White Paper

NetApp: Where Better Backup Is Built-in

By Jason Buffington, Senior Analyst

May 2014

This ESG White Paper was commissioned by NetApp and is distributed under license from ESG.

Contents

Introduction 3
Looking at the Bigger Picture of Storage 3
Challenges to Overcome 4
Technical Data Protection Challenges 4
Storage Consumption Challenges 4
How Many Copies and Versions Do You Need? 4
What You Can Do Instead 5
First, Understand the Business Processes and Needs for Your Data … Then Choose Your Tools 5
To Get Different Results, Do Something Different 6
Step Back and Redefine Your Data Protection Strategy 7
Sometimes, Data Protection Comes Built-in 8
Understanding NetApp Snapshots 8
Snapshots Compared with Backups 9
Integrating Snapshots *and* Backups 10
Other Storage Capabilities That Enable Better Data Protection and More 11
Deduplication 11
Replication 11
But Wait; There's More 11
The Bigger Truth 13

Introduction

As part of its *2014 IT Spending Intentions Survey*, ESG asked IT professionals to identify their most important IT priorities for the next 12 months (see Figure 1).[1] Improving data backup and recovery, along with increasing the use of server virtualization, were among the most commonly cited priorities that the respondents expected would prompt investment by their organizations over the coming year.

Three other notable items in the same IT priority list were managing data growth (25%), regulatory compliance (23%), and business continuity and disaster recovery programs (21%, not shown). Together, these areas of focus paint a picture of IT organizations grasping for a way to get more return from their spinning-disk investments.

Figure 1. Top Ten IT Priorities for 2014

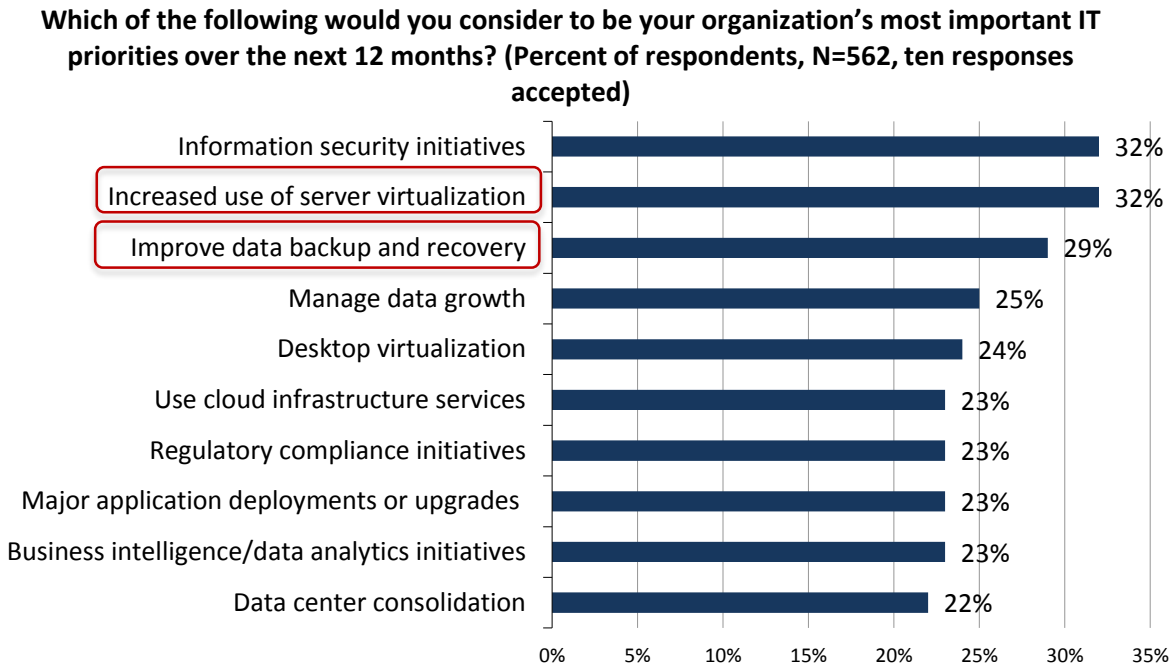

Source: Enterprise Strategy Group, 2014.

Data protection—including backup, snapshots, replication, and BC/DR preparedness—is a vitally important IT function touching all of those priorities in some manner. It is becoming even more crucial as every aspect of business becomes centered on information technology.

Interestingly, however, because backup spending is inescapable, it sometimes receives too little analysis and thought. In other words, it becomes a check-box. If the backup environment works and isn't a source of crisis, then things must be "okay." The reality, though, is that IT decision makers should be taking a *proactive* interest in addressing backup efficiency and cost.

