DNF SECURITY
DYNAMIC NETWORK FACTORY

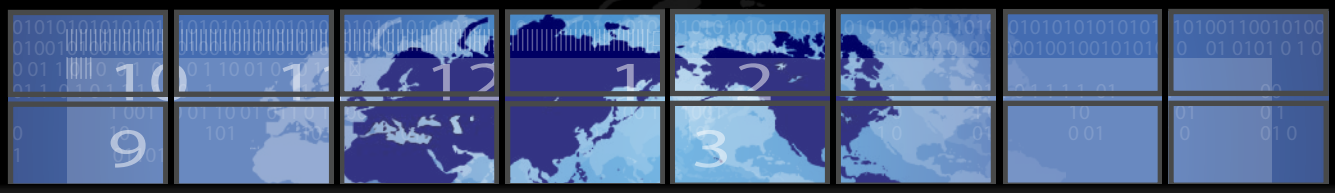

High Definition Data Protection for Reliable Storage Systems

In the early-2000s the Internet boom caused a shift in the IT function, the department focused on implementing a host of projects to increase productivity, standardize their network infrastructure, move to web-based technologies and implement new communications technologies. With the shift to a knowledge-based economy, after the Internet boom, unstructured data in documents, spreadsheets, presentations, web content, intranet sites, archives and backups have grown quickly, over 60%[1] annually in a typical company; this type of information is as essential as critical customer data and business processes.

Typically organizations, after ranking applications and their data on the criticality to business operations, they then design their infrastructure accordingly. For example, a high traffic e-commerce site will spend the bulk of their IT budget on web servers, and data processing applications. A car manufacturer will support their computer aided design systems and other engineering/research and development applications with higher cost systems and infrastructure. Unstructured data like email and user documents tend to be relegated to the lower cost, tier two storage. Backups and archives are typically housed on tier three tapes or disk archives. Disk-based archiving and D2D backup have been implemented by 50% of organizations, and these deployments are increasing year over year.[2] With critical information located on all tiers of storage, data protection is essential and a range of approaches at the disk array and infrastructure level are required to enhance protection and reliability.

Storage Class	Common Applications	Typical Disk Drive(s)
Tier One	• Mission Critical Applications • Operational Data	• Fibre Channel Disks • Enterprise Class SCSI/SAS Disks
Tier Two	• General Storage • File Serving • Email • User Documents	• Midrange SCSI/SAS Disks • SATA Disks • ATA Disks
Tier Three	• Backup • Disk Archives • Permanent Archives	• SATA Disks • ATA Disks • Removable Media • Tape Drives

Early enterprise storage systems relied on heavy duty SCSI disk drives designed for 24/7 operations. These disks had a low rate of failure, since the capacities and rotational speed were much lower than current disk drives. Individual hard disk capacities have increased over 200,000% since their creation in 1956, growing from 5MB to 1,000,000MB (1TB) on a single disk. The first disk drive, IBM's RAMAC, had an areal density of 2,000 bits of data on a square inch; today, Hitachi's commercial disk drives feature an areal density of over 200 Gigabits/square inch[3].

Over the past 10 -15 years two developments have contributed to the growth of hard drive capacity and wider range of deployments: the rapid adoption of personal computers at work and home and the drive to develop lower cost enterprise storage systems by using commodity disk drives.

In the late 90s, companies like Dynamic Network Factory (DNF) began developing lower cost storage systems by using ATA disk technologies. Since ATA disks cost 3-10X less than enterprise disk drives, these systems have been widely adopted in a range of new applications including network attached storage, digital video recording, CCTV infrastructure and well as for critical infrastructure for organizations large and small.

Although ATA disks were very affordable and accessible, these disks were not designed for 24/7 enterprise applications or applications with substantial disk performance demands. ATA based systems were well suited for applications requiring mass storage or organizations looking to reduce capital expenses while increasing overall infrastructure redundancy.

In consumer applications, disks are typically deployed in single disk configurations. Since the disks do not typically work together in a system, minor difference in manufacturing process across batches are insignificant and have limited impact on overall system performance and reliability.

With today's complex storage arrays and systems, 16, 1600 or 16,000 disks may be deployed in a single system. These minor variances add up to a huge impact. Increased drive and system density have impacted overall storage system reliability as more and more data is packed onto a single disk. The chance of experiencing a bad sector or failed drive head is more common than in the past as the system and disk workloads increase. Improving redundancy for all tiers of storage requires a multi-pronged approach combining RAID protection, disk monitoring technologies, careful resource and allocation planning, and advanced integration testing.

