolidated backup proxy-based backup for VMs</i>
<i>VDC</i>	<i>Virtual data center</i>
<i>VDI</i>	<i>Virtual desktop infrastructure</i>
<i>VIO</i>	<i>Virtual I/O</i>
<i>Virtual memory</i>	<i>Operating system or VM extended memory mapped to disk storage</i>
<i>Virtual office</i>	<i>Remote or home office for mobile or remote workers</i>
<i>Virtualization</i>	<i>Abstraction, emulation, aggregation of IT resources</i>
<i>Virtualization</i>	<i>Tools to abstract emulate and aggregate to IT resource management</i>
<i>VLAN</i>	<i>Virtual LAN</i>
<i>VM</i>	<i>Virtual machine or logical partition that emulates a physical machine</i>
<i>VMark</i>	<i>VMware benchmark and comparison utility</i>
<i>VMDK</i>	<i>VMware disk file containing the VM instance</i>
<i>VMFS</i>	<i>VMware file system, stored in a VMDK file</i>
<i>VMotion</i>	<i>VMware tool for migrating a running VM to another physical server</i>
<i>VMware</i>	<i>Virtualization infrastructure solution</i>
<i>VOD</i>	<i>Video on demand</i>
<i>Volume manager</i>	<i>Software that aggregates and abstracts storage for file systems</i>
<i>VPN</i>	<i>Virtual private network</i>
<i>VTL</i>	<i>Virtual tape library based on disk drives emulates tape drives</i>
<i>VTS</i>	<i>Virtual tape system same as a VTL</i>
<i>WAAS</i>	<i>Wide area application services similar to WAFS</i>
<i>WADM</i>	<i>Wide area data management similar to WAFS</i>
<i>WADS</i>	<i>Wide area data services similar to WAFS</i>
<i>WAFS</i>	<i>Wide area file services tools for remote data and application access</i>
<i>WAN</i>	<i>Wide area network</i>
<i>Web 2.0</i>	<i>Second generation Web applications that are two-way or collaborative</i>
<i>WEEE</i>	<i>Waste from Electrical and Electronic Equipment</i>
<i>WGBC</i>	<i>World Green Building Council</i>
<i>WHDI</i>	<i>Wireless high definition transmission</i>
<i>Wide area cluster</i>	<i>Server or storage cluster, also known as grid, spread over a wide area</i>
<i>WiFi</i>	<i>Wireless networking for relative short distances</i>
<i>WiMax</i>	<i>Higher speed, longer distance next generation wireless networking</i>
<i>Windows</i>	<i>Microsoft operating system</i>
<i>Workstation</i>	<i>Desktop PC or laptop computer</i>
<i>WWPN</i>	<i>World wide port name used for power Fibre Channel addressing</i>
<i>x86</i>	<i>Popular hardware instruction set architecture designed by Intel</i>

Xenmotion

VM movement utility for Xen similar to VMware VMotion

XML

Extensible markup language

xWDM

Generic term for Dense Wave Division multiplexing

Source and Curetsy: "The Green and Virtual Data Center" (CRC)

a. Storage Hardware:

StorageIO Comment:

Given the different characteristics and application service requirements for data, different types of storage support online active data along with inactive idle data that vary in performance, availability, capacity, and energy consumption per price point and category of storage. To address the different data access and activity patterns and points in the data life cycle, virtualization provides a means to abstract the different tiers and categories of storage to simplify management and enable the most efficient type of storage to be used for the task at hand.

TIERED STORAGE—BALANCING APPLICATION SERVICE WITH PACE¹⁰ REQUIREMENTS

Tiered storage is an umbrella term and is often referred to by the type of HDD, by the price band, or by the architecture. Tiered storage embraces tiered media, including different types and classes of HDDs, which vary in performance, availability, capacity, and energy usage. Other storage media such as SSDs, magnetic tape, optical and holographic storage devices are also used in tiered storage.

Tiered storage—various types of storage media configured for different levels of performance, availability, capacity, and energy (PACE)—is a means to align the appropriate type of IT resources to a given set of application service requirements. Price bands are a way of categorizing disk storage systems based on price to align with various markets and usage scenarios—for example, consumer, small office/home office, and low-end small to medium-size business (SMB) in a price band of under \$6,000; mid- to high-end SMB in middle price bands into the low \$100,000 range; and small to large enterprise systems ranging from a few hundred thousand dollars to millions of dollars.

¹⁰ Power, Cooling, Floorspace/Footprint, Economics/EHS/Environmental

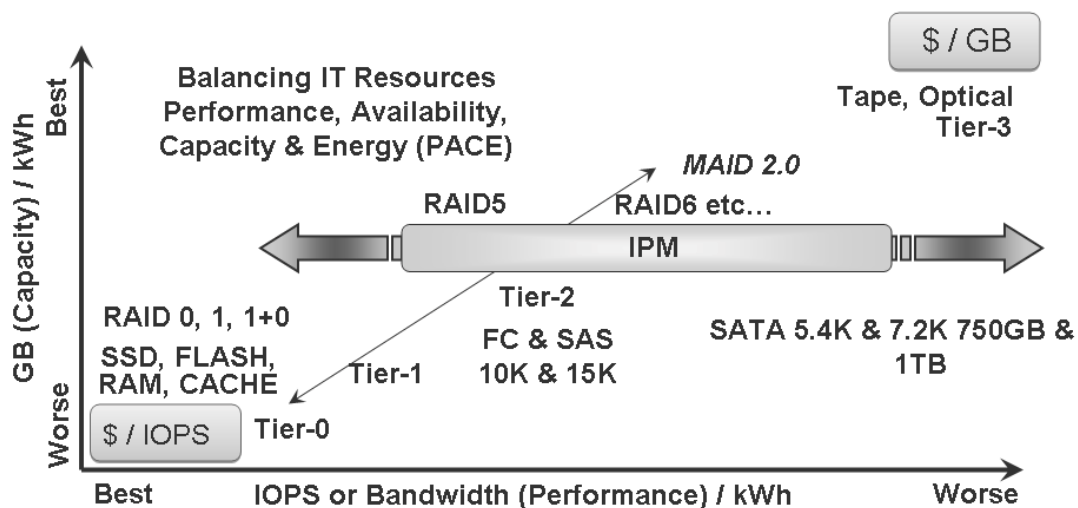


Figure 2 - Balancing Tiered Storage and Energy to Service Needs for Active and Idle States
Source and courtesy: "The Green and Virtual Data Center" (CRC)

Figure 2 shows examples of how tiered storage can be aligned. The lower left portion illustrates the use of high-performance HDDs and applicable RAID configurations to meet Tier 1 service needs measured on a cost-per-transaction basis. The upper right portion shows the other extreme: the most capacity with lowest performance and optimum energy efficiency of offline tape and optical storage. The IT balancing act is to align a given tier of storage to specific application or data needs using PACE resources in an optimal way.

