Simple, Scalable RAID

A scalable stack of drives and built-in RAID controller make Raidion LTX a flexible, off-the-shelf RAID solution

STEVE APIKI

RAID has never been generic. For all its appeal as a low-cost fault-tolerant technology, RAID has often been limited to custom-tailored, platform-specific, monolithic installations. The Raidion LTX from Micropolis takes a new approach: a scalable stack of drives that can start small and grow with storage requirements. What's more, the Raidion looks to the outside world like a single SCSI target, making it a RAID solution for Unix and Mac systems, as well as more commonly targeted NetWare servers.

Raidion LTX is designed to scale readily, and it meets this goal nicely. But like all RAID systems, its primary goals are maximum uptime and bulletproof data protection. Performance tests of the Raidion LTX show that it's a fast system, and I certainly didn't find any problems with data reliability. If there is a downside, it's that the design features of the Raidion that make it so modular (e.g., its cabled backplane) also provide the potential for some minor failure when compared to less flexible RAID systems.

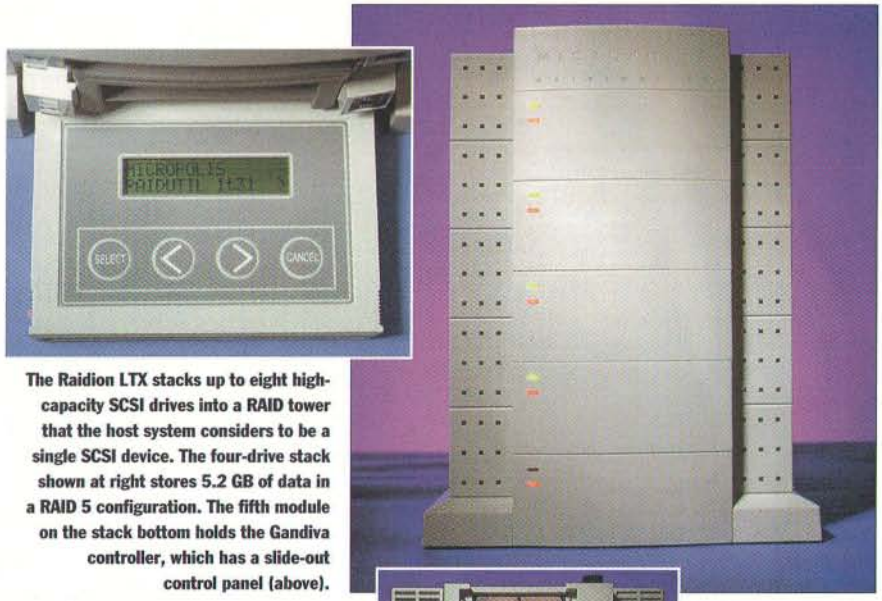

The Raidion LTX stacks up to eight high-capacity SCSI drives into a RAID tower that the host system considers to be a single SCSI device. The four-drive stack shown at right stores 5.2 GB of data in a RAID 5 configuration. The fifth module on the stack bottom holds the Gandiva controller, which has a slide-out control panel (above).

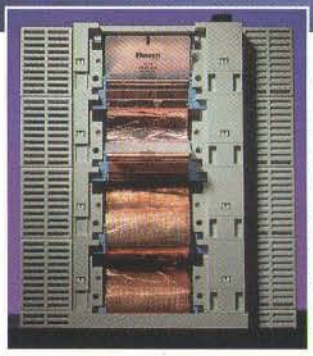

A terminated daisy chain of SCSI ribbon cables (left) connects drive modules, allowing flexible configuration of SCSI channels within a Raidion stack and making it possible to chain multiple Raidion stacks (with up to 28 drives) together via a single Gandiva RAID controller.

Building Blocks

A Raidion LTX array is built from between two and eight drive modules stacked on top of a base containing Micropolis' Gandiva RAID controller. The drive modules consist of a plastic shell that interlocks with the modules above and below, to which the SCSI connectors attach at the rear. The drive unit inside each module can be hot-swapped in and out of the array. To replace drives, you just pop off the module front cover, pull the drive handle to unlock the unit, and slide it out of the shell.

Hot-swappable drive units, complete with individual power supplies, plug into the back plane connector in each module shell.

Each module, including the controller, has its own power supply. As you build up an LTX array, stacking one module atop another, you plug each module into the power outlet sticking up from the module below and hook the new drive into the SCSI chain that runs up the back of the array. The host system connects to the controller module via a Fast-and-Wide SCSI-2 connection at the base of the stack.

Instead of the rigid back plane found in most RAID systems, the Raidion SCSI bus consists of a daisy chain of short ribbon cables that connect at the rear of each module. The Gandiva controller supports up to four independent, synchronous SCSI-2 channels. Ribbon cables run from the Gandiva up through the drives in each channel, ending with a terminated cable stub. Because the Gandiva can support up to seven drives per channel, you can also daisy chain entire arrays by simply running an external SCSI cable from the last drive in the first stack to the first drive in the next. In this manner, you can build a single logical array of up to 28 drives (four stacks) from a single controller.

