Industry Trends and Technology Perspective Solutions Brief

MAID 2.0: Energy Savings without Performance Compromises

Energy Savings for Secondary and Near-line Storage Systems

By Greg Schulz

Founder and Senior Analyst, the StorageIO Group

January 2nd, 2008

Not all Green Storage is created equal. There are business benefits to aligning the most energy efficient storage to meet different data and application requirements. All storage vendors are claiming some form of "green" or some power, cooling and floor space savings. This industry trends and perspectives brief looks at a category of storage to support secondary and near-line or off-line data needs, comparing performance with energy savings, to enable you to be the judge of who is the "greenest" of them all.

Introduction

The advent of cheaper and cheaper compute power has brought with it a tremendous growth in data storage to feed these computers and provide the information organizations need. Until recently, energy efficiency of servers and storage has been less of a concern to IT organizations however this is changing especially as the price of electricity has increased and demand is outpacing the supply of electricity generation and transmission (G&T).

The combination of growing demand for electricity by data centers, density of power usage per square foot, rising energy costs, strained electricity G&T infrastructure and environmental awareness prompted the passage of United States public law 109-431 in 2006. Public law 109-431 instructed the United States Environmental Protection Agency (EPA), part of the Department of Energy (DoE), to report to Congress on the state of IT data centers energy usage in the United States.

In the August 2007 EPA report[1] to Congress, findings included that IT data centers or what is being termed information factories consumed about 61 billion kilowatt hours (kWh[2]) of electricity in 2006 at an approximate cost of about $4.5 billion dollars. Also reported is that IT data centers on average consume 15-20 times or more energy per square foot than compared to a typical office building. Without changes in electricity consumption and improved efficiency, the EPA is estimating that IT data centers power consumption will exceed 100 billion kWh by 2011 further stressing an already strained electrical power G&T infrastructure and increasing previously high energy prices.

Background and issues

While there is a growing environmental impact awareness ("The Greening if IT"), the StorageIO Group, through research and regular discussions with IT personal, has found the more pressing problem facing many IT data centers (approximately 85-90%) are growing bottlenecks and approaching ceilings on available power, cooling and floor space. As IT data centers address their power, cooling and floor space challenges with improved energy efficiency and effectiveness, three main benefits will be realized: helping the environment, reduce power and cooling costs, and enabling sustained application growth to support evolving business information needs.

The major power draw based upon StorageIO research for commonly deployed storage systems are spinning hard disk drives (HDDs) and their enclosures which account for on average 66-75% — controllers account for the balance of electrical power consumption. While it is debatable how much energy in a typical data center is actually consumed by storage (internal to servers and external) along with how much data is active or inactive, what is generally agreed upon is that spinning HDDs require power and one of the different approaches (see www.greendatastorage.com) for addressing this energy usage is to power down inactive HDDs. The challenge, however, is how to do this without adversely impacting application response time or data availability.

The basic premise of MAID (Massive or Monolithic Array of Idle Disks)[3] is to power down HDDs when not in use, similar to a laptop PC or workstation. Unlike a laptop or PC, powering down HDDs in an enterprise storage system presents challenges including negative impact to application performance, loss of availability, potential for HDD failure and damage due to power surges when HDDs are spun back up.

---

1 An analysis by the StorageIO group of the 2007 EPA report to Congress along with links to the full EPA report is located at http://www.storageio.com/Reports/StorageIO_WP_EPA_Report_Aug1407.pdf
2 One kWh is 1,000 watts of energy or the energy usage of a device consuming 1,000 watts per hour
3 Read more about "The Many Faces of MAID" industry trends and perspectives brief also at www.storageio.com

Unlike the HDDs used in high performance storage systems that are designed and optimized to always be spinning, laptop PC HDDs are designed to be powered up and down to conserve power or maximize battery life. There is a debate in the industry around which type of HDDs to use and the merits of powering down HDDs on a regular basis. Consequently, some MAID vendors go to great lengths to explain how and why their implementations are safe and risk averse from power-down effects and potential for premature HDD failure.

Meet MAID 2.0 and Intelligent Power Management (IPM)

MAID enabled devices are evolving from first generation MAID 1.0 where HDDs are either on or off, with associated performance penalties, to a second generation MAID, or MAID 2.0, implementing intelligent power management (IPM). MAID 2.0 leverages IPM to align storage performance and energy consumption to match the applicable level of service being supported. With IPM enabled MAID 2.0, instead of a HDDs being on or off, there can be multiple power-saving modes to balance energy savings with performance and availability needs.

