



White Paper

Improving Economics and Business Workflows by Using a Self-Protecting Data Infrastructure

John Martin
Field Center of Innovation, NetApp

October 2018 | WP-7287

Abstract

Overview of NetApp® Data Fabric platform features and solutions that help enterprises mitigate the risks of data loss, while maintaining an agile IT infrastructure.

TABLE OF CONTENTS

1	Introduction.....	3
2	Business Challenge	3
3	The Problem with Traditional Approaches to Backup and Archive.....	3
3.1	Data Protection Requirements Are Not Aligned with Current Business Outcomes	3
3.2	Most Backup Systems Fail to Satisfy Regulatory Requirements	4
3.3	Tape Is Unsuitable for Long-Term Archiving.....	4
4	Best Practices for Formulating Data Protection Policies.....	5
4.1	Optimize for Both Planned and Unplanned Downtime	5
4.2	Identify Historic and Projected Causes of Data Loss.....	6
4.3	Build Business-Aligned Data Protection Definitions and Requirements	7
4.4	Include Backup Reliability, Administrative Overhead, and Overall Costs in Methodology Selection	11
4.5	Build Recovery Time Objectives and Restore Reliability as Core Requirements	16
4.6	Include Impact of Data Protection Methods on Higher-Level Business Outcomes.....	21
	Conclusion.....	23
	Version History.....	23

LIST OF TABLES

Table 1)	Causes of data loss: perception and reality.	6
Table 2)	Data protection effectiveness by technology type across various failure domains.	7
Table 3)	Typical recovery point objectives.	10
Table 4)	Backup reliability, complexity, and costs.	12
Table 5)	Recovery time objective effectiveness.....	17
Table 6)	Restore effort and reliability.	19
Table 7)	Business effects of backup.	21

LIST OF FIGURES

Figure 1)	NetApp Data Protection Economic Comparison – Cumulative Costs Over Time.	14
Figure 2)	NetApp Data Protection Economic Comparison – Cost Breakdown.....	15

1 Introduction

Despite impressive advances in technology, the unabated, exponential growth of data makes it difficult for enterprises to consistently protect and restore their business-critical information. This situation significantly increases the risk of downtime and data loss, both of which can have a negative impact on revenue, reputation, and legal and regulatory compliance. Additionally, this growth affects alignment with service-level objectives, impedes business agility, and hinders the release of revenue-generating products and services. The comprehensive NetApp Data Fabric platform features solutions that help enterprises mitigate the risks of data loss, while maintaining an agile IT infrastructure that supports key business initiatives.

2 Business Challenge

As data continues to grow, many organizations are finding that the costs and challenges of managing that data grow almost as fast as the data itself. Of these challenges, backup is often cited as the most difficult and costly problem to solve. Traditional backup methods that require whole copies of data to be pulled from storage systems and sent to physical and virtual tape libraries require large and costly infrastructures that act only to mitigate risk and add nothing to top-line revenue. Unfortunately, when their effectiveness is measured objectively, it becomes clear that they often fail in this primary role, because they are not able to meet recovery windows, lack sufficient reliability, and don't work well in a virtualized infrastructure and in cloud contexts.

3 The Problem with Traditional Approaches to Backup and Archive

IT professionals have at their disposal a variety of tools to help businesses achieve their data availability goals, including backup, restoration, replication, and recovery. However, it's crucial to stay focused on the actual goal, which is the availability of the data, and to achieve that goal by using the right set of tools for the specific job. And it's necessary to balance availability with other concepts such as data importance and business criticality, budget, time to deploy, operational capability, and costs of downtime.

Having stated that a data protection and retention solution should be designed from a balanced set of well-defined business requirements, note that this case is rare. Instead, data protection solutions are designed and propagated based on inherited requirements that:

- Were not originally specified in consultation with the business
- Have not been adequately reviewed
- Are not appropriate in the current regulatory, corporate, or technical environment

Unfortunately, backup, like insurance, is an aspect of risk management, and every dollar spent on backup is a dollar that cannot be spent on IT projects that have the potential to improve both top- and bottom-line business outcomes. Therefore, it's important for IT management to ensure that no dollar spent on data protection is wasted on ineffective and outdated strategies.

3.1 Data Protection Requirements Are Not Aligned with Current Business Outcomes

Data protection processes and planning often revolve around the unsettling question "How much can you afford to lose?" This question rarely has a balanced, well-thought-out response. Although it's understandable for the business to want no data to ever be lost under any circumstances and all data created at any time to be kept indefinitely, the costs of doing so are prohibitive.

Rather than asking the difficult questions, many IT organizations fail to engage in any meaningful dialog around data protection with the business and fall into the habit of continuing whatever has gone before, implementing data protection regimes similar to the following:

- Changed data is backed up to tape or a purpose-built backup appliance (PBBA) every night and kept for 2 to 4 weeks.

- All data is backed up to tape or PBBA once a week and kept for 4 to 10 weeks.
- All data is backed up to tape or PBBA once a month and kept for 7 years.

In addition, to protect against site failures, IT might also perform the following procedures:

- Backup tapes or copies of backup tapes are sent off site.
- Mission-critical data is synchronously replicated via dedicated high-speed networks to an alternate data center.

The advantage of these data protection regimes is that they are well known and serve as a general catch-all solution, which is signed off by management as standard IT practice. It's surprising that this practice continues despite numerous examples of its inadequacy. For example, more than 25% of enterprise IT organizations report being either dissatisfied or very dissatisfied with traditional backup infrastructures, and synchronous replication is notorious for instantaneously replicating data corruption across two sites in a variety of data loss events.

3.2 Most Backup Systems Fail to Satisfy Regulatory Requirements

In Australia, Sections 9 and 286 of the Corporations Act 2001 state that the following information needs to be kept: "Financial Records (invoices, receipts, orders for the payment of money, bills of exchange, cheques, promissory notes, vouchers and other documents of prime entry; and such working papers and other documents as are necessary to explain the methods and calculations by which accounts are made up) that correctly record and explain the transactions (including any transactions as trustee) and would enable true and fair financial statements to be prepared."

