

# White Paper

---

## **NetApp Is Accelerating Your Data Protection Strategy to the Clouds**

*By Jason Buffington, Senior Analyst; and Monya Keane, Research Analyst*

January 2015

---

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

## Contents

Introduction .....	3
Apples and Oranges.....	4
How Most People Use Disk, Tape, and Cloud.....	5
Start with Disk.....	5
Tape: In It for the Long Run .....	5
Going to the Cloud.....	5
Where BaaS Fits.....	5
How to Use the Pieces <i>Differently</i> .....	6
A Look at Cloud-extended Data Protection .....	7
Staying Secure .....	7
Efficiency and Effectiveness Go Hand-in-hand.....	7
Risk Mitigation.....	8
NetApp AltaVault.....	8
Extending Storage to the Cloud.....	8
With AltaVault, <i>you can keep your existing backup software</i> if you like it. ....	9
On the Move.....	9
The Bigger Truth .....	9

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.


## Introduction

To IT professionals wrestling with defining a data protection strategy, an abundance of “choice” can be overwhelming. Disk, tape, private cloud, hybrid cloud, backup-as-a-service (BaaS), storage-as-a-service (STaaS), infrastructure-as-a-service (IaaS), and various permutations of those architectures are now available, which can complicate a technology-selection process.

In any case, it is clear that *improving data backup and recovery* is a priority for IT groups in organizations of all sizes, as ESG research shows (see Figure 1).<sup>1</sup> It has consistently appeared as one of the most commonly identified IT priorities by senior IT decision makers participating in ESG’s annual spending intentions survey for several years running, and in midmarket organizations, improving backup and recovery often appears as the most-often cited priority. Evidently, backup is not “solved” yet.

Figure 1. Top Ten IT Priorities for 2015

**Top 10 most important IT priorities over the next 12 months. (Percent of respondents, N=601, ten responses accepted)**



Source: Enterprise Strategy Group, 2015.

There’s also the question of budget. ESG has found that an interest in reducing costs is a major driver of IT strategy. Specifically, *increased use of cloud computing services* and *buying new technology that can deliver a better overall ROI* tied among the most-often-mentioned measures (33%) cited by survey respondents whose organizations are trying to reduce or otherwise contain IT expenditures.<sup>2</sup>

Organizations are willing to spend money on modernizing data protection, but they want to be smart about doing so. They want value for their money. In today’s technology landscape, that desire often leads them to the cloud. In fact, every data protection modernization conversation today really *has to* include the cloud.

Every data protection modernization conversation today really *has to* include the cloud.

<sup>1</sup> Source: ESG Research Report, *2015 IT Spending Intentions Survey*, to be published January 2015.

<sup>2</sup> Ibid.