How does RAID technology improve disk drive reliability?

RAID

RAID (redundant array of independent disks) is a way of storing data across multiple disk dirves to performance or redundancy. This technology places data on multiple disks and the I/O operations can be balanced improve performance.

Key RAID Concepts

RAID techniques mirror or stripe data across multiple disks. Mirroring is copying data on more than one disk. Striping splits data and operations across multiple disks. Lastly, RAID uses error correction to detect, and even fix problems as they arise. Each RAID level may use one or more of these techniques to increase reliability or performance.

Common RAID levels

- **RAID-0:** This level stripes data but does not improve redundancy, offering high performance without fault tolerance.
- **RAID-1:** This level mirrors data across an even number of disks or multiple disk pairs. This level offers improved read performance, since data can be read from multiple disks but offers no write performance benefits but offers absolute fault tolerance for data.

- **RAID-5:** RAID 5 offers a good deal of redundancy by splitting the data across multiple disks for parity. The data is split into chunks written multiple times to multiple disks. Parity information can be used to reconstruct the data in the event a disk is lost due to failure or offline status. RAID 5 requires three disks in an array, but is best suited for multi-user systems where maximum performance is not critical and write operations are less common.
- **RAID-6:** This double parity RAID level protects against multiple drive failures by providing twice the protection than RAID-5, or single parity RAID systems.
- **RAID-10:** This combination of RAID 1 and RAID 0 offers high performance and high redundancy. There are two subtypes: In RAID-0+1, data is organized as stripes across multiple disks, and then the striped disk sets are mirrored. In RAID-1+0, the data is mirrored and the mirrors are striped.
- **RAID-50 (or RAID-5+0):** This type consists of a series of RAID-5 groups and striped in RAID-0 fashion to improve RAID-5 performance without reducing data protection.

RAID improves system reliability by offering redundancy across multiple disks. With RAID technology is designed to work through disk failures by offering protection and the option of changing failed disks in a live system to reduce downtime and maintain data availability.

What causes disk drives to fail?

Disk drives fail for a multitude or reasons. There are predictable and unpredictable failures. Mechanical failures are typically predictable, and monitoring systems can easily alert administrators to the impending failure. Unpredictable failures may be cause by electrical components in individual disk drives, changes in environment, or data corruption. Eliminating as many potential points of failure as possible can help administrators sleep easier, spending more time improving operations not fighting fires.

Mot drive failures are caused by mechanical issues; they tend to be predictable and easily recognized by disk monitoring techniques like SMART. Today's disk drives are largely electronic, managed by tiny microprocessors. Focusing exclusively on mechanical failures is risky, issues with power regulation or on-board processing is not uncovered by disk monitoring techniques and cause disks to fail.

SMART-er Monitoring Techniques

Self-Monitoring, Analysis, and Reporting Technology (SMART) is a process used to alert users or system administrators of impending drive failures. With this information, admins can take preventative measures to protect their vulnerable data: copy the data to an alternate location, move critical systems to another resource, or engage their vendor for technical support. The first version of SMART provided failure prediction by monitoring online hard drive activities, now SMART monitors and prevents failure by attempting to detect and repair sector errors, then tests all data and sectors of drives using offline data to confirm the drives health during its inactivity. These improved techniques have aimed to eliminate mechanical disk failures at the source. Unfortunately this technique doesn't predict and eliminate all causes of disk failure, so organizations are still at risk if they rely on SMART singularly.

Recent independent studies from Goggle and Carnegie Mellon University have concluded that disk drive failure rates are considerably higher than the rates reported by disk drive manufacturers. There is an astonishing 8% annual failure rate for drives that have been in service for two years; which is equivalent to one out of every 12 drives. In addition, 36% of the failed drives did not exhibit a single SMART-monitored

failure.[4] In addition, customers replace disk drives 15 times more often than drive vendors estimated.[5] Hence, disk drive failure predictions are necessary to detect early in order to eliminate marginal disk drives that can cause catastrophic failure in the field. By ensuring your storage systems are reliable from the disk level and beyond, IT organizations can spend less time supporting storage systems and more time improving data management.

