



Storage Efficiency: The Key to Green Storage Operation

(or, It Ain't Easy Bein' Green Storage)

With rising fuel prices and growing environmental concerns, organizations today are more interested than ever in operating efficiently, and that includes making data centers as "green" as possible. While energy consumption of overall data center operations and server-related equipment is generally understood, the story is quite different when it comes to storage systems.

This paper offers practical advice on how to measure storage system efficiency and compares the efficiency numbers of several common storage solutions.

Rob Peglar
Vice President, Technology,
Xiotech Corporation

- **Data Center Energy Efficiency Ratio (DC-EER):** Combination of the above two metrics.²
- **Power Usage Effectiveness (PUE):** Total data center electrical load divided by the IT electrical load; the reciprocal of this is also known as *Data Center Efficiency (DCE)*.³

In addition, the Environmental Protection Agency (EPA) is conducting a nationwide survey of data centers to determine an infrastructure rating technique, similar to the well-known *ENERGY STAR* model used for certain household and industrial appliances.⁴

Recently There Has Been

a surge (no pun intended) of articles and papers published around the topic of “green” IT techniques, in particular surrounding overall data center energy consumption. Clearly, reducing energy consumption while maintaining or improving current service levels (e.g., as defined in SLAs between an organization’s IT function and business functions) is a noble goal. Also, there has been a sharp increase in attention on green data centers from the public sphere, including governmental and advisory/industry bodies.¹

Techniques that apply to overall data center energy consumption and compute-centric functions are well understood. The last hurdle to becoming increasingly green is related to data storage, commonly implemented as disk arrays (often in storage area networks or SANs) and robotic tape libraries.

Energy consumption of storage devices is often misunderstood at best or completely ignored at worst. However, not unlike their compute cousins, storage arrays are best measured using efficiency metrics. Just

Energy consumption of storage devices is often misunderstood at best or completely ignored at worst.

Given this laudable goal, how is one to measure energy consumption and efficiency within a data center and, in particular, in that data center’s use of compute and storage resources—most commonly implemented as servers and/or mainframes?

Some of the more recent overall or compute-centric metrics include the following:

- **Site Infrastructure Energy Efficiency Ratio (SI-EER):** Measures the kilowatts required from the outside (local) utility to deliver one kilowatt of reliable power to the IT equipment.
- **Information Technology Energy Efficiency Ratio (IT-EER):** Measures the compute performance of IT equipment (in particular, compute servers) per embedded watt of required power.

as the IT-EER above measures compute-per-watt, so storage arrays should have standardized metrics, such as capacity per watt and I/O per watt.

However, measuring consumption of energy by data at rest (i.e., data that is being stored but not accessed) is less than meaningful in comparison to metrics such as IT-EER, which assume some compute workload is occurring.

One must measure storage “greenness” using a complete efficiency model, i.e.:

- Energy consumed by storage devices (kW-hour) for data at rest, over a given time period (fixed or, more broadly, the useful lifetime of the data).

- Energy consumed for data being accessed at typical rates (IOPS, GB/sec) as part of a useful compute workload (i.e., one that is normal and customary for a particular data center).
- Energy consumed for data being accessed at maximum rates as part of a synthetic workload (i.e., one generated by synthetic means other than a useful compute workload).

These metrics must also include economic (cost) efficiency, i.e., dollars per metric. Obviously, this calculation is foremost in the minds of many CFOs

popularization of RAID techniques. Although instantaneous power required (watts) is well known for individual disk drives, drive enclosures, and RAID controllers, little or no study has been put into the actual efficiency—under workload—of these devices on a systematic basis.

ISE technology, on the other hand, tightly integrates drive, enclosure, and controller software; pools storage (disk drives) into a singular element that is 20 to 40 times the capacity of a typical disk drive; and distributes much of the processing power and cache into the enclosure. The unique ISE design also

With the advent of ISE technology in 2008 from Xiotech, the world has now entered the era of highly efficient storage.

and CIOs, given the ever-present budgetary impact of data center operations.

significantly reduces vibration and improves cooling. The result is a highly efficient storage system that also provides unmatched reliability and performance.⁵

The Era of Highly Efficient Storage

With the advent of *Intelligent Storage Element (ISE™)* technology in 2008 from Xiotech, the world has now entered the era of highly efficient storage. Previous to ISE, storage arrays were typically designed around variations of a single theme: I/O controllers, often of proprietary design, connected to unintelligent storage (disk) enclosures.

This theme began decades ago with the introduction of such entities as the channel-attached I/O controller, SCSI protocol, and

Efficiency metrics based on common, standardized workloads verify the effectiveness of the ISE design. Two such workloads from the Storage Performance Council (SPC) are examined below:

- **SPC Benchmark 1™ (SPC-1):** Consists of a set of I/O operations designed to demonstrate the performance of a storage subsystem while performing the typical functions of a business-critical application. SPC-1 represents a segment of applications characterized by predominately random I/O operations, and requiring both queries and update operations (e.g., OLTP, database, or mail server applications).⁶
- **SPC Benchmark 2™ (SPC-2):** Represents a segment of applications characterized by predominantly sequential I/O operations, and requiring queries as well as update operations (e.g., large file processing, large database query, and video on demand), including a composite workload.

