

# White Paper

---

## **Multi-layered Protection for Multiple Backup Issues: Hybrid Backup Architectures**

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

**September 2013**

---

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

## Contents

Introduction .....	3
Managing Backups .....	3
Protecting the Edge, not Just the Data Center .....	4
A Hybrid Approach .....	6
Considerations for a Hybrid Backup Architecture .....	6
Additional Key Points .....	7
EMC's Backup Ecosystem .....	9
MozyEnterprise .....	9
Which Technology to Use Where .....	10
ESG Lab Assessment of MozyEnterprise.....	11
The Bigger Truth .....	12

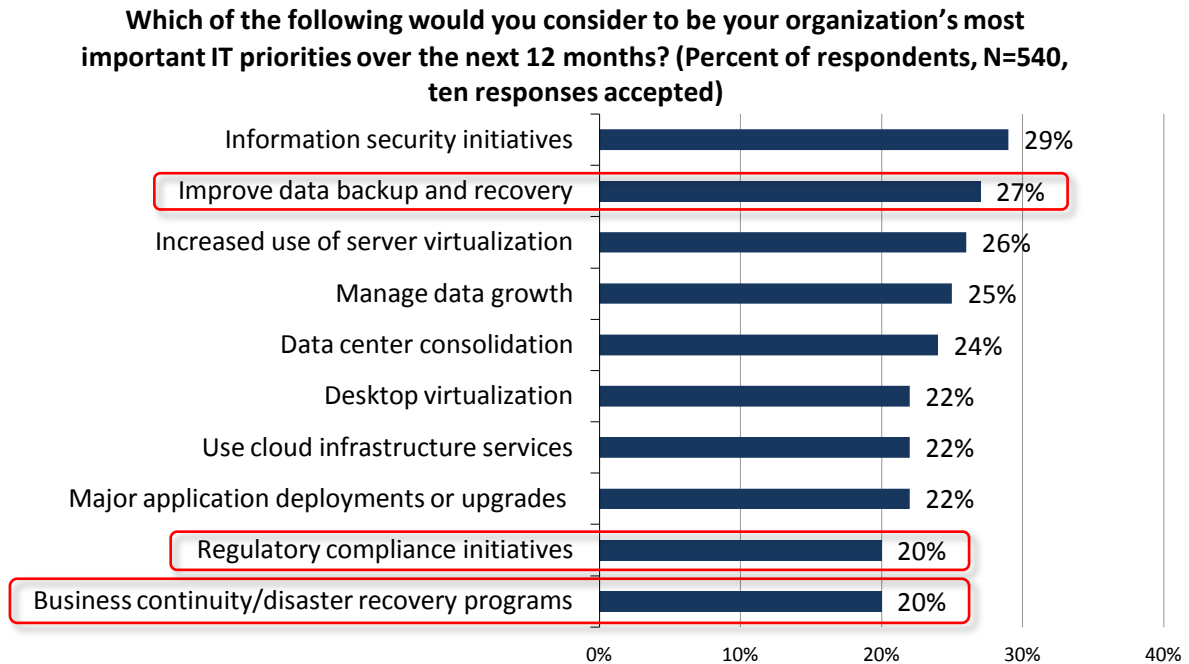
All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from 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

Data protection, a long-time area of concern for IT managers, is holding firmly onto its status as a “forefront” IT priority. According to ESG research, improving data backup and recovery is the second most commonly cited IT priority among respondent organizations for 2013, trailing only information security initiatives (see Figure 1).<sup>1</sup> Two related IT issues—business continuity/disaster recovery and regulatory compliance—also appear on the list of the top-ten most important IT priorities reported by respondents.

The blend of backup, BC/DR, and regulatory compliance can be complicated and can embody both tactical and strategic elements that all contribute to IT’s overall data protection challenge.

Figure 1. Top-ten Most Important IT Spending Priorities for 2013



Source: Enterprise Strategy Group, 2013.