This type of legislation drives the vast majority of the 7-year retention requirement in traditional backup. This requirement is applied as a blanket policy across all data types, regardless of whether or not the data meets the definition of financial records. This situation often results in large amounts of data being kept with little or no business justification.

Unfortunately, even for data that does need to be kept for compliance reasons, the "keep monthly backups for 7 years" policy does not completely satisfy the preceding requirement. For example, consider a spreadsheet that meets the preceding definition of a working paper that was created on the 4th day of the month, used as the basis for a transaction on the 6th day of the month, and then inadvertently deleted or changed on the 9th day. In a typical data protection environment, there is no guarantee that any document of this type will appear in the monthly archives because they are created after the previous month's backup and destroyed before the current month's backup takes place.

A number of other regulations require data to be kept for a certain period of time after a specific event has passed. One example of such a regulation is the Workplace Relations Act, which requires pay slips to be kept for 7 years after employment is terminated. In the case of an employee who has been working for 5 years, the "keep monthly backups for 7 years" retention policy would begin to cause potential noncompliance 2 years after the termination of that employee. As a final complication, the Privacy Act of 1988 states that an organization must take reasonable steps to destroy or permanently redact personal information if it is no longer needed, a legal requirement for which compliance might prove very difficult if traditional tape or PBBA-style backups are the primary method used for data archiving.

3.3 Tape Is Unsuitable for Long-Term Archiving

In addition to the difficulty of expunging data that should no longer be kept, tape is a poor choice for long-term archives for two reasons.

First, in many jurisdictions, the legal requirements for electronic transactions and archiving procedures for electronic records require data retention methods to allow for changes in technology. Tape has a poor track record in meeting this requirement. As a case in point, the LTO specification, introduced in 1999, requires only that an LTO drive be able to read tapes two generations back. With a new LTO generation

being released every 2 to 3 years, it might be only 6 or 7 years before a current-generation device is unable to read data from an archive set.

The second reason tape is unsuitable for long-term archiving is that compared to disk, tape is a relatively delicate contact medium, which degrades with use, can become physically damaged, and is adversely affected by changes in environmental conditions. Data stored on tape can also be lost from exposure to magnetic fields. Therefore, to reliably maintain the integrity of the information, tapes must be periodically refreshed (read and rewritten). Managing refresh cycles for hundreds or thousands of tapes written over many years is a complex and costly task with potentially serious consequences for improper management.

Can Tape Really Keep Data for Long Periods of Time?

Some tape media, such as LTO-5, are often touted as having a 30-year archival life. For a technology that is less than 7 years old, this kind of claim can be relied on only through a fair degree of faith in the vendors' statistical analysis techniques. IT and business management are asked to take this leap of faith while accepting that, unlike disk, there is little or no hard data published for tape about mean time to data loss or annual failure rates under various conditions. Given the high rates of dissatisfaction with tape-based backup, vendor claims of long-term reliability might need to be viewed with skepticism.

Tape Is Only as Good as Its Handlers

One major failing of tape is not the technology itself, but the way in which it is treated. In many cases the staff entrusted with tape management and movement are in entry-level IT positions or are third-party couriers. Even tape media manufacturers acknowledge that the published archival lifetimes are only for tapes kept in optimum operating and storage conditions of 16°C to 25°C, relative humidity 20% to 50%, and no shock or vibration, none of which apply to the courier vans typically used to transport tapes to off-site locations.

In addition, tape drives must themselves be subject to rigorous preventive maintenance. The reason for this requirement is that dirt and debris accumulated on heads, roller guides, and other transport assemblies might get transferred to tape used in a drive that was not well maintained. When these dirty tapes are subsequently used in a good drive, they might transfer those contaminants and degrade a previously clean drive. As the new drive becomes contaminated, a variety of problems can result, including premature wear on the read/write head and crucial parts of the drive transport. This situation leads to an even larger media impact because any new tapes that are used in the drive can also be damaged.

A further caution about using a backup application for long-term archiving is that the media is recorded in a proprietary logical format readable only by the originating application. Backup vendors have been known to discontinue backward-read compatibility for their own logical tape formats, and a change of backup vendors or products from the same vendor might make recovery from archived tapes difficult or impossible. Thus, true long-term archiving would also require archiving the entire backup system, including the computer, recording hardware, software, and multiple copies of the media.

Although many backup systems still use tape as a long-term archive medium, and much of the time this approach works reasonably well, its failures are common enough to be widely known to most IT professionals. To put this fact in perspective, if an organization under a discovery request needed the data residing on a medium that was known to fail and was almost impossible to test for evidence of degradation, how would the risk managers and legal teams feel about having only a pretty good chance of getting the information they needed?

4 Best Practices for Formulating Data Protection Policies

4.1 Optimize for Both Planned and Unplanned Downtime

Traditional data protection strategies focus on reactive response to unplanned events and the associated unplanned downtime. Yet storage and consequent server and application unavailability result from both unplanned and planned downtime. According to some estimates, over 80% of all downtime is planned. A little over 50% of this planned downtime is attributable to database backup, while the maintenance, upgrading, and replacement of application and system software, hardware, and networks typically account for most of the remaining time in this category.

Although much of this planned downtime occurs during nonbusiness hours, changes to the business environment brought about by globalization, online ordering, back-office consolidation, more flexible work practices, and many other factors mean that the number of nonbusiness hours is gradually but inexorably being reduced, while the amount of data that needs to be protected is growing exponentially.

4.2 Identify Historic and Projected Causes of Data Loss

Although outages caused by natural events, terrorist attacks, and utility failures garner significant press coverage, the more mundane day-to-day causes of data loss go unreported and generally unnoticed. Furthermore, a quick Google search on the term “causes of data loss” turns up far more results on what could be more accurately described as “illegal data access,” driven by legislation that mandates public reporting of this class of failure in information security. As NetApp founder Dave Hitz states in his blog, this kind of reporting baffles our risk intuition and results in significant amounts of resources being dedicated to solving problems that might never occur. As a case in point, a 2006 survey of more than 260 IT professionals found that the leading causes of unexpected downtime for databases over the previous year were infrastructure issues, followed by software and database glitches. Although no root cause analysis was reported for the infrastructure outages, one interesting finding from this report was that while most business continuity plans brace for external events that are beyond the direct control of organizations, few if any companies in the survey said that such events contributed to database downtime.