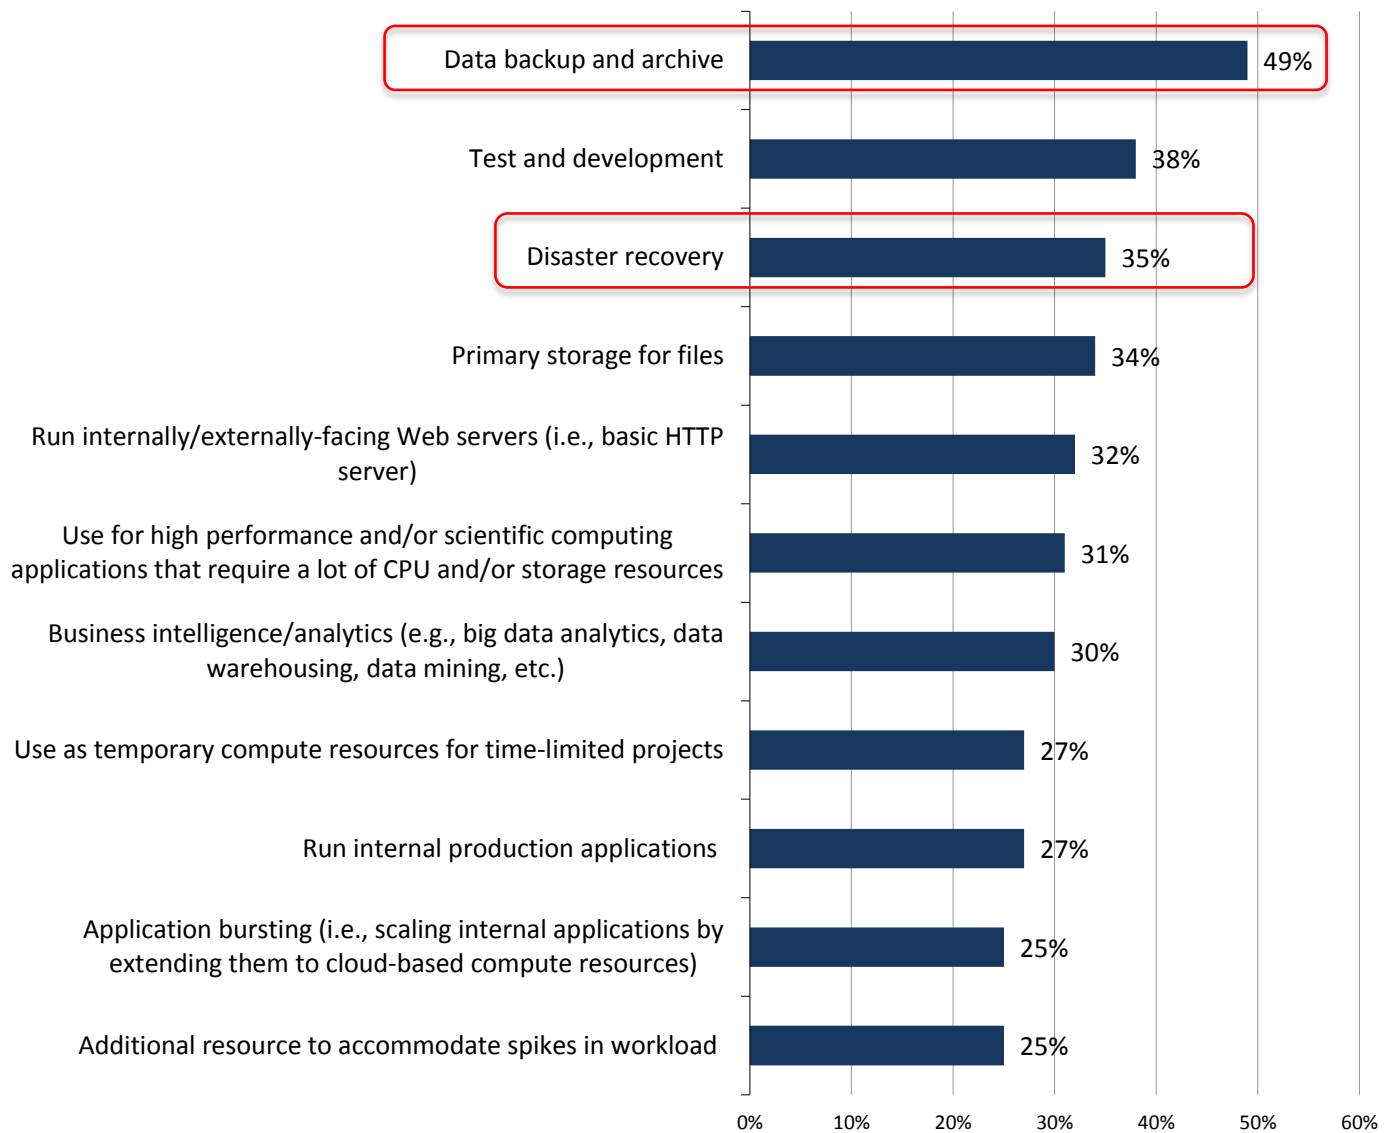
## Apples and Oranges

Cloud backup is a somewhat recent phenomenon. For many organizations planning to modernize legacy onsite data protection architectures, taking advantage of the cloud can mean swapping more than just “an apple for an apple” (i.e., swapping one backup product for another). When it comes to data protection, the better solution might be to swap “an apple for an orange” (i.e., investigate an altogether different, state-of-the-art architecture that has been designed and optimized to leverage the cloud).

