



Technical Brief

NVIDIA RAID Technology
Simplifying the Art of Storage
Optimization



Simplifying the Art of Storage Optimization

A growing number of home computer users are storing large quantities of family videos, photos, games, applications, and other valuable data on their disk drives. Although individual disk drives are increasing their capacity, it is not uncommon for users to use multiple disks to store these media files, thanks in part to the reduced cost of hard drives. But many users do not optimize their use of disks effectively, leaving many gigabytes of usable space open.

Imagine if it were easy to use multiple drives to either add total space or offer data protection. NVIDIA is making this possible by bringing Redundant Array of Independent Disks (RAID) technology—which is used by the world's leading businesses—to the common PC desktop. Industry-standard RAID solutions provide disk striping and mirroring to increase performance and protect valuable data. But many of these solutions require storage management expertise, and can create many headaches for the average home user. An affordable and easy-to-use RAID solution would offer consumers the best performance possible while protecting their personal resources.

However, data storage problems can bring an application to its knees. Bugged-down servers, lost or corrupt files, backup management, and many other storage-related situations generate daily challenges for small businesses that cannot afford an IT staff or team of storage experts. Without an in-depth knowledge of disk technology and RAID configuration techniques, many businesses cannot take advantage of RAID solutions for increasing performance and ensuring high availability of critical information.

The NVIDIA nForce™ media and communication processors (MCPs) feature a unique combination of RAID technology and a powerful but simple-to-use interface, which lets consumers or small businesses apply high-end disk storage management solutions in their computing environments. The integrated NVIDIA RAID technology lowers the total cost of ownership for optimized, highly available data storage solutions.

This paper is an overview of the applicable RAID techniques and the specifics of the NVIDIA RAID solution, and describes the applications and computing environments that can benefit from the NVIDIA RAID solution.

Industry-Defined RAID Techniques

RAID techniques were first published in 1988 by a multivendor consortium—the RAID Advisory Board. RAID techniques were divided into different categories or levels. Originally, RAID levels focused on improving resiliency or data availability. As additional RAID levels were defined, one was introduced for improving performance. For all levels, RAID techniques optimize storage solutions by using multiple disks grouped together that are treated as a single storage resource.

Performance: Disk Striping

RAID level 0 defines a disk striping scheme that improves the disk read and write times for many applications. Data is “striped” across multiple disks in an array. In this manner, reads and writes to multiple drives can be carried out in parallel, improving the overall access time for disk operations (Figure 1).

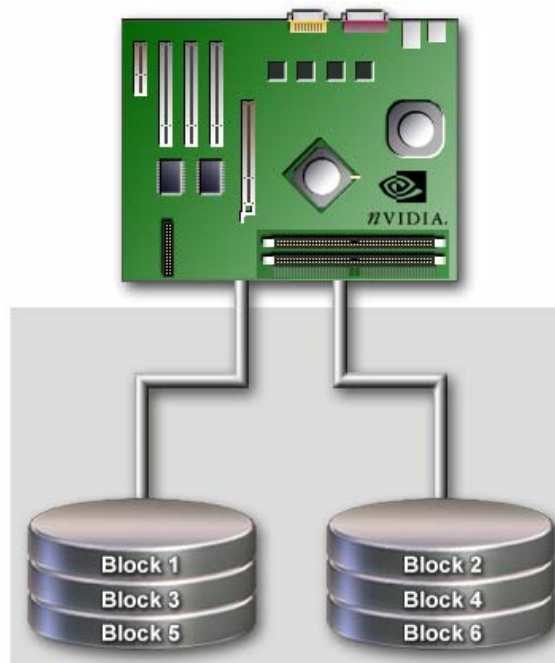


Figure 1. RAID Level 0 Disk Striping

Protection and Simplified Backup: Disk Mirroring

RAID level 1 defines techniques for mirroring data (Figure 2). Every write is actually carried out twice, in parallel. The mirrored, or backup, copy of the data can reside on the same disk or on a second redundant drive in the array. RAID 1 provides a hot-standby copy of data if the active volume or drive is corrupted or becomes unavailable because of a hardware failure. RAID 1 techniques can be applied for high-availability solutions, or as a form of automatic backup that eliminates tedious manual backups to more expensive and less reliable media.