Table 1 shows figures from Kroll Ontrack about the causes of data loss.

Table 1) Causes of data loss: perception and reality.

Cause of Data Loss	Perception	Reality
Hardware or system problem	78%	56%
Human error	11%	26%
Software corruption or problem	7%	9%
Computer viruses	2%	4%
Disaster	1–2%	1–2%

Although the raw data and research methodology behind these figures are not available, it seems reasonable to assume that the data came from Kroll’s customer base. Kroll Ontrack’s data recovery services are heavily oriented toward workstation and laptop users, whose data rarely benefits from RAID-protected highly available disk subsystems. Therefore, it is reasonable to assume that the failure percentages caused by hardware and system problems would be far smaller in enterprise-class data centers, although the added complexity might result in a corresponding increase in the percentages for user error.

A more thorough search of available literature and IT news coverage supports the view that human error is the largest cause of data loss in enterprise environments.

4.3 Build Business-Aligned Data Protection Definitions and Requirements

Without a way of rapidly creating multiple recovery points throughout the day, the largest causes of data loss are not addressed, especially for companies that demand fine-grained recovery point objectives (RPOs). Although synchronous mirroring might satisfy zero RPO for a very small class of data loss events, the majority of failures require recovery by using the legacy bulk copy methods implemented by 99% of data protection specialists. These methods typically have an RPO of between 12 and 24 hours and a recovery time objective (RTO) that can have a significant negative impact on business outcomes.

Identify Failure Domains as the First Step in Creating a Backup Definition

A backup system needs to address numerous failure scenarios. Each one is the result of data loss in a particular failure domain. Although every IT infrastructure has its own unique set of failure domains, from a data protection perspective, these can be broadly categorized as *logical failures* and *physical failures*.

Logical Failures

- **User error.** These errors encompass everything from small failures, such as accidentally deleting or overwriting the wrong file, to major mistakes, such as dropping the wrong table space. In general, these are inadvertent errors due to poor planning.
- **Application failures.** Software bugs or upstream failure of devices might cause corrupted data to be stored.
- **Logical access by actors with malicious intent.** This situation might include external hackers, disgruntled employees, or sophisticated malware.

Physical Failures

- **Hardware storage device or array failure.** These failures might be due to storage administrator-level failures (for example, destroying the wrong LUN), software or hardware failure, or some combination.
- **Site-level failures.** Typically caused by failures in plumbing, power, and air conditioning. A variant of this failure type includes malicious actors who have physical access to the IT equipment in a site.
- **Metro-level failures.** Major weather systems, widespread fires, earthquake, electrical substation, or generator failures.

To prevent data loss in the event of a failure in any of these domains, a copy of that data must be available in a separate failure domain. This situation provides an effective definition for backup. Simply put, backup is a copy of data in a separate failure domain, and each kind of backup protects against specific kinds of failures in specific failure domains. Table 2 outlines the effectiveness of a number of data protection methods for each of the failure domains described earlier.

Table 2) Data protection effectiveness by technology type across various failure domains.

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Local object store with versioning	Yes	Yes	Yes	No	No	No	No
Geo-distributed object store with versioning	Yes	Yes	Yes	Yes	Yes	No	Yes

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Geo-distributed object store with WORM and versioning	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Geo-distributed object store with replication to a separate administrative domain	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Local PBBA	Yes	Yes	Yes	No	No	No	No
Replicated PBBA	Yes	Yes	Yes	Yes	Yes	No	Yes
Cloud backup gateway	Yes	Yes	Yes	Yes	Yes	No	Yes
Crash-consistent local Snapshot™ copies	Yes	Maybe	No	No	No	No	No
Application-aware local Snapshot copies	Yes	Yes	No	No	No	No	No
Metro-replicated Snapshot copies	Yes	Yes	Yes	Yes	No	No	Yes
Geo-replicated Snapshot copies	Yes	Yes	Yes	Yes	Yes	No	Yes
Geo-replicated Snapshot copies with compliance lock	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Site-level synchronous mirroring	No	No	Yes	No	No	No	No
Metro-level synchronous mirroring	No	No	Yes	Yes	No	No	No
Near-line on-site tape backup	Maybe	Yes	Yes	No	No	No	No
Near-line replicated virtual tape backup	Maybe	Yes	Yes	Yes	Yes	No	No
Secured offline on-site tape backup	Yes	Yes	Yes	No	No	Yes	Yes
Secured offline off-site tape backup	Yes	Yes	Yes	Yes	Maybe	Yes	Yes

Why Tape Has Been the Popular Choice and What Comes Next

Tape is often used as the backup medium of choice because it has traditionally had the broadest protection across most if not all failure domains. In particular, a tape that has been handled correctly, is offline, is off site, and is secured in some kind of tamper-proof vault with the optimal environmental conditions is safe from almost every known form of data loss. “Maybe” in Table 2 usually refers to tapes that are kept in the same metropolitan failure domain.

Tape has also been perceived as having the lowest TCO for backup, because the raw media costs in terms of dollars per raw gigabyte are the lowest. However, when combined with software, lifecycle, and handling costs, tape can be one of the most expensive forms of data protection. In contrast, well-configured modern self-protecting data strategies achieve similar or better levels of protection with higher reliability and significantly lower TCO.

The Dangers of On-Site and Online Backup

Tapes and PBBA-based backups that are kept on site and near-line, where tapes and backup sets can be relabeled or erased with relative ease, are subject to administrator error and attack by malicious actors with privileged access. Backups that are kept on site only do not protect against site-level or metro-level failures, even if they are only unavailable for restore. This weakness was noted in the first bomb attack on the World Trade Center, in 1993, when data backups were not actually destroyed, but were sufficiently inaccessible to cause significant business disruption.

