

WHITE PAPER

Enterprise Storage: The Foundation for Application and Data Availability

Sponsored by: NetApp

Brad Nisbet

Noemi Greyzdorf

November 2010

IN THIS WHITE PAPER

This white paper examines the importance of application and system availability and how a growing number of organizations are currently addressing high-availability needs with respect to their storage environments. It explores the importance of application availability and the underlying foundation that reliable disk storage systems provide a much broader set of business requirements and applications. The paper then focuses on deployments of NetApp storage systems and how these deployments have met its customers' expectations with greater than 99.999% availability based on NetApp's rigorous monitoring process.

SITUATION OVERVIEW

Introduction

The Importance of Availability

For a growing number of organizations, having a reliable, highly-available IT solution is critical to success. The importance of IT availability is not relegated to the largest of organizations, nor is it relegated to specific industries. Indeed, availability and uptime will mean different things to different organizations, but the movement toward constant uptime and a 24 x 7 environment continues to permeate small, medium-sized, and large organizations across an expanding number of industries and applications.

Downtime of an organization's IT environment and associated applications can lead to loss of productivity, customers, and current or future business opportunity — all of which can lead to loss of revenue and profit, negative brand recognition, and for public companies, a negative influence on stock prices and market conditions.

The Need for IT Availability Is Pervasive

The growing desire for companies to engage in "real-time" business is pervasive. For many, the reality of real-time business analytics has become a crucial component to leverage business information and increase success. This is particularly true for online retail or other Web-related businesses and services such as reservation systems or news/media. In a highly competitive market, the loss of an online presence directly relates to loss of business.

The need for high availability does not end with Web services but rather touches upon a broad spectrum of industries. Whether it's for applications to support healthcare, higher education, just-in-time manufacturing or order fulfillment - a successful operation increasingly depends on significant uptime to satisfy business needs and a positive perception. High-transaction databases, such as those deployed to support financial institutions, will greatly depend on a highly available IT solution. There is a growing understanding, if not expectation, among mainstream enterprise organizations — in addition to specific vertical industries — that IT will always be ready to serve not only external customers but also internal employees (e.g., file sharing and email access among users within a corporation).

Finally, globalization of business has a tremendous impact on the desire to have 24 x 7 uptime. This is certainly true for large corporations that may already have a global presence, and it is increasingly true for companies of all sizes as the desire to reach untapped global markets expands to a broader set of businesses.

It is important to realize that the overall availability of any IT solution will be influenced by several components, including the application itself and associated software, as well as the hardware components, including servers, networking, and storage. Overall IT availability is predicated on application availability, which is predicated on system availability, including both server and storage systems.

The following sections of this paper discuss the significance of application and data availability within the modern datacenter, the underlying foundation and importance of constant data access, and the role of enterprise storage systems, including those from NetApp.

Application Availability

An overall IT solution consists of several components, but it is most simply expressed as a set of applications and related software running on specific hardware resources such as servers, networking, and storage assets. The applications themselves are fundamentally the lifeblood of the IT solution, enabling an organization to enhance productivity, streamline and automate processes, and leverage business information to generate revenue, improve profitability, and comply with regulatory or business governance.

The relative importance of any one particular application may vary among disparate organizations. For instance, one organization may tolerate an outage in an application, such as email, while another organization may have zero tolerance. Whether large or small, however, organizations of all sizes and industries are evolving toward common requirements around application availability and uptime.

While the types and numbers of applications will vary tremendously among organizations, there is a universal goal for increased uptime with as little downtime as possible, particularly unplanned downtime. Whether organizations deploy common database applications such as Oracle, SAP, or SQL to support ongoing business transactions; email and messaging applications such as Microsoft Exchange; or general file-sharing or directory services, the distinction between "mission-critical" and "mission-trivial" applications continues to blur. CIOs and IT administrators seek to

deploy reliable and available solutions throughout the organization, not just on the front lines.

