White Paper

NetApp RAID-DP™: Dual-Parity RAID 6 Protection Without Compromise

Paul Feresten, NetApp
April 2010 | WP-7005-0410

EXECUTIVE SUMMARY

As disk drive densities continue to increase, and companies adopt SATA disk drives to reduce storage costs, data loss due to concurrent disk drive failure and error in a RAID group has become a growing risk. Single-parity RAID 5 does not protect against such double-failure events. While RAID 1+0 and conventional RAID 6 do provide sufficient protection, each has serious drawbacks—RAID 1+0 requires nearly twice the number of disk drives, and conventional RAID 6 slows performance significantly.

The solution is NetApp RAID-DP, an advanced RAID 6 technology. At a cost equal to or better than that of RAID 5 and with excellent performance, NetApp RAID-DP provides double-parity protection against data loss. In other words, NetApp RAID-DP technology enables “peace of mind” enterprise storage without compromise.

This white paper outlines the issues of disk-related downtime and data loss and compares the effectiveness of available RAID technologies, highlighting the benefits of NetApp RAID-DP.

TABLE OF CONTENTS

1 INTRODUCTION

Businesses today routinely deploy storage systems using 450GB and 600GB Fibre Channel (FC) and Serial Attached SCSI (SAS) disks. Many have already adopted 1TB and 2TB Serial ATA (SATA) drives to reduce storage costs while eagerly awaiting the availability of even larger drives. Aren't higher-capacity and lower-cost disk drives a good thing? Don't they offer a means for enterprises to stay within tight budgets while still accommodating relentless data growth?

The problem is that as disk drives have gotten bigger, their reliability has not improved at the same rate. As a result, the risk of encountering unrecoverable errors that cause data loss has increased proportionately with the larger drives. The risk of data loss is further compounded when lower-cost SATA disks are employed for workload-appropriate applications—the SATA drives use less effective error correction technology than FC drives.

In this setting of high-capacity, low-cost disks, companies face a substantially increased risk of double disk failures or unrecoverable media errors during reconstruction (MEDR) that result in data loss. In fact, NetApp studies show that, in the event of a drive failure, in a RAID 5 group (utilizing 2TB disk drives) results in a more than 33.28% risk of data loss per year. Many enterprises want to aggressively adopt cost-effective SATA drives but the risks of failures, data loss, and downtime have generally restricted their use in single parity RAID 5 environments.

So how can companies take advantage of high-capacity, economical disk drives without risking data loss, application downtime, and business disruption? Certainly not with traditional, single-parity RAID 5 storage, which offers no protection against double disk failures and MEDR. RAID 1+0 is an option—with 1:1 mirroring it offers better fault tolerance from disk failures. But a RAID 1+0 configuration nearly doubles storage costs, effectively negating the cost benefits of the larger, less expensive drives.

This paper makes the case for NetApp RAID-DP, an advanced RAID 6 technology. While all RAID 6 implementations protect against double disk failures and MEDR, most come with a substantial performance penalty that disqualifies them from production storage environments. This paper explains why only NetApp RAID-DP mitigates the risks of RAID 5, the budget impact of RAID 1+0, and the performance issues of other RAID 6 implementations. NetApp RAID-DP is the only RAID 6 solution offering double-parity protection without compromise to budgets, performance, or data protection. NetApp RAID-DP delivers maximum double-disk failure protection, substantially lower performance overhead than competitive RAID 6 implementations, and a cost equal to or less than single-parity RAID 5.

"Hard disk drives do three things. They read, write, and break. That is why RAID is so important. There are storage systems that support hundreds and even thousands of drives, which means statistically drive failures are inevitable. It isn't an issue of whether drives will fail but how often. RAID-DP provides dual-parity protection, ensuring that there is no data loss even if two drives within a RAID group fail. Given today's realities, dual-parity protection is becoming increasingly valuable and in some cases requisite."

