

Online Backup and Recovery of Virtual Server Environments

INTRODUCTION

Information technology (IT) groups are embracing VMware to achieve server consolidation and other benefits. Iron Mountain's LiveVault® online backup and recovery is supported on virtual machines, just as it is on physical machines. Using LiveVault with VMware presents new opportunities to combine LiveVault and VMware capabilities to improve an organization's ability to respond to the range of possible data loss events, from individual file corruption to site disasters. This paper outlines the use of LiveVault and VMware to achieve excellent data protection by leveraging the benefits of each technology.

SUMMARY OF LIVEVAULT BENEFITS ON PHYSICAL SERVERS

LiveVault Online Backup and Recovery is offered as a subscription service on a monthly fee basis and as licensed software for large enterprises that wish to operate their own back-end vaults. In this paper LiveVault refers to the subscription service in which Iron Mountain provides the back-end infrastructure. The features and benefits are common to both offerings. For readers that may not be familiar with LiveVault, this section introduces some of the features, all of which continue to be benefits when deployed in a virtual environment:

- **Efficient and immediate off-site protection.** For each backup, LiveVault uses patented technology to identify, compress, encrypt and transmit over the Internet just delta blocks of new and changed data. With LiveVault's continuous protection, you can have full backup versions every 15 minutes – all stored off-site.
- **No more tapes.** Tapes require human labor, time and attention. LiveVault eliminates this.
- **Automatic backups.** No more daily labor.
- **Alerts and monitors.** LiveVault is a true service. If there is a problem or potential problem, you are alerted. No daily oversight is required.
- **Integrated open file and database backups.** LiveVault uses a consistent approach to protect SQL, Exchange, Oracle, and other databases.
- **Selectable retention.** Backed up data, at the server, directory or file level, is retained for a specified period of time, 30 days, 1 year or 7 years.
- **Extensive security and access control features.** LiveVault includes a state-of-the-art suite of security features, far more than just data encryption.
- **Optional on-site TurboRestore™ Appliance.** An optional on-site TurboRestore Appliance keeps recent backups locally, thus allowing fast restores over a LAN.
- **Web user interface.** Management of LiveVault at all sites and servers can be done through a single web portal that does not require any special software or VPN connections.
- **Single, global vendor.** Iron Mountain owns and develops the LiveVault technology, operates extremely secure data vaulting facilities in the U.S., Canada, Europe and the Far East, and provides customer support.