Backup and recovery was also tied as the most frequently reported IT priority in ESG’s 2012 IT Spending Intentions Survey,<sup>2</sup> and it has been one of the top five priorities cited by respondents annually in the survey since 2010, further underscoring the trend: Organizations have been struggling with deficiencies in their backup and recovery environments for a long time.

## Managing Backups

Improving the data backup process calls for solutions to old *and* new challenges—considering that today’s IT environments are bigger because of rampant data growth and are more sophisticated because of advances in server virtualization. The IT industry also has seen increases in the use of cloud-based backup/restore services and diversification in the associated data protection tools. And ESG research shows that public cloud computing services are definitely affecting IT strategy: Seventy-eight percent of respondents to one ESG survey report that cloud-hosted apps or cloud-based compute/storage services will have a significant impact—or at least will factor into—their storage-related strategic planning over the next five years.<sup>3</sup>

<sup>1</sup> Source: ESG Research Report, [2013 IT Spending Intentions Survey](#), January 2013. Regulatory compliance initiatives and business continuity/disaster recovery programs were tied at 20% with business intelligence/data analytics initiatives, improve collaboration capabilities, and deploying applications on or for new mobile devices (not shown).

<sup>2</sup> Source: ESG Research Report, [2012 IT Spending Intentions Survey](#), January 2012.

<sup>3</sup> Source: ESG Research Report, [2012 Storage Market Survey](#), November 2012.

So, IT managers must assess the new technologies that are available. Specifically, they should familiarize themselves with a wider range of solutions than they might have pursued even just two years ago.

## Protecting the Edge, not Just the Data Center

A traditional IT department would place its highest backup priorities on the data center-centric environment. And when IT managers think about “fixing the backup problem,” many of them still tend to focus on bytes contained on the corporate raised floor. But the data most in need of improved protection these days is often housed within the organization’s branch offices.

That reality reflects a larger paradigm shift underway—a shift from device-centric or data center-centric computing to *people-centric* computing. Today’s end-users are not tethered to a corporate desktop. They are not limited to accessing data via VPN. Instead, they are employing multiple devices throughout the day to maximize their productivity, and they expect to enjoy the same basic read/write/edit functionality across all those devices. Perhaps the most pervasive change in IT (ever) is this people-centric approach to service/data delivery, regardless of their device or location.

Clearly, a bridge between “data protection” and “data accessibility” is developing. Data protection is evolving and expanding beyond just *backing up* data to encompass giving *access* to data wherever and whenever end-users want it—with IT retaining responsibility for protecting and securing it all.

Perhaps the most pervasive change in IT (ever) is this people-centric approach to service/data delivery, regardless of the worker’s device or location.

This is a particularly important shift for **branch offices**, where data protection can become almost an afterthought from a strategic perspective. For example, IT may back up information stored in the data center every two hours but decide that nightly backups are “good enough” for remote locations. This decision doesn’t result from an organization believing branch office data is unimportant. Rather, the IT managers deem that nightly backups—or even less frequent backups—are sufficient because they don’t think they can do any better.

### WHAT ABOUT MOBILE PHONES AND CONSUMER TABLETS?

Although many people may say that these devices are productivity assets, the reality is that most of them are *data consumption-devices* and not *creation devices*, meaning that they typically do not have unique data to be backed up.

Most mobile OS vendors offer a rudimentary built-in backup for the purpose of restoring settings after losing or replacing a device, but their data tends to come from other cloud services, including online file stores, mail servers, and media services. Because of this reality, most reasonable data protection strategies do not address consumption-oriented endpoint devices.

**Endpoint devices** also are giving rise to new data access-related needs. Corporate bring-your-own-device (BYOD) initiatives for laptops and other creation devices (tablets/desktops) are breaking many longstanding IT-oversight rules, and IT teams are grappling with how to enable and simplify data access through these devices. IT has been especially challenged in backing up mobile users’ computing platforms. The problem is, if laptops and desktops residing at the edge of an IT environment *aren’t* backed up via a corporate-IT-delivered process, the organization can’t count on having that data should a primary data-loss event occur.