Mark Peters
Senior Analyst, Enterprise Strategy Group

2 THE CHALLENGE: USING HIGH-CAPACITY, LOW-COST DISK DRIVES WITHOUT INCREASING THE RISK OF DATA LOSS AND DOWNTIME

Data loss and downtime cost businesses billions of dollars each year, so avoidance of both must be a primary objective of every IT manager. But what's the best strategy for leveraging new, cost-saving disk technologies without reaching a tipping point when it comes to reliability? Or adequately protecting petabytes of critical corporate data stored on high-density disks?

2.1 DRIVE CAPACITIES GROW FASTER THAN ERROR CORRECTION CAPABILITIES

First, consider the realities of modern disk architectures—while disk capacities double roughly every 18 months, error correction technologies do not improve at comparable rates. Drive-vendor specifications have noted relatively constant uncorrectable error rates (UERs) and mean time between failure (MTBF) statistics across drive generations. The harsh truth is that there has been far too little improvement in error handling and failure rates from one generation of disk technology to the next.

An example helps clarify the significance. Assume the published error rate is one in every 1014 bits transferred. If the error rate does not change, the denser the disk becomes, the higher the probability that the disk will encounter an unrecoverable error. During normal operations, such errors pose no problem since RAID parity can easily recover the data to avoid data loss. However, in the event of a disk failure, the risks increase dramatically for some RAID implementations (such as RAID 5). In this case, a drive failure forces the RAID group to enter "reconstruction," a state in which the data from the failed drive is recalculated and written onto an available spare drive. Not only is parity protection not available during the reconstruction phase, but the odds of hitting an unrecoverable error go up substantially because of the large amount of data being read from the remaining disks in the array. As FC and SAS drives exceed the 450GB mark and SATA drives grow to 2TB capacities, the risk of data loss due to unrecoverable MEDR will become even greater with conventional protection such as RAID 5.

2.2 SATA DRIVES ARE LESS RELIABLE, BUT ECONOMICALLY ATTRACTIVE

SATA drives work exceptionally well as secondary and near-line storage solutions. But unadorned with the same level of enterprise-class data protection technology inherent to high-end Fibre Channel drives, SATA drives tend to fail more frequently. It should not be news that SATA-based disk drives are not the most reliable choice. Originally designed for the desktop environment, SATA drives were never meant to spin 24x7. But enterprise customers have found it hard to ignore the attractive economies of SATA and low-cost FC and SAS drives.

2.3 SCARY MATH: SATA BASED RAID 5 GROUP= 33.28% CHANCE OF DATA LOSS

Deploying a high-density SATA solution with single-parity RAID protection (such as RAID 5) increases the risk of serious data loss and downtime. In fact, analysis shows that, in the event of a drive failure, utilizing a SATA RAID 5 group (using 2TB disk drives) can mean a 33.28% chance of data loss per year. Figure 1 illustrates this risk and how RAID-DP can dramatically reduce the probability of data loss.