Looking at the Bigger Picture of Storage

Traditional backup is actually on a lot of people's minds (judging by the research and spending trends above), but they shouldn't view it in a vacuum. For example, **disaster recovery** as a strategic capability also appears as one of the reported IT spending priorities for enterprises and midmarket organizations alike,[2] and both traditional backup and BC/DR are about creating and storing multiple copies in case of crisis. But other operational reasons also exist to create and access multiple copies of the same data—reasons outside of data protection and arguably even more critical to business effectiveness.

[1] Source: ESG Research Report, *2014 IT Spending Intentions Survey*, February 2014.
[2] Source: Ibid.

Those multiple copies may be:

- Copies of data for reporting that do not affect production workloads.
- Copies of data for application development and testing of patches and upgrades.
- Copies of data for customer support to recreate and diagnose problem scenarios.

Such data-heavy activities are key business enablers. But the cost of such enablement can be exorbitant storage consumption—so much so that a new discussion topic in any data protection conversation should be on "Copy Data Management" to determine how many copies the organization really needs (and where it needs them). All of those copies may not make anyone's top-ten storage investment list, but they do consume appreciable amounts of primary and near-line storage beyond what backup and BC/DR consume in secondary and tertiary capacity. What does it all point to? Much like the advice to "work smarter instead of working harder," IT needs *smarter* storage instead of more storage.

Challenges to Overcome

Considering today's chronically tight budgets, organizations would not be making storage investments tied to data protection initiatives if IT were not experiencing technical and business challenges justifying those expenditures.

Technical Data Protection Challenges

IT organizations and backup solution vendors are struggling to meet the protection and recovery demands of today's IT infrastructures. ESG research in 2012 revealed an average success rate of only 86% for traditional backup job completion within backup windows, according to respondents focused on applications and databases, and the success rate was just one percentage point higher for respondents focused on data protection.[3] Most organizations do have workarounds, so a survey indicating that roughly one-seventh of backup jobs fail might not signal a calamity. However, the situation is cause for investigation because it implies that many IT environments' backup solutions may not be living up to expectations.

When the lens of a recovery time objective (RTO) or recovery point objective (RPO) service-level agreement is applied, the picture looks worse. Survey respondents reported having even less success getting recovery jobs to meet RTOs or RPOs defined by SLAs—they missed the mark 18% of the time.

Collectively, these numbers represent resources operating ineffectively and talented people wasting cycles dealing with a chronic problem.

Storage Consumption Challenges

Imagine an environment that has exactly 1TB of data in **production storage**.

In a traditional backup scenario, one might generate 4TB to 6TB in **secondary storage** stemming from multiple full and incremental/differential backup copies. Deduplication obviously mitigates part of this amount.

For disaster recovery, another 1TB to 1.5TB will be consumed in **tertiary storage** at another location, along with the replication bandwidth required for transmission.

ESG data shows that both production storage and secondary/tertiary storage are growing at nearly the same rate,[4] a situation that continues to create pressure for CapEx and OpEx associated with storage.

How Many Copies and Versions Do You Need?

One typically starts making copies of data with the on-premises backup solution, the offsite disaster recovery environment, and the tape-based retention system.

In addition to copies created for data protection, odds are that demand exists for copies to support **reporting and**

[3] Source: ESG Research Report, *Trends in Data Protection Modernization*, August 2012.
[4] Source: ESG Research Report, *Backup and Archiving Convergence Trends*, April 2014.

analytics, **development and testing**, and so on. Such copies have always proliferated, thanks to users who create multiple copies of individual MS Office files through versioning and file sharing, or due to organizations running secondary copies of applications and data to support different functions and business units.

The ability to move and copy data is an obvious benefit of a modern IT infrastructure. But it is also a source of many direct and indirect costs—increasing the need for hardware and making management more complex. (In fact, with snapshotting, an opportunity exists to rethink data-centric practices.)

You *do* need a range of copies and versions of your data. Often, you're in a situation in which your organization can't afford not to have them. But supporting them is often burdensome.

What You Can Do Instead

Business and operational concerns must drive storage strategies. By making sure the *proper* concerns are driving decisions, IT can begin to build a storage infrastructure that is more responsive, more efficient, and better attuned to organizational needs. This smarter approach to IT delivery will also reduce cost and complexity.

First, Understand the Business Processes and Needs for Your Data … Then Choose Your Tools

Unfortunately, many IT managers determine their recovery options based on the backup tools they have in place. (The same can be said for how they replicate and test.) Basically, people tend to assess their business capabilities based on the tools or components already within their reach instead of establishing their goals first, and then investigating the most viable methods to achieve them.