That kind of swap from legacy products to new data protection architectures and solutions that are cloud-enabled seems to be happening at many organizations that are modernizing their data protection environments. Protection is proving to be an extremely popular use case for cloud IT. In fact, ESG recently conducted research in which current data backup and archive and disaster recovery were two of the most commonly cited cloud infrastructure services being leveraged by current IaaS users (see Figure 2).<sup>3</sup>

Figure 2. Cloud Infrastructure Service Use Cases

**For which of the following purposes does/did your organization use cloud infrastructure services?  
(Percent of respondents, N=327, multiple responses accepted)**



Source: Enterprise Strategy Group, 2015.

<sup>3</sup> Source: ESG Research Report, *2015 Public Cloud Computing Trends*, to be published March 2015.

## How Most People Use Disk, Tape, and Cloud

IT organizations now have disk, tape, and the cloud at their disposal. One good strategy is to combine the platforms to create a hybrid architecture. Its exact makeup would depend on an organization's requirements for retention, distance-based protection, and on-premises recovery agility.

### Start with Disk

In a data protection architecture, disk should be the primary means of ingest and the first source for recovery. In other words, the first backup copy you make should go onto an onsite disk-based backup target, regardless of whether you later send the data to tape, cloud, or another disk. Disk:

- Helps keep backup windows as short as possible.
- Is ideal for instant VM and any type of recovery, and it is best able to support stringent RTOs and SLAs.
- Supports data deduplication for efficiency, helping to keep data growth under some control.
- Is flexible for restoration efforts, thanks to its random access-based design.

### Tape: In It for the Long Run

Thanks to tape's durability and portability, it is an excellent medium for archival retention. Of course, some labor is involved—tapes must be mounted, dismounted, stored, shipped, labeled, and so on. Buying more cartridges, drives, and libraries will occasionally be necessary. If your data-retention periods are lengthy (i.e., decades), the tape media and associated tape technologies you use today may turn out to be obsolete before that retention period is up—so retrieving the data may become difficult. But for many organizations, the scaling advantages, high reliability levels, and current LTFs-related data-access capabilities of tape may outweigh its costs.

### Going to the Cloud

The cloud offers many possibilities. But organizations may struggle to put the cloud to use for data protection in part because of the confusion caused by the many BaaS and disk-to-cloud (D2C) offerings on the market.

D2C involves sending a copy of data from a primary production server *right* to the cloud for protection. Now, backing up data to the cloud probably isn't going to be a problem; restoring data from the cloud can be. Recovery may not happen fast enough to suit an IT organization operating under today's tight SLAs. For that reason, ESG recommends that production data requiring protection go first to an on-premises disk-based backup target, and then to the cloud, which serves as the tertiary or long-term retention container.

### Where BaaS Fits

Certainly, backup-as-a-service has a place in data protection, but it also presents challenges:

- **Some solutions were born as consumer-oriented products** or continue to be marketed for consumer "bring your own device" (BYOD) access or file sharing. They may not be optimized to handle high data volumes; they may be unable to ensure quick recovery, or they may provide insufficient security and lack the IT oversight needed for regulatory/policy compliance.
- **BaaS can lock you in.** Your data may not be portable from one BaaS vendor to another—that's a problem if you wish to switch providers to save money or avoid vendor lock-in. If your data was stored in a BaaS provider's proprietary format, moving it to a new BaaS vendor could involve a lot of time and expense, if it's even possible. Also, depending on your industry's regulatory requirements, changing BaaS vendors prior to the end of a legally required data retention period could be considered a breach of custodianship.