The Challenge with Crash-Consistent Snapshot Copies

In contrast, crash-consistent NetApp Snapshot™ copies are generally not seen as “real” backups because they protect against just one, or arguably two, failure scenarios. The caution against application-level protection is due to the lack of guarantee that a crash-consistent backup of an application or database will be recoverable. That said, database technology is sufficiently robust that the probability of successful recovery from a crash-consistent backup is very high. In the rare case of a “torn page” or similar failure, a rollback to a previous crash-consistent recovery point is likely to work well. Given the frequency with which crash-consistent Snapshot backups can be made, a recovery is almost guaranteed. This recovery usually happens with recovery point objectives and recovery time objectives that are superior to those provided by any other data protection method.

Application-Consistent Snapshot Copies

To resolve the issue of crash consistency, application and database vendors often provide capabilities to create application-aware or application-consistent Snapshot copies. For example, SQL Server includes the Microsoft SQL Writer Process, which provides an interface that allows a consistent copy of the database to be made either by using a Snapshot copy or by transferring copies of the database tables to a separate device. For many years, Oracle has provided interfaces, such as the “alter tablespace begin backup” command, to facilitate similar capabilities. In a modern data center, most applications have a similar ability to quiesce datasets for backup. When these interfaces are used by products such as NetApp SnapCenter®, a completely consistent backup of the application can be made and restored. Furthermore, database maintenance tasks such as consistency checks and log truncation can be scheduled with minimal impact to the application by allowing I/O-intensive tasks to be offloaded to a secondary host.

Site Failure Is a Bigger Problem Than Hardware Reliability

Because modern storage arrays have mature software and redundant hardware components, including high-speed erasure coding, data scrubbing, proactive identification of lost writes, and end-to-end data integrity features, the site in which the array sits is now often a more common cause of failure than the array itself. For this reason, replication of data to multiple physical locations is considered a best practice. This replication is usually done in the metro failure domain to a site that contains sufficient compute capability for it to act as a disaster recovery site in the case of site-level failure. In addition, many companies are looking to the major cloud vendors to create a third copy, or “data bunker,” to host their data outside of the metro failure domain. This approach allows off-site replication to continue, even after site-level failure has caused failover to a disaster recovery facility.

This combination of Snapshot copies or versioning plus replication and cloud integration forms the core of the NetApp self-protecting data infrastructure approach. It offers a more cost-effective and superior level of protection against more failure scenarios than any other solution.

The Limitations of Depending on Only Synchronous Mirroring

Synchronous-level mirroring in and between sites provides protection for physical-level events. However, it typically does not provide any protection to logical-level or metro-level failures. It is included here only because of its common use as a requirement for the protection of mission-critical data where the tolerance for data loss is low. However, this approach often gives a false sense of comfort because it fails to protect against the most common forms of massive data loss, such as administrator error and crypto-locker malware.

Solving Multiple Problems with Geographically Distributed Object Stores

In the case of geographically distributed object stores, the concepts of primary and secondary sites are blurred, and data is physically placed in multiple locations on an eventually consistent basis by using a combination of replication and erasure coding. With a NetApp StorageGRID® object-based storage solution, the locations and level of redundancy are determined by a metadata-driven policy at ingest and can change over time as the data’s value moves through its lifecycle on an object level of granularity.

Design Solutions Based on Recovery Point Objectives

Data protection design often centers around the uncomfortable question “How much data can be lost?” The answer, expressed in a period of time from when the data loss event occurred, is called the recovery point objective (RPO). Table 3 outlines some typical values for data loss for each of the failure domains and backup technologies.

Table 3) Typical recovery point objectives.

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Local object store with versioning	0	0	0	Undef	Undef	Undef	Undef
Geo-distributed object store with versioning	0	0	0	5 min.	5 min.	Undef	5 min.
Geo-distributed object store with WORM and versioning	0	0	0	5 min.	5 min.	0	5 min.
Geo-distributed object store with replication to a separately secured administrative domain	0	0	0	5 min.	5 min.	12 hr.	0
Local PBBA	12 hr.	12 hr.	12 hr.	Undef	Undef	Undef	Undef
Replicated PBBA	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	Undef	12 hr.
Cloud backup gateway	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	Undef	12 hr.

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Crash-consistent local Snapshot copies	1 hr.	<4 hr.	Undef	Undef	Undef	Undef	Undef
Application-aware local Snapshot copies	12 hr.	12 hr.	Undef	Undef	Undef	Undef	Undef
Metro-replicated Snapshot copies	1 hr.	<4 hr.	1 hr.	1 hr.	Undef	Undef	Undef
Geo-replicated Snapshot copies	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	Undef	12 hr.
Site-level synchronous mirroring	Undef	Undef	0	Undef	Undef	Undef	Undef
Metro-level synchronous mirroring	Undef	Undef	0	0	Undef	Undef	0
Near-line on-site tape backup	12 hr.	12 hr.	12 hr.	Undef	Undef	Undef	Undef
Secured offline on-site tape backup	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	Undef
Secured offline off-site tape backup	24+ hr.	24+ hr.	24+ hr.	24+ hr.	24+ hr.	24+ hr.	24+ hr.

Table 3 shows the superiority of the self-protecting data infrastructure, where approaches based on versioning and Snapshot copies are combined with replication, allowing much more frequent backups and very fine-grained RPOs.

Table 3 assumes that tape is moved off-site in the metropolitan area, which explains its vulnerability. It also assumes that tape backups are performed once a day and that there is a 12-hour elapsed period between the time the backup is complete and the time the data loss event occurs. Note that even though synchronous replication is shown as having zero RPO, many environments use this approach in conjunction with tape to protect against logical-level failure, leading to an RPO for most data loss scenarios that can be measured in hours.

4.4 Include Backup Reliability, Administrative Overhead, and Overall Costs in Selecting a Methodology

Not only must there be a design objective to ensure that data loss is kept to a minimum, but also the process and technology involved must reliably execute the policies that are designed to meet those objectives.

Assess and Mitigate Limitations of Existing Expertise and Processes to Address Issues of Reliability and Coverage

Some problems other than downtime itself stem from data protection systems that depend on manual configuration and management. Highly skilled personnel with multiple skillsets are required to manage, configure, and optimize the performance of large, distributed data protection infrastructures.