Data protection and recovery—and all the associated opportunity costs and direct costs—is too important to be driven by things as trivial as the status quo, historical processes, or what is "doable" using the current toolset. Instead, organizations should (1) decide how they want to do *restore* (e.g., whether they need granular object, multi-site, whole-VM restore, etc.), (2) define the associated business-driven SLAs, and then (3) move ahead to design a data protection strategy based on those needs.

Start with the restore-related decision, the purpose of secondary copy(ies), the use-cases, etc. Understand the processes and insights already developed. Likewise, make sure you are familiar with other business processes needing support. Have authentic conversations covering business stakeholders, application/workload owners, and anyone else responsible for related data protection and data management capabilities.

After you have those conversations, you should be able to determine which data is the most important, how rapidly it needs to be recovered in the event of a problem, where additional views into the data should be provided, and how those views will be utilized. At this point, you are going to be much better equipped to make realistic choices.

Along the Way, You Might Learn Something

During your data protection and data management conversations with workload owners and business stakeholders, look for opportunities to optimize your existing processes. The conversations with the business stakeholders not only help you to ensure continued service to them during a crisis, but also will almost certainly spur additional conversations about where some operational processes/procedures can be optimized *now*.[5] You might discover:

- Production servers (often "just VMs"[6]) that have been clandestinely launched within a department and never backed up or monitored.
- Server resources that are so critical, they really should be clustered or have some other means of high availability implemented.

[5] Source: ESG *Technical Optimist.com* blog, "The best part of BCDR planning is what you get before the disaster," Jason Buffington, November 2011.

[6] Source: ESG *Technical Optimist.com* blog, "What to Look for in Virtualization Protection in 2014," Jason Buffington, November 2013.

- Servers that are still running "in case they're needed" but haven't been utilized in a while. These servers might be running on expensive hardware that could be virtualized, or they may be enjoying maintenance-agreement attention that could be downgraded.

It would be nice if these conversations were organic in all companies. But the reality is that instead of pausing to gain a consensus on goals and develop strategies to achieve them, the conversations don't occur—simply due to the hectic, tactical nature of IT. Reassessing one's data protection or data management strategy will often bring important topics to light. The conversations also will spur dialogues that promote understanding between business stakeholders and their IT counterparts. And that understanding can yield big empathy/cooperation benefits.

To Get Different Results, Do Something Different

If you do the same thing over and over again, then you should expect the same results over and over again, too.

No silver bullet can make problems with backup and restore disappear. IT has been wrestling with the problem for decades. When it comes to backup, although workloads will evolve and technologies will innovate, if you back up data in the way you've always done it, then you should expect those less-than-ideal 82% to 87% success rates revealed by ESG's research.

To achieve success rates or recovery SLAs better than what traditional backups can deliver, you will likely need to consider "something different" from the methods used in the past for data protection and recovery. If those methods include adding snapshots or replication to your backup strategy, you are not alone (see Figure 2).[7]