Another dimension of tiered storage is tiered access, meaning the type of storage I/O interface and protocol or access method used for storing and retrieving data—for example, high-speed 8-Gb Fibre Channel (8GFC) and 10-Gb Fibre Channel over Ethernet (FCoE) versus older and slower 4GFC or low-cost 1-Gb Ethernet (1GbE), or high-performance 10GbE-based iSCSI for shared storage access, or serial attached SCSI (SAS) and Serial ATA (SATA) for direct attached storage (DAS), or shared storage between a pair of clustered servers.

Additional examples of tiered access include file- or network attached storage (NAS)-based access using network file system (NFS) or Windows-based Common Internet File System (CIFS) file sharing, among others. Different categories of storage systems combine various tiered storage media with tiered access and tiered data protection. For example, tiered data protection includes local and remote mirroring (also known as replication), in different RAID levels, point-in-time (pit) copies or snapshots, and other forms of securing and maintaining data integrity and meet data protection RTO and RPO requirements.

TIERED STORAGE SYSTEM ARCHITECTURES

Tiered storage solutions, also known as different storage system architectures, include high-end cache-centric or monolithic frame-based systems typically found in upper-price-band or mid-range and clustered storage systems. Definitions of these terms have become interface-, protocol-, and host application-independent, whereas in the past there were clear lines of delineation between different storage system architectures, similar to the traditional lines of demarcation for various types of servers.

Differences used to exist between block and file (NAS)-based solutions or enterprise (mainframe) and open systems or modular and monolithic systems, but now the packaging and features/functions are blurred. All of these systems are characterized by their ability to scale in terms of performance, capacity, availability, physical size, functionality, and connectivity. For example, these systems support a mix of IBM mainframe FICON and ESCON connectivity along with open-system Fibre Channel, iSCSI, and NAS access, natively or via optional gateways, routers, or protocol converters.

In addition to supporting both open-system and mainframe servers natively, high-end cache-centric storage systems, as their name implies, have very large amounts of cache to boost performance and support advanced feature functionality. Some systems support over 1,000 HDDs including ultra-fast SSD-based devices, fast Fibre Channel HDDs, and lower-cost and high-capacity fat SATA or Fibre Channel HDDs. While smaller mid-range storage systems can in some cases rival the performance of cache-centric systems while drawing less power, an advantage of the larger storage systems can be to reduce the number of storage systems to manage for large-scale environments.

Mid-range and modular storage systems span from the upper end of price bands 6 and 7 for enterprise solutions down to price band 1 for low-end SMB-based storage solutions. The characteristics of mid-range and modular storage systems are the presence of one or two storage controllers (also known as nodes), storage processors or heads, and some amount of cache that can be mirrored to the partner controller (when two controllers exist). Dual controllers can be active/passive, with one controller doing useful work and the other in standby mode in the event of a controller failure.

Mid-range and modular controllers attach to some amount of storage, usually with the ability to support a mix of high-performance fast HDDs and slow, large-capacity HDDs, to implement tiered storage in a box. The controllers rely on less cache instead of cache-centric solutions, although some scenarios that leverage fast processors and RAID algorithms can rival the performance of larger, more expensive

cache-centric systems. As of this writing, modular or mid-range storage systems max out in the 250- to 400-HDD range, but by the time you read this, those numbers are expected to increase along with performance.

TIERED STORAGE MEDIA OR DEVICES

Tiered storage mediums or devices, often generically referred to as tiered storage, include different types of magnetic HDDs such as fast, high-performance Fibre Channel and SAS and lower-performing, high-capacity SATA, SAS, and Fibre Channel HDDs. Other types of tiered storage devices that can be configured into storage solutions include magnetic tape, optical media (CDs, DVDs, magneto-optical) and semiconductor disk drives (SSDs).

Given that data storage spans categories from active online and primary data to offline and infrequently accessed archive data, different types of storage media addressing different value propositions can be found in a single storage solution. For example, to address high-performance active data, the emphasis is on work per unit of energy at a given cost, physical, and capacity footprint. On the other hand, for offline or secondary data not requiring performance, the focus shifts from energy efficiency (doing more work per unit of energy) to capacity density per cost, unit of energy, and physical footprint.

Compare different tiered storage media based on what applications and types of data access they will be supporting while considering cost and physical footprint. Also consider the performance, availability, capacity, and effective energy efficiency for the usage case, such as active or idle data.

As an example, a current-generation 146-GB, 15,500 (15.5K)-RPM, 4-Gb Fibre Channel or SAS 3.5-inch HDD consumes the same, if not less, power than a 750-GB, 7,200 (7.2K)-RPM SATA or SAS 3.5-inch HDD. For active online data, the 15.5K-RPM HDD delivers more performance per unit of energy than the larger-capacity SATA HDD. However, for capacity-intensive applications that do not need high performance, the SATA drive has better density per unit of energy in the same physical footprint as the faster 146-GB HDD. Which drive to use depends on the application; increasingly, a mix of high-speed FC or SAS drives is configured in a storage system with some number of lower-performing, high-capacity or FAT HDDs for a tiered storage solution in a box.

SOLID-STATE DEVICES (SSDS)

The need for more effective I/O performance is linked to the decades old, and still growing, gap between server and storage performance, where the performance of

HDD storage has not kept up with the decrease in cost and increase in reliability and capacity compared to improvements in server processing power.

Reducing energy consumption is important for many IT data centers. Although there is discussion about reducing energy by doing less work or powering down storage to reduce energy use, the trend is toward doing more with less power per unit of work. This includes intelligent power management when power consumption can be reduced without compromising application performance or availability as well as doing more IOPS or bandwidth per watt of energy.

FLASH is relatively low-cost and persistent memory that does not lose its content when power is turned off. USB thumb drives are a common example. DDR/RAM is dynamic memory that is very fast but is not persistent, and data is lost when power is removed. DDR/RAM is also more expensive than FLASH. Hybrid approaches combine FLASH for persistency, high capacity, and low cost with DDR/RAM for performance.