Of course, IT managers also have to figure out what to do about all the laptops and desktops at headquarters. Again, they must assess when it would be best to back up those devices to the on-premises data center, or conversely, when it would be preferable to back them up to a cloud service (for example, when data accessibility beyond the firewall is necessary or when employees take laptops home to do after-hours work).

A stationary desktop at a corporate location, tethered to the corporate network, obviously will be backed up per policy—which could include IT-managed backups to headquarters servers or cloud-based protection of desktops in remote offices.

Laptops access data in a broader variety of ways that will affect one's choices in protecting them:

- If the laptop user is headquarters based and only accesses server data (e-mail and files), then the laptop's cached copies of the data may not require protection—because the server copy of the data is perhaps more “backup-able.”
- But if that same headquarters user is creating new data on that device and doesn't have a server synchronization capability, then an IT delivered (on-prem or cloud-based) backup solution is equally reasonable.
- For users who have to connect remotely to everything (corporate services or cloud providers), cloud-based protection often makes the most sense, though some corporate cultures will drive those backups to corporate servers for alternative user access or confidentiality purposes.

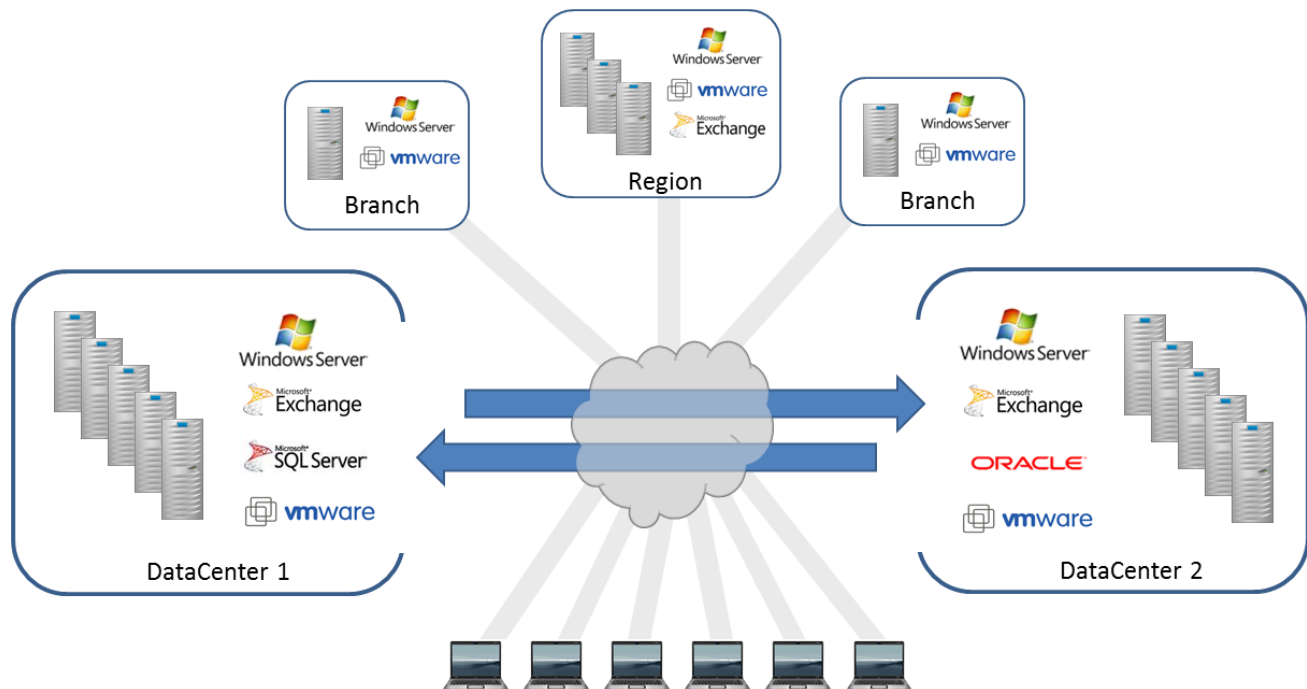
Other variations exist as well, but (1) the types of data being accessed (and their synchronizeability), as well as (2) the predictable proximity of the laptop to a backup server versus a cloud connection, will often drive the default decision process for how endpoint computing devices get protected. And those policies will almost always have exceptions based on cultural/corporate mandates, even more than IT technical rationales.