Figure 2. Supplementing Virtualization Protection Mechanisms Beyond Backup

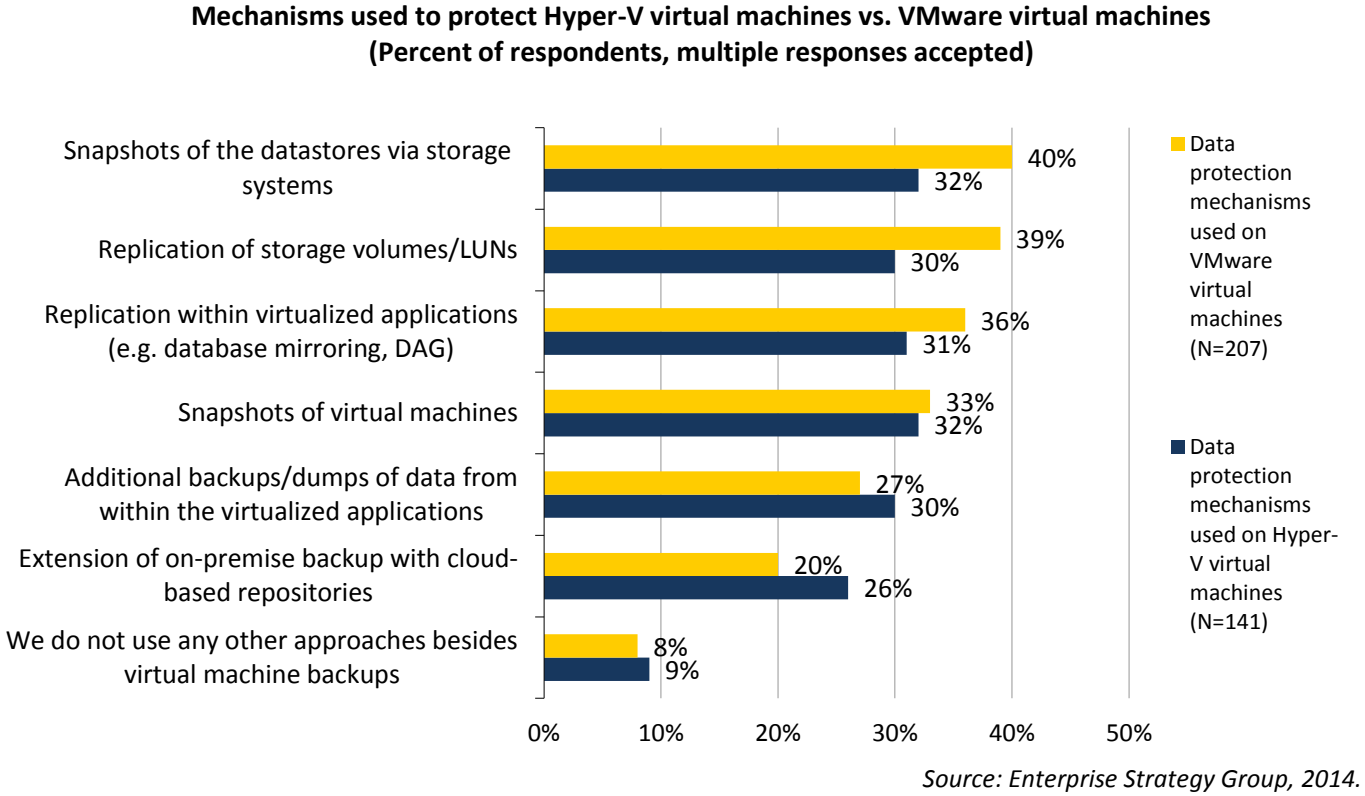

Source: Enterprise Strategy Group, 2014.

As Figure 2 shows, less than 10% of respondents who were protecting highly virtualized environments relied upon VM backups alone. More than 90% of respondents used some combination of snapshots or replication to ensure the agility of virtual machines and the services running within them.

[7] Source: ESG Research Report, *Trends for Protecting Highly Virtualized and Private Cloud Environments*, June 2013.

However, it is important to recognize that although snapshots and replications can provide more frequent operational restore points and better restore-time SLAs, they aren't the be-all and end-all solution to the problem. Not all snapshots or replication solutions are the same, and snapshots and replication don't entirely replace traditional backup from a long-term retention perspective.

The real takeaway is that you should be choosing your data protection and data management tools based on what the business requirements are. Some restore mechanisms deliver sub-minute restore and offer multi-site recovery. Not all snapshots are practical if you require snapshots throughout the working day because they may cause a performance impact. It is important to research what is available and to be open to new methods of protection, replication, and distribution based on how you want to leverage the data, not on what you presume your current solution can offer.

When you decide that snapshots or other storage efficiency technologies will enable potentially better backup and restore results, you will likely discover that the same technologies will enable other business scenarios, too. Start with this question: "*How do I want to store data, and what outcome(s) do I want?*" You want production data and support for backup and offsite data. You need access to data for functions such as development and test, sales demos, and customer service efforts. And you need to deal with the all-encompassing issue of data growth. You have valid reasons for storing data or versions of data. So, the question simply becomes how to do it more effectively.

IN THIS ESG PAPER, the diverse sub-categories of "data protection" include:

Backup—The making of secondary or tertiary copies of data on non-production systems, using traditional file/application processes, as a means of preserving previous points in time, often where length of retention or application awareness are paramount.

Snapshots—The leveraging of underlying storage technologies to capture previous points in time within a storage array or file system, where recovery time within the original systems is most important.

Replication—The duplication of data over distance between disparate but logically linked systems.

Other aspects of "data protection" beyond these three categories exist—and that is the point of this paper.

Step Back and Redefine Your Data Protection Strategy

To summarize the guidance above:

- Start by understanding the operational goals and recovery requirements of the applications and the business stakeholders—including a financial assessment of the importance of the data/systems—to understand "the cost of the problem" before considering "the price of a solution."
- Then assess which data protection mechanism(s) such as backup, snapshots, and replication are most appropriate to achieve those recovery goals—but recognize that multiple methods may not always require multiple vendors or even multiple management experiences.

