

## A Short Primer: Active-Active RAID and High-Availability



With the introduction of VTrak™ active-active, dual controller RAID from Promise Technology to the Apple market, come several new questions for prospective purchasers. These include:

- How does active-active RAID operation differ from Xserve RAID?
- What are the tradeoffs?
- Is an active-active system a high-availability storage system?

### Dual Controller, Active-Active RAID Operation vs. Xserve RAID

As the names suggest, a single controller-RAID architecture like that used for Xserve RAID has only one controller per attached storage volume(s). If any component in the *storage to host data path* fails (including controller, attached Fibre Channel fabric/links, or host controller ports serviced), access to RAID data will stop. This will ultimately cause connected host systems to cease operating.

With two controllers (C1, C2) designed for active-active operation, failure of a *single controller* triggers takeover by the remaining controller (an operation called failover), ensuring continuing access to all data stored on the RAID.



Technically speaking, the VTrak RAID from Promise uses two “active-passive” controllers that are in standby for the other until the other fails, when logic causes the remaining controller to perform the failover. Depending on what fails, failover can cause a pause in system operation.

This contrasts with “active-active” controllers used in high-end, enterprise arrays that mirror controller data needed for takeover in battery-backed cache for faster access upon failover and data integrity of the array.

### Tradeoffs

A dual-controller array like VTrak provides a higher level of **access** availability than a single controller array, since failure of a single controller will not stop system operation. Why then, are single-controller RAID systems even used? One common reason is cost of an additional controller. Another is added operational complexity — a single controller RAID optimized for an application (like Xserve RAID for video), is often faster and simpler to manage.

In case you think that failover is always instantaneous, it's not. When it comes to computer operations, data transfer is never a continuous process, and failover logic must be able to differentiate between legitimate cessation of data transfers and those caused by device failure. Sometimes, programming bugs in not just the controller, but drivers or device firmware, cause failover incompatibilities between combinations of devices that testing can't easily detect.

Thus, when it comes to choices about availability, there's no such thing as free lunch: a higher level of **access availability** can certainly be achieved, but it often comes with added management complexity, performance compromises, and added cost.

Simplifying these and other SAN complexities without performance tradeoffs is one of the reasons Vicom developed the purpose-built Vmirror SAN appliance with RAID-independent, active-active multi-path failover, which is used not only for managing Apple storage, but also multi-platform enterprise storage environments.

### Is Active-Active RAID a High-Availability Storage System?

The simple answer is no — while a dual controller, active-active RAID will increase the level of storage **access** compared to a single controller RAID system, it is **not** a high-availability storage system in storage industry terms, where operational continuity (no downtime or operating interruptions) is the principal concern. According to the IEEE, a *high availability system*:

*“Is a system design protocol and associated implementation that ensures a certain **absolute degree of operational continuity** during a given measurement period.*

*‘Availability’ refers to the ability of the user community to access the system, whether to submit new work, update or alter existing work, or collect the results of previous work. If a user cannot access the system, it is said to be unavailable. Generally, the term downtime is used to refer to periods when a system is unavailable.”*

Active-active controller design will guard against controller or data path failures, but it will not ensure operational continuity during failures. Only a fully redundant, mirrored storage system can provide this, and is the primary reason behind Vicom’s Vmirror SAN appliance.


RAID, of course, is designed to protect against single-disk failures. But failure of a disk requires rebuilding of the RAID array, which in turn, degrades performance during re-build period — typically, 1-2 days for most RAID capacities used for Apple video or IT applications. Degraded performance for performance-intensive applications like video playout or work in large collaborative workgroups usually interrupts operational continuity.

Moreover, during this period of degraded operation, loss of a second disk will cause catastrophic loss of data. Using a non-redundant RAID system for business-critical operations also creates maintenance problems because it is difficult to take the storage system out of service.

For these reasons, redundant, high-availability storage systems, and not dual-controller RAID arrays, are used near-universally for enterprise and business-critical applications. Of course, higher availability comes at a higher cost. A mid-range solution will cost \$200,000-\$300,000, not including service. For these applications, however, the cost of downtime or loss of data is many times more.

### How to Build a High-Availability Storage System with VTrak RAID

For VTrak deployment into business-critical environments, we recommend using the 4Gbit, Vmirror 4G with VTrak to create a redundant, high-availability storage solution as illustrated below.



Typically, organizations that choose to implement these configurations have operating needs where interrupted performance or loss of data will have a severe impact on the business. Examples include broadcast, video-post production, advertising, and large printing operations.

It’s true that high-availability Vmirror solutions require twice as much RAID storage, thus costing more than VTrak-only storage. But keep in mind that Apple-based systems offer tremendous value. With fully redundant data paths, storage, and active-active logic from RAID through the entire SAN fabric, this combination represents an high-availability storage solution comparable to mid-range high-availability storage solutions at 4-5 times the price.