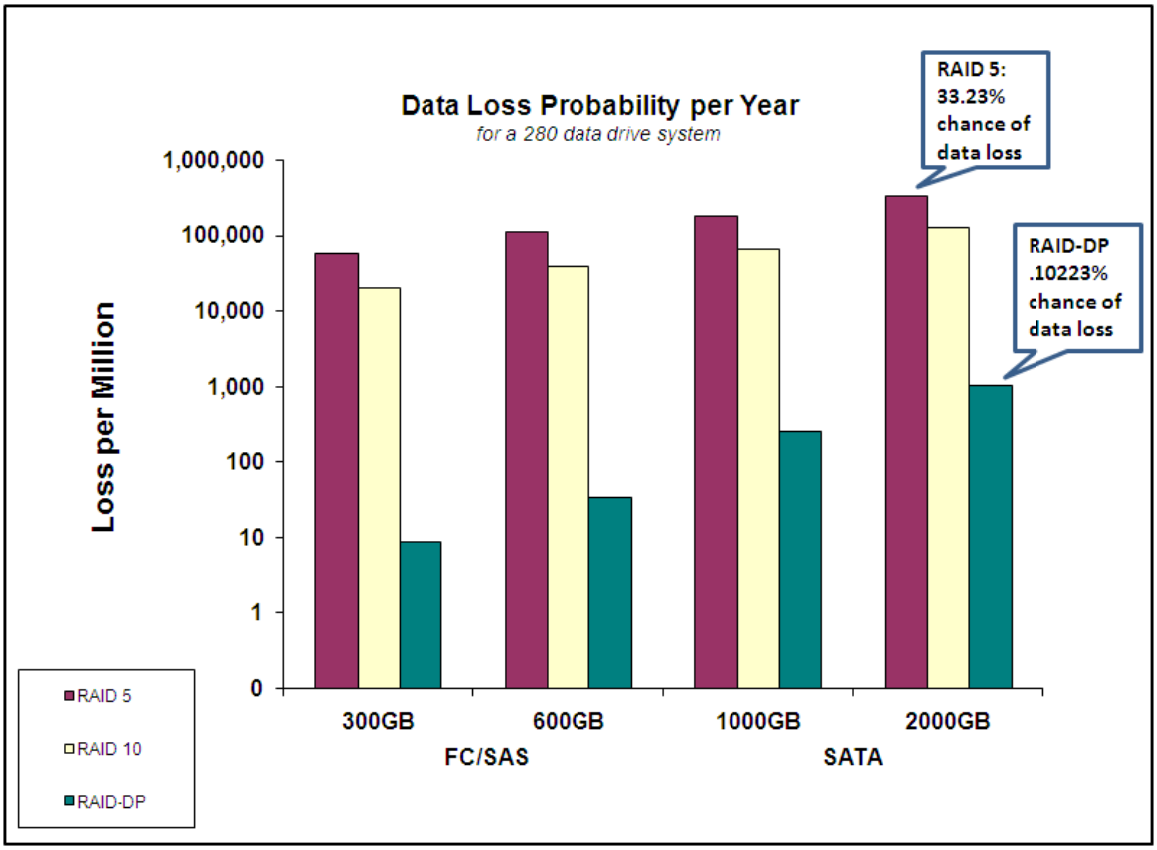

Figure 1) NetApp RAID-DP: Hundreds of times better than RAID 5.

When one disk fails, RAID recreates data from the remaining disks in the array onto a spare disk. However, significant increases in disk capacities have resulted in much longer reconstruct times for data lost on the failed disk. And using a larger SATA drive makes reconstruction times even longer—the more data, the longer the reconstruction, and the longer the reconstruction, the greater the chance of another disk failure or uncorrectable error occurring.

The most alarming impact of this reality is the loss of mission-critical data when a second disk failure or uncorrectable error event occurs in a single-parity RAID configuration. To make matters worse, consider that the probability was calculated using industry-standard statistics based on, among other things, drive failure rates. Typically, given standard drive failure rates, one would expect a very low percent-annual-failure rate. But what happens when the deployment expands to 100 or 1,000 disks? The more disks in the deployment, the more the probability starts to catch up. When double failures occur in a RAID 5 configuration, the enterprise can face serious data loss, downtime, and the financial consequences associated with interrupted business operations.

3 THE PROS AND CONS OF RAID 5, RAID 1+0, AND CONVENTIONAL RAID 6

Clearly, with the odds so stacked against enterprise-class reliability, businesses need to implement data protection technology that works more effectively with modern disk architectures. All of the major storage solutions vendors offer RAID technologies: EMC, Hitachi Data Systems (HDS), HP, IBM, and Sun each offer RAID 1+0 and RAID 5. HDS, HP, and Sun also offer RAID 6 options. Do these RAID systems help protect your data? Some do, but all force compromises on either cost or performance.

3.1 RAID 5: GETTING RISKY, PLUS COST ISSUES

"A form of parity RAID in which the disks operate independently, the data stripe size is no smaller than the exported block size, and parity check data is distributed across the RAID array's disks."[1]

RAID 5 implementations use single parity to protect against single disk failure. Single disk failure is the issue. Traditional single-parity RAID offers adequate protection against a single failure event. The caveat is that no other disk failure or uncorrectable media error can occur while reconstruction is still in progress. During normal operation, if there is an unrecoverable error and the event is a read error, then recreating data from parity occurs almost instantaneously, and the array remains online. However, if a disk in a RAID 5 group fails, then all its data has to be recreated, and the array remains in a vulnerable degraded mode (without parity protection) until data has been reconstructed onto a spare disk. For a high-density or slow disk, that could be a very long time.