There is a myth that SSD is only for databases and that SSD does not work with files. The reality is that in the past, given the cost of DRAM-based solutions, specific database tables or files, indices, log or journal files, or other transient performance-intensive data were put on SSDs. If the database were small enough or the budget large enough, the entire database may have been put on SSDs.

More SSDs are not in use because of the perceived cost. The thought has been that SSDs in general costs too much compared to HDDs. When compared strictly on a cost per gigabyte or terabyte basis, HDDs are cheaper. However, if compared on the ability to process I/Os and the number of HDDs, interfaces, controllers, and enclosures necessary to achieve the same level of IOPS or bandwidth or transaction or useful work, then SSDs may be more cost-effective for a given capacity usage case. The downside to RAM compared to HDD is that electrical power is needed to preserve data.

RAM SSDs have, over the years, addressed data persistence issues with battery-backed cache or in-the-cabinet uninterruptible power supply (UPS) devices to maintain power to memory when primary power is turned off. SSDs have also combined battery backup with internal HDDs, where the HDDs are either stand-alone, mirrored, or parity-protected and powered by a battery to enable DRAM to be flushed (destaged) to the HDDs in the event of a power failure or shutdown. While DRAM-based SSDs can exhibit significant performance advantages over HDD-based systems, SSDs still require electrical power for internal HDDs, DRAM, battery (charger), and controllers.

FAST OR HIGH-PERFORMANCE AND FAT-CAPACITY HDDS

As a technology, magnetic HDDs are over 50 years old, and they have evolved significantly over that time, increasing usable capacity, performance, and availability while reducing physical footprint, power consumption, and cost. The mainstays of data center storage solutions today are based on 3.5-inch high-performance enterprise and high-capacity desktop HDDs along with emerging small-form-factor 2.5-inch high-performance enterprise HDDs.

With a growing focus on “green” storage and addressing power, cooling, and floor space issues, a popular trend is to consolidate data from multiple smaller HDDs onto a larger-capacity HDD to boost storage capacity versus energy usage for a given density ratio. For idle or inactive data, consolidating storage is an approach to addressing PCFE issues; however, for active data, using a high-performance drive to do more work using fewer HDDs is also a form of energy efficiency. As seen in Table 1, each successive generation of HDDs had improved energy usage.

Capacity	73 GB	73 GB	73 GB	146 GB	300 GB	500 GB	750 GB	1.5 TB
Speed (RPM)	15.4K	15.4K	15.5K	15.5K	15.5K	7.2K	7.2K	7.2K
Interface	2GFC	4GFC	4GFC	4GFC	4GFC	SATA	SATA	SATA
Active watts/hour	18.7	15.8	15.2	17.44	21.04	16.34	17.7	17.7
Capacity increase	N/A	N/A	N/A	2×	4×	6.5×	10×	20×

Table 1 - Balancing Performance, Availability, Capacity, and Energy across Different HDDs
 Source and courtesy: “The Green and Virtual Data Center” (CRC)

How much energy is needed to power 100 TB of storage? The answer depends on the size of the storage system (e.g., price band), if performance or capacity optimized, the category or type of tiered storage medium being used, how it is configured, and, if comparing raw versus usable RAID-protected storage, which RAID level plus storage is being compared on an effective basis using a compression or de-duplication ratio.

In Table 1, active watts/hour represents an average burdened configuration, that is, the HDD itself plus associated power supplies, cooling, and enclosure electronics per disk for active running mode. Lower power consumption can be expected during low-power or idle modes as well as for an individual disk drive minus any enclosure or packaging. In general, 100 TB of high-performance storage will require more power and therefore cooling capacity than 100 TB of low-cost, high-capacity disk-based storage. Similarly, 100 TB of high-capacity disk-based storage will consume more power than 100 TB of magnetic tape-based storage.

As an example, a mid-range, mid-price storage system with redundant RAID controllers and 192 750-GByte, 7,200-RPM or 5,400-RPM SATA HDDs in a single cabinet could consume about 52,560 kWh of power per year (not including cooling).

Assuming an annual energy cost of 20 cents per kilowatt-hour, which factors cost to cool the 100 TB of storage along with applicable energy surcharges or perhaps even carbon taxes being passed on by an electric utility, the total cost would be about \$10,500.

This works out to about 39.4 tons (this will vary by specific location and type of energy being used) of CO₂ emissions per year or the equivalent of about five to eight automobiles (which, of course, would vary depending on the type of automobile, miles driven per year, and type of fuel being used for the vehicle). To put this even more into perspective, for the 100 TB in this scenario, at 20 cents, the cost to power a terabyte of storage is about a penny an hour.

The previous is a middle-of-the-road example, with tape-based solutions being much lower in cost and high-end, large-enterprise storage systems with many fast HDDs being more expensive. “Your mileage will vary” is a key phrase, as different vendors with similar configurations and will have different power consumption levels, based on the overall energy efficiency of their systems independent of data footprint reduction or other enhancements capabilities.

Large-capacity drives certainly store more data at less power per gigabyte. This comes, however, at the expense of reduced performance, which can be aggregated due to density. Table 1 shows several different types of HDDs, ranging from 73 GB to 1.5 TB. The power consumption values shown in Table 1 may be higher than those normally found in HDD manufacturer specification sheets because they include infrastructure overhead. Overhead, or burden, includes the power used by the HDD as well as the HDD’s share of the enclosure and associated cooling.

It is easy to compare drives on a power-per-gigabyte basis, but it is also important to consider the drive in the context of how it will be used. Look at efficiency and how power is used with respect to how the storage is being used. That is, if using storage for active data, look at how much work can be done per watt of energy such as IOPS per watt, bandwidth per watt for sequential or video streams per watt of energy. If the data is inactive or idle, then consider the energy required to support a given capacity density while keeping in mind that unless it is for deep or time-delayed access, some amount of performance will be needed.

For those who think that the magnetic HDD is now dead, in actuality, just as disk is helping to keep magnetic tape around, SSD (both DRAM and FLASH) will help take some performance pressure off HDDs so that they can be leveraged in more efficient and economical ways, similar to what disk is to tape today. While magnetic HDDs continue to decline in price per capacity, FLASH price per gigabyte is declining at a faster rate, which makes storage using SSDs a very complementary technology

pairing to balance performance, availability, capacity, and energy across different application tiers.