The point is that in an era when backup to ensure business continuity and regulatory compliance is so vitally important, a failure to protect all of an organization's data *frequently and thoroughly* is unacceptable, regardless of whether that data resides within headquarters, a branch office, or a laptop.

As Figure 2 makes clear, a modern data protection infrastructure encompasses more than a single data center. It's more customary to see one large enterprise with at least two data centers replicating one another bi-directionally, leveraging advanced technologies for backup, fast recovery, and extended retention.

Expanding beyond that core environment, it is also common to find multiple remote and branch offices with many desktops and laptops that all demand a suitable backup approach.

*Figure 2. Data Needing Protection Extends Far Beyond the Data Center*



Source: Enterprise Strategy Group, 2013.

IT frequently attacks the branch-office backup issue by (1) implementing backup to corporate headquarters, or (2) installing a self-contained local solution such as a small appliance or extra disk. Both methods can potentially lead to problems:

- When data moves between thin pipes in the **back-to-headquarters approach**, WAN traffic congestion can result, and it can get bad enough to affect production environments.
- The **local backup approach** presents management problems. Responsibility for tending to the environment typically falls on the shoulders of an onsite non-IT manager who has other jobs to do besides looking after backups. If local backups are managed remotely, that means someone in IT has the disagreeable task of monitoring potentially hundreds of sites to determine which were properly backed up and whether any backup sessions need to be restarted.

## A Hybrid Approach

An alternative deserving consideration is the **hybrid approach**. With it, the data center's backup operations are handled through local and replicated copies as is customary. The remote sites have the flexibility of choosing:

- WAN-based backups (if fat pipes and good management tools are available).
- Internet-based backups (putting day-to-day responsibility into the hands of an external cloud services provider, with high-level oversight by the IT organization).

In a hybrid backup architecture, some parts of the enterprise's data-protection strategy are delivered using on-premises, data center-centric technologies. Other parts are delivered via cloud-based technologies.

This decentralized data protection strategy is not synonymous with unmanaged backup, even though remote-office data isn't being copied "home" to a back-end set of monolithic systems. In-house IT personnel still manage the backup services even if they don't directly fulfill them. (In fact, IT managers should manage the backup policy and strategy regardless of whether backup lives on-premises or at a service provider's facility.)

The hybrid model means an organization does not need to make "either-or" data protection decisions but can instead use centralized and distributed backup approaches where they make sense—for example, centralized data protection could provide the foundation for compliance, while distributed backup could add agility in dealing with remote data.

## Considerations for a Hybrid Backup Architecture

When describing hybrid backup considerations, it's important to distinguish between backup-as-a-service (BaaS) and online file synchronization. They are both cloud services, and they do share a number of characteristics. *But they serve different purposes.*

Consumers and even prosumers (tech-savvy people who manage business-grade data) often use consumer-oriented online file synchronization offerings as if those offerings were BaaS solutions. It's apparent why, given the technologies' surface-level similarities: Both place client software on a user's Internet-connected machine, provide dynamic updates typically at the byte or block level, and are purchased on a consumption-based model. But BaaS and online file synchronization each should be used as intended:

- BaaS offers cloud-based protection of secondary copies of data over time.
- Online file synchronization enables access to files via a user's multiple devices and/or facilitates collaboration among multiple users of a primary copy of a file.

BaaS is about retention of data. OFS is about enabling productivity. They are not the same. BaaS and OFS may look similar, be parts of a holistic endpoint strategy, and may be delivered through the same technology stack, but they are not the same, period.

