Industry Trends and Technology Perspective Brief

"The Many Faces of MAID Storage Technology"

A look at how Massive or Monolithic Array of Idle Disks (MAID) is evolving and being implemented in various storage solutions

By Greg Schulz

Founder and Senior Analyst, the StorageIO Group

December 11th, 2007

There is a growing awareness of the need to more effectively and efficiently utilize available or already constrained power, cooling and floor space resources in IT data centers. For some IT environments the focus is on reducing energy consumption to save budget dollars while for other environments reducing energy consumption means freeing up available energy and cooling capacity to power growth of additional servers, storage and networking devices. There can also be a focus on reducing energy related emissions and green house gases. This paper looks at MAID technology in general along with some current implementations, benefits, caveats and what to consider when evaluating MAID enabled technology to support primary, secondary and off-line data storage requirements.

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

Introduction

Generally speaking, in the U.S. IT data centers accounted for only 1.5% of all consumed electricity in 2006[1] with approximately ½ of that (.75% of total consumed power) being used for cooling of IT equipment. The balance of the power used by data centers is consumed primarily by servers, followed by storage and networks. In addition to controllers, the major power draws of storage systems are the spinning hard disk drives (HDDs) that are used for storing data.

The StorageIO Group has prepared and will be publishing additional reports pertaining to data center power, cooling, floor space, performance and environmental (or green storage) topics. Sign up for RSS feeds to receive notification of these reports at www.storageio.com.

While it is debatable how much electricity in a typical data center is actually used by storage (internal to servers and external) along with how much data is active or in-active, what is generally agreed upon is that spinning HDDs require power and one of many different approaches (see www.greendatastorage.com) for addressing energy usage is to power down inactive HDDs while avoiding adverse effects to application response time or loss of availability.

Another energy saving approaches is the movement of inactive data to magnetic tape or optical media and emerging removable HDD (RHDD). Energy usage for storage can also be optimized by shifting high performance active data to semiconductor storage devices (SSD) or more energy efficient storage that delivers more performance per watt of energy consumed (transactions per watt, IOPS per watt, bandwidth per watt). The balance of this brief will look at MAID technology with other energy saving technologies discussed in other StorageIO reports and briefs.

Background and issues

Similar to a laptop or PC workstation with energy saving modes, an approach to save on energy consumption by larger storage systems is to power HDDs down when not in use. That is the basic premise of Massive or Monolithic or Misunderstood Array of Idle or Inactive Disks (MAID). Unlike a laptop or PC, simply powering down HDDs when not in use in a typical storage system presents many challenges including negative impact to application performance, loss of availability, potential for HDD failure and power surges when HDDs are spun back up.

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 on power or maximize battery life. There is 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 explaining how and why their implementations are safe and risk averse from powering down different types of HDDs.

MAID evolution and early RAID DejaVu

Monolithic purpose built MAID devices are similar to what we saw about 15-17 years ago with the early first generation RAID implementations such as those from SF2 (South San Francisco Forklift) company whose patents are now owned by EMC. First generation RAID devices were very large and bulky utilizing purpose built components with limited functionality and poor performance. Current generation RAID enabled technologies commonly have multiple RAID levels to align performance and data protection to specific data and application service requirements leveraging technology improvement curves (processor, memory, I/O adapters and HDDs) to boost performance, lower cost and optimize

[1] Learn more about IT data center power, cooling, green and related topics in the StorageIO Group Industry Trends and Perspectives white paper analysis and synopsis of the EPA Report to congress found at www.storageio.com

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

power consumption. First generation MAID devices are similar to first generation RAID devices in that early MAID devices are the pioneers shifting focus from theoretical white paper and white board discussions to technology that is field and customer deployable.