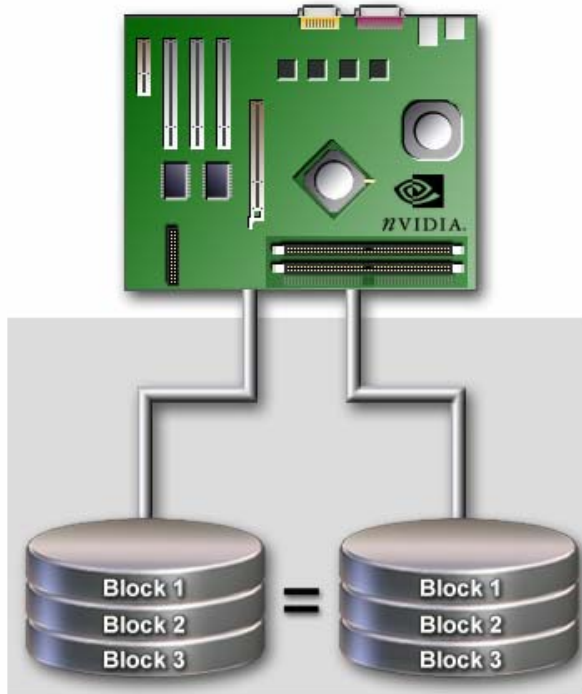


Figure 2. RAID Level 1 Disk Mirroring

Combining the Two

RAID 0 drives can be mirrored using RAID 1 techniques. A RAID 0+1 solution (Figure 3) provides improved performance, plus resiliency.

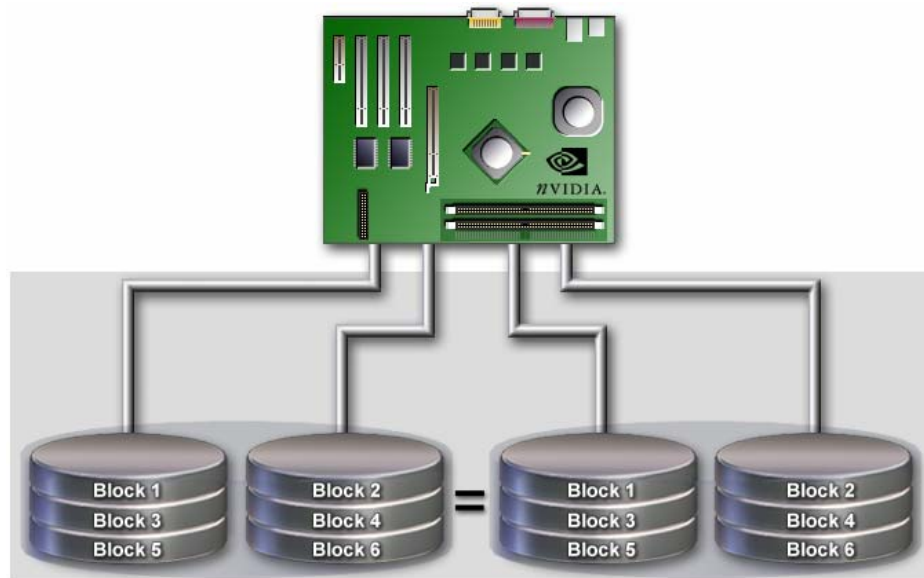


Figure 3. RAID Level 0+1 Disk Striping and Mirroring

NVIDIA RAID

NVIDIA RAID technology implements standard RAID 0, RAID 1, and RAID 0+1 techniques to maximize storage assets. NVIDIA RAID technology additionally introduces many innovations that simplify and optimize the management of RAID features and disk resources.

The NVIDIA User Interface

Using the intuitive user interface, anyone—including users without RAID experience—can apply and manage the NVIDIA RAID technology. Point-and-click operations lead users through the steps to specify which disks to configure in a RAID array, turn on striping, and build mirroring volumes. Configuration changes can be quickly and easily carried out at any time using the same interface (Figure 4).



Figure 4. NVIDIA's Interface: Easy Setup and Management of Disk Storage Resources (Mirroring)

Cross-Controller RAID

Unlike competitive RAID solutions, the NVIDIA solution uniquely supports both Serial ATA (SATA) and parallel ATA disk devices within a single RAID array. Users need not know the semantics of each disk drive—individual setup differences are transparent to the user. A single setup procedure applies to all drives, shortening the learning curve for users who have a variety of disk drives in their computing environment. This flexibility also makes it possible to use all the storage resources and lets users easily expand the number of drives in a particular array.

Booting from a RAID drive

The NVIDIA RAID solution fully supports the use of a RAID drive for loading the operating system at power-up. This means that all available disk drives can be included in the RAID array for optimal performance and protection of all storage.

On-the-Fly Rebuilds

If a disk fails, RAID 1 techniques allow continuous operation by taking advantage of the mirrored data copy in the array. The NVIDIA RAID solution goes a step further and lets a user rebuild a new mirrored copy for the data while the system is up and running, without disrupting user and application access to the data set. Rebuilding on the fly eliminates down time and maximizes protection for critical storage resources.

