

The Data Protection Advantages of iStor Intelligent RAID Management

by Gene Nagle
Director of Technical Marketing
iStor Networks, Inc.

Only iStor Networks offers a combination of three vital technologies that work together to offer one of the industry's highest levels of data protection in a popularly priced iSCSI data storage system.

- **Down-Drive** is a form of preventative data migration that provides an automated background process for copying off data from a drive that is going bad before an actual failure.
- **Micro-Rebuild** takes care of the tendency of SATA drives to occasionally go off-line briefly, avoiding a full drive rebuild when this is not necessary.
- **RAID Protected Virtual Volumes (RPV²)** offers volume-level RAID configuration for significant reductions in rebuild time should a drive failure occur

If the growth of your data requires you to look for additional robust storage with low total cost of ownership, you should read this short paper to understand the extra protection offered by these iStor technologies. You will see why most iStor iS325 users can choose the higher performance and capacity of RAID 5 and do not need to make the sacrifices necessary to go to RAID 6.

Avoiding the RAID Rebuild Process

RAID with parity, commonly known as RAID 5, has become the standard for basic data protection in disk arrays. Most storage arrays create RAID sets at the drive level; that is, the user chooses some or all of the drives in the array, and a RAID 5 set is created using that entire set of drives. When one of the drives in a RAID 5 array fails and is replaced, with either a "hot spare" that is already in the array or a newly inserted drive, the new drive must be rebuilt from the other drives using the parity mechanism. With today's high capacity SATA drives this can be a lengthy process, and it is a potentially risky one since if a second drive should fail during that rebuild time, the user's data is lost. This is why iStor has developed two important data protection technologies that help to avoid ever having to do a rebuild.

The iStor intelligent RAID system monitors each disk drive, keeping track of all uncorrectable bad blocks and other information provided by the drive's internal Self-Monitoring, Analysis and Reporting Technology (SMART). To enhance this drive monitoring process still further, the user may set up scheduled media scans or parity scans during times when the storage array is not busy. Users of the iS325 are encouraged to take advantage of this background scanning option since bad sectors are likely to be discovered much sooner with drive scanning than during normal drive operations.

When the number of uncorrectable bad blocks reported by a drive reaches a pre-defined level, or other indicators from the SMART system reach a critical threshold, the iStor intelligent software identifies the drive as potentially going bad and issues an email alert. The user can then request that a **Down-Drive** process be started which moves data off the failing drive in advance of an actual failure occurring. This is a background block-by-block copy process rather than a RAID rebuild process using parity, so it requires far less in the way of system resources and thus a minimal performance penalty. The copying process is done to those blocks that are actually in volumes on the identified drive, not all blocks on the drive, and when it is complete, the offending drive is automatically marked as bad by the iStor software and will no longer be used.

"Among popular iSCSI disk arrays with the lowest cost of ownership, only iStor has RPV² Technology, for managing RAID data protection at the volume level."



“... most iStor iS325 users ... do not need to make the sacrifices necessary to go to RAID 6.”

A second intelligent feature of the iStor iS325 that is designed to keep RAID rebuilds to an absolute minimum is called **Micro-Rebuild**. Many Serial ATA (SATA) drives at times go into a brief unresponsive state, and ordinary RAID systems often mistake that for a drive failure, with the result that when the drive comes back on-line a complete rebuild of the drive is triggered. On the other hand, when a drive becomes unresponsive in the iS325, the iStor intelligent

RAID management system automatically creates a log file that keeps track of all writes destined for that drive. If the drive becomes responsive again within five minutes, the iS325 performs a Micro-Rebuild of just those writes in the log file and then resumes normal operation. This totally avoids the overhead and performance degradation caused by unnecessarily rebuilding disks that are still viable.

How Ordinary RAID 5 Works

As noted above, the majority of disk array systems have RAID sets created at the drive level; that is they span multiple whole disk drives. A RAID set is then carved up into partitions of various sizes, sometimes called volumes, based on the requirements of the user's applications. Should one of the disk drives fail, the entire drive needs to be rebuilt, even including space on the drive that was not used by any partition. Large SATA drives, such as those with capacities of 750 Gigabytes (GB) can take a significant length of time to completely rebuild from the parity protection on the rest of the drives, and if a second drive should fail during that time the user's data is lost.

Even though the likelihood of two disks failing at close to the same time is quite remote, the above risk of data loss is the reason that some disk array buyers feel the need for double parity, also known as RAID 6. RAID 6 provides protection from the failure of two drives at once, but at a cost: throughput is often significantly compromised by RAID 6, and this is especially apparent on writes to disk. Double parity also requires that disk space otherwise used for data must be relinquished for additional parity space. In the case of an eight-drive RAID set created from 750 GB drives, RAID 5 provides the user 5.25 terabytes (TB) of useable space, whereas RAID 6 only provides 4.5 TB. Note also that RAID 6 requires a minimum of four drives, versus three for RAID 5.

RPV² is no Ordinary RAID Technology

To minimize the time needed to rebuild a RAID 5 set should a disk drive failure occur, iStor has developed an advanced form of storage virtualization in the iS325 storage system called RAID Protected Virtual Volumes or RPV². With RPV², each application (user created volume) becomes its own "RAID set".

For example, if the user needs 200 GB for one of the applications whose data will be stored on the iS325, he or she creates a RAID 5 volume to hold that data by using the Create Volume function on the iS325. The iStor system will use several disk drives for this RAID set (selected automatically or by the user), placing what is known as an Extent on each of several drives for optimum performance. Moreover, the automatic volume creation process is programmed to always use those drives with the most free space when a new volume is created, thus as much as possible leveling the space usage across all available drives.