For example, a storage system can implement MAID Level-0 (no real energy savings, no impact to performance) for active data. For less active data, an administrator can choose a user selectable setting to transition the storage system to MAID Level-1 were power is reduced by retracting HDD read/write heads. For even better power savings, a HDD or RAID group or some other unit of storage granularity could be put into a MAID Level-2 mode where the speed of the drive platters is reduced. For the best power savings, the administrator could select MAID Level-3 where a HDD or the RAID group is completely powered down or put into a suspended standby sleep mode.

A current example of MAID 2.0 is the SATABeast™ from Nexsan which has implemented MAID 2.0 and intelligent power management with AutoMAID™ technology on their standard SATABeast™ storage systems. This system has multiple power saving modes to intelligently align different levels of energy usage to different quality of service (QoS) level needs. For example, a Nexsan SATABeast™ can be configured to support high performance data movement (reading and writing) during peak periods (MAID Level 0) with the user selectable options of MAID Level-1, Level-2 and Level-3 set for various periods or tiers of inactivity.

What this means is after a period of inactivity (user selectable), the storage system could automatically take the first power saving step of unloading the read/write heads, reducing the power consumption to MAID Level-1 and 21% or more power savings. But, the HDDs can very rapidly return to active I/O duty without incurring a time delay or power spike associated with HDD spin up found in first generation MAID implementations. Additional energy savings can be achieved after a longer inactive period by putting the HDDs into MAID Level-2, 38% or more energy savings. At this level, HDDs can respond to I/O requests in seconds without a power spike. Further energy savings can be realized with increased inactivity by putting the HDDs into MAID Level-3, increasing energy savings to over 56%. MAID Level-3 also eliminates the power surge associated with first generation MAID by leveraging intelligent power management and sequenced HDD spin-up when re-activating in-active HDDs.

Secondary Storage System Power and Performance Comparison

To compare and contrast on-line secondary, near-line and off-line storage solutions, a raw (unformatted) and native (no compression or de-duplication) capacity of 1.3PByte is used in Chart-1. Chart-1 graphically shows a synopsis of the information found in Table-1 comparing annual energy costs for MAID Level-0 (no energy savings) and MAID Level-3 (idle or standby) mode if supported. For example, the EMC CLARiiON™ and NetApp R200 do not support an energy savings mode, HDS AMS1000 supports MAID 1.0 (HDDs in a RAID group are on or off), and Nexsan SATABeast™ supports MAID 2.0 (variable energy savings and performance capabilities).

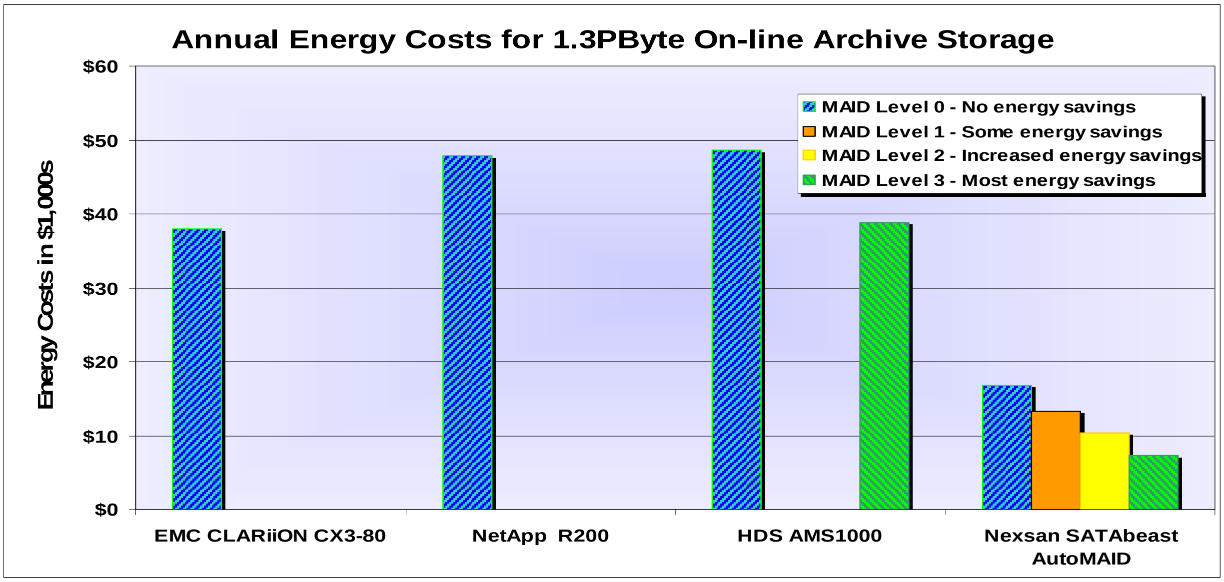

Chart-1: Power Costs for Disk Based On-line Secondary, Near-line and Off-line Archive Storage