The Role of Reliable Storage Systems to Support Application & Data Availability

The reliance on resilient storage systems has always been an important component to ensure timely access to data and availability of applications. However, the role of storage is increasingly critical as enterprise organizations adopt evolving datacenter environments to address a growing set of business and operational demands.

New Datacenter Environments

Driven from a new set of business realities in the datacenter, organizations are accelerating their efforts to reduce and optimize capital and operational IT spending. These organizations are seeking new approaches and technologies that can help reduce or possibly eliminate the amount of IT equipment needed while at the same time move to a more predictable operational environment.

Ultimately, the quest to ease management and deliver more efficient, faster and flexible IT services to internal business units and external customers results in next-generation datacenter solutions incorporating the following elements:

- ☒ **Virtualized Environments** – server virtualization represents the massive effort to reduce capital spending and increase operational efficiency.
- ☒ **Shared Storage Environments** – shared, networked, and virtualized storage to support a virtualized server environment and to capitalize on improved management and utilization efficiencies.
- ☒ **Multiple Applications and Multi-Tenancy** – virtualization and shared storage provide a foundation to manage multiple applications across multiple users or business units within a common environment.
- ☒ **Maturing Cloud Environments** – an increasing effort to adopt technologies and build infrastructure that can deliver a flexible, scalable IT delivery model based on internal and/or external, third-party service offerings. For public service providers looking to build infrastructure to support cloud-based offerings, the elements of virtualization, shared storage, and multi-tenancy become crucial.

From a storage perspective, datacenter environments that adopt the elements described above will need to deploy systems and solutions that provide reliable access to stored information in a virtualized, shared and increasingly multi-tenant environment.

Access to Data Is Crucial

Just as application availability is the foundation for a successful IT deployment, so too is data availability the foundation for a successful and reliable application deployment. For any application, the preservation of and continuous access to data is critical. The loss of data or access to it means downtime for most applications. This has a direct impact, which places utmost importance on the availability of both primary and

secondary storage systems that are deployed to store, provide access to, manage, and protect business data.

Integration of Storage with Application

Storage systems that are optimized to work closely with specific applications provide an excellent customer experience for a variety of reasons. First, the ease of use and management of a storage system in supporting a specific application creates an improved customer experience. The complexities of allocating, administering, and managing storage resources to support a specific application can be greatly reduced. For database applications such as Oracle or SQL, the working relationship between the database and storage administrators can be dramatically improved if the storage system is tightly integrated with the performance and capacity characteristics required by the database administrator to support the application.

Second, the tight integration and optimization of a storage system to support a specific application can greatly reduce the risk that the storage system may fail in supporting that application. This has tremendous impact and further increases the reliability and availability characteristics of that storage system to contribute to application availability and ultimately the success of the organization.

Customers among organizations of all sizes have come to realize the importance and benefit of having storage systems that can be optimized to support their important applications. In addition, customers have grown to expect availability on these storage systems to be extremely high, as this system availability is critical to support the data access and ultimately the application availability that is critical to a successful organization. This expectation has driven many customers to explore options based not only on typical metrics around performance, capacity, scalability, and ability to integrate with specific applications but also on the high availability and reliability of any given storage system.

The NetApp Solution

System Availability: Nothing New for NetApp

One such company that is addressing the high-availability needs for users among a wide variety of applications across many industries is NetApp. From conception of the product and initial production deployments in 1992, generations of NetApp disk storage systems have been designed to provide highly reliable data access to applications. NetApp's legacy of developing solutions that are easy to deploy and manage, combined with compelling high-availability characteristics, has created an impressive base of loyal customers who rely on NetApp storage to serve as a reliable data foundation for application availability.

