



“ For over a decade Coolspirit have been supplying the UK’s top organisations with storage products and solutions so be assured we will meet your requirements head on.

It’s all about getting things right first time, quickly and simply! ”

**Damon Robertson**  
Coolspirit Ltd

**Our address**

24 The Bridge Business Centre  
Beresford Way  
Chesterfield  
S41 9FG

**Get in touch**

Call us on: 01246 454222  
Email us: [web@coolspirit.co.uk](mailto:web@coolspirit.co.uk)  
Find us: [View location map](#)  
Web: [www.coolspirit.co.uk](http://www.coolspirit.co.uk)

**Office hours**

mon - thurs 8:30am - 5:30pm  
fri 8:30am - 5pm  
sat - sun Closed

“ Boost your storage buying power...  
use ours! ”

Buy with confidence from  
Coolspirit your authorised  
Dot Hill Partner





white paper

**A CASE FOR VIRTUAL  
RAID ADAPTERS**

# Beyond Software RAID





## Table of Contents

1. Abstract .....	3
2. Storage Configurations.....	4
3. RAID Implementation.....	4
4. Software RAID.....	4-5
5. Hardware RAID Adapters.....	6
6. Virtual RAID Adapters.....	7-8
7. Summary.....	9
8. About Dot Hill Systems Corp.....	10

## Table of Figures

1. Storage Configurations.....	4
2. Enterprise Class RAID for entry level servers and workstations.....	7
3. RAID Implementation Comparison.....	9

## Beyond Software RAID; A Case for Virtual RAID Adapters

### **Abstract**

A Virtual RAID Adapter (VRA) is a unique and new class of pure, host-based software RAID solution. It combines the key advantages of pure software RAID with many of the key advantages of hardware RAID adapters. Unlike other host-based offerings, VRA solutions are not tied to any single vendor's disk I/O controller. This makes VRAs universally capable of working on most mainstream SATA or SAS I/O devices. The universal nature of VRA technology allows OEMs to easily migrate to new hardware platforms while maintaining their same look and feel and without having to change RAID vendors.

---

Small to large users alike have used RAID for more than 20 years to protect against data loss due to hardware or system malfunction. RAID has been used to improve system read/write performance as well. An acronym for **Redundant Array of Independent Disks** (originally **Redundant Array of Inexpensive Disks**) RAID is a computer data storage scheme that can divide and replicate data among multiple hard disk drives and can be used to increase data reliability or increase input/output performance. While this is not a tutorial on RAID technology, the most common forms of RAID are:

- RAID 0 (Data Striping)
- RAID 1 (Data Mirroring)
- RAID 5 (Block Striping w/ Striped Parity)
- RAID 6 (Block Striping w/ Multiple Striped Parity)
- RAID 10 (or RAID 1+0 – Mirrored Striped Array)
- RAID 50 (or RAID 5+0 – RAID 5 Striped Array)

Traditionally, RAID has been implemented in the controller of an external disk array or up on an application server via an add-in RAID adapter card. With faster processors and the multi-core processors of today, RAID now runs efficiently on an application server in concert with the operating system; this is usually referred to as “software RAID.” Regardless where it runs, it is important to understand that in essence all RAID is done in software. The difference is where the code gets executed; external disk array controller, host RAID adapter card or host CPU. The storage can be shared or dedicated but it takes a microprocessor and memory somewhere to run the RAID function – i.e., the software<sup>1</sup>. Moreover, one type is not inherently more reliable than the other as it depends upon the product, feature sets and robustness of the solutions. Each has its advantages and disadvantages depending on the end-user environment, application requirements, budget and the problem being solved.

---




<sup>1</sup> Although for array-based controllers and host-base RAID adapters, the code is often referred to as firmware, but it is software none-the-less

## Storage Configurations

RAID does not drive storage decisions rather storage configuration decisions are a function of capacity requirements and expected growth, performance requirements, data sharing requirements, data protection requirements over and above RAID, security requirements, budget constraints and more. Once the storage configuration decision has been made (whether SAN-based shared storage or dedicated application-specific storage) RAID implementation discussions can commence. The more demanding the requirements the more RAID implementation leans towards external disk arrays at the traditional top of the food chain.