Table-1 shows more detail on the various metrics for on-line secondary, near-line and off-line storage including disk and tape based solutions supporting 1.3PByte of raw (unformatted) storage.

The assumed kWh power cost (cents) = 0.08 <-- Enter energy cost here in cents per kWh

	Non MAID Capable		MAID 1.0 or 2.0 Capable Storage			Tape Based Storage		CAS Object Storage	
Configuration for 1.3PByte raw storage capacity	EMC CLARiiON CX3-80	NetApp R200	HDS AMS1000	Nexsan SATAbeast AutoMAID	Copan Revolution 300 MAID	Sun SL8500 LTO4 Tape	HP ESL 720e LTO4 Tape	EMC Centera 4LP	Nexsan Assureon
Number of storage systems required	4	8	4	32	2	1	2	16	32
Floor space footprint (cabinets or racks)	12	8	8	4	2	1	2	16	4
Total number disk or tape drives	1,800	2,688	1,680	1,344	1,792	32	32	1,792	1,344
Total raw non formatted capacity (Tbytes)	1,350	1,344	1,296	1,344	1,344	1,158	1,072	1,344	1,344
Total hourly bandwidth performance (Tbyte/hr)	16.14	7.20	21.89	80.64	10.41	13.82	13.82		
MAID Level 0 = Normal active disks (no saving) kWh									
MAID Level 0 - Annual Power (no cooling) kWh	474,792	598,764	606,893	210,240	128,299	18,273	20,709	560,640	236,520
MAID Level 0 - Annual energy costs ($1,000s)	$38	$48	$49	$17	$10	$1	$2	$45	$19
MAID Level 1 = Park disk read/write heads									
MAID Level 1 - Annual Power (no cooling) kWh	-	-	-	165,389	-	-	-	-	191,669
MAID Level 1 - Annual energy costs ($1,000s)	-	-	-	$13	-	-	-	-	$15
MAID Level 2 = Reduce disk RPM speed									
MAID Level 2 - Annual Power (no cooling) kWh	-	-	-	128,667	-	-	-	-	154,947
MAID Level 2 - Annual energy costs ($1,000s)	-	-	-	$10	-	-	-	-	$12
MAID Level 3 = Standby (sleep) power									
MAID Level 3 - Annual Power (no cooling) kWh	-	-	485,514	91,384	62,652	12,264	9,496	-	117,664
MAID Level 3 - Annual energy costs ($1,000s)	-	-	$39	$7	$5	$1	$1	-	$9
MAID Level-0 - kWh per Tbyte capacity (raw space)	0.040	0.051	0.053	0.018	0.011	0.002	0.002	0.048	0.020
MAID Level-1 - kWh per Tbyte capacity (raw space)	-	-	-	0.014	-	-	-	-	0.016
MAID Level-2 - kWh per Tbyte capacity (raw space)	-	-	-	0.011	-	-	-	-	0.013
MAID Level-3 - kWh per Tbyte capacity (raw space)	-	-	0.043	0.008	0.005	0.001	0.001	-	0.010
MAID Level-0 kWh used per Tbyte/hr bandwidth	3.358	9.493	3.165	0.298	1.406	0.151	0.171	-	-

Table-1: On-line Secondary, Near-line and Off-line Archive storage for 1.3PByte of Storage Capacity[4]

[4] Annualized energy costs are based on 8 cents per kWh using publicly available vendor specifications

Chart-2 below displays the performance in TByte per hour bandwidth of on-line secondary and near-line storage systems showing that a system with MAID 2.0 (Nexsan SATAbeast™ with AutoMAID™) has improved power savings with little or no impact to storage and application performance.