MAGNETIC TAPE, OPTICAL, AND HOLOGRAPHIC STORAGE

For applications and environments that need the lowest-energy-consuming storage and where response time or application performance are not required, for example, offline storage, magnetic tape remains a good option and companion to HDD-based online and near-line storage systems. Magnetic tape has been around for decades in different formats and with various characteristics.

While utilizing similar techniques and basic principles, over time magnetic tape has evolved along the lines of other technologies, with increased densities, reduced physical size, better reliability, faster performance, easier handling, and lower costs.

Tape today is being challenged by lower-cost, high-capacity disk-based storage such as SATA devices to perform disk-to-disk (D2D) and disk-to-disk-to-tape (D2D2T) backups and archives. This is not to say that tape is dead, as it continues to be the most cost-effective medium for long-term and offline data retention, however tape is finding new uses¹¹ for long-term data retention given its low cost and energy efficiency for off-line data. For example, tape cartridges are capable of storing over 1 TB of native, uncompressed data. Alternative media include magneto-optical media (MOs), CDs, DVDs, and emerging holographic storage.

INTELLIGENT POWER MANAGEMENT AND MAID 2.0

Intelligent power management (IPM), also called adaptive power management (APM) and adaptive voltage scaling (AVS), applies to how electrical power consumption and, consequently, heat generation can be varied depending on usage patterns. Similar to a laptop or PC workstation with energy-saving modes, one way to save on energy consumption by larger storage systems is to power down HDDs when they are not in use. That is the basic premise of MAID, which stands for Massive (or Monolithic or Misunderstood) Array of Idle (or Inactive) Disks.

Given performance or other service requirements, not all storage or data applications are appropriate for MAID. Look at second generation MAID- and IPM-enabled storage solutions to determine the granularity and flexibility options. For example, do all of the disks in a storage system have to be under MAID management, or can some disks be included while others are excluded? Similarly, are only SATA HDDs supported, or is a mix of SATA, SAS, or even Fibre Channel disks with IPM capabilities and different MAID level granularities at the drive, RAID, or volume group level possible?

¹¹ See Tapes Changing Role In Data Protection: <http://www.devx.com/ibmsolutioncenter/Article/42045>

BALANCING PACE¹² TO ADDRESS PCFE¹³ ISSUES WITH TIERED STORAGE

Using tiered storage and aligning the most appropriate storage technology to the task or application is part of an overall approach to address power, cooling, floor-space, and environmental (PCFE) issues while balancing usage with application service-level requirements. Table 2 shows different tiers of storage independent of specific architectures (for example, cache-centric or monolithic frame based; modular, distributed, or clustered) or tiered access methods (DAS, SAN, NAS, or CAS), contrasting performance, availability, capacity, and energy consumption as well as relative cost.

	<i>Tier 0 Very High Performance</i>	<i>Tier 1 Performance and Capacity</i>	<i>Tier 2 Capacity and Low Cost</i>	<i>Tier-3 High Capacity and Low Cost</i>
<i>Uses</i>	<i>Transaction logs and journal files, paging files, look-up and meta-data files, very active database tables or indices</i>	<i>Active online files, databases, email and file servers, video serving needing performance and storage capacity</i>	<i>Home directories, file serving, Web 2.0, data backups and snapshots, bulk data storage needing large capacity at low cost</i>	<i>Monthly or yearly full backups, long-term archives or data retention with accessibility traded for cost or power savings</i>
<i>Comparison</i>	<i>Dollar per IOPs IOPS or activity per watt of energy and given data protection level</i>	<i>Activity per watt of energy and capacity density and given data protection level</i>	<i>Capacity density per energy used with performance for active data at protection level</i>	<i>Capacity density per energy used with bandwidth when accessed at protection level</i>
<i>Attributes</i>	<i>Low capacity and high performance with very low power consumption: DDR/RAM, FLASH, or some combination</i>	<i>Primary active data requiring availability and performance: 10K- or 15K-RPM 2.5- or 3.5-inch FC, SCSI, and SAS HDDs</i>	<i>Low cost, high density: 5.4K- or 7.2K-RPM SAS, SATA, or FC HDDs with capacities in excess of 1 TB</i>	<i>Low cost and high capacity or "FAT" 5.4K- or 7.2K-RPM SAS, SATA, or FC HDDs; magnetic tape and optical media</i>
<i>Examples</i>	<i>Cache, s, SSD (FLASH, RAM).</i>	<i>Enterprise and mid-range arrays</i>	<i>Bulk and IPM-based storage</i>	<i>Tape libraries, MAID/IPM or optical storage, removable HDDs</i>

Table 2 - Comparison of Tiers of Storage and Service Characteristics in the Same Footprint
Source and courtesy: "The Green and Virtual Data Center" (CRC)

Alignment of the most appropriate tier of storage to application needs is an effective technique, along with others such as data footprint reduction, to address data center bottlenecks without continuing the vicious cycle of sparse storage allocation and later consolidation.

¹² Performance, Availability, Capacity, Energy/Economic

¹³ Power, Cooling, Floorspace/Footprint, Energy/EHS/Economic

From a power and performance perspective, SSD provides very good IOPS per watt or bandwidth per watt of energy used and capacity per watt. There is, however, a trade-off of cost and capacity. As an example, 1 TB of usable FLASH-based SSD can fit into 2 rack units height (2U) in a standard 19-inch rack consuming 125 W or .125 kWh of power capable of delivering 100s of MB/second bandwidth performance. The FLASH SSD is comparable to two other storage solutions: storage-centric, high-capacity, low-cost HDDs; or high-performance, medium-capacity, disk-based storage systems.

Even with the continuing drop in prices of DDR/RAM and FLASH-based SSDs and increasing capacity (Figure 3) for most IT data centers and applications there will continue to be a need to leverage tiered storage, including HDD based storage systems. This means, for instance, that a balance of SSDs for low-latency, high-I/O or performance hotspots along with storage-capacity, high-performance HDDs in the 146-GB and 300-GB, 15.5K-RPM class are a good fit with 500-GB, 750-GB, and 1-TB HDDs for storage capacity-centric workloads.