There are 3 fundamental storage configurations:

(Figure 1)

<b>Internal Storage</b>	<b>JBOD Storage</b>	<b>Network Attached Storage</b>
 <ul style="list-style-type: none"><li>• Integrated server &amp; storage</li><li>• 2-8 internal drives</li><li>• RAID function in the host; typically software</li><li>• Storage dedicate to the host (i.e., DAS)</li></ul>	 <ul style="list-style-type: none"><li>• JBOD storage external to server</li><li>• Internal storage may be coupled w/ JBOD storage or dedicated to boot image</li><li>• RAID function in the host; software or HBA</li><li>• Storage dedicated to the host (i.e., DAS)</li></ul>	 <ul style="list-style-type: none"><li>• Server and storage array network attached via switch</li><li>• RAID function in the storage array controller</li><li>• Storage shared among multiple hosts (i.e., SAN)</li></ul>

## RAID Implementation

Internal and internal + JBOD storage configurations typically leverage host-based software or hardware adapters for RAID protection. External JOB storage configurations typically leverage host-based hardware adapters for RAID protection while external RAID disk arrays have dedicated controllers with embedded RAID firmware. In particular it is with internal storage and internal + JBOD expansion storage configurations where software versus hardware RAID comparison is relevant. Even here there are different implementations of RAID. Software RAID can either run entirely on the host with no additional hardware, or it can be hardware-assisted. Hardware-assisted RAID leverages a plug-in host bus adapter (HBA) with a RAID BIOS or an additional RAID BIOS integrated with the motherboard.

## Software RAID

A simple way to describe software RAID is that the RAID tasks run on the CPU of your computer system. Therefore, the RAID code utilizes the CPU's calculating power as well as memory and internal bus resources. The code that provides the RAID features runs on the system CPU, sharing the computing power with the operating system and all the associated applications. Software RAID uses hard disk drives which are attached to the computer system via a built-in I/O interface or a processor-less host bus adapter (HBA). The RAID function becomes active as soon as the operating system has loaded the RAID driver software, and is typically limited to RAID 0, 1 and 10 options.



## Advantages

The salient advantages of software RAID include:

- Lower total cost of ownership (TCO)
  - Lower Acquisition Costs – Eliminate the extra cost of the hardware RAID adapter
  - Lower Operating Costs – Eliminate the associated power and cooling costs
  - Lower Support Costs – Commonality of RAID stack (i.e., 1 versus multiple RAID stacks) reduces training and supports costs
- Cost effective in that software RAID utilizes the host resources and I/O components of the motherboard. This maximizes the investment of the motherboard and makes good use of idle resources
  - Uses CPU cycles, internal bus, main memory and onboard I/O ports
  - A CPU and memory upgrade can be less expensive and more effective than a hardware controller
- Scaled performance – Overall RAID speed increases as host CPU count (multi-thread, multi-core) increases, following current market trends.
- Server slot management – PCI slots can be used for other function (e.g., SAS tape adapter, FC HBA)

## Disadvantages

Key disadvantages of software RAID include:

- Consumes and places additional load on host resources (i.e., CPU, memory and internal bus)
- No boot protection – While some pure software RAID solutions have boot protection, many do not. For those that don't, disk failure prior to loading the operating system during the boot process can result in an inoperable system.
- RAID1/5 Reliability – For RAID1 and RAID5 (see RAID5 write hole) running on a host with no NVRAM, system power failures or hardware or software errors occurring for in-progress write operations can result in mirror or parity consistency issues, resulting in corrupt data.
- Write performance – Software RAID is always in write-through mode which extends the write-commit back to the host (vs. write-back cache).
- Limited expansion capability
  - 10s not 100s of disk drives
  - The more drives the more burden on host resources

## **Hardware RAID Adapters**

Hardware RAID adapters have their own processor and memory to run the RAID application. In essence it is an entire computer unto itself dedicated to the RAID application, and acts as a conduit between the server and the array off-loading this task from the host system. They are separate RAID adapter cards that are installed in available server expansion slots. A RAID adapter manages the system disk drives and then presents the logical drives to the server. The server sees the array or arrays as just one or more very fast hard disks; the RAID is completely hidden from the host machine. Hardware RAID adapters usually offer a full range of RAID options (i.e., 0, 1, 5, 6, 10, 50).