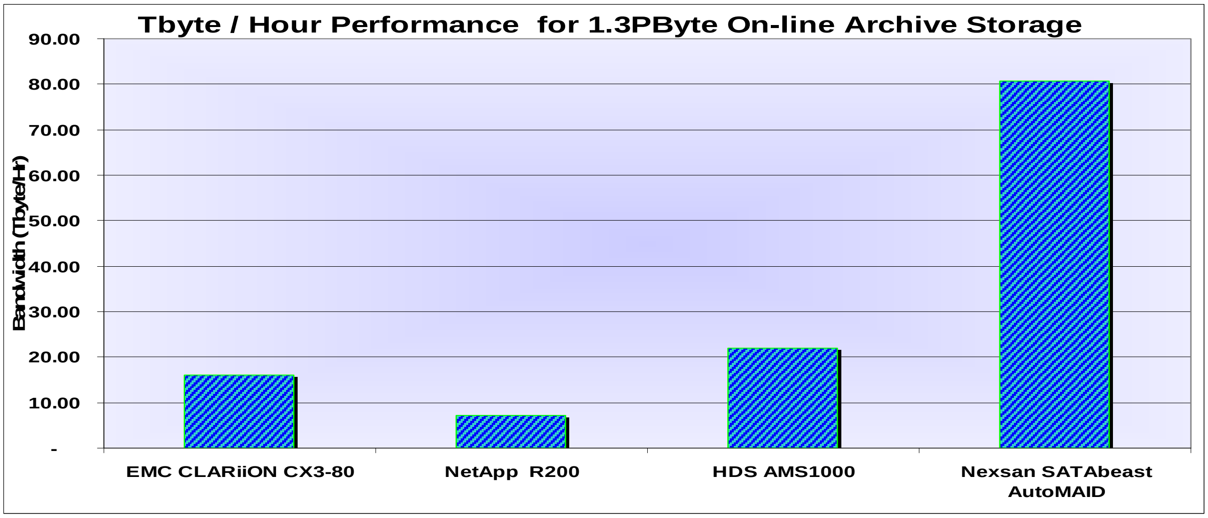

Chart-2: Performance for Disk Based On-line Archive Storage

Comparing Chart-1 (power consumption) and Chart-2 (bandwidth performance) for on-line disk-based secondary and near-line storage systems, the Nexsan MAID 2.0 enabled storage system has 3.5X to 11X the performance with 65% improved power savings ($32,000) with MAID Level-0 compared to other solutions in Table-1. At MAID Level-3, Nexsan exhibits an 81% ($32,000) improvement compared to the HDS AMS 1000 with Power Savings Storage feature enabled.

