Achieving Five-9's Availability: The Dot Hill Systems Commitment to Reliability

INTRODUCTION

Organizations take extraordinary measures to maximize the availability of on-line data, as well as to minimize or even eliminate the loss of any valuable data. RAID (redundant array of independent disks) configurations supplemented with data replication or snapshots all help achieve both goals. These measures are both more effective and less expensive, however, when they are built upon a reliable storage area network or virtualized storage array foundation.

When data in the storage area network (SAN) are unavailable, so too are the applications that must access this data. During such outages, work grinds to a halt. Many employees are forced to stop what they are doing, only to get immediately on the phone to the help desk to find out what's wrong. Indeed, the help desk staff may be among the few employees who are still able to perform their regular duties until normalcy is restored.

This is why any downtime of mission-critical applications is expensive. While figures vary by the type and size of the organization, some industries, such as energy and telecommunications, report losses of revenue from $33,000 to $47,000 for every minute of downtime, according to a study by Meta Group (subsequently acquired by Gartner).

Availability is measured by percentage of uptime. Five-9's of availability translates into just over 5 minutes of downtime annually; four-9's means about 53 minutes of downtime. These are statistical averages, of course, but the additional 48 minutes of annual average uptime can readily justify the relatively modest additional investment required for greater redundancy in the SAN infrastructure supporting mission-critical applications.

Dot Hill Systems understands the importance of high availability for its customers. Throughout its 27-year history, Dot Hill has built a reputation for exceptional quality in both hardware and software. This commitment has resulted in an industry-leading high availability, which is embodied in the very name of Dot Hill's product family: AssuredSAN™.

This paper's remaining content is divided into two sections followed by a brief conclusion. The section on Designing for High Availability describes the measures Dot Hill takes during the engineering design phase to ensure high reliability and availability. The measures taken to validate system reliability are then covered in the section on Proof of Success with Rigorous Analysis and Field Data.

DESIGNING FOR HIGH AVAILABILITY

High availability is achieved through a combination of three design elements:

- High reliability (measured by the Mean Time Between Failures or MTBF) of the system and its several subsystems;
- Redundant subsystems to eliminate as many single points of failure as possible; and
- Rapid repair of any failure (measured by Mean Time to Repair or MTTR) by using Field Replaceable Units (FRUs) for all critical subsystems.

The following equation for availability demonstrates the vital role of serviceability in the system's design. Maximum availability can be achieved only by minimizing the time it takes to affect a repair, which is reduced significantly by using FRUs.

$$\text{Availability} = \frac{\text{MTBF}}{\text{MTBF} + \text{MTTR}}$$

To achieve maximum availability, Dot Hill designs for reliability and serviceability, as well as for manufacturability; the specific aspects of each are described in turn here.

DESIGN FOR RELIABILITY & SERVICEABILITY (DFRS)

Designing hardware for high reliability and serviceability involves both the system and its several subsystems. To achieve high availability at the system level, Dot Hill integrates reliability into the design process in several ways. The first and most obvious is the use of storage device (disk drive) redundancy with RAID configurations (RAID 1, 3, 5, 6, 10 and 50) and dual power supplies, which each includes its own fan to prevent over-heating (and thereby, accelerated component failures). Even higher availability is achieved by using redundant controllers. By eliminating single points of failure in these critical subsystems, the system itself continues to operate normally during a failure of any single FRU. While such a failure does factor into the subsystem's MTBF (its rated reliability), it does not diminish the availability of the system itself.

Dot Hill's AssuredSAN architecture features full redundancy for every subsystem requiring a significant number of active components. The mechanical chassis itself cannot be redundant, of course, and there is a single midplane that performs the simple function of connecting the redundant controllers to the redundant disk drives. The midplane has minimal active components, however, and Dot Hill selects these for the highest possible reliability. The result is an extraordinarily high MTBF for the chassis and its midplane, and therefore, virtually no impact on system availability.

To enhance system serviceability for the shortest possible MTTR, Dot Hill utilizes the two complementary design techniques. The first is the use of a modular chassis with Field Replaceable Units (FRUs). The ability to swap out a confirmed failed subsystem quickly and easily minimizes the time it takes to repair an installed system and restore it to full operation. By utilizing such a modular design, which provides convenient access to all subsystems, Dot Hill's AssuredSAN products can be maintained seamlessly with minimal or no disruption in service during most repairs.