A list of common drive technology issues:

Type	Symptom	Cause
On Time Spin-Up	Drive Not Ready	Stiction Problems
Signal Integrity	Signal to Noise Ratios (SNR)	Media Errors
Remaps	Data Mismatch	Read/write Errors
Random Seeks	Data and Track-Following Problems	Zone Boundary Issues
Reset	Power Spike	Data & Control Crosstalk
Random Power-Down/ Snoozing/ Low Activity State	Out of Tolerance (for use in RAID)	Temperate Tolerance Issues

These issues do not necessarily impact disks in day to day use or cause immediate failure, but they reduce the overall lifespan of the disk drive.

There are four major classes of data corruption:

- **Firmware Errors:** These errors are usually a result of incorrect settings in the BIOS or corruption of the firmware that controls the hard disk drive. The drive may be undetectable at startup or misnamed because of the corruption of service data.

- **Hardware faults**: This is the most common class of data corruption; faults are caused by the failure or malfunction of the hardware components. Faulty hardware is typically caused by mechanical wear and tear and the deterioration of the components over time. Permanent disk damage is caused by disk head crashes or bad sectors, and hardware malfunctions can lead to corrupted data being processed or stored.

- **Human-induced errors:** People can introduce data corruption accidentally. When power to a disk is lost while the system is active, data may be saved erratically, causing unreadable, unusable corrupted data.

- **Silent corruption:** In these cases data is written to a drive then changed without a warning or error. The system continues to operate normally until the block or sector is needed. This can cause a catastrophe ranging from the loss of a file, database or file system. The typical causes of silent corruption include firmware bugs or changes, hardware errors, software bugs, bit error rates or noise from transport protocols.

Silent corruption is the hardest cause of drive failure to diagnose and recognize. Its elusiveness also leaves administrator more vulnerable to data loss and drive failures. To eliminate failure caused by silent corruption, early detection is key. Standard manufacturing quality control and testing efforts fail to detect silent corruption, and storage systems utilizing standard commercial disk drives are vulnerable to this problem. By prequalifying disk drives in a real world application enthronement, where large I/O operations are performed and validates, silent corruption can be detected and eliminated in the lab and discovered before your data is involved.

How does drive qualification reduce field drive failures?

Storage system manufacturers do not develop hard drives, and the disk drives used do not vary significantly form vendor to vendor. Vendors can ensure reliable systems land in the field by using extensive qualification process to normalize initial build quality. These processes are key in delivering mission-critical, high-availability systems, where predictable performance and uptime are critical. The qualification process can eliminate that exhibit signal quality problems, data mismatch errors, temperature and power regulation issues, or other conditions that typical manufacturer's testing processes do not uncover. DNF has developed an extensive qualification process, High Definition Disk Certification, for all disk drives that has reduced our field failure rates by 50%, and reduces annual drive failure rate to 1%, 80% less than Google's reported annual failure rate.[6]

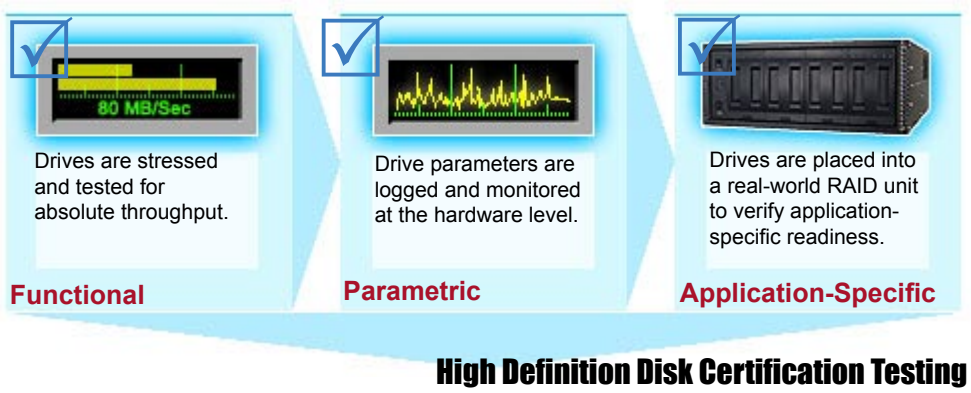

How does the High Definition Disk Certification Process work?

The 120 hour certification process qualifies commercial disk drives with a rigorous certification process weeding out drives not well suited for storage systems or RAID arrays. There are three testing processes that cover functional testing, parametric testing and application-specific testing.