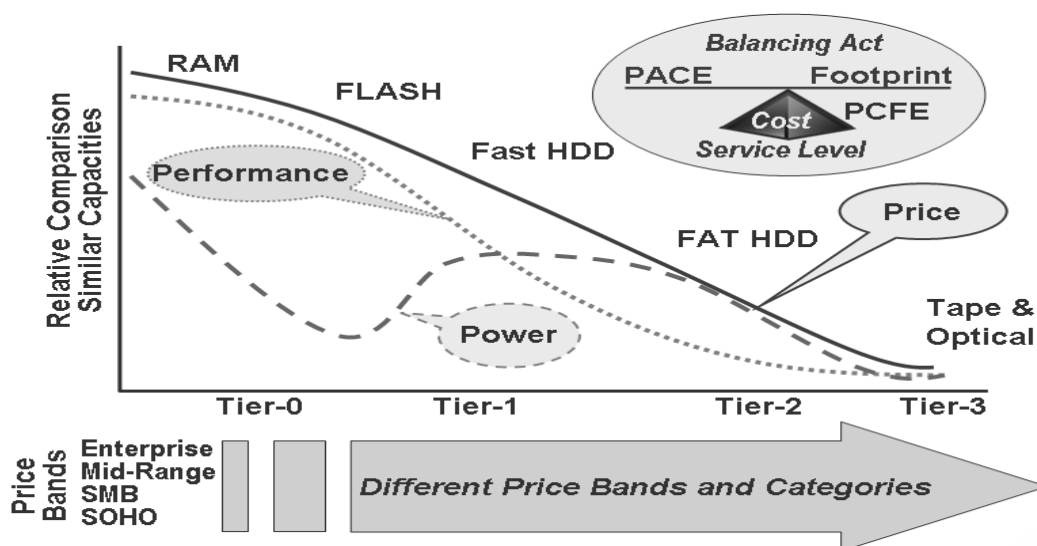


Figure 3 - Tiered Storage Options
 Source and courtesy: "The Green and Virtual Data Center" (CRC)

For active storage scenarios that do not require the ultralow latency of SSD but need high performance and large amounts of affordable capacity, energy-efficient 15K-RPM Fibre Channel and SAS HDDs provide a good balance between activity per watt, such as IOPS per watt and bandwidth per watt, and capacity, as long as the entire capacity of the drive is used to house active data. For dormant data and ultra large storage capacity environments with a tolerance for low performance, larger-capacity 750-GB and 1.5-TB "fat" HDDs that trade I/O performance for greater storage capacity provide a good capacity per watt of energy.

Another variable to consider is how the storage system is configured in terms of RAID level for performance, availability, and capacity. RAID levels affect energy consumption based on the number of HDDs or SSD (FLASH or DDR/RAM) modules being used. Ultimately, the right balance of PACE should be weighed with other decision and design criteria, including vendor and technology preferences, to address various PCFE issues.

1. Storage Media: The physical material that stores data. Storage Media used in Enterprise Storage applications may be electrical (e.g. solid state drive), magnetic (hard disk drive, tape drive), or optical (optical disc drive).

StorageIO comment:

In addition to the preceding material, storage media should include disk drives such as magnetic hard disk drive (HDD), solid state disk (SSD) drive, tape drive as well as other mediums including optical or emerging holographic. These drives and/or their removable media (e.g. removable hard disk drive, FLASH drive, tape or optical) are incorporated as part of overall storage systems. Some of the components such as HDDs of individual SSDs may lend themselves to a sub-category of Energy Star® given their inclusion in servers, desktop/notebooks and other devices in addition to in data center storage systems.

2. Storage Product: A system composed of integrated storage controllers, storage media and software that provide data storage services to one or more Computer Servers and/or other devices. While Storage Products may contain one or more embedded processors, these processors do not execute user-supplied software applications but may execute data-specific applications (e.g. data replication, backup utilities, data compression, install agents, etc.). This definition is specifically intended to exclude aggregating storage elements such as RAID array subsystems, robotic tape libraries, filers, and file servers.

StorageIO comment:

Why are the exclusions listed? RAID array controllers, tape libraries, file servers (aka filers or NAS devices) are all essential and commonly deployed data center storage products. However, if the intention is to eliminate appliances or servers running aggregation or virtualization software that consolidates data center storage systems, that could be a different matter and simply one of clarification and discussion. See preceding discussions as well as those of definitions, terms and taxonomies from other industry sources to help clarify what should be included and is typically deployed in data center environments.

3. Storage Controller: TBD

StorageIO comment:

This makes for an interesting discussion point in that some storage systems, storage solutions or products utilize custom or purpose built controllers for performing RAID and other storage functions while other vendors leverage general purpose commercial servers including Energy Star® compliant models for their controllers. A storage controller, aka storage processor, Node, NAS head, appliance is a key and essential element of any storage system from block to NAS to object to VTLs among others and thus should be further clarified.

4. Storage Product Family: A group of Storage Product configurations where every configuration includes base components with the same or similar technical specifications and power specifications.

5. Blade System: A system composed of both a Blade Chassis and one or more removable Blade Servers or Blade Storage units. Blade Systems are designed as a scalable solution to efficiently package and operate multiple Computer Servers or Storage units in a single enclosure, and are designed for technicians to be able to easily add or replace hot-swappable boards in the field.

StorageIO comment:

Some storage solutions are being deployed today leveraging off the shelf, commercial blade system or blade center or blade servers running specialized software to create a storage solution such as clustered NAS or bulk storage for near-line, on-line or off-line applications. There is an interesting notion or possibility that an underlying component for example the blade server, plus disk enclosures, HBAs or NICs, as well as unifying software combined could be an Energy Star® for data center candidate as well as some of the comments also being Energy Star® such as a server.

6. Blade Chassis: An enclosure containing shared resources for the operation of Blade Servers and Blade Storage units. These resources may include power supply(s) for power conversion, shared storage, and hardware for DC power distribution, thermal management, system management, and network services. A Blade Chassis features multiple slots which can be populated with blades of different types.

StorageIO comment:

Some storage solutions are being deployed today leveraging off the shelf, commercial blade system or blade center or blade servers running specialized software to create a storage solution such as clustered NAS or bulk storage for near-line, on-line or off-line applications. There is an interesting notion or possibility that an underlying component for example the blade server, plus disk enclosures, HBAs or NICs, as well as unifying software combined could be an Energy Star® for data center candidate as well as some of the components also being Energy Star® such as a server.

7. Blade Storage: A storage-specific element that relies on shared resources (e.g., power supply, cooling, etc.) for operation. Blade Storage units are designed to be installed in a Blade Chassis, are hot-swappable and are incapable of operating independent of the chassis.