But there is more[8] to consider when redefining a modern data protection strategy:

- Plan for a multi-platform infrastructure in which servers/services will run on a combination of physical servers, VMware-hosted VMs, and Hyper-V-hosted VMs—and ensure that your storage solution and your backup capabilities accommodate all three platforms.
- Similarly, plan for a hybrid data preservation system that should almost certainly start with agile production and secondary disk, but likely also include tape for retention and/or cloud storage for resiliency (disaster-recovery-as-a-service) or as a tape-alternative for preservation.
- Deduplication and other storage optimizations are no longer optional. Every part of an IT infrastructure must be reconsidered with efficiency in mind. In the case of backups, simply storing the final copy better

[8] Source: ESG *TechnicalOptimist.com* blog, "8 Suggestions for Every Data Protection Strategy," Jason Buffington, February 2014.

isn't enough—the primary data needs to be stored more efficiently, and that optimization should not be undone as data moves throughout its lifecycle into secondary and tertiary copies.

- And most importantly, recognize that modernization does not equate to merely uplifting your current components or capabilities. It is highly likely that if you have not reassessed your data protection strategy in the last 24 months, by doing so you will reduce your CapEx and your OpEx. And more than anything else, you will increase your data's agility by investing in new methods.

Sometimes, Data Protection Comes Built-in

In contrast to always thinking about protection "after" production, newer approaches for data protection are often best accomplished by starting at the production systems themselves.

A few companies, NetApp being a notable example, have ventured into this technological maelstrom with particular effectiveness. Founded in 1992 with a fundamental focus on storage, NetApp today focuses its offerings to meet a range of needs among a growing range of organizations. NetApp's storage efficiency, snapshotting, and replication technologies offer a *different* approach to data protection and data management—an approach adapted to virtualization and data growth.

Understanding NetApp Snapshots

As Figure 2 showed, many approaches for protecting data are already designed to fit the needs of individual organizations (hopefully based on an understanding of their unique business requirements). Snapshots, for example, represent an excellent way to recover quickly without a lot of overhead or infrastructure.

One of NetApp's underlying technologies is its Snapshot software (see Figure 3). Because not all snapshots work the same way, the key to appreciating how NetApp's snapshots might fit into a data protection strategy centers on understanding the nature of the disk blocks and pointers underlying NetApp's file system. Those elements are the foundation for the NetApp snapshot.

Figure 3. How NetApp Snapshots and Pointers Work

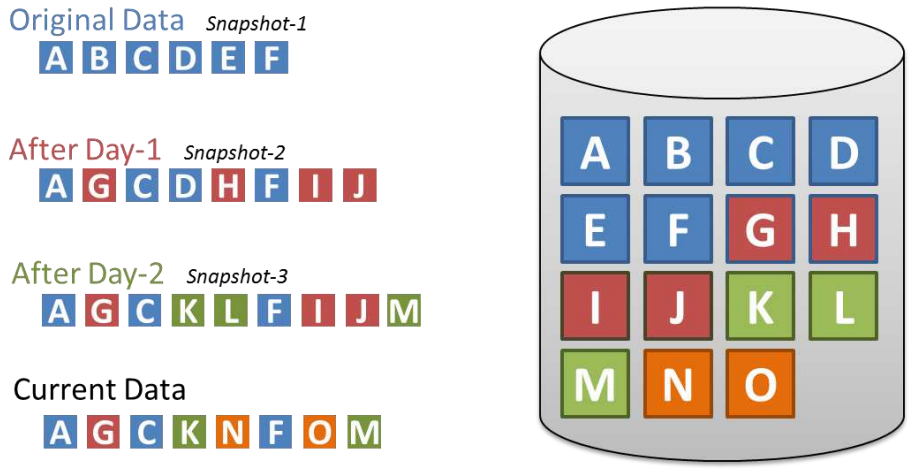

Source: Enterprise Strategy Group, 2014.

In the very simple example above, the left side of the figure shows four iterations of a data file composed of six to nine blocks.

1. The original file was made up of blocks ABCDEF. A snapshot of the original data would simply comprise six pointers to the original six blocks of data.
2. After Day 1, data blocks B and E were replaced with data blocks G and H, and data blocks I and J were added. A snapshot created at this point lists the pointers to the remaining blocks (ACDF) along with pointers to the new blocks (GHIJ).

3. After Day 2, more replacements occur, and more new data is appended—with another snapshot created.
4. Now, the file comprises blocks AGCKNFOM.

Presuming that a daily snapshot schedule is being enforced, each of those points in time would be immediately restorable or "retrievable." If the storage system needed to be reverted to Day 1, the snapshot (i.e., the pointers to the eight blocks AGCDHFIJ) would be used to present the data as of that point in time.

All of this activity occurs within the NetApp storage system, so the snapshots (lists of pointers) are derived without affecting production server performance or consuming additional space (other than the very minor catalog of pointers).