SPC-1 metrics include the following:

- **SPC-1 IOPS™ (per Disk):** Measured I/O per second (IOPS) divided by the count of configured disk drives.
- **SPC-1 Price-Performance – \$ per IOPS:** Published price of the system divided by the measured IOPS.

SPC-2 metrics include the following:

- **SPC-2 MB/sec™:** Measured throughput of the three specific workloads and the composite workload.
- **SPC-2 Price-Performance – \$ per MB/sec:** Published price of the system divided by the measured throughput.

unit cost of the storage resource necessary to sustain the workload.

SPC Results for ISE

Emprise 5000 from Xiotech is a compact, self-enclosed virtualized storage system, designed around a single ISE (one enclosure per Emprise 5000 system). This solution was tested in 2008, and its SPC-1 and SPC-2 metrics have been published⁷:

- 294 SPC-1 IOPS per disk
- \$3.53 per SPC-1 IOPS
- 645.60 SPC-2 MB/sec
- \$32.25 per SPC-2 MB/sec

A rating of power efficiency should also include metrics such as watts and kilowatt-\$ per IOPS and MB/sec.

A rating of power efficiency should also include the following metrics, (lower is more efficient):

- Watts (W) per IOPS per disk
- W per MB/sec per disk
- Kilowatt (kW)-\$ per IOPS
- kW-\$ per MB/sec

Note: If the cost per kW-hour of energy is known, the latter two figures can be measured in pure cost terms, i.e., \$ per kWh per \$/MB/sec, or the cost of providing the energy resource per the

Likewise, the maximum instantaneous power (watts) required by the system has also been published:

- 600 watts per Emprise 5000 system

Therefore, the efficiency metrics for Emprise 5000 are as follows:

- 2.04 watts per enclosure per IOPS per disk (600/1/294)
- 2.11 kW-\$ per enclosure per IOPS (.600*1*\$3.53)
- 0.93 watts per enclosure per MB/sec (600/1/645.60)
- 19.35 kW-\$ per enclosure per MB/sec (.600*1*\$32.25)

The I/O efficiency ratings of Xiotech's Emprise 5000 system compared to the two popular modular arrays are significant:

- For I/O efficiency (W/enclosure/IOPS/disk), Xiotech's Emprise 5000 earned a rating of 2.04 versus 2.88 (EMC) and 2.63 (HP).
- For economic efficiency (kW-\$/IOPS/enclosure), the Emprise 5000 rating was 2.11 versus 9.67 (EMC) or 10.57 (HP)—representing a four- to five-fold improvement.

Comparison to Nonefficient Storage

There are a plethora of inefficient storage arrays available today; as stated above, most of the inefficiency lies in their design. For comparative purposes, published SPC results for a sampling of popular storage arrays are shown in the table below, with their corresponding efficiency metrics.

The I/O efficiency ratings of Xiotech's Emprise 5000 system compared to the two popular monolithic arrays are even more significant:

- For I/O efficiency, the Emprise 5000 rating of 2.04 is nearly 50 percent better than HDS's 2.92 rating and nearly a seven-fold improvement compared with the 3PAR array's 13.95.
- For economic efficiency, Emprise 5000 earned a 2.11 versus 21.49 (3PAR) or 10.03 (HDS)—representing a nearly five- and ten-fold improvement respectively.

Metric	Popular Modular Arrays		Popular Monolithic Arrays		ISE-Based Array
	EMC CX3-40 ⁸	HP EVA 2C12D ⁸ (mirrored cache)	3PAR Inserv S800 ⁸	HDS USP-V ⁸	Xiotech Emprise 5000 ⁷
Storage Controllers	2 (SPE)	2	8	5 (integrated frames)	1 (integrated ISE)
Drive Enclosures	11 (DAEs)	12	24 (40 drives each, 960 total disks)	N/A (1024 disks, 64 total enclosures)	N/A (1 ISE enclosure with 2 DataPacs)
SPC-1 IOPS™ (per Disk)	162	168	104	195	294
SPC-1 Price Performance (\$ per SPC-1 IOPS)	\$20.72	\$23.88	\$14.81	\$17.61	\$3.53
SPC-2 MB/sec™	Not published	Not published	Not published	Not published	645.60
SPC-2 Price Performance (\$ per SPC-1 MB/sec)	Not published	Not published	Not published	Not published	\$32.25
Total Watts (# controllers*watts per controller) + (# disk enclosures*watts per disk enclosure)	5,135 ((2*230)+(11*425))	5,312 ((2*400)+(12*376))	34,824 ((8*924)+(24*1,143))	36,480 (7,488*5)	600 (600*1)
Watts per Enclosure per IOPS per Disk (watts/enclosures/IOPS/disks)	2.88 (5,135/11/162)	2.63 (5,312/12/168)	13.95 (34,824/24/104)	2.92 (36,480/64/195)	2.04 (600/1/294)
kW-\$ per Enclosure per IOPS (kilowatts*\$ per IOPS/enclosures)	9.67 (5.135*\$20.72/11)	10.57 (5.312*\$23.88/12)	21.49 (34,824*\$14.81/24)	10.03 (36,480*\$17.61/64)	2.11 (.600*\$3.53/1)
Watts per Enclosure per MB/sec per Disk	Not published	Not published	Not published	Not published	0.93
kW - \$ per Enclosure per MB/sec	Not published	Not published	Not published	Not published	19.35

Conclusion

Energy consumption and efficiency within a data center and effective use of compute and storage resources is fast becoming one of the most critical areas in IT. As such, what began decades ago with such entities as the channel-attached I/O controller, SCSI protocol, and RAID techniques has now been put to further scrutiny, and found wanting.