Because RAID 5 cannot protect against double disk or media failures, most vendors recommend it only for non-mission-critical corporate data or in some limited business-critical applications where cost is a constraint. But be aware that RAID 5 still does not come cheaply. Many vendors limit RAID group size to 3+1 or 5+1. That means 17% to 33% of the disk budget will go to parity disks.

3.2 RAID 1+0: MOST COSTLY OF ALL RAID AND STILL SOME RISK

A stripe of mirrors.

RAID 1+0 technology uses mirroring to protect against double disk failures. It is generally considered the gold standard for mission-critical applications—but it comes at the price of gold as well. RAID mirroring requires purchasing twice the disk capacity needed to store data. To emphasize, that means twice the up-front acquisition costs and twice as much hardware to service and maintain—think two times the disk capacity, two times the disk enclosures, two times the physical rack space, two times the power costs, two times the cooling costs, and so on.

For such a high price tag, one might expect complete protection against disk or media failure. But RAID 1+0 does not deliver in that respect. While RAID 1+0 protects against many double disk or media failure scenarios, gaps remain. In the event that both disks on the same mirror fail, the result is data loss. To clarify, when one disk fails in a mirror, the only way to recover from that failure is to copy all of the data from the mirror onto an available spare disk. This copy period is a protracted, vulnerable time during which the odds of an unrecoverable error increase proportionately with the number of bits transferred. To be fair, some RAID 1+0 vendors can offer an add-on synchronous mirroring product that eliminates the problem of data loss during this copy period. But such protection adds to the overall price tag—the synchronous mirroring license alone can cost hundreds of thousands of dollars. In addition, the synchronous mirroring requires a completely separate array as a destination. All told, a RAID 1+0 solution with synchronous mirroring can easily exceed several million dollars in cost.

3.3 CONVENTIONAL RAID 6: EASIER ON THE WALLET, BUT CAN YOU AFFORD A 33% PERFORMANCE HIT?

"Any form of RAID that can continue to execute read and write requests to all of a RAID array's virtual disks in the presence of any two concurrent disk failures. Several methods, including dual check data computations (parity and Reed Solomon), orthogonal dual parity check data, and diagonal parity have been used to implement RAID Level 6."

Conventional RAID 6 implementations use dual parity to protect against double disk failures. RAID 6 technology provides excellent data protection without the extreme cost of RAID 1+0 solutions. But standard RAID 6 implementations compromise performance. Hitachi Data Systems is one of the few big-name storage vendors to offer RAID 6. According to the company's own published statistics, an HDS RAID 6 solution suffers up to a 33% performance hit.

[1] Storage Networking Industry Association; www.snia.org/education/dictionary/r.

3.4 THE DESIRED SOLUTION: DOUBLE-FAILURE PROTECTION, MINIMAL PERFORMANCE OVERHEAD, COST-EFFICIENT CAPACITY UTILIZATION

So what would characterize the desired RAID solution for business and mission-critical enterprise application environments? Obviously the solution must protect against double disk or media failures. Without that level of protection, companies with large deployments of high-capacity disks and/or lower-duty-cycle products such as SATA are at risk of data loss and downtime. Secondly, the RAID implementation should ensure high levels of capacity utilization—otherwise the acquisition, installation, and maintenance costs can be prohibitive for many organizations. Finally, the RAID technology should incur minimal performance overhead so that applications can maintain a high level of predictable performance.

4 NETAPP RAID-DP TECHNOLOGY: BEST-IN-CLASS DOUBLE-FAILURE PROTECTION WITHOUT COMPROMISE

NetApp RAID-DP technology uniquely satisfies key requirements of double-failure protection, efficient capacity utilization, and minimal performance overhead. NetApp RAID-DP is a dual-parity, RAID 6—as defined by the Storage Networking Industry Association (SNIA)—implementation. At the most basic layer, RAID-DP offers a dual-parity configuration for greatly enhanced data protection against any two disk failure events occurring in the same RAID group.