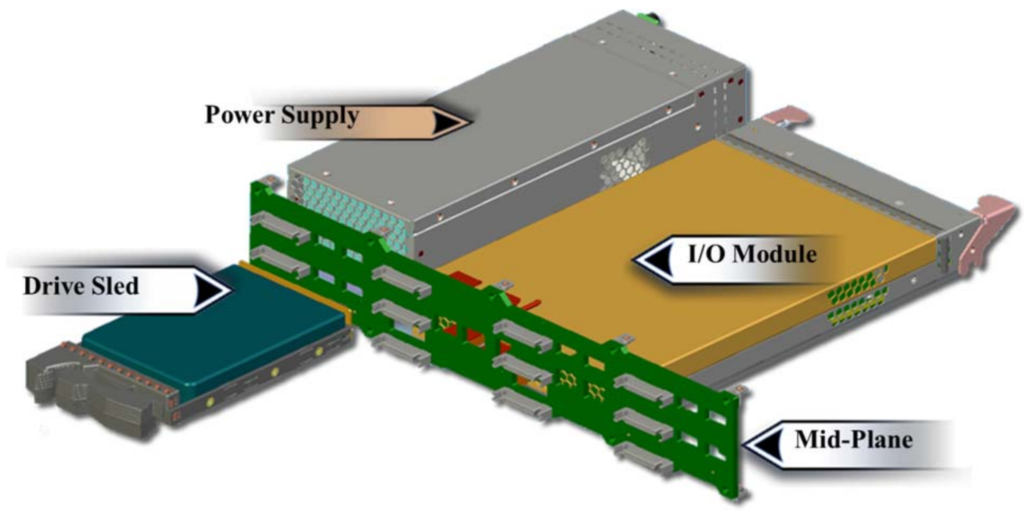

Dot Hill's mechanical design enables the power supply and fan, I/O module and controller, and disk drives all to be serviced quickly as hot-swappable Field Replaceable Units (FRUs). Being able to replace redundant FRUs while the system is fully operational further enhances availability. Note how the power supply and I/O module are accessible from the rear of the chassis, while the disk drives are accessed from the front. Note also that redundant FRUs are not shown here to identify the individual subsystems more clearly.

The second serviceability technique is immediate notification of any failure. The longer it takes to detect a failure, the longer it will take to repair; this is rather obvious. Time is of the essence for another reason, however: The failure of a redundant subsystem creates, in effect, a temporary single point of failure that increases the risk of a system-level outage. For this reason, the firmware in all Dot Hill systems is designed to detect, isolate and confirm any failure, initiate a failover to a redundant subsystem, and provide immediate notification. The actual "messaging" of the notification can also be configured to match operational procedures to ensure that on-duty staff is properly and quickly notified.

At the FRU or subsystem level, Dot Hill utilizes four separate design techniques to maximize the MTBF of each, while at the same time also maximizing the inclusion of leading-edge SAN features. The first is reducing the part count. Because any individual part can fail, the fewer there are, the higher the inherent reliability of the subsystem. Dot Hill engineers endeavor, therefore, to minimize the parts required on all printed circuit boards and other subsystem FRUs.

The second technique is to use only high quality parts. Higher quality parts cost more, of course, but their superior performance and longer service lives normally contribute to a lower total cost of ownership in the long-run. Despite the higher per-part cost, minimizing the part count, while concurrently enhancing functionality, helps to improve the overall price/performance of a highly-reliable design. For these reasons, Dot Hill utilizes only the highest quality parts available from reputable suppliers.

The third technique involves the de-rating of selected parts. Operating any part or component at or near its rated capacities inevitably shortens its useful service life. For critical parts, Dot Hill selects only those that will be able to operate at approximately 50% of their maximum allowable specifications for voltage, power and/or current. This can substantially increase the service life, and therefore, the MTBF of the subsystem.

The fourth technique is unique to Dot Hill: designing for software reliability. In modern designs, software reliability is just as important as hardware reliability, and in some ways even more important. The reason is: Software bugs (including those in firmware) that cause downtime normally take significantly longer to resolve than the more obvious hardware failures. Bugs are often dependent upon system state (the set of circumstances leading up to the failure), making them difficult to reproduce and isolate quickly, and any patch or update must be tested before it can be released. Both add considerably to the MTTR for software failures, thereby adversely impacting on system serviceability and availability.