Functional Testing:
Functional testing, more commonly known as stress testing or "burn-in," places units in an artificial environment and forces them to perform as fast as possible for an extended period of time. Drives are put through a range of read/write situations, including worst-case scenario alternating fast 0/1 writes on the dense inner tracks of a drive. However, functional testing is not comprehensive, and only eliminates drives that perform outside of the specified range.

Parametric Testing:
This testing section uses a real-time RTOSS operating system, to monitor raw drive behavior(hardware and interface)l are monitored and logged. to eliminate anomalies in seek time, read error rate, voltage.

Application Testing:
Application Testing is one of the most crucial features of the certification procedure. Individual drives typically fall in line with manufacturer's tolerance levels, but individual variations are unacceptable in system or array deployments; each drive must conform with its teammates in the application.

As data is striped across a high-performance RAID array, each drive accepts an exact percentage of the write load. When the next write is striped on the array, each drive must have completed its previous write task for the RAID to work at full speed. If a single drive fail to complete its previous write in time, the entire array will slow down to accommodate it. If the performance anomalies are serious enough, drives may flicker offline and threaten the array's integrity. Application testing is the last test; it takes the winners, tests them in unison, in a RAID application, to ensure the disks are working well together. This offers improved performance and reliability for real applications like streaming media, high I/O databases, IP surveillance

or multi-purpose disk arrays, over using any set of commercial disks. RAID Certification qualifies drives to perform in real-world applications -- not just lab tests. This increases field reliability for any system using qualified commercial disk drives, SATA, SCSI, ATA, SAS or fibre channel.

How should I design a fault tolerant environment?

Our checklist for multi-faceted data protection:

- ☐ Look for redundant components in each array that are field replaceable: hard drives, power supplies, cooling fans and so on
- ☐ Make sure your system offers automatic notifications for any failed components or unusual circumstances
- ☐ Look for hardware RAID protection for the standard RAID levels: 0, 1, 5, 6, 10, 50 (remember software bugs are one of the main causes of silent data corruption)
- ☐ Deploy a hot-spare (or multiple spares) in each array no matter what RAID level you use
- ☐ Use array based disk monitoring tools to predict mechanical disk failures
- ☐ Look for storage vendors that qualify commercial drives, before using them in your system, drives straight from the distributor are not well suited for your production environment
- ☐ Create an internal baseline for data protection. Define your required uptime and critical systems
- ☐ Create redundancy outside of your array with SAN mirroring, clustering or remote sites

(Footnotes)

[1] IDC, "Worldwide Storage 2008 Top 10 Predictions: New Paradigms," Doc # 209796, December 2007.

[2] InfoStor, "User support for D2D rising…" http://www.infostor.com/articles/article_display.cfm?Section=ARCHI&C=Newst&ARTICLE_ID=232817, July 2005.

[3] Hitachi, "Hitachi Achieves Nanotechnology Milestone For Quadrupling Terabyte Hard Drive," News Release, October 15, 2007.

[4] _Google Labs, "Failure Trends in a Large Disk Drive Population," Research, February 2007.

[5] ComputerWorld, "Disk Drive Failures 15 Times What Vendors Say," Doc # 9012066, March 2007.

[6] DNF internal research from technical support and RMA data from 2002-2006.

About Dynamic Network Factory

Dynamic Network Factory (DNF) takes pride in its innovative spirit, engineering excellence, and broad product line. DNF's six business units focus on specific vertical markets and technologies to cover the business and government technology space. Since 1998, DNF has focused on delivering storage solutions from direct attached storage for small business to enterprise applications, networked storage, iSCSI and file servers in capacities from 1TB-10PB. In 2006, its acquisition of iSCSI pioneer StoneFly Networks expanded DNF's storage portfolio to include enterprise class iSCSI and storage virtualization technology. With the expertise of DNF Storage, DNF Systems and StoneFly, Inc, DNF Security leverages innovative storage and server technologies for mission critical IP surveillance solutions. For more information visit **www.DNFsecurity.com**.

DNF SECURITY
DYNAMIC NETWORK FACTORY

www.DNFsecurity.com
toll free: 800.947.4742
fax: 510.265.1565
email: sales@DNFcorp.com

Corporate Headquarters
21353 Cabot Boulevard
Hayward, CA 94545
main: 510.265.1616