Unfortunately, these individuals are becoming more difficult to hire and retain at the same time that traditional data protection regimes are forcing them to perform most of their implementation and troubleshooting when most people would prefer to be in their beds or with their families. Although the professionalism of data protection specialists is typically high, the pressures of late nights, increasing workloads, and decreasing resources often lead to increased staff turnover with the consequent rise in data loss caused by human error.

Because of this fact, there is generally an inverse relationship between the amount of administrator involvement in the setup, tuning, and ongoing operation of the overall data resiliency architecture and improved outcomes across a broad range of technology choices, regardless of how much money is spent. (See Table 4.)

Table 4) Backup reliability, complexity, and costs.

	Reliability	Admin Burden	\$ per Protected Terabyte
Local object store with file versioning	Very high	Very low	\$
Geo-distributed object store with file versioning	Very high	Very low	\$\$
Geo-distributed object store with WORM and file versioning	Very high	Very low	\$\$
Geo-distributed object store with replication to a separately secured administrative domain	High	Low	\$\$\$
Local PBBA	Medium	High	\$\$\$\$\$
Replicated PBBA	Medium	High	\$\$\$\$\$
Cloud backup gateway	Medium	Medium	\$\$\$
Crash-consistent local Snapshot copies	Very high	Very low	\$
Application-aware local Snapshot copies	High	Medium	\$\$
Metro-replicated Snapshot copies	High	Low	\$\$\$
Geo-replicated Snapshot copies	High	Medium	\$\$\$
Site-level synchronous mirroring	Very high	Low	\$\$
Metro-level synchronous mirroring	Very high	Low	\$\$\$\$
Near-line on-site tape backup	Medium	Medium	\$\$\$

	Reliability	Admin Burden	\$ per Protected Terabyte
Secured offline on-site tape backup	Low	High	\$\$\$
Secured offline on-site tape backup	Very low	Very high	\$\$\$

Although these reliability and administrative burden factors might be debatable, they are based on the author’s 15 years of experience with multiple backup solutions from multiple backup vendors. Some of this consideration includes the duty of care shown by numerous tape off-siting companies and the practice of using relatively inexperienced IT staff to physically handle the tape. This practice lowers the overall reliability of any data protection strategy, which depends on the physical movement and handling of backup media.

The other major challenges of traditional bulk copy methods of data protection stem from trying to copy exponentially increasing amounts of data on ever-increasing disk and SSD capacities even though the throughput speeds of those disks and SSDs have not increased correspondingly. Overcoming this data density issue has led to increasingly complex engineering, which is difficult to maintain, test, and troubleshoot. For example, the cycle time to test a change in a full backup policy is in excess of 24 hours; often the test can be run only once a week. In contrast, a change in a versioning-based and replication-based data protection infrastructure can be tested in an hour. The superiority of data protection methods based on Snapshot copies, versioning, and replication derives from their maturity and elegance in how they solve the major challenges of data protection.

Assessing Costs of Various Data Protection and Resiliency Architectures

The relative costs of traditional backup compared to Snapshot copies and replication are the result of a variety of causes, including:

- Per-terabyte costs of backup software, which typically starts at \$10,000 per protected terabyte of storage before discounting. Even with aggressive discounts on licenses and support, the 3-year cost is typically more than \$US6,000 per terabyte.
- Additional infrastructure costs for data movers, high-speed storage for backup indexes, and dedicated network infrastructure.
- Opportunity costs and other hidden impacts of scheduled downtime for backup operations.

Quantified Backup Costs for Traditional Structured and Unstructured Data

The Evaluator Group quantified some of these costs at <https://www.evaluatorgroup.com/data-protection/>, which includes the interactive calculator illustrated in Figures 1 and 2.

Figure 1) NetApp data protection economic comparison: cumulative costs over time.