BaaS can lock you in, and a wholesale migration to BaaS may be overkill. Don't presume that to leverage the cloud, you have to change your entire backup infrastructure.

- **A wholesale migration to BaaS may be overkill.** A BaaS solution is markedly different from a traditional on-premises protection environment. Moving to pure BaaS requires replacing all agents in all servers to be backed up. It requires recreating job schedules. And it requires retraining staff. The term for this approach is “throwing the baby out with the bathwater.” Don’t presume that to leverage the cloud, you have to change your entire backup infrastructure.

That said, BaaS has a legitimate place in data protection, and its market growth and popularity for backup, archiving, and DR isn’t surprising. For example, BaaS is excellent for protecting the devices of BYOD workers. And it’s worthy of investigation by any organization that doesn’t operate under heavy regulatory pressure or have heavy demands for fast data restoration. Expect BaaS to continue to grow and thrive.

## How to Use the Pieces Differently

We can rely on two certainties: (1) On-prem disk is the “best, first” backup tier, and (2) tactical backup and strategic BC/DR are prevalent use cases for the cloud. But beyond those certainties, a few more nuanced decisions still need to be made.

For example, the preferred architecture—disk-to-disk-to-cloud (D2D2C)—can still encompass different “flavors” or permutations, each with varying benefits and concerns:

- Some onsite backup servers can replicate data to a second instance of the server operating in the cloud, almost as if you had installed a server at your own “Site A,” and another at your own “Site B.” These solutions can be a challenge due to where they can restore the data to, as well as how fast.
- Some traditional backup hardware vendors are partnering with certain cloud repositories. Using this approach, you might point your on-premises appliance to that vendor’s repository for replication, but it requires that the cloud provider hosts the largest of that storage array vendor’s platforms and then uses multi-tenant controls to segment its subscribers.

Each of those approaches has advantages *and* limits:

- If your existing backup software is functioning adequately but the overall solution isn’t achieving your broader range of tactical and strategic goals, you may be better off keeping the software and changing the target solution that the software is utilizing. Do it right, and you’ll still have an on-premises disk target, and your backup software will remain in place. All the installed agents, scheduled jobs, and staff training you’ve deployed and customized over time will stay the same.
- If you aren’t committed to your backup software, you can change it. But be wary of any data protection technology that locks you in and thus may prevent you from pursuing future strategic-level data protection enhancements.

If your existing backup software is functioning adequately but the overall solution isn’t achieving a broader range of goals, you may be better off keeping the software and changing out the target solution that the software is utilizing.

No matter which cloud strategy “flavor” you employ, a D2D2C approach will give you the best of both worlds. You’ll benefit from agile restores, and you’ll enjoy the cost-reduction benefits of deduplication and compression that on-premises disk excels at. By maintaining a remote copy, you’ll also ensure data survivability for DR. And because the capital expenditures for cloud protection are lower than they are for onsite disk, you may be able to keep several years of compressed, deduped archival data in the cloud—perhaps a volume of data that would have been cost-prohibitive to keep using just local storage.

## A Look at Cloud-extended Data Protection

To evaluate a cloud-extended data protection solution, consider your fundamental needs for security, efficiency, effectiveness, and risk mitigation.

### Staying Secure

According to ESG research, security is the most-cited concern for organizations that hadn't yet extended their data protection efforts to the cloud. However, organizations that already *are* using the cloud for data protection tell a different story, with 37% of respondents citing improved security as a benefit of using cloud-based data protection services—in fact, improved security was their most-often-mentioned benefit.<sup>4</sup> Cloud security proves to be better than people expect in large part because security is a strategic necessity for cloud vendors who can't afford to be lax.