StorageIO comment:

Some storage solutions are being deployed today leveraging off the shelf, commercial blade system or blade center or blade servers running specialized software to create a storage solution such as clustered NAS or bulk storage for near-line, on-line or off-line applications. There is an interesting notion or possibility that an underlying components for example the blade server, plus disk enclosures, HBAs or NICs, as well as unifying software combined could be an Energy Star® for data center candidate as well as some of the components also being Energy Star® such as a server.

8. I/O Device: Provides data input and output capability to the Storage Product from other networked devices. Examples of I/O Devices include: Ethernet devices, InfiniBand devices, external RAID/SAS controllers and Fibre Channel devices.

StorageIO comment:

This should include block, file (NAS), object, tape and virtual tape as well as various interfaces including those listed above as well as Fibre Channel over Ethernet, iSCSI, NFS and CIFS, HTTP and cloud protocols including SOAP, Bit-Torrent and DICOM among others. What probably needs to be clarified here is if this is referring to the I/O ports and adapters used by servers to communicate with storage systems, or the storage controllers, or the back-end storage devices (e.g. media). Refer to SNIA and other industry and vendor specific taxonomies for additional perspectives.

StorageIO comment:

Where would storage software, both software implemented as firmware, microcode, layered software as part of a storage system as well as optional storage software running on appliances or on application servers be included in this model?

b. Storage Characteristics:

1. Capacity: The sum of the raw unformatted, uncompressed capacity of all Storage Media installed in a Storage Product.

StorageIO comment:

Capacity is an interesting and contentious topic as does it mean native, raw, formatted, allocated, compressed, optimized via dedupe or other techniques? Is the reference towards the raw capacity before RAID or data protection applied or after? Does it refer to file system raw or device raw?

2. **Direct-connected:** Storage designed to be under the control of a single host, or multiple hosts, in a non-shared environment.

StorageIO comment:

Direct-connected aka DAS can imply storage directly attached to a single, or, multiple servers when using a shared interface however without using a SAN or LAN switch. Direct-connected can be shared or non shared storage, it could be SAS, Fibre Channel or others without using a SAN or LAN switch to create a network as in a point to point connection¹⁴. A common misperception is that DAS is internal dedicated storage when in reality it can be external shared SAS, Fibre Channel or even iSCSI storage connected in a point to point topology (e.g. without a network) to one or more servers (e.g. shared storage). Some data center storage is SAN or NAS connected while other storage may be shared, or, simply external DAS using SAS or some other interface.

3. **Network-connected:** Storage designed to be connected to a host via a network connection (e.g., Ethernet, InfiniBand, and Fibre Channel).

StorageIO comment:

Networked-connected (SAN or NAS) implies some type of network including Ethernet for iSCSI or NAS, Fibre Channel, Fibre Channel over Ethernet and InfiniBand among others. Some data center storage is SAN or NAS connected while other storage may be shared, or, simply external DAS using SAS or some other interface.

c. Other Data Center Hardware:

1. **Computer Server:** A computer that provides services and manages networked resources for client devices, e.g., desktop computers, notebook computers, thin clients, wireless devices, PDAs, IP telephones, other Computer Servers and other networked devices. Computer Servers are designed to respond to requests and are primarily accessed via network connections, and not through direct user input devices such as a keyboard, mouse, etc.

2. **Blade Server:** A Computer Server consisting of, at minimum, a processor and system memory that relies on shared resources (e.g., power supplies, cooling, etc.) for operation. Blade Servers are designed to be installed in a Blade Chassis, are hot-swappable, and are incapable of operating independent of the chassis.

StorageIO Comment:

While blade servers are primarily intended for supporting general purpose operating systems and applications including database, transaction processing, web hosting or

¹⁴ Refer to "Resilient Storage Networks – Desinging Scalable Flexible Data Infrastructure" (Elsevier) – SNIA Endorsed Book

Also shown in Figure 5 is the performance in terabytes per hour, to show how different solutions vary in their ability to ingest and store data at a baseline configuration. From this baseline it is possible to apply various compressions and de-duplication strategies to make additional comparisons, as well as various RAID levels and RAID group sizes and numbers of hot spare disk drives.

1. **Idle:** An operational state in which the operating system and other software have completed loading and the Storage Product is capable of completing workload “external” transactions, but no active workload transactions are requested or pending by the system.

StorageIO Comment:
See preceding discussion.

2. **Active:** TBD

StorageIO Comment:
See preceding discussion.

3. **Maximum:** TBD

4. **Full Load:** TBD

d. Questions for Discussion:

1. How are Active, Idle, Maximum, and Full Load states defined in the industry?

StorageIO Comment:
See preceding discussion.

2. What are the critical factors in determining if a Storage Product is idle?

StorageIO Comment:
See preceding discussion.

Are there other accepted terms for these states?

StorageIO Comment:
See preceding discussion.

3. Are there other Operational States specific to Storage Products that will need to be defined in this specification?

What portion of the enterprise storage market has the greatest need for an ENERGY STAR® label to help customers identify the most efficient products? In what market segments can EPA expect to garner the most energy savings?

Building Block #3: Energy Efficiency Criteria and Test Procedures

a. **Purpose:** Once it is determined which products will be included in the ENERGY STAR® specification, the next step is to identify metrics for energy efficiency performance. Metrics may address the efficiency of key components (e.g., power supplies), operational modes (e.g., Idle state), and/or whole system energy efficiency. Efficiency metrics must be supported by generally accepted test procedures. Further, while one efficiency metric across a broad range of products is preferable, EPA will evaluate the need to develop unique requirements for segments of the market if it can be shown that key product functions or purposes require additional energy. In other words, EPA does not intend to develop unique energy efficiency criteria for different storage products based only on differences in the underlying technology.

The efficiency metric(s) will ultimately be used to determine ENERGY STAR® qualification. The ENERGY STAR® program strives to set requirements such that 25% of the products available in the market at the time the specification becomes effective will be able to qualify.

StorageIO comment:

While whole system efficiency or effectiveness is applicable to what IT data centers customers are focusing on or would be buying, keeping the notion that Energy Star® for data center storage also has value to vendors, vars or other systems integrators. Consequently components that may make up an overall solution may need to be separately handled. For example a disk drive enclosure shelf that includes slots for some number of disk drives, power supplies and cooling fans, control cards and other items that are commonly leveraged by OEMs and other vars for solutions.