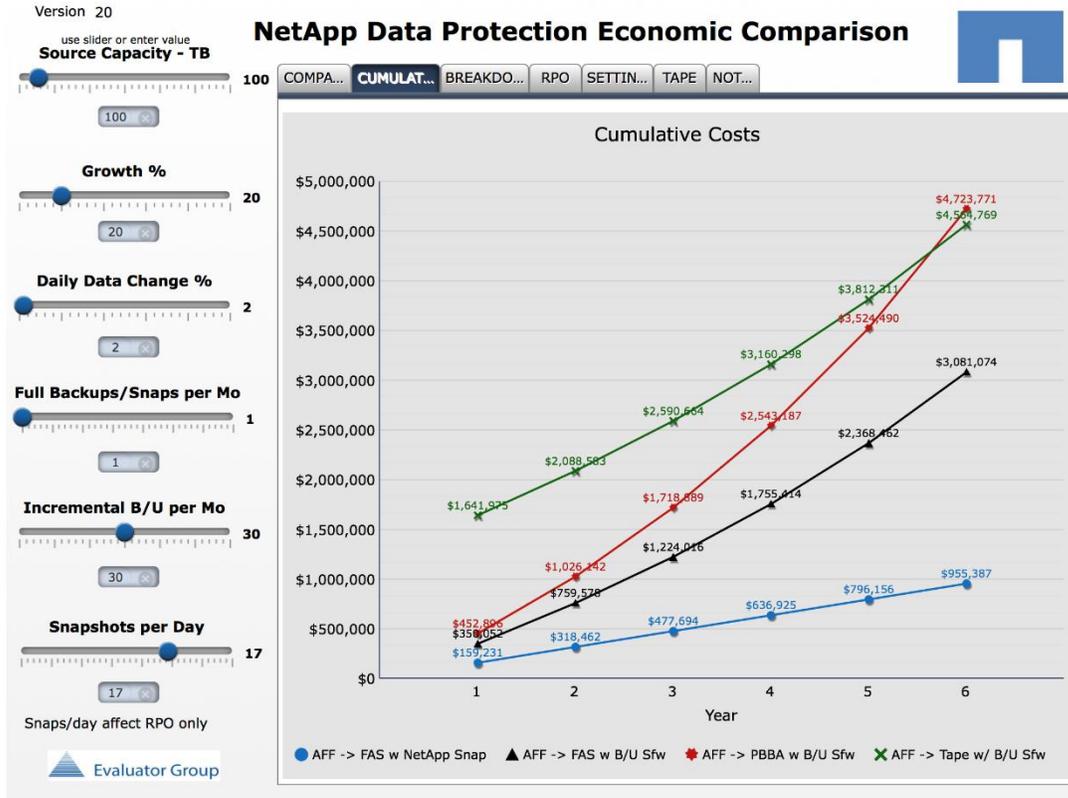
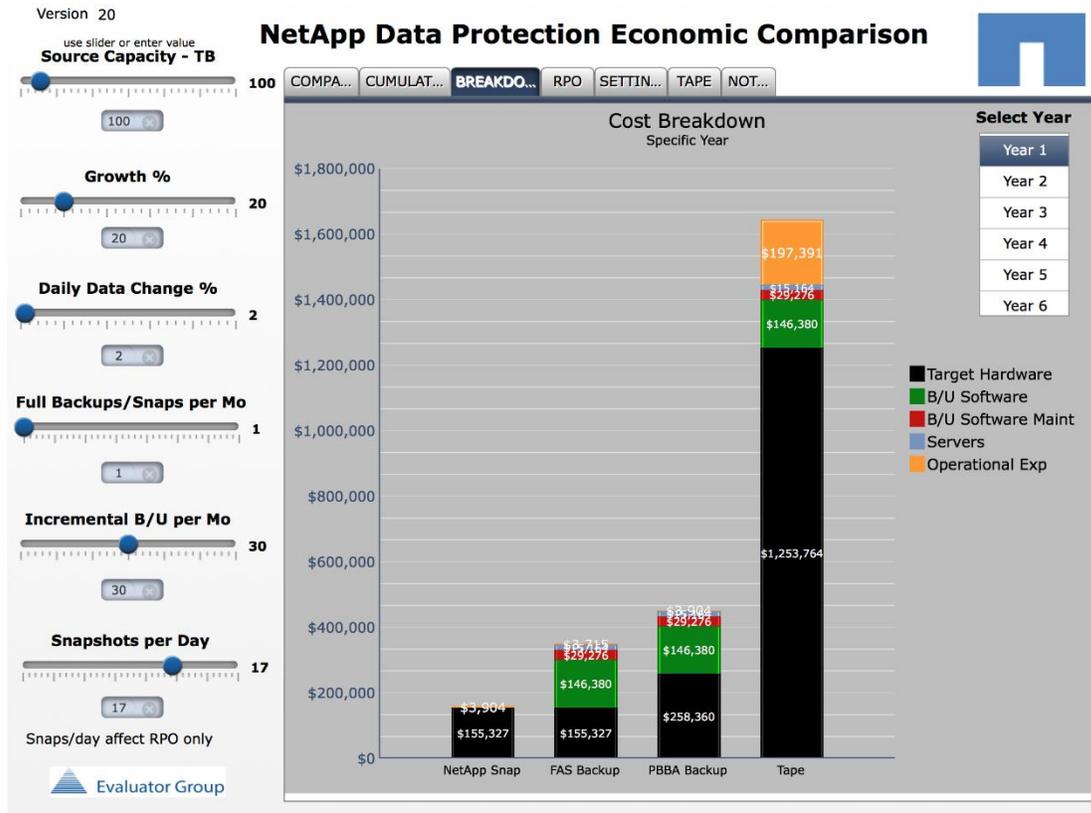


Figure 2) NetApp data protection economic comparison: cost breakdown.



Backup Costs for Large-Scale Unstructured Data in Distributed Object Stores

Snapshot copies and replication create breakthrough improvements in both cost efficiencies and reliability for data stored in traditional POSIX-style data containers. Object stores such as StorageGRID take this practice one step further by building multisite resiliency into the policies of the grid, while still holding data at a very low cost per terabyte.

For example, if a single copy of 1TB of data in an object store at a single site cost \$1,089 (the full international US\$ list price of a small StorageGRID implementation, including 3 years' service and support), extending this implementation to two sites would simply double this cost to \$2,178, while providing full resiliency across site failure and automatic zero RPO recovery. Extending this implementation to a third site and adding a further 20% to hold additional versions would increase this cost to \$3,920/TiB.

This calculation shows that even at full list price, StorageGRID is approximately half of the discounted software cost of a traditional backup system. Using the \$7,453 difference in costs of backing up 100TiB using traditional backup compared to using just Snapshot copies and replication, the Evaluator Group model shows that StorageGRID, at worst case, costs almost \$4,000 less per terabyte than traditional backup software. These cost benefits can be improved further by eliminating the costs of backup target capacity, data movers, and other ancillary hardware required to run a bulk copy backup architecture.

Furthermore, as this data ages, the previous versions could be deleted, removing the 20% additional overhead, and the data could move from three replicas to an 8+2 erasure-coding configuration, resulting in similar levels of redundancy and resiliency, while reducing the price to just \$1,362/TiB. This cost reduction happens based on lifecycle policies, while still keeping the data in a form that can be accessed immediately by applications, without needing to run through a restore process first, with only a moderate penalty to overall throughput and time to first byte. Depending on policies, this new access could allow the object to be automatically reinstated back to the beginning of the lifecycle, distributing multiple copies back into the grid and improving the throughput and time to first byte for that object.

4.5 Build Recovery Time Objectives and Restore Reliability as Core Requirements

Ultimately, disruption to a business from a data loss event is due to the time it takes to restore the data from the backup medium, along with the time it takes to recover the lost or missing data. In many cases, the lost data is written off or is reentered over time. However, the business impact caused by a loss of access to information systems during the recovery from a data loss event is the most painful experience of current data protection systems and causes the most dissatisfaction. The time that elapses between the data loss event and the eventual restoration of business services is another key performance indicator of the effectiveness of a data protection solution; it is referred to as the recovery time objective (RTO). Table 5 outlines some example RTO values, although they can vary widely, depending on the business processes that surround a data loss event.