When you're evaluating cloud vendors, though, you should still be careful in one regard: How does the provider and your software/hardware vendors help ensure that no one else can gain access to your data? Ask this question directly during the vendor evaluation process. The response you're looking for should address the quality of encryption for data at rest in storage and data in flight while moving to or from the container.

It shouldn't be a difficult question for the vendor to answer. Again, in most cases, encryption and physical control are actually better than with many of the pure on-premises legacy systems organizations are using today.

### Efficiency and Effectiveness Go Hand-in-hand

ESG research shows that among the top five reasons surveyed IT organizations would add or replace a data protection solution, the two most often mentioned reasons related to efficiency, and the third through fifth reasons related to effectiveness (see Figure 3).<sup>5</sup> Specifically, the top two reasons centered on economics, reflecting OpEx and CapEx budgetary concerns. The next three revolved around various challenges with legacy solutions.

Figure 3. Top Five Reasons Organizations Would Replace Their Current Backup Solution/Vendor



Source: Enterprise Strategy Group, 2015.

Efficiency and effectiveness are directly linked to ROI (a top consideration for justifying IT investments to the business management team), and they include factors such as the time and labor needed for training, retraining, maintenance, and vendor-relationship issues that do not center on the equipment itself.<sup>6</sup> Taken together, efficiency and effectiveness improvements can cut thousands of dollars a month from an IT budget.

Efficiency and effectiveness also are tied to specific parts of the architecture. Efficiency comes from utilization of the cloud. Effectiveness, for most environments, often requires local disk to ensure high-performance ingest and fast recovery agility.

<sup>4</sup> Source: ESG Research Report, [Data Protection-as-a-service \(DPaaS\) Trends](#), September 2013.

<sup>5</sup> Source: ESG Research Report, [Trends in Data Protection Modernization](#), August 2012.

<sup>6</sup> Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), to be published in January 2015.

## Risk Mitigation

The third essential need in a data protection architecture is risk mitigation. IT modernization demands change, and change can bring risk. With the cloud, risk can come from several directions:

- Vendor longevity—A web search for cloud-based solution providers brings up many, many vendors. Months (or maybe a few years) from now, will all of them still exist? Will the one you choose exist?
- Lock-in—If you want to change vendors for cost control or other reasons, can you expect a smooth transfer? Some BaaS providers store their data in proprietary formats that cannot easily be exported for re-hosting with a different provider.
- Unsafe exposure—If you’re migrating from one on-premises backup solution to another, will you be running two architectures in parallel for a period of time? You may need to add or change agents mid-migration to accommodate a new backup job. Suddenly, you’re exposed.

Tough questions are worth asking, even if there are no simple answers. The better you understand your own protection goals, the more comprehensive and accurate your vendor assessments will be.