NetApp's current line of disk storage systems consists of the following product series:

- FAS2000 Series** — systems scaling to 272TB, for entry-level or distributed midsize enterprises

- ☒ **FAS3100 & FAS3200 Series** — systems scaling to 1920TB, for greater capacity, availability, and performance supporting business and virtualized applications at enterprises
- ☒ **FAS6000 & FAS6200 Series** — systems scaling to 2880TB, for demanding datacenter applications and virtualized environments at enterprises and service providers

All NetApp systems run on the same operating system, Data ONTAP. The vast majority of FAS systems deployed run on version 7G of Data ONTAP. However the company also ships a scale-out, clustered version of the operating system with Data ONTAP 8. The clustered version of Data ONTAP uses global namespaces to provide continuous data availability when moving data within the cluster. Further, all FAS storage systems offer multiprotocol functionality, including simultaneous connectivity via NAS (e.g. CIFS and NFS) or SAN block protocols (e.g. Fibre Channel and iSCSI).

In addition, NetApp offers an expansive suite of software tools that enhances the customer experience in a variety of ways. Along with providing ease of deployment and administration, NetApp's wide range of available software features - such as Snapshot, FlexVol, and FlexClone - arms administrators with a powerful set of advanced data management tools designed to enable them to optimize the storage, management, and retrieval of business information.

What's Behind NetApp's High Availability?

From the initial architecture conceived in 1992 through today's products, NetApp storage systems are designed from the ground up to ensure high system availability. Several components contribute to the high-availability characteristics of NetApp's storage system solutions, fundamentally comprising hardware and software elements. While the hardware side generally consists of leveraging well-tested, industry-standard and redundant components, the software capabilities are equally important to provide data access and ultimately application uptime even during certain hardware failures.

Hardware

From the beginning, NetApp FAS systems have been designed to maximize uptime and protect against hardware failure. This is true of the controller units (diskless "gateways") that house the intelligence of the system as well as the storage subsystem (the disks and associated enclosures such as the DS4243 and DS2246 SAS-based storage enclosures.) In addition, NetApp systems leverage industry-standard components, such as the latest 64-bit, multi-core processors and well-tested SAS, SATA, and Fibre Channel hard disk drives and solid state disk storage, all of which are proven in numerous systems and environments.

Plus, the ability to deploy FAS storage controllers in pairs to create a high-availability solution protects against system downtime due to hardware failure. This hardware redundancy is facilitated through advanced software functionality, which is described in the following sections, along with other software features that contribute to the overall high-availability characteristics of NetApp storage solutions.

Software

To complement the highly reliable hardware architecture, NetApp has consistently devoted a significant amount of attention to develop sophisticated software features to minimize downtime and ensure continuous data access even during certain hardware failures or other data corruption risks. NetApp software features that address uptime and data availability can be categorized into the following, broad categories:

Software features inherent to the Data ONTAP operating system:

- RAID-DP (the NetApp implementation of RAID 6) to protect against data loss associated with double disk failures within a RAID group
- Checksums for each 4Kb block of written data to protect against firmware bugs or other hardware/data path problems
- Automatic detection of and protection against degrading/failing drives
- WAFL's ability to detect and prevent "lost writes"
- Snapshot technology designed to efficiently provide point-in-time copies of the storage system, which sets the stage for data recovery and other disaster recovery scenarios
- RAID Reconstruction Prioritization to ensure that doubly-degraded RAID groups (those at immediate risk of data loss) are reconstructed before singly-degraded RAID groups (those not at immediate risk of data loss).
- Multipath HA to provide secondary path connections to all storage shelves attached to the system
- IOM3 and IOM6 shelf module (SAS Expander) to provide direct disk access such that all disks are isolated from each other so disk errors don't propagate.

Software that addresses planned or unplanned downtime:

- SyncMirror and SnapRestore – these optional software features extend the value of Snapshot technology by providing robust synchronous mirroring and data recovery.
- HA Pair Controller – a data service deployed on a dual-controller FAS storage system that enables an automatic, active-active failover environment. This software protects against planned downtime should one or both storage systems need to be taken down for repair or maintenance. It also protects against controller or storage failure in a single path HA configuration.
- MetroCluster – provides an extra layer of isolation and protection for an HA Pair Controller configuration. The systems can be separated across distance in the same building or between sites located miles apart to protect against

data loss in the case of disruption. It includes synchronous mirroring to ensure zero Recovery Point Objectives (RPO). The solution also complements VMware vSphere HA and FT for an end-to-end continuous availability solution.