Another example would be individual disk drives that are used by various vendors. This could be a low hanging fruit opportunity for Energy Star® and the industry in general to create a subset of Energy Star® compliant disk drives that could then be leveraged and parlayed by both Energy Star® for data center storage, as well as Energy Star® for servers where internal dedicated disk drives are used, as well as in Energy Star® compliant notebooks or laptop computers.

Notes on Value Added Resellers (VAR):

In some cases, Storage Products may be shipped from the Original Equipment Manufacturer (OEM) to a VAR that then configures the device for sale to the end user.

It is important to EPA that a product qualified for ENERGY STAR® by an OEM, processed by a VAR, and ultimately installed and used at a customer site continues to meet ENERGY STAR® requirements. The ENERGY STAR® program for Computer Servers addressed this situation in the following manner:

In order for the VAR to sell the Computer Server as ENERGY STAR® qualified under the OEM brand name, one of two conditions must be met:

a. The end configuration sold by the VAR must have been originally qualified by the OEM; or

b. In the case that the end configuration has not been qualified by the OEM, the VAR must become an ENERGY STAR® partner, and test and qualify the configuration. OEM partners selling Computer Servers to VARs must provide the VAR with a list of qualified configurations for that model, using approved components, which have been initially qualified and reported to EPA by the OEM Partner.

b. Initial Approach: EPA intends to adopt or develop a test procedure to measure the energy consumed by a Storage Product while performing a realistic and representative workload. This workload will likely include power measurements under both Idle and Active conditions. EPA would like to adopt an industry standard procedure, but if no suitable procedure exists EPA will work with stakeholders to develop one.

StorageIO Comment:

Benchmarks, particularly industry-accepted workloads that are applicable to specific applications and environments, can be a useful gauge of how different solutions compare; however, they are not a substitute for actually testing and simulating production application and workload conditions. Take benchmarks with a grain of salt when considering their applicability to your environment, application characteristics and workloads.

Once a test procedure has been identified, EPA will begin an effort to collect and analyze test data from product tests performed by manufacturers. This data collection is critical to the success of the program, since the test data will be used to inform the development of the final ENERGY STAR® performance criteria.

EPA understands that there are many software-based approaches to improving the energy efficiency of storage products. The benefits of virtualization, data de-duplication, and other software-based data management techniques are well documented. These software solutions, perhaps even more so than the hardware itself, are heavily customized for specific customers and applications. Achieving maximum efficiency gains is highly dependent upon proper software architecture, implementation, operation, and maintenance by individual users. A key objective of the ENERGY STAR®

5. It has been noted that the SPECsfs2008 benchmark might be an appropriate workload for NAS devices and is currently used by many NAS vendors. Is the benchmark suitable for more than NAS?

StorageIO comment:

Assuming that a Energy Star® for data center storage is to be encompassing and not just for high-end block based storage, there needs to be considered and including of and for file based storage systems. Network attached storage (NAS) which includes NFS and/or CIFS file sharing represents one of the fast growing segments of storage.

The popularity of rich media and Internet-based applications has resulted in explosive growth of unstructured file data that requires new and more scalable storage solutions. Applications such as video pre- and postproduction processing, animation rendering, on-demand video and audio, social networking websites, and digitalization of data from cell phones, personal digital assistants (PDAs) and other sources have increased burdens on storage performance and capacity. Unstructured data includes spreadsheets, PowerPoint presentations, slide decks, Adobe PDF and Microsoft Word documents, Web pages, and video and audio JPEG, MP3, and MP4 files.

The diversity of rich media and Internet applications ranges from many small files with various access patterns to more traditional large video stream access. Consequently, storage systems, in order to scale with stability to support Internet and Web 2.0 applications, will need to support variable performance characteristics from small random access of meta-data or individual files to larger streaming video sequences. Data growth rates range from the low double digits to high double or triple digits as more data is generated, more copies of data are made and, more data is stored for longer periods of time.

While structured data in the form of databases continues to grow, for most environments and applications, it is semi structured email data and unstructured file data that created the biggest data footprint impact and subsequent bottlenecks. Unstructured data has varying input/output (I/O) characteristics that change over time, such as data that starts out with a lot of activity, and then goes idle for a time before extensive reads, as in the case of a video or audio file becoming popular on a media, entertainment, social networking, or company-sponsored website. As another example, usually, when a development or research project is completed, the data or intellectual property is archived or migrated to lower-cost, lower-performance bulk storage until it is needed again for further research or sequel projects.

Consequently to support un-structured data, NAS needs to be considered as a significant piece of data center storage. SPEC based benchmarks are a defacto standard for file serving measurements in the industry today. In addition to NAS or file based data and storage systems, there are also traditional block based storage that often are used for supporting server based file systems as well as databases, transaction processing, financial and other applications. Object based storage is another consideration that is evolving and thus should be considered in a strategy.

6. Are the power supply test procedures, test methods and levels from the Computer Server specification applicable to Storage? Is there data available that support these levels? What modifications might have to be made to address any of the unique characteristics of storage power supplies?

StorageIO comment:

They serve as a basis and in particular for storage systems that are based on using industry standard servers as their underlying controller hardware as is the case with some mid-range block storage systems, many virtual tape libraries (VTLs), dedupe appliances or other storage appliances. Special consideration also needs to be taken for storage enclosures, or, the disk drive shelves and the work being done by the Storage Bridge Bay (SBB) group among others as well as factoring in work by 80 plus.

7. What are the typical loading ranges for Enterprise Storage power supplies? What are the typical redundancy configurations of these systems?

StorageIO comment:

Typically at least N+1 or minimum of dual power supplies for components including controllers (or appliances) and disk drive shelves.

8. Would a specification for net power loss or efficiency over the load range be appropriate for Enterprise Storage?

9. Does the VAR sales channel play an important role in the Enterprise Storage market? Would the above requirements make clear how to qualify storage products through third-party sales channels?

StorageIO comment:

Yes to varying degrees as there is a perception that VARs or channel only sell to low-end yet in reality, many data centers buy via a channel or VAR as a means of fulfillment. Many manufactures are increasingly relying on the channel for both touch/pull as well as for fulfillment. Even some government entities buy data center storage via Var or resellers (e.g. the channel).