### **Advantages**

The salient advantages of hardware RAID adapters include:

- Improved host system and I/O performance – The demand on host system resources (CPU, memory and internal bus) are mitigated by offloading the RAID function to a hardware RAID adapter. This is especially important when one or more CPU intensive applications are running, or as in a virtualized server environment where the aggregation of multiple physical servers into multiple virtual servers running on one physical server places significant demands on the internal CPU, memory and internal bus resources.
- Boot protection – The boot image can be RAID protected prior to loading the operating system mitigating corruption should a disk failure occur during the boot process.
- RAID1/5 Reliability – Battery-backed write cache ensures mirror and parity consistency in the event of system failure for in-progress writes.
- Improved write performance – With a cache memory, write-back caching can be enabled speeding the write-commit back to the host; as opposed to write-through mode which extends the write-commit back to the host.
- Scalability – The number of drives supported is typically much greater than with software RAID.
- Advanced Functions – Hardware RAID adapters typically serve as a platform for advanced functions such as online RAID level migration (RLM), on-line capacity expansion (OCE) and SAS expansion.

### **Disadvantages**

Key disadvantages of hardware RAID adapters include:

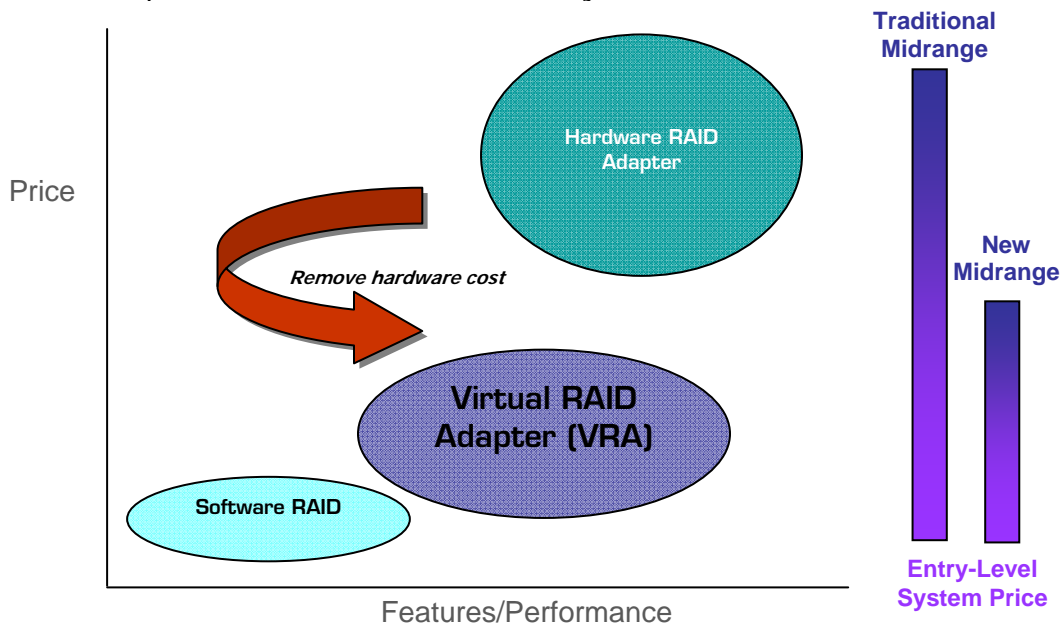
- Higher system cost
  - Acquisition Costs – Extra cost of the hardware component
  - Operating Costs – Extra power and cooling costs of the hardware component
- No controller spanning – Hardware RAID adapters can only manage disks dedicated to that controller. This can result in multiple RAID stacks (one for core logic SATA I/O ports and one for the SAS IOC ports) and multiple storage pools managed independently.
- Fixed performance – I/O performance may not scale as host processing power increases
- Fixed feature set
  - May have to purchase functionality you don't want
  - Will have to buy a new HBA for upgraded or new functionality