Table 5) Recovery time objective effectiveness.

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Local object store with file versioning	5 min.	5 min.	0	Undef	Undef	Undef	Undef
Geo-distributed object store with file versioning	5 min.	5 min.	0	0	0	Undef	0
Geo-distributed object store with WORM and file versioning	5 min.	5 min.	0	0	0	0	0
Geo-distributed object store with replication to a separately secured administrative domain	5 min.	5 min.	0	0	0	<12 hr.	0
Local PBBA	12 hr.	12 hr.	1 wk.	Undef	Undef	Undef	Undef
Replicated PBBA	12 hr.	12 hr.	1 wk.	1 mo.	1 mo.	Undef	36 hr.
Cloud backup gateway	12 hr.	12 hr.	1 wk.	1 wk.	1 wk.	Undef	36 hr.
Crash-consistent local Snapshot copies	5 min.	12 hr.	Undef	Undef	Undef	Undef	Undef
Application-aware local Snapshot copies	12 hr.	12 hr.	Undef	Undef	Undef	Undef	Undef
Metro-replicated Snapshot copies	12 hr.	12 hr.	12 hr.	12 hr.	Undef	Undef	Undef
Geo-replicated Snapshot copies	12 hr.	12 hr.	12 hr.	12 hr.	12 hr.	Undef	12 hr.
Site-level synchronous mirroring	Undef	Undef	0	Undef	Undef	Undef	Undef
Metro-level synchronous mirroring	Undef	Undef	5 min.	5 min.	Undef	Undef	0
Near-line on-site tape backup	12 hr.	12 hr.	1 wk.	Undef	Undef	Undef	Undef
Secured offline on-site tape backup	12 hr.	12 hr.	1 wk.	Undef	Undef	Undef	Undef
Secured offline off-site tape backup	48 hr.	48 hr.	1 wk.	1 mo.	3 mo.	3 mo.	3 mo.

Superiority of Versioning Plus Replication

Again, as with RPOs, the combination of incremental versioning (using Snapshot copies or object versions) and replication provides the best levels of service for the most common data loss events.

Superiority of Intrasite and Intersite Active-Active Configurations

With active-active configurations spread across multiple physical array failure domains, synchronous replication offers capabilities for recovering from array-level failures automatically, with no noticeable interruption to the application (zero RTO). For site-level failures where zero RTO is also possible, levels of automation can be achieved through mechanisms such as a third-party witness. This level of functionality is a core aspect of a modern distributed object store. For more traditional POSIX and block-based datastores, Table 5 assumes that a business process that declares a disaster needs to be invoked before the service is restored at the DR site.

Evidence for Long RTO Times Using Traditional Data Protection Strategies

Similarly, the 48-hour RTO for off-site tape assumes business processing delays and transportation delays for tape and full site recoveries outside of the cloud of one month or more and reflects the time it would take to procure replacement hardware and also reconfigure a large amount of the physical infrastructure, in addition to restoring data. For a real-world example of how long it takes to recover data from tape or PBBA-based infrastructure for a major SAN outage, refer to appendix A of https://www.ato.gov.au/uploadedFiles/Content/CR/downloads/js39322_ATO-systems-report_w.pdf, which demonstrates how a well-managed response was able to restore a minimum viable environment in 12 days, and it took 4 weeks to recover the IT environment to a business-as-usual state. Note that the ATO outage was for only a partial failure of a centralized storage device, not a complete site failure.

Improving Outcomes by Combining Automation with Minimized Data Movement

Recovery from data loss events using NetApp versioning and replication-based technology such as StorageGRID, SnapManager®, and MetroCluster™ is designed to not only reduce by 95% the amount of data that needs to be transferred, but also to automate and simplify business processes. Furthermore, although backup reliability is often measured daily, the importance of reliability for restore is arguably more important, but it is rarely tested or quantified due to the limitations of legacy bulk copy techniques. Versioning and replication techniques are much easier to test for conformance, and therefore they are good candidates for automation.

Although recovery from array-level events over long-distance networks might take significantly longer than 2 days, replication-based solutions also have an option to use physically transportable media when necessary. Taking this fact into account, for each recovery scenario, even after factoring in the engineering considerations about the time it takes to move the data, the primary determining factors of RTOs are usually the processes involved in restoring business services after a data loss event. This fact underpins the importance of automation and testability in determining end-to-end recovery time.

Restore Effort and Reliability

Table 6 compares the difficulty of restore and its relationship to the reliability of various data protection strategies.

Table 6) Restore effort and reliability.

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Local object store with file versioning	Low	Low	0	Undef	Undef	Undef	Undef
Geo-distributed object store with file versioning	Low	Low	0	Low	0	Undef	0
Geo-distributed object store with WORM and file versioning	Low	Low	0	0	0	0	0
Geo-distributed object store with replication to a separately secured administrative domain	Low	Low	0	0	0	Medium to high	Low
Local PBBA	Medium	High	High	Undef	Undef	Undef	Undef
Replicated PBBA	Medium	High	High	Extreme	Extreme	Undef	Extreme
Cloud backup gateway	Medium	Medium	High	Extreme	Extreme	Undef	Extreme
Crash-consistent local Snapshot copies	Low	Medium	Undef	Undef	Undef	Undef	Undef
Application-aware local Snapshot copies	Low	Medium	Undef	Undef	Undef	Undef	Undef
Metro-replicated Snapshot copies	Low	Medium	Medium	Medium	Undef	Undef	Undef
Geo-replicated Snapshot copies	Low	Medium	Medium	High	High	Undef	Medium
Site-level synchronous mirroring	Undef	Undef	0	Undef	Undef	Undef	Undef

	User	Application	Array	Site	Metro	Malicious Actor with Privileged Logical Access	Malicious Actor with Site Access
Metro-level synchronous mirroring	Undef	Undef	0	0	Undef	Undef	0
Near-line on-site tape backup	Medium	High	Extreme	Undef	Undef	Undef	Undef
Secured offline on-site tape backup	High	High	Extreme	Undef	Undef	Undef	Undef
Secured offline off-site tape backup	High	High	Extreme	Extreme	Extreme	Extreme	Extreme

Some organizations are required by regulation or business to exercise their disaster recovery plan and prove that they can restore data in a prescribed interval. Data protection methods that do not provide easy ways to verify that the recovery methods work not only incur significant operational costs, they are also exposed to additional risks of downtime associated with recovery of data over the top of live production systems. Furthermore, the reliability of the restore is inversely proportional to the effort involved. The greater the complexity, the less likely it is to work.