- ❑ Alternate Controller Path (ACP) recovers an I/O module during certain fault scenarios by reset or power cycle. ACP is an out-of-band management path so control over the shelf does not require the same communication path as the data.
- ❑ Additional options through the NetApp Integrated Data Protection offerings include SnapVault for disk-to-disk backup and recovery, SnapMirror for Disaster Recovery, SnapLock for Compliance as well as MetroCluster for continuous data availability (listed above).

In addition to the specific software features listed above, the NetApp unified storage architecture enables higher levels of availability by easing the deployment and administration of NetApp FAS storage systems via a single storage operating platform (Data ONTAP). By reducing the complexities associated with deploying storage assets with disparate operating platforms, the unified storage architecture approach from NetApp allows greater confidence to administer and manage data among multiple systems and provide continuous access among a variety of applications.

Further, to complement NetApp's ability to provide reliable and effective file-sharing capabilities, the company's suite of application-specific software modules, in the form of SnapManager, helps to optimize the integration between application and storage. This further reduces the possibility of a storage- or data-related failure that is particular to the specific application. NetApp has developed SnapManager modules for the following business applications:

- ☒ Oracle
- ☒ SAP
- ☒ SQL Server
- ☒ Microsoft Exchange
- ☒ Microsoft SharePoint Server

NetApp Monitors System Availability

NetApp has taken steps to monitor the reliability and uptime of its FAS disk storage systems across a significant portion of its installed base. In an effort to provide remote serviceability and event alerting to as large a customer population as possible, the company monitors all NetApp FAS storage systems for system availability and downtime for all customers that participate in a service plan. The service plan can be administered directly by NetApp or through a NetApp Global Services partner.

AutoSupport Monitoring

To monitor system availability, NetApp uses direct remote monitoring and its automated "dial home" service called AutoSupport (also known as ASUP). About 80% of all NetApp FAS storage systems are deployed with AutoSupport enabled and in use, which provides a significant portion of the NetApp customer population with direct, real-time event monitoring. The percentage of systems utilizing AutoSupport is greater among higher-end systems such as the FAS6000 series or large FAS3100 systems versus midrange or entry-level systems such as smaller FAS3100 or FAS2000 products.

NetApp evaluates the availability of its disk storage systems that are serviced by NetApp or service partners for total possible runtime, downtime, and uptime. Runtime hours are the total number of hours (in a given time period) possible for a population of a particular array model to be up and running properly. NetApp tracks the total hours of unplanned downtime resulting from the following possible causes:

- Any hardware failure leading to an unplanned outage caused by failure of either the storage controllers (e.g., storage processors, network cards) or the disk subsystem (e.g., disk drives, enclosures, cabling, power and cooling components)
- Any software failure leading to an unplanned outage caused by the Data ONTAP operating system, the WAFL file system, or any of the associated NetApp software products deployed for data protection and data management (e.g., SyncMirror, SnapRestore, MetroCluster, SnapManager, FlexVol, FlexClone)
- Any planned outages resulting from a recognized failure (i.e., the downtime associated with retroactively resolving an identified system failure)

As part of the availability calculation, NetApp monitors but does not include failures induced by the following:

- Datacenter or site power failures due to energy provider or natural disaster
- Datacenter or site environmental failures (temperature falls above or below specifications)
- Operator-induced downtime that is not associated with a recognized system failure (i.e., planned outages, hardware upgrades, or other maintenance)

Based on the preceding criteria, the total amount of downtime is calculated over the measurable time period and subtracted from the total runtime, yielding the total uptime for the defined population of storage systems. Availability is then calculated by producing the ratio of total uptime to total runtime.