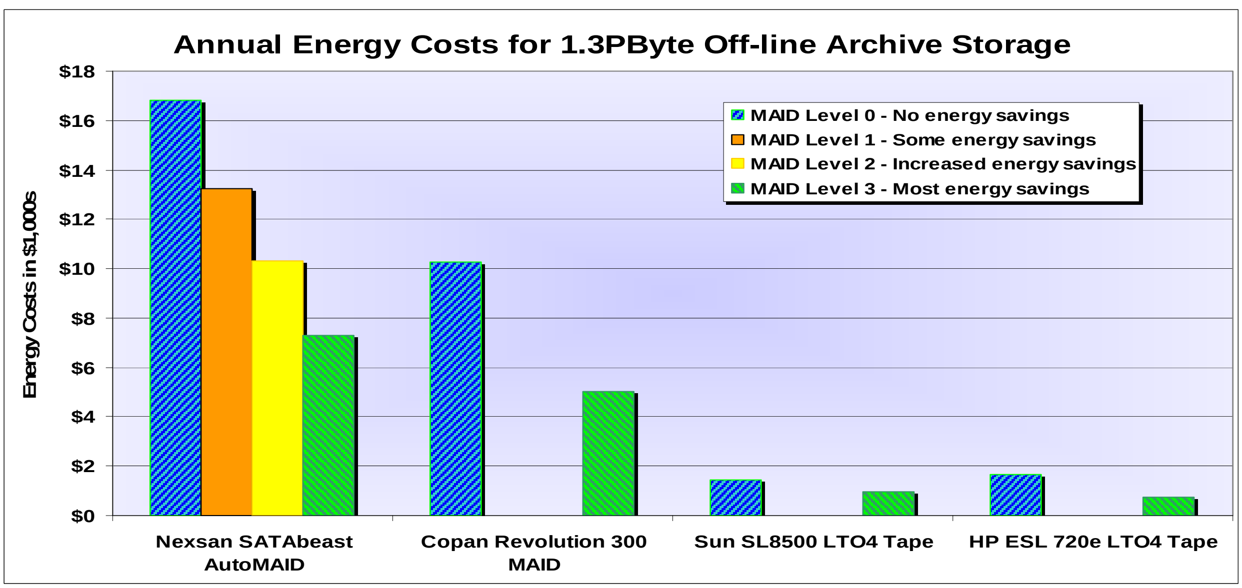

Chart-3: Power Costs for Off-line Archive Disk and Tape Based Storage

Comparing off-line disk and tape based secondary storage (Charts-3 and Chart-4), solutions like the Sun LTO4 Tape based system has better power savings versus the best on-line or off-line disk-based solutions. From a performance view (Chart-4), MAID 2.0 enabled storage is best for near-line or off-line bandwidth solutions with tape being a fit for some off-line applications to meet recovery time objectives.

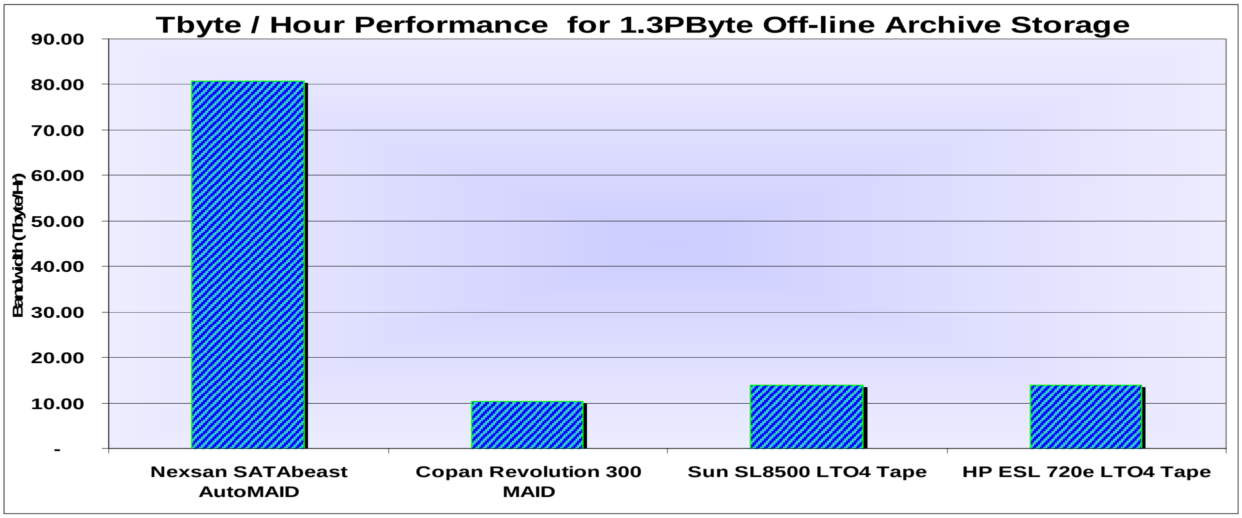

Chart-4: Performance comparison between Off-line Archive Disk and Tape Based Storage

For applications and environments that need the lowest energy consuming storage and where response time or application performance are not required, for example off-line storage, magnetic tape remains a good option and companion to HDD based on-line and near-line storage systems.

First generation MAID solutions, like the Copan Revolution, exhibit better energy consumption compared to traditional HDD on-line primary high performance storage systems however at the expense of application performance. Given that Copan spins a maximum of 25% of HDDs at any given time, application performance is compromised for example during large scale data restoration when bandwidth is needed to meet recovery time objectives. Similarly, when compared on an energy savings and performance basis, MAID 2.0 enabled HDD based storage systems shine as does magnetic tape demonstrating viability for long term off-line storage needs.

Taking archive or near-line storage to the next level with Content Addressable Storage (CAS), the EMC Centera™ appears very energy efficient at 125 watts per G4 LP (low power) node with 3TByte raw storage capacity. When, however, configured into a 32[5] node cluster, the Centera™ consumes the same, if not more, watts of energy than a comparably configured EMC CLARiiON™. Depending on the data protection mode, for example mirroring or parity, an EMC CLARiiON™ with the same number and type of HDDs may have even better power and performance characteristics when compared to a Centera™ on a raw storage basis. A more apples-to-apples comparison would be EMC Centera™ compared to the Nexsan Assureon™ object- based content addressable storage solutions. This CAS system is based on MAID 2.0 storage and has over 2X advantage in power savings at MAID Level-0. Further, it has MAID Levels 1, 2, and 3 where the Centera™ lacks similar energy saving modes. Similarly, the more applicable EMC CLARiiON™ comparison would be against the HDS AMS1000, NetApp R200 and Nexsan SATABeast™.