In response to the challenge, Xiotech has ended the era of inefficient storage with its Intelligent Storage Element (ISE) technology that tightly integrates drive, enclosure, and controller software; pools storage (disk drives) into a singular element that is 20 to 40 times the capacity of a typical disk drive; and distributes much of the processing power and cache into the enclosure. The unique ISE design also significantly reduces vibration and improves cooling. The result is a highly efficient storage system that also provides unmatched reliability and performance.

References

- ¹ Brill, Kenneth. 2007. *Opinion of the Institute: EPA Report Should Spur Industry-Wide Green Data Center Movement*. Uptime Institute (www.uptimeinstitute.org)
- ² Brill, Kenneth. 2006. *High Density Computing: The Path Forward 2006*. Uptime Institute (www.uptimeinstitute.org).
- ³ The Green Grid. 2007. *Green Grid Metrics: Describing Data Center Power Efficiency*. (www.thegreengrid.org).
- ⁴ U.S. Environmental Protection Agency (EPA) and U.S. Department of Energy (DOE). 2008. Data Collection Initiative to Develop an ENERGY STAR Rating for Data Centers (www.energystar.gov and www.energystar.gov/index.cfm?c=prod_development.server_efficiency#data_collection)
- ⁵ Xiotech Corporation. 2008. *ISE, The New Foundation of Data Storage*, www.isepower.com
- ⁶ Storage Performance Council. 2006. SPC Benchmark 1 (SPC-1) Specification (www.storageperformance.org/specs).
- ⁷ Storage Performance Council. 2008. SPC-1 Benchmark Results Report #A00064 – Xiotech Emprise 5000 (www.storageperformance.org/results/benchmark_results_spc1#a00064); SPC-2 Benchmark Results Report #B600031 – Xiotech Emprise 5000 (www.storageperformance.org/results/benchmark_results_spc2#b00031).
- ⁸ Storage Performance Council. SPC-1 Benchmark Results (http://www.storageperformance.org/results/benchmark_results_spc1).
 - Report #A00059 – EMC CLARiiON CX3 Model 40 (2008)
 - Report #A00007 – hp StorageWorks Enterprise Virtual Array Model 2C12D (2003)
 - Report #A00027 – 3PAR InServ S800 X-Series (2004)
 - Report #A00054 – Hitachi Universal Storage Platform V (2007)

Important Notice

By accepting, reviewing, or using this document, you, as the recipient, agree to be bound by the terms of this notice. Information in this document is subject to change without notice. Names and data used in examples are fictitious unless otherwise noted. No association with any real company, organization, product, domain name, email address, logo, person, place, or event is intended or should be inferred. Xiotech and/or its licensors may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights covering the subject matter in this document. The configuration(s) tested or described in this document may not be the only available solution(s). This document is not intended (nor may it be construed) as an endorsement of any

product(s) tested, as a determination of product quality or correctness, or as assurance of compliance with any legal, regulatory, or other requirements. This document provides no warranty of any kind, including, but not limited to, any express, statutory, or implied warranties, whether this document is considered alone or in addition to any product warranty (limited warranties for Xiotech products are stated in separate documentation accompanying or relating to each product). No direct or indirect damages or any remedy of any kind shall be recoverable by the recipient or any third party for any claim or loss in any way arising or alleged to arise from or as a result of this document, whether considered alone or in addition to any other claim.

Xiotech, Magnitude 3D, Magnitude, and TimeScale are registered trademarks of Xiotech Corporation. Emprise, ISE, GeoRAID, Dimensional Storage Cluster, and DataScale are trademarks of Xiotech Corporation. SPC Benchmark 1, SPC Benchmark 2, SPC-1 IOPS, and SPC-2 MB/sec are trademarks of the Storage Performance Council. All other trademarks or service marks are the property of their respective owners. The reproduction of any trademark or service mark, registered or otherwise, belonging to any third party appears in this document for the purposes of identification of the goods or services of that third party only. No license of any kind is provided by Xiotech and/or its licensors to any party because of this document. No right or title in any mark or other proprietary interest is transferred by Xiotech and/or its licensors to any party because of this document.

© 2008 Xiotech Corporation. All rights reserved. PN 070606-0708



Corporate Headquarters:

6455 Flying Cloud Drive
Eden Prairie, MN 55344-3305
1.866.472.6764
www.xiotech.com