So, what should organizations look for in a hybrid backup architecture? The goal should be to find one with the following characteristics:

- It provides and ensures data protection and recoverability.
- It is manageable by the IT organization.
- It supports remote office/branch-based personnel and highly mobile workers.
- It provides data protection and data access securely, both to ensure that important files are backed up and to proactively help employees boost their productivity by offering them “anytime/anywhere” access to their files. In other words, it conforms to the *people-centric* computing ideal.
- It offers enterprise-grade security, including a cloud component that is a hardened, mature service (ISO certified, Safe Harbor compliant, with a successfully completed SSAE 16 audit, etc.).

### Additional Key Points

It’s certainly worth emphasizing that a great hybrid architecture doesn’t just prevent drops in productivity; it helps *increase* productivity. Specifically, a solid hybrid backup architecture should:


- **Relieve branch office staff and remote/mobile workers from the burden of managing backups.** That task goes to the IT administrators who are better suited for it.
- **Avoid invasive management techniques** when it comes to backing up endpoint devices. IT groups may wish to install monitoring agents and backup apps on those devices, as they generally do with corporate-owned IT assets. But BYOD users typically won’t tolerate that “heavy-handed” IT desire.

Accordingly, organizations should look for a consumer-style hybrid backup architecture—perhaps one in which an app store (or whatever installation mechanism users are accustomed to) installs a lightweight backup agent that doesn’t require VPN or otherwise hinder the device.

- **Address business-centric requirements** while providing consumer-like functionality. For example, the technology should give IT administrators visibility to determine whether backups are happening, and it should give those admins the ability to wipe a lost or stolen device remotely.
- **Allow IT to provide a cloud-based backup function that appeals to end-users.** An ESG survey of non-IT personnel (see Figure 3) identified cloud-based backup and storage as the second most common service that end-users purchase without IT’s knowledge or consent.<sup>4</sup>

---

<sup>4</sup> Source: ESG Research Report, [2013 IT Spending Intentions Survey](#), January 2013.

**Figure 3. IT Evaluation and Purchase Decisions Made Outside the Purview/Approval of the IT Organization**

Source: Enterprise Strategy Group, 2013.

The upshot: If IT fails to deliver a backup solution that non-technical end-users like, then those users may “acquire cloud backup” on their own—including force-fitting consumer-oriented applications meant for online file synchronizing and collaboration into serving as their “backup” app.

But the IT organization can purchase BaaS and thus be part of the solution, instead of being perceived by end-users as part of the problem. Then for example, when an end-user leaves the company, IT retains access to the business data on that user’s endpoint devices. When an end-user breaks or loses a device, IT can help the user recover the data. IT *can’t* play those roles without knowing how or whether endpoint device backups are occurring.

With all of the challenges related to a distributed collection of devices, a single data protection strategy may be composed of multiple data protection products. But that does not mean that one’s data protection execution should be fragmented or conducted within isolated silos.

Instead, IT organizations should know that portfolios of technologies that are complementary and usable in a cohesive fashion do exist. One such set of offerings is from EMC.



## EMC's Backup Ecosystem

So now what? IT teams have more data than ever to back up. Organizations as a whole are now filled with highly mobile end-users. And remote offices are increasingly vulnerable unless they are being included as part of a solid, secure backup scheme. The [EMC Backup Recovery Systems \(BRS\)](#) division is trying to address those challenges, offering products and services to protect data on platforms ranging from the largest enterprise arrays to the smallest client-side devices. This backup solutions portfolio provides key building blocks that an IT organization can use to deploy a comprehensive data protection architecture.

Notably, one member of the EMC data protection family, [MozyEnterprise](#), enables a hybrid approach that EMC has not had until now. The MozyEnterprise technology, sometimes in conjunction with Avamar (see sidebar), enables IT departments to more completely protect critical data on laptops, desktops, and servers, including those situated in remote and branch offices.

### MozyEnterprise

Mozy by EMC launched as a standalone BaaS provider in 2005, with EMC acquiring the company in 2007 and transitioning the business into its BRS division in 2013.