4.1 DIFFERENTIATING FEATURES FOR SPEED AND IMPROVED DATA PROTECTION

Several important implementation features differentiate NetApp RAID-DP technology from other systems. If a double disk failure occurs, RAID-DP automatically raises the priority of the reconstruction process so that the recovery completes faster. A second uniqueness is that RAID-DP takes advantage of the likelihood that in a double disk failure, one disk failed some time before the second failure event and at least some of the information was already recreated with traditional row parity. NetApp RAID-DP automatically adjusts for this by starting recovery at the point where two elements are missing from the second disk failure.

4.2 NEGLIGIBLE PERFORMANCE OVERHEAD

NetApp RAID-DP technology introduces minimal overhead—typically less than 2% compared to NetApp RAID 4 single-parity implementations. Tables 1 and 2 compare RAID performance in HDS RAID 6 and NetApp RAID-DP implementations. Table 1 data is normalized RAID 5 and RAID 6 performance data extracted from the HDS white paper entitled “Using RAID-6 with Hitachi TagmaStore™ Storage for Improved Data Protection.” In this comparison table, HDS RAID 5 (7D+1P) provides a baseline of 100%.

Table 2 data derives from NetApp RAID level performance testing. In the NetApp example, the default RAID group size for RAID 4 (7D+1P) provides the baseline of 100%.

Comparing the performance results from each table indicates a clear performance advantage with NetApp RAID-DP technology. While the HDS RAID 6 implementation suffers from a 33% random write performance penalty, the NetApp RAID-DP implementation introduces a negligible 1% to 2% overhead across all read/write operations.

Table 1) HDS TagmaStore RAID level performance comparisons.

RAID Level	Random Read/ Sequential Read	Sequential Write	Random Write
RAID 5 (7D+1P)	100%	100%	100%
RAID 6 (6D+2P)	100%	85.7%	66.7%

Table 2) NetApp RAID level performance comparisons.

RAID Level	Random Read/ Sequential Read	Sequential Write	Random Write
RAID 4 (7D+1P)	100%	100%	100%
RAID-DP (14D+2P)	99%	99%	98%

Combining the inherent efficiency of the NetApp WAFL® (Write Anywhere File Layout) file system with the RAID-DP striping algorithm, NetApp accomplishes clear advantages with minimal performance difference between single-parity and dual-parity RAID. In particular, RAID-DP achieves virtually equivalent random write performance as NetApp single-parity RAID 4. Practically speaking, that means that NetApp RAID-DP implementations can be deployed in demanding enterprise environments with virtually zero impact on application performance.

4.3 THE BUSINESS MATH: BETTER PROTECTION THAN RAID 5, HALF THE COST OF RAID 1+0, AND INDUSTRY-LEADING RAID 6 PERFORMANCE

NetApp mathematical models highlight the extreme advantages of NetApp RAID-DP technology as compared to alternative RAID implementations. Consider the following statistics and comparisons:

- **Maximum data protection**. With NetApp RAID-DP, the chance of data loss as a result of a double disk failure is hundreds of times less likely than in competitive RAID 5 configurations. While RAID 1+0 offers much better data protection than RAID 5, it still presents a risk of data loss in the event of double mirrored disk failures. RAID-DP offers 100% double disk failure protection at half the cost of RAID 1+0.
- **Lowest cost**. RAID 5 implementations often limit RAID group size to 3+1 or 5+1 (which represents a 17% to 25% parity cost overhead). RAID 1+0 requires 1+1 (a 50% overhead). In contrast, NetApp supports RAID group sizes of up to 28 (26+2) for a low 7% capacity overhead. For maximum data protection, RAID-DP offers the industry's most cost-efficient RAID scheme for maximum capacity utilization without compromise.
- **Uncompromising performance**. Competitive dual-parity technologies are often perceived as incurring a substantial performance penalty. This is reflected by publications from Hitachi Data Systems that indicates a 33% performance impact with HDS RAID 6 technology. NetApp RAID-DP incurs virtually zero performance penalty compared to single-parity RAID.