More impressively, imagine that the snapshot schedule is not once daily, but once every 15 minutes, with an ability to unobtrusively create immediately restorable data with fine granularity reflecting very short time periods.

In its most simplistic form, a snapshot of a given point in time retains the pointers to the blocks within its file system that represent how data looked at the time the snapshot was created. This structure results in a couple of possible restore scenarios:

- Reverting back an entire file system from a snapshot can be done nearly instantaneously by simply referring back to an earlier snapshot that references an older set of pointers.
- If the proper file system and metadata is maintained, it is possible to selectively restore a file or folder by simply recalling those blocks—often with the intent of restoring them elsewhere, which can be as fast as a file copy within the same disk or system.

Snapshots Compared with Backups

Snapshots of data on primary systems do not replace traditional backups. They augment traditional backups as part of a broader, more agile data protection strategy. (This is due to the restoration agility that snapshots enable.) Still, they offer value. For instance, they can quickly recreate what happened only a few moments before and thus restore you to a desired previous state. They are about the recent past—days, hours, or minutes ago.

Snapshots can be thought of, in some sense, as an availability solution almost as much as a data protection capability (due to the recovery speed). Snapshots allow you to simply "put things back" to the way they were a short time earlier. Backups, conversely, reside on a separate storage device for longer periods of time to enable a different kind of recovery.

Snapshots within production storage systems are a useful adjunct to backup. Not too many years ago, partisans had sparked a passionate debate that implied an either/or choice: You were either going to do snapshots or do backups.

Imagine a dialogue[9] between two ardent IT experts, both of them trying to persuade an application owner to use their data protection capability:

> *Storage Administrator: Use snapshots because they are **fast**, and they don't consume server **CPU**.*
>
> *Application Owner: They are fast, but your storage is too far (logically) from my application, and **storage doesn't understand what my application needs** before you make that restoration point. Oh, and I can't afford to keep snapshots around for weeks or months.*
>
> *Backup Administrator: Do backups. **They understand applications** via agents, and they can retain your data for **years** on tape.*
>
> *Application Owner: You do understand my application, but **your restore isn't nearly as fast**. And your agents are okay, but they consume CPU, sometimes lots of it, which can slow my app down.*

The two sides of that imaginary debate can become as vehement as politics or religion. But it all comes back to avoiding letting your method dictate your outcome. Choose your desired outcome(s) first, then let that choice

[9] Source: ESG *Technical Optimist* blog, "Snapshots vs. Backups—a great debate, no longer," Jason Buffington, October 2012.

determine your method(s) to achieve that goal. If you want near-instantaneous recovery, then you want snapshots. If you want application consistency and long-term retention, then you want backup.

However, it isn't an "either/or" choice. Often, the right answer is to do *both*.

NetApp's approach exemplifies "reaching across party lines," which includes offering the ability to leverage snapshots within the primary/production storage systems and within the secondary storage systems for even more agility (see Figure 4).

Figure 4. Leveraging Primary and Secondary Snapshots Within a Comprehensive Data Protection Strategy

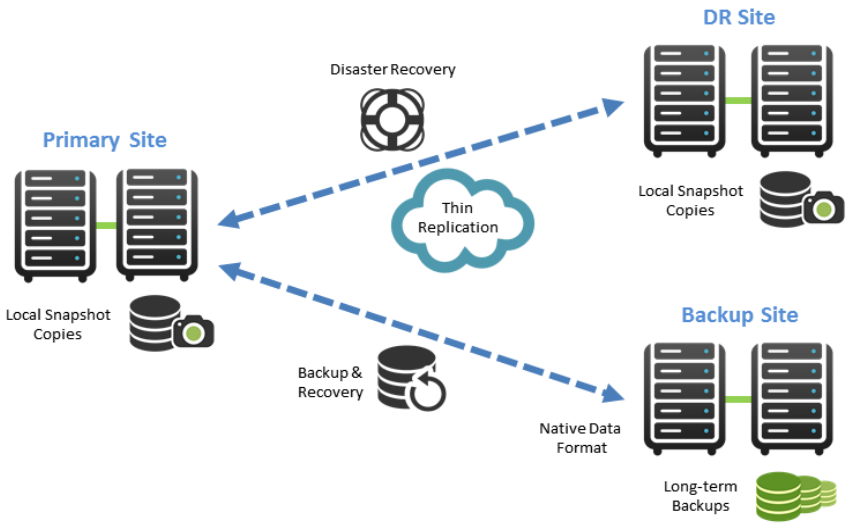

Source: Enterprise Strategy Group, from material supplied by NetApp, 2014.