Building Block #4: Information and Management Requirements

a. Purpose:

EPA is interested in developing tools to facilitate the efficient design and operation of data centers. Improvements in the energy efficiency of Storage Products will lead to even greater efficiencies in data center design and operation, due to reductions in infrastructure requirements. To this end, EPA intends to take the Standard Information Reporting requirements and Data Measurement and Output requirements from the ENERGY STAR® Computer Servers program and implement them in the Enterprise Storage specification. These requirements are intended to encourage proper capacity planning and the procurement of the most efficient equipment for a particular end-use scenario.

- The Standard Information Reporting requirements specify that manufacturers publish a Power and Performance Data Sheet for all ENERGY STAR® qualified products. The data sheet provides information on the energy performance, advanced power saving features, and thermal characteristics of the product in a standard and accessible format. Templates for the Power and Performance Data Sheet for Computer Servers can be found on the ENERGY STAR® website for Computer Servers at www.energystar.gov/NewSpecs. The Data Measurement and Output requirements specify that ENERGY STAR® qualified Storage Products have the ability to measure and self-report power consumption, utilization and input air temperature in an open, accessible format to interface with third-party management software. This information is intended to give data center operators more visibility into the real-time performance of their storage hardware to inform power management decisions.

b. Initial Approach:

Following are excerpts from the Standard Information Reporting and Data Measurement and Output requirements from the Version 1.0 Computer Server specification, with minor modifications for Storage Products.

- ***Standard Information Reporting Requirements***
 - *Partners must provide a standardized Power and Performance Data Sheet with each ENERGY STAR® qualified Storage Product. This information must be posted on the Partner's Web site where information on the qualified model, or qualified configurations, is posted.*
 - *Partners shall also provide links to a detailed power consumption calculator and to information on the power consumption of different system configurations.*

- *Partners are encouraged to use the standard data sheet template, but may also create their own template provided that it is identical in format and design as the referenced template, and has been approved by EPA.*
- *Each Power and Performance Data Sheet must include the following information:*
 - *Model name/number, SKU and/or Configuration ID;*
 - *System characteristics (form factor, max storage capacity, power specifications, etc.);*
 - *System configuration;*
 - *Power data for Idle and full load, estimated kWh/year, link to power calculator (where available);*
 - *Additional power and performance data for at least one benchmark chosen by the Partner;*
 - *Available and enabled power saving features (e.g., power management);*
 - *Information on the power measurement and reporting capabilities of the device; and*
 - *Select thermal information from the ASHRAE thermal report;*

• **Data Measurement and Output Requirements**

- *All Storage Products must have the ability to provide data on request over network connections on input power consumption in watts, inlet air temperature, and utilization.*
- *A service processor, embedded power or thermal meter (or other out-of-band technology shipped with the device), or preinstalled operating system may be used to collect and disseminate data.*
- *Data must be made available in a published or user accessible format so as to be readable by third-party, non-proprietary management systems.*
- *When an open and universally available standard protocol becomes available to report and collect data, manufacturers should incorporate the universal standard into their systems.*
- *Measurement Accuracy:*
 - *Input Power: $\pm 5\%$ or ± 5 watts accuracy, whichever is greater.*
 - *Utilization measurements: TBD*
 - *Inlet air temperature measurements: $\pm 3^\circ$ C.*
- *Sampling Requirements: Data must be provided as a rolling average with a period of less than 30 seconds.*

c. Questions for Discussion:

1. What, if any, characteristics of the above requirements may not be appropriate for Storage Equipment? What data reporting capabilities do typical Storage Equipment currently possess? Are there industry trends towards the inclusion of these reporting capabilities?

2. What additional information specific to Enterprise Storage might be required for the Power and Performance Data Sheet?

StorageIO comment:

A standardized format including peak or circuit breaker power in Amps and KVA along with typical or nominal energy usage as well as expected cooling requirements in Btus. This should be readily accessible both in human readable format on a label or tag on a device as well as in product documentation, on-line as well as via sense code inquiry status information from a storage system.

3. How is utilization defined for Storage Equipment? What utilization information would be helpful to managers for procuring the proper equipment and integrating those systems into their data centers?

4. Are there any industry accepted communication protocols available for measuring utilization, temperature, etc?

StorageIO comment:

SNMP MIBs and vendor specific APIs.

End of EPA Energy Star® for Enterprise Storage Content

Conclusion

IT organizations and data centers are realizing that in addition to conserving power and avoiding unneeded power usage, addressing time-sensitive applications with performance enhancements can lead to energy efficiency. Various techniques and existing technologies can be leveraged to either reduce or support growth with current power and cooling capabilities, as well as with supplemental capabilities. There are also several new and emerging technologies to be aware of and consider. An Energy Star® for enterprise or data center storage makes sense and is applicable to help customers, as well as manufactures and VARs align the most applicable and effective storage related technology to the task and usage case at hand.

These range from more energy-efficient power supplies, storage systems, and disk drive components to performance enhancements to get more work done and store more data per unit of energy consumed in a given footprint. Improvements are also being made in measuring and reporting tools to provide timely feedback on energy usage and enable tuning of cooling resources.

Action and takeaway points include the following.

- Energy avoidance can be accomplished by powering down storage.
- Energy efficiency can be accomplished by using tiered storage to meet different needs.
- Measure and compare storage based on idle and active workload conditions.
- Storage efficiency metrics include IOPS or bandwidth per watt for active data.
- Storage capacity per watt per footprint and cost are a metrics for inactive data.

Bottom line, An Energy Star® for data center storage would be a good thing assuming that such a specification includes and address the various types of storage from on-line activity doing work while storing information, to off-line at rest while storing information and the various modes in-between. The approach needs to factor in both hardware and software functionality without being too overall encompassing so as to be useless or unable to be implemented. Above all, the specification must be of benefit to the buying organizations providing insight into both efficiency while doing work, while storing information and providing a return on investment.

About StorageIO and the Author

Greg Schulz is founder of StorageIO, an IT industry analyst consultancy firm and author of the books *The Green and Virtual Data Center* (CRC) and *Resilient Storage Network* (Elsevier). Learn more at www.storageio.com, www.storageioblog.com and on twitter @storageio.

All trademarks are the property of their respective companies and owners. The Server and StorageIO (StorageIO) Group makes no expressed or implied warranties in this document relating to the use or operation of the products and techniques described herein. StorageIO in no event shall be liable for any indirect, inconsequential, special, incidental or other damages arising out of or associated with any aspect of this document, its use, reliance upon the information, recommendations, or inadvertent errors contained herein. Information, opinions and recommendations made by StorageIO are based upon public information believed to be accurate, reliable, and subject to change.