4.4 PROVEN IN THE MARKETPLACE: STANDARD ON ALL NETAPP SYSTEMS

Integrated and included as a standard component of the NetApp Data ONTAP® operating system and available across the entire product line, NetApp RAID-DP adds no cost or special hardware requirements to NetApp solutions. Because RAID-DP delivers significant data protection advantages without compromising performance or introducing cost penalties, NetApp recommends RAID-DP as the standard RAID configuration for all NetApp storage deployments. For mission-critical applications that require maximum system fault tolerance beyond double disk failure protection, NetApp SyncMirror® software can be added to RAID-DP for ultimate storage resiliency. As illustrated in Figure 2 below, SyncMirror and RAID-DP are key components of the NetApp storage resiliency product family.

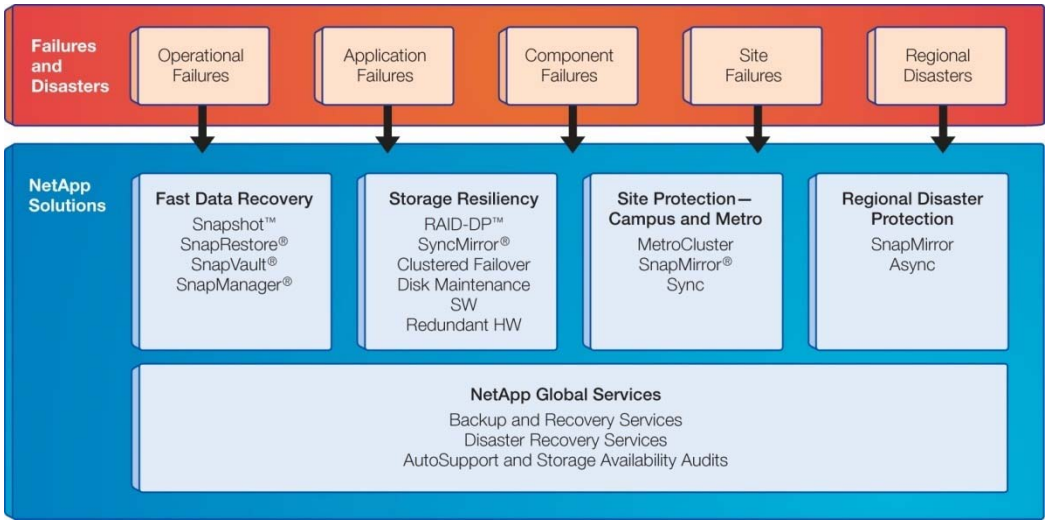

Figure 2) Storage resiliency is an important component of NetApp availability and disaster recovery solutions.

Unlike competitive vendors that lack a RAID 6 solution or have seen limited acceptance of the technology because of performance issues, NetApp has successfully proliferated RAID-DP across the full spectrum of applications, industries, company sizes, and geographies.

5 SUMMARY

5.1 DATA PROTECTION WITHOUT COMPROMISE

NetApp RAID-DP technology offers true RAID 6 dual-parity protection without compromise to budgets, performance, or data integrity. Alternative RAID 5, RAID 1+0, and even competitive RAID 6 offerings force tradeoffs that limit applicability of advances in new drive technologies. In contrast, NetApp RAID-DP extends protection against disk failure to enable confident utilization of high-density and/or lower-duty-cycle SATA for all workload environments, ranging from near-line secondary to mission-critical primary storage.

With data loss and downtime costing $1 million to $6 million per hour in lost revenue, maximum data protection is expected as an inherent part of any storage system deployed in business-critical or mission-critical application environments. Enterprise customers cannot afford either the degraded performance of competitive RAID 6 implementations or the crippling cost penalties of RAID 1+0. As the RAID 6 industry leader with thousands of proven installations, NetApp RAID-DP delivers enterprise-class data protection without compromise.