

**Category:** Storage Technology.

**Subject:** Defining MAID – “Massive Array of Idle Disk”.

**Author:** Bill Mottram

The concept of MAID technology is the brainchild of a research team from the University of Colorado who hypothesized that MAID (massive array of idle disks) would be a storage structure that would deliver the density of tape with the performance similar to that of disk and with a very small power envelope.<sup>1</sup>

SNIA Definition – “MAID, a storage system comprising of a **massive array of (idle) disk drives** that are powered down individually or in groups when not required.”<sup>2</sup>

The motivation for such an architecture was to deliver a solution that exploited relatively inexpensive SATA disk technology to create a commercially viable, enterprise class, mass storage solution that exhibited much of the access performance and data integrity characteristics of a disk array but with the economics of a tape library. The sweet spot for this technology, and where it will deliver the most benefit, is in the storage and management of persistent data that is, infrequently accessed data (low IOP’s), data that will rarely if ever be changed, but data that is serving applications that need faster access to individual files than magnetic tape can deliver. Adding to the sweet spot characteristics is high energy efficiency.

But not all “MAID” labeled solutions are created equal and the purpose of this brief is to highlight some the key characteristics and provide some wisdom to those considering the value of MAID or are in the MAID acquisition thought process.

**Purpose built architectures:**

The following table contrasts the design principles of traditional, transactional storage with an architecture specifically designed for long term storage of persistent data.<sup>3</sup>

<u>Transactional Array</u>	<u>Persistent Data Array</u>
<ul style="list-style-type: none"><li>• Very low latency response, high IOP’s</li><li>• Read-write-Modify Model</li><li>• Optimized to provide access to data at most if not all times.</li><li>• Exhibits temporal or spatial locality cacheable</li><li>• Optimized to small-grained data access</li><li>• Data retention tends to be short lived.</li><li>• Design principles above drive cost upwards</li></ul>	<ul style="list-style-type: none"><li>• Once created, data is rarely accessed or modified, often immutable</li><li>• Low temporal access locality, caching tends to be a wasted expense.</li><li>• Only a small portion of the data repository is accessed at any one time.</li><li>• Often access is to medium to large grained data</li><li>• Long term data retention</li><li>• Eliminating unnecessary sophistication controls design and production cost.</li></ul>

<sup>1</sup> The Case for Massive Array of Idle Disks (MAID): Dennis Colarelli, Dirk Grunwald and Michael Neufeld, Dept of Computer Science, University of Colorado, Boulder. January 7<sup>th</sup>, 2002.

<sup>2</sup> The Dictionary of Storage Networking Technology, Storage Networking Industry Association (SNIA), 2005/2006.

<sup>3</sup> Persistent Data Storage Architecture: COPAN Systems, September, 2006

## What are the key characteristics of a MAID Solution that drive user benefit?

A MAID system suitable for an enterprise data center application is identified by unique design characteristics that are not noticeable in traditional array design.

1. Very high drive packaging density.
  - *The fact that in a MAID architecture the number of drives that can spin at any one time is limited. This allows extremely dense packaging not possible in conventional architectures. Example, a single COPAN frame can support up to 896 drives while the new EMC Infiniflex 10000 aka Hulk tops out at 300 drives and DataDirect Networks S2A Storage Scalar at 600 drives per rack. Greater storage density delivers floor space savings, elimination of aging technology through consolidation and the opportunity to reallocate more expensive storage to more appropriate use and potentially delaying an expensive purchase.*
2. The number of drives that can spin at any one time is limited and will not exceed 50% of the total number of drives installed.
  - *This is the original definition as presented by the University of Colorado researchers. COPAN currently limits this number to 25%. The key being that not all drives spin at any one time reducing maximum power requirements, reduced heat generation which in turn reduces the necessary cooling infrastructure and aides in the elimination of rotational vibration issues*
3. The power available in the Cabinet will not support all drives spinning at any one time.
  - *Limited power budget drives power efficiencies and prevents any misguided attempt to power up all drives.*
4. The component count, power supplies, power converters, fans etc will be significantly less than traditional architectures.
  - *Reduced component count should equate to a less cost and improved system reliability. However only actual reliability data will confirm how well theory performs when reduced to practice.*
5. Access to data on drives that are powered down will be 15 seconds or greater. Longer than expected from a traditional disk subsystem but significantly less than off board tape.
  - *Applications must be MAID aware. A request for data on a powered down LUN will experience a delayed response back to the application and if the application is not MAID aware the delay may trigger a time out or recovery action. Applications must be MAID aware.*
6. To meet data center, enterprise class expectations the solution should have embedded data and device integrity checking and self healing capabilities.
  - *The value and usefulness of data lives much longer after its creation and its initial period of activity. Corporate governance and government compliance regulations are causing data to be stored for increasing longer periods of time and it is this accretive process that is fueling the explosive data growth issue. Any solution that targets the storage of long term data must be architected to ensure the integrity and availability of the data when requested. The more automated the processes the better.*

What is not a characteristic of a MAID solution is the increasingly common feature of drive spin down, occasionally referred to as sleepy drives, a software feature added to traditional array architectures

Spin down is a compromise approach enabling vendors to add this feature to existing array architectures. It is better classified as a power efficiency feature that will deliver 25% to 60% of the power savings possible from a MAID implementation. Companies who offer a spin down option include Fujitsu, HDS, NEC, DataDirect Networks, EMC, Nexsan with Pillar Data promising a future deliverable.

A broader discussion of MAID and its relevance in the data center is presented in the white paper “Defining MAID – “Massive Array of Idle disk”. A discussion of its relevance in the data center.”

email any comments regarding this brief to -

**[bill@verdictusassociates.com](mailto:bill@verdictusassociates.com)**

Verdictus Associates Inc  
Verdictus Associates, a *pragmatic* consultancy focused on data storage technology, management and information technology marketing.

Verdictus Associates Inc  
14 Via Sierra Grande,  
Manitou springs,  
Colorado 80829  
303 588 6674

[www.verdictusassociates.com](http://www.verdictusassociates.com)

Blog:[www.storagetopics.com](http://www.storagetopics.com)