To maximize software reliability, Dot Hill monitors the improvement in the Mean Time to Discovery (MTTD) of bugs to assess the maturity of all software and firmware during development. It is important to note that MTTD is not an industry metric, but is instead a software maturity metric created by Dot Hill as part of the company's commitment to quality and reliability. All designs must have a sufficiently high and stable MTTD before being finalized.

The design is released to manufacturing only after passing three comprehensive tests. The Engineering Verification Test (EVT) and the Design Verification Test (DVT) ensure that the system and/or subsystem(s) fully satisfy all design specifications, including those for high reliability of both the hardware and software. These tests also confirm that marginal variations in parts from component suppliers will not compromise system reliability over the product's useful life of a minimum of 10 years. The Reliability Demonstration Test (RDT) is a separate and rigorous evaluation of the final production hardware that verifies its calculated reliability, availability and serviceability (covered below). Whereas some vendors use only a few samples in a fairly short demonstration test, Dot Hill's RDT uses 18 to 20 fully-configured systems in a 13-week test that must produce zero hardware failures to pass.

Design for Manufacturability (DFM)

To ensure high availability, Dot Hill's Quality & Reliability Assurance (QRA) Group establishes a comprehensive set of assurance controls in parallel with the engineering team's design effort. These controls ensure that the design facilitates manufacturing best practices for maintaining high quality and reliability with high yields and minimal defects. Dot Hill's DFM process also evaluates all components and suppliers for quality assurance using a strict set of qualifications to ensure minimal failure rates.

All hardware is then manufactured under similar controls for the production process itself, which must comply with two separate Dot Hill specifications: Quality Requirements for Mechanical Parts; and Supplier Production Part Approval Process Qualification. Additional best practices during manufacturing include burn-in of components, subsystems and full systems to minimize infant mortality failures, and Ongoing Reliability Testing. For the ORT, random samples are selected every week for a 4-week test to ensure that the manufacturing process is yielding the desired levels of quality and reliability by not compromising the design's inherent capabilities.

Proof of Success with Rigorous Analysis and Field Data

The reliability of any system can be assessed using a "bottom-up" analysis of its several component parts. While these analyses can be remarkably accurate, especially when conducted in accordance with proven methodologies, it is also prudent to validate these calculated assessments with actual data from units in production use at customer sites. Dot Hill does both.

Reliability, Availability & Serviceability (RAS) Analysis

Dot Hill's RAS methodology models the design's reliability, availability and serviceability at both the system and subsystem levels, and also provides estimates of RAID data protection levels. Component-level Belcore MTBF predictions at 25ºC and 40ºC are included for all FRUs as a foundation for the analysis, and all analyses are performed in accordance with Telcordia's SR-NWT-000332 Reliability Prediction Procedure for Electronic Equipment. This Belcore/Telcordia prediction methodology assumes a series model with the MTBF based on any random hardware failure (including in redundant subsystems), regardless of whether or not it causes a system-level outage.

Per the RAS analysis methodology, all FRUs are assembled via reliability block diagram (RBD) methods (shown in the figure below), and Monte Carlo simulations are applied to account for different data protection levels (Vdisks/LUNs for RAID5 N+1 and RAID6 N+2) and different MTTRs; for example: eight hours ("Gold") and 16 hours for ("Bronze") service restore times for any FRU or software failure. The analysis assumes an average 24 second failover time to the redundant controller per Dot Hill internal test data.