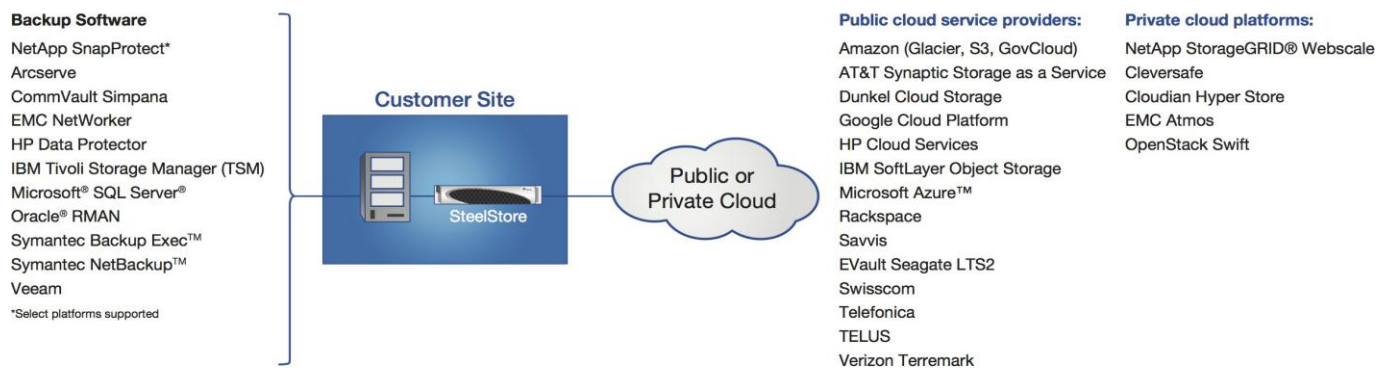
## NetApp AltaVault

One vendor ESG believes is worth assessing is [NetApp](#), a Sunnyvale, California-based company that has been innovating around storage and data protection solutions for many years. In 2014, NetApp acquired the AltaVault® (formerly SteelStore) technology from [Riverbed Technology](#), which first made its name in WAN optimization—pushing massive volumes of data through thin wires efficiently.

## Extending Storage to the Cloud

NetApp® [AltaVault](#), NetApp’s cloud-enabled backup gateway appliance, builds on the WAN-optimized pedigree from its Riverbed roots and carries forward the idea of optimizing WAN throughput to significantly innovate core IT scenarios that NetApp leads in; namely data protection. More specifically, the AltaVault solution delivers an on-premises storage solution that transparently extends its capacity through a variety of public cloud providers (see Figure 4).

Figure 4. AltaVault Cloud-integrated Storage Topology Diagram



Source: NetApp, 2015.



## **With AltaVault, you can keep your existing backup software if you like it.**

Therefore, you can keep all of the:

- Agents you've already installed on your production servers.
- Scheduled backup jobs you've already configured.
- Staff-training efforts you've already expended effort on.

Without disruption or extra effort, and with less risk, you continue to be “recovery agile” because you still have enterprise-class, cloud-integrated, local disk for recovery.

You're not going directly to the cloud and thus exposing your organization to the risk of missed recovery SLAs. Instead, you have fast onsite disk within the AltaVault appliance. And that appliance can be virtual or physical, running on-premises or perhaps in a cloud itself.

Using AltaVault cloud-integrated storage, you can choose to use whichever public or private cloud you prefer, and easily switch later if needed. The only thing end-users may see is that some files retrieve a few milliseconds faster than others because the “slower” ones are coming to them from the cloud.

## **On the Move**

If you're upgrading your data protection architecture, looking to install a better compression and deduplication appliance or backup storage container, or modernizing from tape-based to disk-based backup, AltaVault is a logical consideration. Metricon, who started using AltaVault while it was still being developed by Riverbed, reduced data transport times to the cloud by 75%. Another AltaVault customer, Spot Trading, reduced data footprints by 85%.<sup>7</sup> This technology offers the onsite compression and deduplication benefits organizations need today, and it provides a very advantageous hybrid local-disk-to-remote-cloud-extended platform. You—and your backup software—will both be happy.

## **The Bigger Truth**

The cloud should be part of every data protection discussion these days. The question is *how* you will use the cloud, not *whether* you will use it. If you want to upgrade your data protection, be aware that many ways exist to do so. Changing your existing backup software may be necessary ... *or it may not be*. If you like your backup software, then you might be better off leaving it in place.

After all, implementing new backup software will be more disruptive than simply extending your current solution using cloud-integrated storage. You'll need to weigh your needs against what's available to you in the market and look closely not just at vendor offerings, but at the vendors themselves.

If you are exploring using the cloud in a D2D2C protection architecture, consider NetApp AltaVault, a solution highly worthy of investigation.

---

<sup>7</sup> Sources: [SPOT TRADING IMPLEMENTS RIVERBED WHITEWATER CLOUD STORAGE APPLIANCE TO OPTIMIZE THEIR DATA STORAGE PROCESS](#) and [Riverbed SteelStore Helps Australian Home Builder Metricon Back Up Data to the Cloud and Reduce Its Storage Volume by 80%](#).



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](http://www.esg-global.com)