Because of its simplicity and ease of use, Mozy by EMC is still believed by some people to be a consumer data backup and protection service. However, the Mozy team has been hardening the offering for years, and in particular since 2010, it has been emphasizing its ability to support extremely large enterprises.

The fact is, Mozy has had an enterprise offering in the market since 2008. This division of EMC is selling to enterprises in earnest, and a great deal of Mozy's business now comes from upper-midmarket and enterprise-scale corporate accounts.

The MozyEnterprise product should not be misconstrued as an SMB or midmarket solution that "grew" into one capable of supporting a large enterprise. MozyEnterprise supports backup of both corporate-owned devices and servers and personal-device data, and it is being used by customers as a remote office/branch office backup solution as well.

Mozy by EMC is also a player in the online file synchronization market, launching an offering called Sync in 2013. Online file synchronization falls outside the primary focus of this paper and is not a replacement for a backup service in any case. But Sync can serve as a complementary component of a broader data-management strategy.

Figure 4 illustrates how an EMC-centric hybrid backup architecture might look. NetWorker, Data Domain, and Avamar offer protection in the data center, as a combination of products and services support remote/branch offices.

### EMC OFFERINGS FOR ON-PREMISES AND CLOUD-BASED DATA PROTECTION

**EMC Data Domain**—The EMC Data Domain product line consists of deduplication appliances with software intelligence to let EMC customers reduce redundant backup and archive data, which can help to speed up backup sessions. EMC says its Data Domain dedupe systems provide up to 31TB/hr throughput.

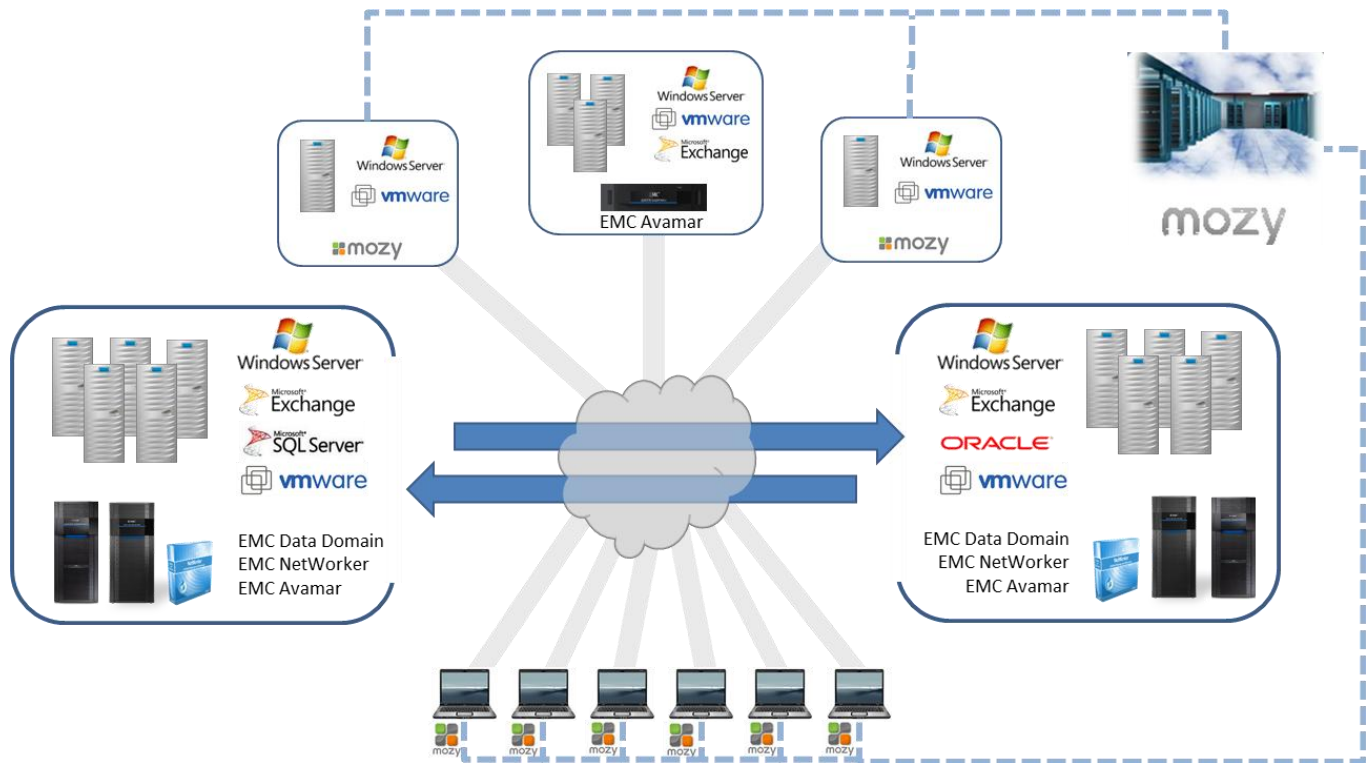
**EMC Avamar**—These backup appliances are optimized (predominantly) for virtualized IT environments. The product set includes backup solutions for remote offices and network-attached storage devices, EMC Avamar Application Modules integrate with enterprise applications, providing backup capabilities for Exchange, Oracle, and SQL Server environments, among others.