A MAID 2.0 system like the Nexsan SATABeast™ with AutoMAID™ provides variable energy savings and performance. For example, assuming that the HDDs have been configured accordingly, after 20

[5] For a 32 node Centera™ cluster, 28 nodes are configured as storage nodes with 4 nodes as access nodes

minutes of no activity, HDDs can achieve a 21% energy savings by disabling read/write heads (MAID Level-1), then after another 10 minutes slow the HDDs down (MAID Level-2) reducing power by 38% and then after another 5 minutes put the HDDs into MAID Level-3 mode (idle standby) for an energy savings of over 55%. This implementation of MAID 2.0 allows for rapid spin up and return to active service of HDDs from MAID Level-3 without completely powering them down avoiding excessive drive wear, tear and pre-mature failures. When in suspended or sleep mode, HDDs automatically wake up for periodic surface scan and data integrity checks; for example once a week for a few hours determined by user configuration. For active systems, AutoMAID™ will not be invoked until drives go idle based on user selected criteria based on the time and level of power savings desired to meet application QoS levels.

Calculating Your Own Annual Energy Savings

You can calculate your own energy savings by using the values in Table-2. First, select the energy costs for your location and the number of kWh required to power your device for an hour. For example, if you have a storage device that consumes 100 kWh per hour of power and the average energy costs for your location is 8 cents per kWh; your annual energy cost would be $70,100 dollars. Then, select the AutoMAID™ level you desire, for this example MAID Level-3 for 56% savings. Apply this savings factor to your annual energy costs in the first step and determine your first year's savings, 56% x $70,100 = $39,256 in power savings. Do not forget to include energy savings from reduced cooling costs in your calculation. This will increase your energy savings beyond the $39,256 used in this example.

Hourly Power Usage	5 cents per kWh	8 cents per kWh	10 cents per kWh	12 cents per kWh	15 cents per kWh	20 cents per kWh
1 kWh[6]	$438	$701	$806	$1,051	$1,314	$1,752
10 kWh	$4,380	$7,010	$8,060	$10,510	$13,140	$17,520
100 kWh	$43,800	$70,100	$80,600	$105,100	$131,400	$175,200

Table-2: Annual costs for various levels of energy consumption

Tips and general comments

A storage system with MAID technology needs to be able to adapt to changing workloads, workflows and dynamic application service level requirements. If your applications and data requirements are for long term use, meaning reading of data, look for storage density (capacity), performance (how much data can be accessed, backed up or recovered in a given amount of time), energy efficiency (reduced power and cooling costs or enable more performance), on-going data integrity checks and safeguards as features enabled by MAID technology enabled storage.

General comments and recommendations include:

- ✓ Look beyond simple yes or no support, asking vendors how, when and where they support MAID
- ✓ For backup and recovery, bandwidth is important in order to recovery data in a timely manner
- ✓ Align MAID technology, including MAID level, to the applicable tier and application needs
- ✓ Understand the performance tradeoffs of using MAID as an alternative to traditional disk and tape
- ✓ For additional energy savings, look into leveraging data footprint reduction technologies[7]

---

[6] Annual kWh computed as 365 x 24 x kWh used per storage system

[7] Refer to the StorageIO Group Industry Trends and Perspectives report "The Business Benefits of Data Footprint Reduction" at www.storageio.com and www.greendatastorage.com

Closing comments

Not all storage or data applications lend themselves to MAID given performance or other service requirements. However, if you are looking at secondary, on-line archive, or near-line storage, some of the key attributes are high density storage (or capacity optimized), energy efficiency with intelligent power management enabled with MAID 2.0 (reduce power and cooling costs) and affordability.

MAID continues to evolve with many different faces and implementations and additional vendors preparing to support MAID 2.0, intelligent power management, data footprint reduction and other energy saving or storage optimization techniques. As with other storage and management tools, it is important too understand the capabilities of the technologies as well as where and how they can address various issues to minimize surprises and disappointments from deploying the wrong technology or technique when seeking energy savings without application compromises.

About the author

Greg Schulz is founder and senior analyst of the StorageIO Group and author of the book *Resilient Storage Networks — Designing Flexible Scalable Data Infrastructures* (Elsevier Digital Press).

www.storageio.com P.O. Box 2026 Stillwater, MN 55082 651-275-1563 info@storageio.com