If one were to overly simplify the "snapshots versus backups" debate, the discussion would boil down to "agility to rapidly recover" (snapshots on primary storage) versus "longer-term retainability and granularity" (backups from secondary media). Neither generalization is completely accurate. But they allow for a third scenario—snapshots within secondary systems that might be invoked on a different schedule (thus enabling a different retention tier) while enabling the kind of agile, fast recovery operation shown in Figure 4.

Integrating Snapshots *and* Backups

Happily, we are past the "snapshots *or* backups" discussion and can recognize the complementary benefits of each.

To provide its customers with a broad range of data protection capabilities, NetApp offers solutions of its own and partners with traditional software-based backup vendors in a "better-together" approach. In particular, NetApp offers SnapProtect, which combines snapshot copy creation and replication using a single management interface and catalog to create, track, and restore snapshot copies from disk or tape. The product uses agents to create application-aware snapshots across applications and virtual environments.

Although NetApp's solutions are based primarily on snapshot abilities, SnapProtect does provide traditional backup to complement those snapshots. Storing multiple space-efficient snapshot copies on secondary storage or even tape as a "backup" is one of NetApp's core value propositions. (Admins can set up schedules for SnapVault to retain snapshots for very long periods of time.) And, in a truly better-together scenario, SnapProtect catalogues the snapshots' data to support a unified restore experience.

SnapProtect reflects NetApp's understanding that its customers need both long-term backups and rapidly recoverable snapshots. It also demonstrates NetApp's willingness to deliver to the marketplace better-together solutions (with, for example, IBM, Catalogic, or Symantec) in addition to offering solutions under its own NetApp banner.

Other Storage Capabilities That Enable Better Data Protection and More

Looking at the broader picture—beyond "just" backups and snapshots—it is clear that NetApp's architecture provides NetApp customers with additional capabilities stemming from its approach to storing data.

Deduplication

The same storage engine that uses pointers to blocks can, with some additional intelligence, discern whether the storage holds multiple copies of the same blocks of data. Similar to avoiding keeping multiple copies of data between snapshots, the deduplication capabilities of the file system determine whether redundant blocks are being sent to a volume (for example, when two users store the same file in their home directories or when backup software is writing yet another copy of unchanged data to its store). When duplicates are identified, NetApp filers can reduce storage consumption by simply updating the pointers for both instances to a common block, not consuming space with what would have been duplicate information. Having the deduplication discernment occur at the block level rather than the file level results in more identification of duplicate files (or partial file fragments) that can then be eliminated, resulting in more space savings.

Replication

As described in Figure 2, snapshots are not the only mechanisms IT managers are considering for augmenting traditional backups. Replication, such as what NetApp offers in its SnapMirror technology (usually for disaster recovery or higher availability scenarios), is also receiving new consideration by leveraging *storage systems'* ability to replicate (instead of application- or backup-based replication).

How Deduplication and Replication Enable Better Data Protection

By combining the intelligence of deduplication and replication, additional data protection scenarios are enabled. If blocks that are about to be replicated are determined to reside on the target platform already, then those blocks aren't transmitted. Instead, relevant pointers are updated without network impact or storage consumption.

But Wait; There's More

Other built-in features also either enhance or broaden one's data protection capabilities:

- ***Compression that works well with block deduplication:*** Data blocks that are compressed on the primary storage remain compressed during network transmission, resulting in more effective data protection and data movement. Compression works seamlessly with block deduplication because the blocks themselves are compressed.
- ***Storage resiliency:*** Enabling data protection isn't simply about offering multiple copies throughout an enterprise. With so much deduplication, compression, and other reduction, it is vitally important to ensure that data is not lost within a storage system. To ensure that drive failures or other internal issues do not compromise data, NetApp provides a storage architecture called RAID-DP, whereby dual-parity drives ensure that even hard drive failures are mitigated without the expense of disk mirroring or the performance impact of commodity RAID mechanisms.
- ***High availability:*** No discussion of a broad data protection strategy is complete without considering the resiliency of the storage appliance(s) overall. NetApp Clustered Data ONTAP enables nondisruptive operations, allowing admins to perform critical tasks without business interruption. For example, storage controllers can be replaced without disruption and without moving data. The ability to assign, promote, and retire storage resources dynamically improves service levels over the lifespans of applications.
- ***Business continuity:*** Synchronous replication with MetroCluster (see Figure 5) adds a business continuity component that protects data beyond the data center with zero data loss and zero downtime. MetroCluster combines array-based clustering with synchronous replication across city or metro areas up

to 200 kilometers. And because it is part of Data ONTAP, it doesn't require an external device to manage it. After MetroCluster is set up, it no longer requires ongoing configuration whenever an admin adds or removes LUNs or volumes. Both SAN and NAS protocols are supported, as are deduplication, compression, and all the other data protection features. As Figure 5 shows, MetroCluster uses many of the other storage-enabled data protection augmentations to traditional backup to create underpinnings not only for "failover," but also for ongoing business continuity. After all, BC is the ultimate goal of all data protection mechanisms—to ensure accessibility of data and systems in an effort to maintain and bolster corporate productivity.