**EMC NetWorker**—With its 20-year-plus track record, EMC NetWorker is one of the company's most established data protection products. This enterprise backup software provides wide-ranging platform support as it automates backup and recovery across an organization's IT environment. Customers may add data deduplication capabilities to NetWorker through integration with Avamar client services and the Data Domain Boost product.

**EMC Data Protection Advisor (DPA)**—This offering monitors backup status across multiple products. It supports Avamar and NetWorker as well as non-EMC products including those from CA (ARCserve), Veeam, CommVault, and Symantec (NetBackup). DPA also monitors replication products such as EMC RecoverPoint on an EMC VMAX array.

For example, Data Domain paired with Avamar or NetWorker could suit the remote locations that boast fat pipes. Sites with less WAN bandwidth might opt for internet (cloud) backup with Mozy, which would also make mobile device data accessible for backup. The cloud portion of the architecture would still be centrally manageable by IT through a Mozy administrative console.

Figure 4. A Hypothetical Hybrid Backup Architecture Based on the EMC BRS Ecosystem



Source: Enterprise Strategy Group, 2013.

## Which Technology to Use Where

Using the architecture diagram of Figure 4 as an example, one can glean a few ideas about when to use each part of EMC's data protection portfolio:

- **For larger data centers' server workloads**, EMC's Data Protection Suite (including NetWorker, Avamar, and Data Protection Advisor) paired with Data Domain protection storage seems the obvious choice.
- **For field personnel with laptops**, Mozy's laptop protection and online file synchronization capabilities appear most appropriate.
- **For smaller offices** (see top of Figure 4), the answer isn't as apparent, due to the overlapping capabilities of MozyEnterprise's server protection capabilities and the downsized traditional EMC offerings, including a small Avamar appliance and/or a small Data Domain appliance (both of which can initially help with protecting the data, then can replicate it to a larger appliance at the data center). In fact, within a wholly VMware-powered small office, vSphere Data Protector Advanced (VDPA) may also be suitable because it is built upon Avamar's technologies.

As with most data protection decisions, start with your recovery goals in mind. In this case, consider who has to have access to the data and who needs to be involved in the recovery process.

- If recoverability is primarily for the original remote users, the MozyEnterprise solution may be better.
- If recoverability is primarily for headquarters teammates after a crisis, a data center-centric model may appear better, but even then, those users could also be granted access to recover from the Mozy cloud.

Not clear yet? Then consider the intranet architecture.

- If all remote offices' Internet traffic must pass through a data center (intranet hub and spoke), then perhaps data center-centric backup makes sense.
- If remote offices' Internet traffic goes directly to and from those offices, then a cloud solution such as Mozy reduces that intranet traffic—optimizing access to headquarters resources that would otherwise be competing for bandwidth.