NetApp currently monitors over 100,000 deployed storage systems on a monthly basis (those with active AutoSupport service), spanning systems across the FAS6000, FAS3100, FAS3000, and FAS2000 series, as well as several older products still deployed throughout the installed base. To be included in the calculation, the system must be reporting on AutoSupport for at least one month. In

addition, the calculation includes those systems running on Data ONTAP versions 7.2 or greater. Further, the calculation excludes downtime not associated with case creation (e.g., planned downtime resulting from test case scenarios).

IDC has found NetApp's methodology to be rigorous and acceptable for measuring NetApp storage system availability in terms of percentage of uptime over a given time period. This sound method for monitoring and the associated analysis have led NetApp to determine that its monitored storage systems have achieved greater than 99.999% uptime as measured on a rolling basis between December 2009 and September 2010.

Storage Availability Audit

In addition to constant monitoring of system availability and other system characteristics, NetApp provides other services for AutoSupport customers, such as a reoccurring Storage Availability Audit. This automated summary report leverages the information obtained from the AutoSupport tracking data and allows customers to review the availability characteristics of their NetApp storage systems, enabling them to make adjustments or to identify potential issues before they occur. With this service, NetApp is not only identifying best practices around capacity and performance utilization and system availability but also proactively helping customers to optimize those best practices. A quarterly Storage Availability Audit is included with NetApp SupportEdge Standard service, and onsite physical inspections are provided with NetApp SupportEdge Premium service.

CHALLENGES/OPPORTUNITIES

NetApp's FAS storage systems with its built-in and value-add resiliency features help customers meet application uptime requirements by providing a solid foundation for data access to critical applications. As the company strives to build upon its reputation for providing a positive customer experience, it should consider the following:

- ☒ Availability of storage systems across all suppliers in the market will continue to evolve. Over the long term, high availability will continue to move toward becoming a "checklist" item, if not an assumed requirement. While NetApp has established a meaningful focus in this area, particularly around the ability to recover from hardware failure and provide data and application availability, the company will need to continue the evolution of its total storage solution availability and differentiate its offerings from those of its competitors as solutions become more sophisticated.
- ☒ As NetApp continues to expand its strategic partnerships with other technology suppliers to pursue deployments of IT that span elements of servers, storage, networking and virtualization, it will need to ensure that the high-availability characteristics of its own solutions are not diffused by other components. This will be an increasing challenge shared among all storage suppliers as partnerships expand to address "holistic" IT deployments in the datacenter.

- ☒ With the continuing evolution of NetApp's Data ONTAP operating system, the company must ensure that similar, if not better, levels of availability are achieved across an environment of multiple storage systems working together to provide increased levels of scalability.

CONCLUSION

IT has grown to become a vital component within successful organizations. While the numbers and types of application deployments may vary among a wide variety of organizational sizes and verticals, the reliance on IT solutions is causing availability and uptime to become paramount across the board.

In particular, the persistent availability of the applications is top of mind for many IT administrators and CIOs, as these applications are at the heart of maintaining a successful IT deployment. This desire for high application availability is dependent on continuous access to the supporting data, which, for a growing number of organizations, resides on networked disk storage systems.

The importance, therefore, of highly available networked storage systems becomes a crucial component. By focusing on both hardware and software aspects, NetApp delivers highly available disk storage solutions that are crucial in today's highly virtualized environments supporting multiple applications across multiple tenants. The company not only has designed a highly reliable hardware architecture but also, more importantly, has developed a family of software products and operating system features that are specifically aimed at providing very high levels of system availability.

In addition, NetApp's ability to tighten the relationship between popular applications and the supporting storage has enabled customers to deploy effective and reliable data access to many important database applications or file-sharing activities, thus effectively contributing to the ever crucial application uptime.

With NetApp's rigorous approach to measuring system availability and the ability to leverage that information with frequent Storage Availability Audits, the company continually builds upon strong relationships within its customer base to improve the quality and reliability of storage deployments, not to mention improving the customer experience. With current system uptime measured by NetApp to be over 99.999%, customers can be assured that NetApp disk storage solutions can play a vital role in supporting the application and data availability needed for success.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2010 IDC. Reproduction without written permission is completely forbidden.