Copan is usually considered the first vendor to ship a purpose built monolithic MAID device and installed for use at a customer environment. Their early installations are used as a benchmark for newer products and technologies. Given that Copan spins a maximum of 25% off HDDs at any given time, large energy savings compared to on-line primary storage can be achieved. This comes, however, at the expense of negative performance compared to primary on-line and high performance storage. When compared to tape, the Copan MAID approach utilizes more power; however, it also enables an individual file to be accessed faster. The Copan device is designed for storing infrequently used data where performance (IOPS or bandwidth) is not a concern and response time needs to be better than accessing a single file from tape. The sweet spot for Copan MAID is applications or data that do not require performance in terms of IOPS or bandwidth and that need faster access to individual files compared to magnetic tape with associated energy savings. However, for sensitive time and large bandwidth operations including DR or full database restore, tape or regular disk can be a better approach.

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 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 storage service and is being implemented in traditional storage systems. With IPM enabled MAID, instead of a HDD being either on or off, there can be multiple power saving modes to balance energy consumption or savings with performance and service level availability needs.

For example, a storage system can implement MAID Level 0 (no energy savings, no impact to performance) for active data and then based on a timer or other user selectable setting, transition to a MAID Level 1 reducing power by retracting HDD read/write heads and slowing the drive down from 7,200 RPM to a lower speed. For even lower power savings, a HDD or RAID group or some other unit of storage granularity could be put into a MAID Level-2 standby mode or further power savings by going into a MAID Level-3 idle or completely powered down mode.

Storage vendors with MAID functionality in addition to Copan include Fujitsu (Eternus™), HDS (AMS™ series), NEC and Nexsan. Some vendors support MAID granularity on a RAID group- for example, HDS on their AMS™ series report potential energy savings of 20% by powering down an entire RAID group when not in use. By implementing at a RAID group granularity HDS is able to power down both SATA and higher performance Fibre Channel HDDs without impacting performance of other HDDs in other RAID groups. Fujitsu on their Eternus™ series of storage systems support enablement of MAID on selected groups of HDDs and storage pools essentially enabling primary on-line high performance storage and secondary low power MAID enabled storage tiers in the same storage systems. Nexsan has implemented what they refer to as AutoMAID™ technology that enables MAID 2.0 functionality on their standard SATABeast™ storage systems with multiple power saving modes to align to different tiers and 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) similar to a laptop PC or workstation. What this means is that after a user selectable period of time of inactivity, take a first power saving step of slowing the HDDs down (MAID Level 1) to achieve a 30% or greater energy saving, yet the HDDs can very rapidly return to active I/O duty without incurring a time delay or power spike associated with

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

spinning up the HDDs. Additional energy savings are achieved after a longer in-active time when HDDs are put into idle (MAID Level 2) mode or MAID Level 3 mode powered down mode.

Another vendor, Hitachi Global Storage Technologies (HGST), not to be confused with Hitachi Data Systems (HDS), is shipping HDDs with multiple power saving modes to implement IPM and enable MAID 2.0 functionality on general purpose storage systems. Other vendors, including NEC, have deployed MAID into their existing storage system with other vendors making announcements or hinting on their future MAID 2.0 capable plans and alternative approaches to managing storage power and cooling issues.

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 of data (reading and accessing), 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 increased 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; ask vendors how, when and where they support MAID
- ✓ 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
- ✓ Look beyond energy savings and factor in cost of acquisition and site prep along with TCO and ROI
- ✓ If using MAID for backup investigate how many TByte/hour or MB/sec the system can support
- ✓ For additional energy savings, look into leveraging data footprint reduction technologies[2]

Closing comments

Given performance or other service requirements, not all storage or data applications lend themselves to MAID. Like RAID, MAID continues to evolve, consisting of many different faces and implementations with 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, understand capabilities of the technologies, where and how they can address various issues to minimize surprises and disappointment from deploying the wrong technology or technique. Learn more and sign up for RSS feeds about approaches, technologies and trends pertaining to data storage performance, power, cooling, floor space and environmental topics including white papers, articles, webcasts, presentations and tips at www.storageio.com and www.greendatastorage.com.

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).

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