In addition to the problems for tapes held for long-term storage, recovering large amounts of data from tape-based archives is particularly unreliable for most organizations. The durability of that data should not be relied on for more than 3 or 4 years without a significant investment in quality control and correction processes, along with maintaining duplicate tape media.

The major difference is in the recoverability of an application or database from a crash-consistent backup, which might require more than one attempt to find a recoverable backup. Geo-level backups are shown as being slightly less reliable because of the increased complexity of global networks and the potential impact of latency and increased recovery times on restore processes.

4.6 Include Impact of Data Protection Methods on Higher-Level Business Outcomes

Finally, a number of other concerns about data protection are independent of the failure domains, as summarized in Table 7.

Table 7) Business effects of backup.

	Cost per Recovery Point	Backup Window	Impact During Backup	Verifiability	Cloud/Multitenant	Centralized Management	Support for Continuous Delivery
Local object store with file versioning	Lowest	Instant	0	Excellent	Easy	Best	Excellent
Geo-distributed object store	Very low	Instant	0	Excellent	Easy	Best	Excellent
Geo-distributed object store with WORM and file versioning	Very low	Instant	0	Excellent	Easy	Best	Excellent
Geo-distributed object store with replication to a separately secured administrative domain	Low	Instant	0	Excellent	Easy	Good	Excellent
Local PBBA	Moderate	Long	High	Fair	Hard	Very Good	Inhibits
Replicated PBBA	High	Long	High	Fair	Hard	Very Good	Inhibits
Cloud backup gateway	Low	Long	High	Fair	Hard	Very Good	Inhibits
Crash-consistent local Snapshot copies	Very low	Instant	0	Excellent	Easy	Very Good	Excellent
Application-aware local Snapshot copies	Low	Very short	Low	Excellent	Easy	Good	Excellent

	Cost per Recovery Point	Backup Window	Impact During Backup	Verifiability	Cloud/Multitenant	Centralized Management	Support for Continuous Delivery
Metro-replicated Snapshot copies	Moderate	Short	Low	Excellent	Easy	Very Good	Very good
Geo-replicated Snapshot copies	Low	Medium	Low	Fair	Easy	Very Good	OK
Site-level synchronous mirroring	High	Instant	Very low	Fair	Easy	Good	Neutral
Metro-level synchronous mirroring	Very high	Instant	Low	Fair	Easy	Good	Neutral
Near-line on-site tape backup	Moderate	Long	High	Fair	Hard	Very Good	Inhibits
Secured offline on-site tape backup	Very low	Long	High	Poor	Hard	Good	Inhibits
Secured off-site offline tape backup	Very low	Long	High	Very poor	Hard	Fair	Inhibits

Cost Effectiveness of Snapshot and Object Versioned Backups

Even though Snapshot and object versioned backups consume space on primary storage, their density allows them to compete effectively with tape on a cost per recovery point for their primary use cases. When only changed blocks or objects are stored, the effective deduplication ratios compared to regular full backups are often well in excess of 25:1, especially for crash-consistent backups, where 100+ backups might be kept for any given day. Because fewer database and application-consistent backups are kept, the effective cost per recovery point is higher for these kinds of backup. This approach also reflects the additional costs of administration and software licenses.

Multitenancy

Multitenancy has less to do with the underlying differences between bulk copy and versioning with replication approaches than it has to do with the evolution of the management tools.

Most enterprise backup infrastructures are built around central control of the backup from a centralized console with a dedicated backup team. These have shared tape libraries and large centralized indexes and are built around a recovery workflow that involves coordination between the backup team and the application owners. This approach is necessary because the backup and recovery infrastructure is a shared resource with a historically limited number of tape drives and potential contention for single-threaded backup media. In this environment, a recovery request from one application owner could conceivably conflict with the resources required for a backup or recovery request from another.

The management of backup and recovery based on Snapshot copies and replication initially focused on the application owners themselves, providing tools designed specifically for self-service. This approach was possible because the impact of performing a backup or restore on other application owners was negligible. In addition, the Snapshot copy itself has self-contained metadata, which significantly reduces and in most cases eliminates the need for a centralized backup index to facilitate the recovery process.

The self-contained or encapsulated nature of the self-protecting data infrastructure versioning-based backup approaches lends itself to fine-grained role-based access control and self-service. However, given the delegated model, it can make centralized control and monitoring of application-based Snapshot copies more complex than the inherently centralized system maintained by traditional backup approaches.

Continuous Delivery

The ability to support continuous delivery requires that the backup system can be extensively automated and that it responds at least as rapidly as the build process requires:

- Backups complete in minutes
- Restores complete in minutes
- Supports multitenancy
- Is easily tested and verified
- Has programmer-friendly APIs

All of these are core attributes of the NetApp approach to integrated data protection and are a foundational part of the NetApp agile data infrastructure offerings.

Conclusion

Although traditional data protection technology and policies provide a good measure of protection from data loss, they face increasing challenges in an environment of extreme scale. Furthermore, as data increases, the impact on the environment of creating full copies of data prevents the kinds of continuous testing that are driven by requirements of agile methodologies and business models. NetApp integrated data protection effectively protects against the most likely causes of data loss and is easily extended to cover almost any imaginable failure scenario. It is founded on a proven set of technologies that protect data in highly virtualized environments at petabyte scale and leverage the strengths of traditional data protection strategies where appropriate.

Version History

Version	Date	Document Version History
Version 1.0	July 2018	Second working draft
Version 1.1	October 2018	Sixth working draft: includes cloud backup and object stores. Minor edits for table color and terminology consistency

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2018 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.