The key benefit of RPV² technology is that individual disk drives in the iStor iS325 normally have only a fraction of their space used for volumes. So if a drive should fail, only the volumes that have extents on that drive will be rebuilt by using space available on other drives, and/or on a hot spare drive. Furthermore, since the iS325 intelligently searches for space on physical disks that are not currently part of the volumes to be rebuilt, the rebuild process is distributed across many drives, or "virtual spares." This in turn allows multiple affected volumes to be rebuilt simultaneously. So the rebuild process proceeds quite rapidly compared to rebuilding an entire single drive, and the time that the system is exposed to risk of a second drive failure is minimized.

RPV² Technology, then, reduces the time of rebuild, and for many users makes the use of RAID 6 and its accompanying sacrifices of performance and capacity completely unnecessary. Another important benefit is that the volumes of a failed drive can be rebuilt without a hot spare drive in the

“Micro-Rebuild... avoids the overhead and performance degradation caused by unnecessarily rebuilding disks that are still viable.”



system. This is possible when the unused space of all other physical drives not part of the volumes being rebuilt is sufficient to hold the volume extents from the failed drive. Although one or more disks can be set aside as hot spares if desired, the user can have all disks contribute to the performance and capacity of the array by maintaining enough space across the system as a whole for virtual spares.

RPV² also allows each volume to be created with the quality of service (QoS) that is best suited for that application; for example, RAID 10 mirroring can be used instead of RAID 5, for a volume whose data is even more critical (requires more protection and performance), and the RAID 10 volume can utilize the same drives as some other RAID 5 volumes. In addition, RPV² allows an iS325 system to have drives of dissimilar capacities. With ordinary RAID systems, RAID sets are limited by the size of the smallest drive in the set, whereas with RPV² the extra capacity of the larger drives can be fully utilized.

Figure 1 provides an example of volume management with RPV² in the iS325, where three volumes of different RAID types are configured in a nine-drive system that has various size drives in it. For this illustration, the cylinders represent physical drives, and each symbol represents an extent that is 100 Gigabytes in size. The light gray circles make up a large 500 GB RAID 5 volume with 400 GB available for data. The dark gray hexagons make up a RAID 1 mirrored volume with a net capacity of 100GB. While the black squares are the two parts of a 200GB RAID 0 striped volume.




Figure 1. Sample of Volume Management with RPV²

The Intelligent Way to Protect Your Data

The table below provides a summary of the data protection features of iStor's iSCSI Controller Technology along with benefits that our users derive from these features.

FEATURES	BENEFITS
Micro-Rebuild	Ability to rebuild a drive incrementally in response to a brief anomaly, avoid complete rebuild.
Drive Monitoring	Early warning of possible drive issues.
User Scheduled Drive Scanning	Enhances the early warning capabilities of drive monitoring.
Down-Drive	Moves data off a failing drive in the background, then takes it off line.
RPV ² : RAID type determined at volume level	Quality of Service determined for each individual volume.
RPV ² : Drive Space Utilization Management	Keeps space usage and wear on the drives in the array more even.
RPV ² : Multiple virtual spares	Faster rebuild of failed drives, no need for dedicated hot spares

Among popular iSCSI disk arrays with the lowest cost of ownership, only iStor has RPV² Technology, for managing RAID data protection at the volume level. RPV² and other important intelligence features designed into the iStor iS325, including **Down-Drive** and **Micro-Rebuild**, make possible one of the industry's highest levels of data protection. And these features allow the user to choose the higher performance and useable capacity of RAID 5 instead of slower and more expensive RAID 6 while still minimizing the risk of possible data loss.

Gene Nagle is the Director of Technical Marketing at iStor Networks, Inc.





iStor Networks, Inc.®

Corporate Headquarters

iStor Networks, Inc.
7595 Irvine Center Drive Suite 100
Irvine, CA 92618
Phone: (949) 753-8999
Toll Free: (888) 98iStor
(888) 984-7867
Fax: (949) 753-1068
Email: info@istor.com

Worldwide Sales Information
Phone: (949) 753-8999
Fax: (949) 753-1068
Email: sales@istor.com

Customer Support
Phone: (949) 753-8999 x159
Toll Free: (888) 98istor x159
(888) 984-7867 x159
Email: support@istor.com

Sales Management

Leo Kameya
Director of Sales - Americas
Office: (949) 753-8999 x181
Fax: (949) 753-1068
Mobile: (760) 707-6247
Corporate: (949) 753-8999 x181
Email: lkameya@istor.com

Peter Cmaylo
VP, NA & EMEA Sales & Business Development
Office: (508) 543-7966
Fax: (508) 543-7954
Mobile: (508) 208-8835
Corporate: (949) 753-8999 x105
Email: cmaylo@istor.com

International Offices

Europe

Tim Beck
Director of Sales EMEA
7 North Avenue
Goring By Sea
Worthing
BN12 4DA
United Kingdom
Office: +44 (0) 1903-502800
Fax: +44 (0) 5600-492410
Mobile: +44 (0) 7939-310522
Email: tbeck@istor.com

Asia Pacific

Dennis Lin
VP, Asia Sales and Business Development
7F., No. 46, Lane 10, Kee Hu Road, Nei Hu,
Taipei(114), Taiwan, R.O.C.
Office: +886-2-2799-9198 x8228
Mobile: +886-9-8771-9526
Fax: +886-2-2799-9398
Email: dlin@istor.com