Although this scheme provides unparalleled flexibility, it is also the source of two potential problems. First, because a single power supply runs the controller module, it is a single point of failure on which the entire array depends. Second, large Raidion LTX configurations rely on a host of interconnected cables that must be installed carefully and kept securely attached. To this end, the cable connections within stacks are designed to fit tightly and attach at right angles to the cable plane, making it virtually impossible for them to work loose accidentally. They're also housed in an enclosed space accessible by removable panels. Further reducing the chance of problems and making installation easier is the fact that three- and four-drive Raidion units come as preassembled stacks.

In addition to preconfigured stacks, Micropolis also sells all the parts (and excellent documentation) for building your own arrays. I tested a preconfigured four-drive unit; hardware setup didn't involve much

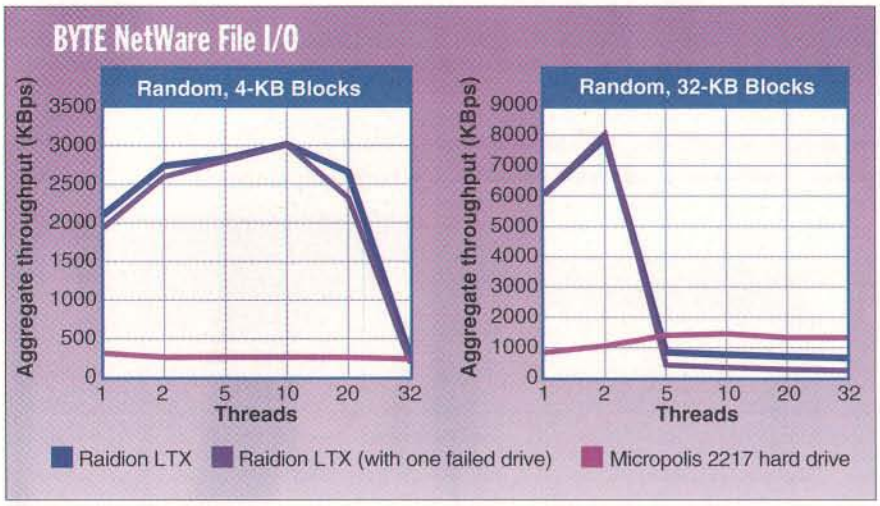

Raidion LTX performance on the random portion of BYTE's NetWare File I/O tests. Each graph shows aggregate throughput as additional tasks (test threads) are added. Both figures show results for a single Micropolis 2217 1.7-GB hard drive for comparison. With small block accesses, (4 KB) the Gandiva controller's over 6 MB of cache keeps the Raidion operating at high-throughput levels until about 20 tasks are running, when throughput begins to degrade rapidly. With large block sizes (8 KB) the throughput is impressive, but once the cache is overwhelmed, the array's performance drops below that of a single drive. Both figures show only a small performance drop when operating with a drive down. The test does not take advantage of NetWare software caching.

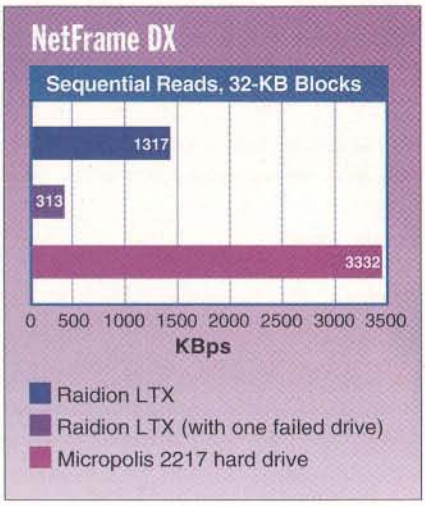

Differences between a fully functional and an impaired array (with one of the four drives down) are more dramatic on sequential-access tests, where overhead from recalculating missing data takes its toll. The early LTX firmware I tested (in OLTP, not Multimedia configuration) did not implement a read-ahead cache, which Micropolis plans to add. As a result, the array did worse on this test with sequential I/O than a single hard drive (one of the Micropolis 2217 drives used in the array).

more than plugging in the SCSI cable. Micropolis sells both a wide-to-wide and a wide-to-narrow cable for attaching the drive to SCSI-2 host adapters, but connecting to a Mac isn't so easy: You need either a SCSI-2 add-in card or a (hard-to-find) converter from Mac DB-25 SCSI to SCSI-2. Micropolis says that it recommends NuBus SCSI-2 cards for performance reasons but is working with cable manufacturers to find a more readily available solution for standard Mac SCSI.