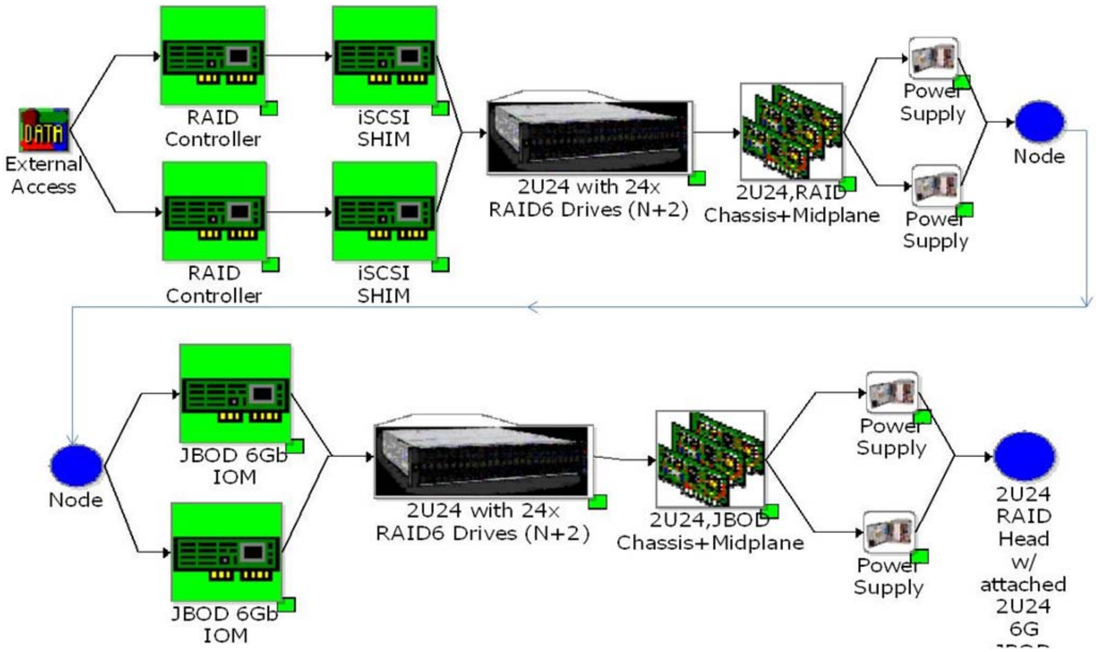

Depicted in this reliability block diagram is the configuration of the AssuredSAN 3000 Series system modeled here, which includes two iSCSI 4-port 2U24 RAID chasses with dual power supplies/fans, and dual RAID controllers (each with an attached JBOD with dual power supplies/fans and dual JIOMs). Each RAID or JBOD chassis contains 24 x 300GB SAS disk drives (RAID5 and RAID6), resulting in two different configurations.

A summary of the availability results from the RAS analysis are shown in the following table:

	RAID5 Dual Controllers	RAID6 Dual Controllers
"Gold" 8 Hour Service Restore Time	99.999516%	99.999520%
"Bronze" 16 Hour Service Restore Time	99.999024%	99.999043%

These statistical availability rates translate into the following average expected minutes of downtime per year:

	RAID5 Dual Controllers	RAID6 Dual Controllers
"Gold" 8 Hour Service Restore Time	2.54	2.52
"Bronze" 16 Hour Service Restore Time	5.13	5.03

Proof of Success with Actual Field Data

Dot Hill works primarily with major vendors in the storage industry as an original equipment manufacturer (OEM). Under this business model it is the major vendors, and not Dot Hill, who perform the field maintenance of the installed base of SAN systems. These vendors all utilize the results of Dot Hill's RAS Analysis to help maintain an inventory of spare FRUs, and all also track actual failures in the field. While this data are regularly shared with Dot Hill for quality control purposes, most vendors consider the data to be confidential.

One vendor that takes particular pride in providing high availability for its customers, however, has agreed to permit Dot Hill to disclose aggregate results from its "Uptime Meter" used to track availability continuously and in real-time. The vendor utilizes Dot Hill's AssuredSAN 3000 platform rebranded as a private label product. The Uptime Meter has confirmed Dot Hill's claim of five-9's availability over time from actual performance across its entire installed base of systems.

Conclusion

As demonstrated by the many measures outlined in this white paper, Dot Hill is committed to delivering the highest possible availability across our entire product line. Dot Hill also takes great pride in the fact that this commitment to quality and reliability has enabled us to deliver "carrier-class" five-9's of availability in products with "enterprise-class" price/performance.

No other company delivers such high availability at such a competitive price. But don't take our word for it. Compare for yourself Dot Hill's high availability with any other product in its class. Evaluate the vendor's commitment to quality and reliability by its design, manufacturing and testing processes. Look at the bottom line results. Talk to colleagues about their experience. We are confident that if you do, you too will come to realize what many others already know: Dot Hill Systems delivers industry-leading high reliability and availability.

To learn more about how your organization can benefit from high-availability SAN solutions from Dot Hill, visit us on the Web at www.dothill.com or call us at 303-845-3200.