Spare Disk Allocation

RAID 1 techniques let users designate spare drives that can be configured as hot standbys, protecting arrays in the event of a disk failure. A *shared spare* can protect multiple arrays of drives, or a *dedicated spare* can serve as a hot standby for a particular drive array. The spare disk feature, which offers protection that goes beyond mirroring data, has traditionally been limited to high-end RAID systems. The affordable NVIDIA RAID solution brings this high-end capability to PC desktops. A dedicated spare can take over for a failed disk until the repair is completed—giving support teams the flexibility to choose convenient repair times.

NVIDIA Stability and Reliability

NVIDIA innovations have been characterized with exceptional stability and reliability. The Company's hallmarks of industry-leading engineering and quality control have also been applied in the development of the NVIDIA RAID technology.

To maximize stability, the NVIDIA hardware supports spare disk allocation solution, and incorporates features that contribute to exceptional disk performance and measurable increases in uptime:

- Hardware acceleration

The NVIDIA RAID technology spans hardware, driver software, and the

NVIDIA user interface. The combination of hardware acceleration and optimized software architectures adds up to superior performance.

- ❑ Hot-plug support for SATA drives
Some NVIDIA solutions¹ let users take advantage of this drive capability to avoid powering down an array for swapping out a failed drive.

Applications

RAID solutions deliver several categories of benefits to various applications:

- ❑ Performance optimization
Disk striping techniques result in higher throughput to and from an array of disk drives.
- ❑ Data protection
By mirroring data to a second drive, the probability of data loss because of drive failures is zero.
- ❑ Data availability
NVIDIA RAID technology implements RAID techniques in a manner that minimizes down time for individual drives and the entire array.
- ❑ Automated backups
For users that do not have the staff or expertise to perform regular backups and manage data restores in the event of failures, RAID techniques offer a cost-effective alternative. The NVIDIA RAID solution uniquely simplifies the process with an intuitive interface for setting up and managing RAID arrays.

Applications and users that can benefit from NVIDIA RAID solutions are identified in the following paragraphs.

Archiving Videos and Photos: Performance, Data Protection

Personal archiving of family videos and photos is a growing application area for home computing. As families opt for the convenience of online storage, editing, and file sharing with friends and relatives, the storage requirements have expanded. The requirements now include adequate performance when sharing large files, and protection (backup) for priceless family history. NVIDIA RAID 0+1 technology addresses all these requirements, without requiring the home user to become a storage expert.

Multiple drives can be easily configured to mirror (or back up) critical data and support striping for enhanced performance to avoid dropped frames during video editing operations. As disk storage continues to drop in price, this alternative has become a very practical solution within home computing environments.

¹ Not all NVIDIA nForce solutions support this feature. Please check product technical specifications for details.

Architectural: Data Protection, Automated Backups

Small engineering firms typically cannot afford a dedicated IT support team, but must somehow ensure the availability and protection of valuable engineering drawings and project files. NVIDIA RAID technology provides an affordable and easy-to-use solution with immediate benefits to these types of small businesses.

- ❑ RAID 1 data mirroring ensures the *availability* of stored files.
- ❑ Mirroring also provides an *automatic backup* copy, eliminating the need for time-consuming copying of files to expensive tapes or other media.
- ❑ RAID 0 striping can be implemented to *shorten the access times* for large engineering files, and to improve productivity, especially for projects that require collaborative efforts and file sharing.

Manufacturing: Data Availability

Any production line incurs a cost (lost revenue) when operation is disrupted. The PCs controlling the line must have constant access to critical program and data files to avoid down time. RAID 1 technology can be applied in these types of mission-critical situations to affordably gain high availability and protection for critical storage resources.

Multithreaded Client/Server Applications: Performance

Multithreading, a popular programming technique for large software applications, puts an increased burden on a storage system. Multiple active program paths generate additional disk accesses and can quickly bog down a single drive. Similarly, many server applications result in overly taxed PC clients, plus increased loads on the desktop storage devices. RAID techniques, by striping data across multiple drives, can increase the throughput of a desktop storage solution and ensure the performance required for today's complex applications.

Conclusion

The NVIDIA RAID technology, a standard feature in the latest NVIDIA nForce MCPs, provides a cost-effective foundation for reliable, highly available, and optimized data storage solutions. With an easy-to-use interface, the NVIDIA solution removes the barriers associated with RAID storage management in the past. NVIDIA RAID technology shifts the scale back in the direction of reliability and stability, and lowers the total cost of ownership for optimized, highly available data storage solutions.



Notice

ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, and NVIDIA nForce are trademarks or registered trademarks of NVIDIA Corporation. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2004 by NVIDIA Corporation. All rights reserved.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com