After hardware setup, you need to configure and format the array. Gandiva supports RAID levels 0, 1, and 5. You can set various stripe sizes and set up hot spares (spare drives that automatically come online to replace a failed drive), so you have a variety of configuration options.

Configuration, drive recovery, and other tasks are handled either through a menu-driven front panel or with software utilities. The panel pops out from the Gandiva controller module providing a two-line LCD and four input keys. I used the DOS and NetWare utilities, but preferred the front panel (which you can use while the array isn't attached to any host system). Full access to maintenance functions through the front panel makes the Raidion LTX a truly host-independent system.

Raidion LTX 5.2 Building Blocks

General specifications	
Array capacity	5.2 GB
Test configuration	4 drives, RAID 5
Drives/stack	2–8
Interface	Wide Fast SCSI-2
RAID controller	
Model	Gandiva
RAID levels supported	0, 1, 5
Maximum drives/controller	28
Processor	33-MHz R3051
Cache memory	8 MB
Drive units tested	
Model	Micropolis 2217
Capacity	1.76 GB
Seek time	10 ms
Latency	5.56 ms
External transfer rate	10 MBps
Power system	one power supply/drive, one power supply/controller
Software	configuration utilities for: DOS, NetWare, Macintosh, Sun OS/Motif[1], HP-UX[1], SCO/Motif[1], UNIX Ware/Motif[1], AIX/Motif[1]

[1] Motif-based configuration utilities are planned, but were not available at press time.

About the Product

Raidion LTX 5.2.....$13,470
Host interface cable....$125
Micropolis Corp.
21211 Nordhoff St.
Chatsworth, CA 91311
(800) 395-3748
(818) 709-3333
fax: (818) 701-2809
Circle 1001 on Inquiry Card.

Array Performance

With a fast hardware RAID controller like the Gandiva, a RAID-5 configured Raidion array benefits from increased performance as well as fault tolerance. In addition to a fast 32-bit RISC processor for performing RAID parity calculations, the controller has 8 MB of memory. In the test configuration, the Gandiva operating system takes up about 1.5 MB of that 8 MB of on-board RAM, leaving the balance available as a data cache.

Micropolis will offer the Raidion in two firmware versions, one for OLTP (on-line transaction processing) and the other for multimedia. It will also support a customized video-on-demand configuration through resellers. At this writing, Micropolis offers only the OLTP configuration that I tested. The OLTP firmware maintains a write-through transaction cache but does not support read-ahead caching, which is in the works for OLTP. As a result, random-access performance was much better than sequential access. My speed tests focused on random data access.

I ran BYTE's NetWare File I/O test with the Raidion LTX, the Raidion LTX with one failed drive (pulled), and a stand-alone 1.7-GB Micropolis 2217, which is the drive used in the test array (see the figure "BYTE NetWare File I/O"). The test measures throughput as access threads (simulating multiple users) are added. I ran these tests with a stripe size of eight 512-byte

blocks. With small access sizes (see the figure, "BYTE NetWare File I/O"—Random, 4-KB Blocks), the array peaks at just over 3-MBps throughput, much better than the single drive. With large access sizes (see the same figure—Random, 32-KB Blocks) aggregate throughput goes even higher (over 8 MBps), but once the cache is exceeded, performance drops off rapidly. In general, Raidion LTX showed excellent performance; the real lesson from the 32-KB test is that you must tune the stripe size to accommodate the size of typical data requests if you expect peak performance. In that test, block size was much larger than stripe size.

The random-throughput figures show little loss in speed when a drive is lost. This can be critical if you must maintain a certain performance level even under failure conditions. However, sequential performance, as measured by NetFrame's DX benchmark, does fall a bit off the pace when a drive goes down (see the figure "NetFrame DX").

The last critical performance metric is the time it takes to rebuild the array after failure. This is the window of vulnerability during which a second failure would be disastrous. It took just over 90 minutes (unloaded) for the Raidion to swap in a hot spare in a three-drive array.

How Raidion Stacks Up

There are some changes in the works for Raidion LTX: The 1.76-GB Micropolis 2217 drives will be replaced by larger 2.1-GB 4221 drives. Another option will be 9-GB drives. More important, the release of multimedia-tuned firmware and read-ahead caching should provide a big performance boost for sequentially oriented applications.

I found the Raidion LTX reliable, fast, and unbelievably easy to configure and maintain. Raidion's scalability and flexibility makes it a natural for smaller sites with critical data (like my consulting office) as well as larger, traditional RAID installations. At a price per megabyte of about $2.50, Raidion LTX stacks up reasonably against other fault-tolerant options, and larger companies with mixed-platform networks will save money from reduced spares. ■

Steve Apiki is a BYTE contributing editor and senior developer at Appropriate Solutions, Inc., a Peterborough, NH-based consulting firm specializing in multiplatform development. You can reach him on BIX at apiki or on the Internet at apiki@apsol.com.