But what about the capital expense and operational requirements of the extra backup hardware compared with a cloud solution?

- The centrally manageable but cloud-powered Mozy solution may trump even the distributed NetWorker/Avamar presumptions that some data center IT professionals are used to.
- Depending on the organization's financial models (including the ability to chargeback, preference of depreciation of assets versus operational line-item charges, etc.), it still could go either way.

**That is the point:** *By adding the MozyEnterprise cloud capabilities to the EMC data center portfolio, EMC is giving their customers choices—with flexibility to design a data protection strategy that fits their unique requirements.*

With that in mind, ESG took a closer look at the MozyEnterprise capabilities in consideration of how IT organizations might add it to their enterprise data protection strategy.

## ESG Lab Assessment of MozyEnterprise

ESG Lab recently conducted an assessment of MozyEnterprise with a special focus on claims related to ease of use and installation, security/trust capabilities, and efficiency for hybrid backups.

Some highlights from the ESG Lab Report:<sup>5</sup>

- ☑ MozyEnterprise demonstrated easy-to-manage online backup and restore.
- ☑ Installation was extremely simple, and the option to separate installation from activation made it easy to deploy at scale.
- ☑ ESG Lab validated that MozyEnterprise offers the highest levels of security with redundant operations; Reed-Solomon erasure coding; state-of-the-art, fully protected NOCs and data centers; and military-grade encryption.
- ☑ A key MozyEnterprise feature is the ability for customers to maintain their own encryption keys with no access by the cloud provider. MozyEnterprise includes this privacy and security option that ensures Mozy has no capability whatsoever to decrypt user data, even under penalty of law.
- ☑ ESG Lab validated that IT can retain backup configuration and enterprise-wide control while Mozy takes care of managing backup storage. This makes it a highly scalable and affordable backup solution.
- ☑ In ESG Lab testing, MozyEnterprise minimized the amount of data transmitted by performing only block-level incremental backups after the initial full backup, and by using pointers to data already stored.
- ☑ The 2xProtect local drive option demonstrated speedy backup and restore, enabling fast, efficient return to productivity as well as minimizing network usage.

In consideration of a hybrid architecture composed of data centers, regional offices, and endpoints that contain critical and unique corporate data, the ESG Lab team found the MozyEnterprise technology suitable for the kinds of data protection that businesses of all sizes are looking for. (Some customers use Mozy to protect tens of thousands of seats.) It might complement the data center protection capabilities that the broader EMC BRS portfolio can offer through EMC Data Domain, EMC Avamar, and EMC NetWorker, as Figure 4 showed.

---

<sup>5</sup> Source: ESG Lab Report, [MozyEnterprise: Secure, Efficient, Cloud-based Backup](#), August 2013.

## The Bigger Truth

When it comes to data protection, a major goal is centralized *management*, not necessarily a centralized repository. And when you pursue this goal, you may have to put some preconceived ideas to the test.

It might seem that anchoring an enterprise backup strategy entirely in the data center is desirable, but it may not be the best technical decision. Considering the shift toward leveraging the cloud for IT, a hybrid approach can be preferable: Such an approach offers management control while leveraging the cloud's economies. It enables—even empowers—increased productivity as part of a “people-centric computing” paradigm; it facilitates users' online file synchronization, and it supports the protection of remote and branch office business data very effectively.

IT managers may have to discard traditional ways of doing business to address today's complex enterprise backup needs. Endpoint devices are going to the cloud even as data center protection remains mostly an internal function. Remote office/branch office backup represents a gray area in which no single solution is likely to fit every distributed location. With a hybrid model, backup data from branches can go to the cloud or to an array in the data center; both of those options are always open.

The nature of backup is changing. Nevertheless, the IT department can maintain control. IT can choose the solution, fund it through a corporate budget, and definitely remain a part of that solution from a data-recovery perspective.

In that manner, the backup strategy remains within the purview of IT, even if the repositories exist in the cloud. That's the beauty of hybrid backup.



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)