## Virtual RAID Adapters

With advances in processors, memory and bus technologies a new form of RAID is emerging which bridges the gap between pure software RAID and hardware RAID adapters. Virtual RAID adapters (VRAs) combine the key advantages of pure software RAID with many of the key advantages of hardware RAID adapters. For example, the VRA retains the cost advantages of pure software RAID but can be a platform for advanced functions such as RAID 5, online RAID level migration (RLM), on-line capacity expansion (OCE) and SAS expansion. It can also solve the problem of boot protection by having an Option-ROM integrated with the system BIOS for bring up during the boot process.

Figure 2

*Enterprise class RAID for volume entry level servers and workstations*



## Advantages

Virtual RAID adapters blend the advantages of pure software RAID with many of the advantages of hardware RAID adapters.

- Lower total cost of ownership (TCO)
  - Lower Acquisition Costs – Eliminate the extra cost of the hardware RAID adapter
  - Lower Operating Costs – Eliminate the associated power and cooling costs
  - Lower Support Costs – Commonality of RAID stack (i.e., 1 versus multiple RAID stacks) reduces training and supports costs
- Cost effective in that software RAID utilizes the host resources and I/O components of the motherboard. This maximizes the investment of the motherboard and makes good use of idle resources
  - Uses CPU cycles, internal bus, main memory and onboard I/O ports
  - CPU and memory upgrade may be less expensive and more effective than a hardware controller

- I/O Performance – For applications requiring write-through I/O, the RAID function is performed by the host CPU. Since there are fewer components to flow through and the CPU processor and internal memory bandwidth typically exceeds that of a hardware RAID adapter, virtual RAID adapter I/O performance can exceed that of the plug-in HBA.
- RAID1/5 Reliability – The touched region logging (TRL) feature of Virtual RAID Adapters ensures mirror and parity consistency in the event of system failure for in-progress writes.
- Scaled performance – Overall RAID speed increases as host CPU count (multi-thread, multi-core) increases, following current market trends.
- Server slot management – PCI slots can be used for other function (e.g., SAS tape adapter, FC HBA)
- Boot protection – With an Option ROM integrated with the system BIOS, the boot image can be RAID protected prior to loading the operating system mitigating corruption should a disk failure occur during the boot process.
- Scalability – With SAS expansion support, the number of drives supported is much greater than with pure software RAID, approaching the capabilities of hardware RAID adapters.
- Controller spanning – There is an additional level of virtualization as all drives, whether attached to core logic SATA I/O ports or SAS IOC ports, can be managed as a single storage pool with LUNS provisioned across the storage pool not just from within a limited subset.
- Silicon independence and freedom of choice – Freedom of choice to migrate between different vendors' motherboard platforms without having to change RAID software stacks. Virtual RAID technology allows users to combine different disk I/O ports from multiple vendors' controllers and treat them as one large virtual RAID adapter utilizing a unique controller-spanning technology, creating even greater flexibility for end-users.
- Advanced Functions – Virtual RAID adapters serve as a platform for advanced functions such as online RAID level migration (RLM), on-line capacity expansion (OCE) and SAS expansion.

### **Disadvantages**

While Virtual RAID adapters mitigate the many of the disadvantages of pure software RAID, some disadvantages remain.

- Consumes and places additional load on host resources (i.e., CPU, memory and internal bus)
- Write performance –Software RAID is always in write-through mode. This extends the write-commit back to the host (vs. write-back cache). For applications that can leverage write-back cache, I/O performance of a hardware RAID adapter will typically exceed that of a virtual RAID adapter.