Figure 5. NetApp MetroCluster for Availability Enablement Within a Data Protection Strategy

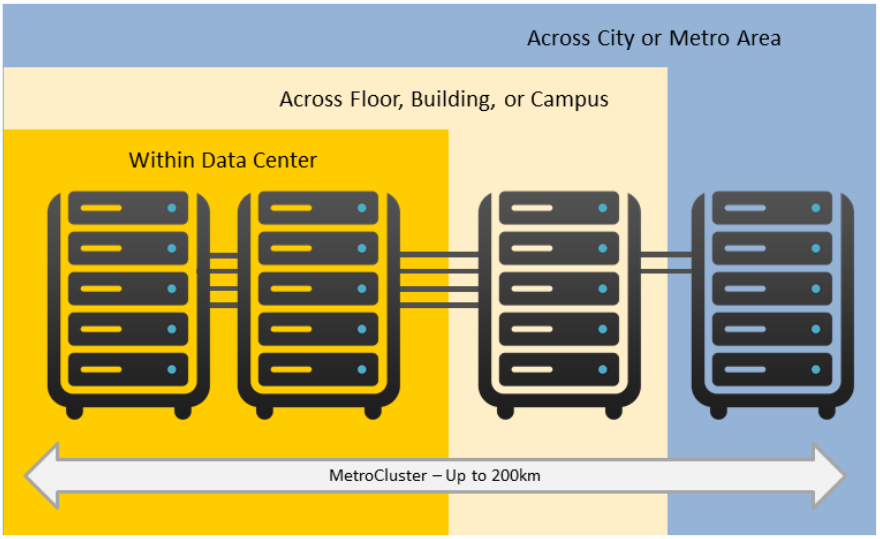

Source: Enterprise Strategy Group, from material supplied by NetApp, 2014.

- ***Cloning:*** NetApp FlexClone software enables true cloning—instant replication of data files, LUNs, and volumes without requiring additional storage space at the time of creation (see Figure 6). Each cloned file, LUN, or volume is a transparent, virtual copy of your data that can be used for essential enterprise operations such as testing and bug fixing, platform and upgrade checks, multiple simulations against large data sets, remote office testing and staging, and provisioning of server and desktop images.

Figure 6. Space-saving FlexClone

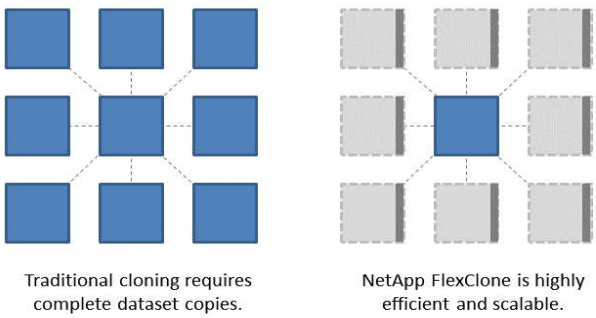

Source: Enterprise Strategy Group, from material supplied by NetApp, 2014.

The Bigger Truth

You can breathe a sigh of relief. The "religious wars" between backup/recovery and snapshotting zealots are over. They have both won, and it turns out that they were both right. Alone, neither approach meets the needs of IT today. But in a better-together world—as exemplified by NetApp's range of data protection and data management options—organizations can aim for highly efficient, highly reliable storage for production, backup, and recovery.

Think about what IT outcomes you need, then envision the smartest way to achieve those outcomes. In the case of data protection, decision makers should focus on *how they want to restore*, which will help them determine *which methods of protection to use* (not vice versa).

That same approach should be taken when one is faced with the larger perspective of storage management. Strategies related to storage (how much, what kind of tiering, etc.) should be driven by business value and the goals for how the data needs to be utilized, not by some arbitrary notion of how storage should be deployed or by what already exists in-house.

Just as most folks' preference in mainstream IT is storage that provides "unified" capabilities that were previously debated as "SAN versus NAS," today's environments need to embrace the better-together(ness) of backup *plus* snapshots. This combination means you don't just protect your data "better," you protect your data "smarter" by first understanding the business processes, then optimizing throughout the data protection and data management lifecycles (not just the redundant repositories).

ESG expects to continue seeing not only better integration between "snapshots plus backups," but also more maturity—where production platforms contribute more to overall data protection enablement. It is something that NetApp has already been doing for quite a while.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com