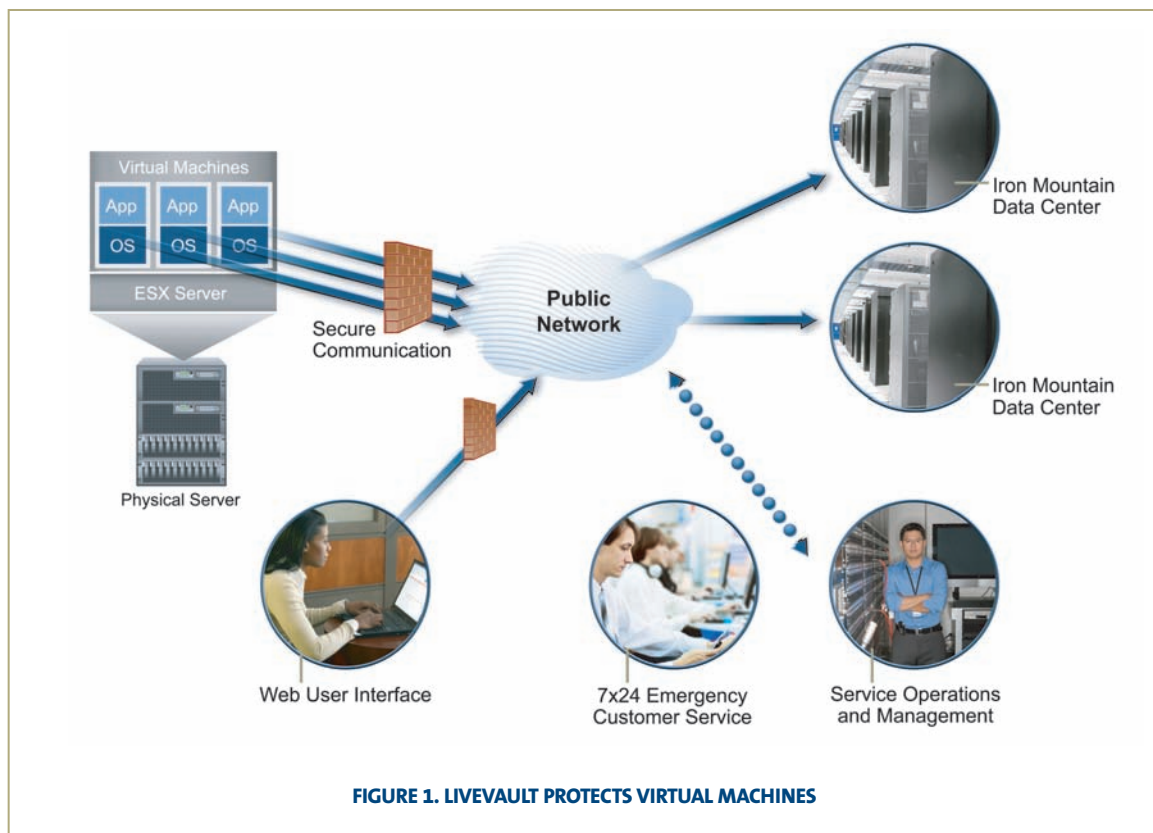


FIGURE 1. LIVEVAULT PROTECTS VIRTUAL MACHINES

LIVEVAULT DEPLOYMENT ON VMWARE

When used with VMware, the lightweight LiveVault agent is installed in each virtual operating system, not in the ESX host environment, as shown in Figure 1. This may be counter intuitive for someone who is thinking purely about disaster recovery. Why does not LiveVault backup the underlying virtual machine implementation at the ESX host level? At this level each virtual machine is just a few files that would need to be protected. While offering protection at the host level would require fewer instances of the LiveVault agent software, there are many reasons why having a LiveVault agent on each protected virtual machine is a good design:

- **Selective backup policies.** Having a LiveVault agent on each virtual server allows the LiveVault web interface to show the administrator the native file system structure (directories, file names and attributes) so that he/she can specify what should be backed up, and what should be excluded. Furthermore, the administrator can create policies to govern the retention and backup schedules. This flexibility allows the administrator to keep the backup cost consistent with the data protection requirements.
- **Selective restores.** In a similar fashion, having a LiveVault agent on each virtual server allows an administrator to restore just one or a few files into a working virtual machine.
- **Open file and database handling.** The LiveVault agent's lightweight protection of open, actively changing files and databases operates in the virtual world just as it does in the physical world.

- **No concern about virtual machine migration.** VMware's Distributed Resource Scheduler (DRS) allows virtual machines to easily migrate from one physical host to another. By having the LiveVault agent as part of the virtual machine environment, LiveVault's backup and restore activity is unaffected by (and unaware of) virtual machine migrations.
- **Simple installs, no added cost, no maintenance.** Installing the LiveVault agent requires only 5 minutes and one reboot. This is a one time event. With the LiveVault subscription service there is no fee for the LiveVault agent, so installing one or a dozen agents does not affect the cost of using LiveVault. After installation, software updates are automatic so management of multiple LiveVault agents is not an issue.

VMWARE CONSOLIDATED BACKUP

VMware offers an add-on called Consolidated Backup. This allows "copies" or snapshots of the virtual machines within an ESX cluster to be instantiated on a designated "proxy" server outside of the ESX environment (that shares the same SAN disk system as the other virtual machines in the cluster). The designated machine exists as a place for backup software to run, and thus the backup software can offload the backup process to another physical server outside of the VMware Infrastructure via SAN-based snapshots. Use of Consolidated Backup with LiveVault would not be a good idea (and is not supported by Iron Mountain). Over 90% of LiveVault's customers use LiveVault's continuous backup schedule to capture up to 96 separate backups per day. This would not be practical if done via Consolidated Backup. As noted above, the LiveVault agents are unobtrusive, maintenance-free, and do not deal with hardware such as tape drives, so there is no down side to running the agents directly in the virtual machines being protected.

SYNERGY STATEMENT

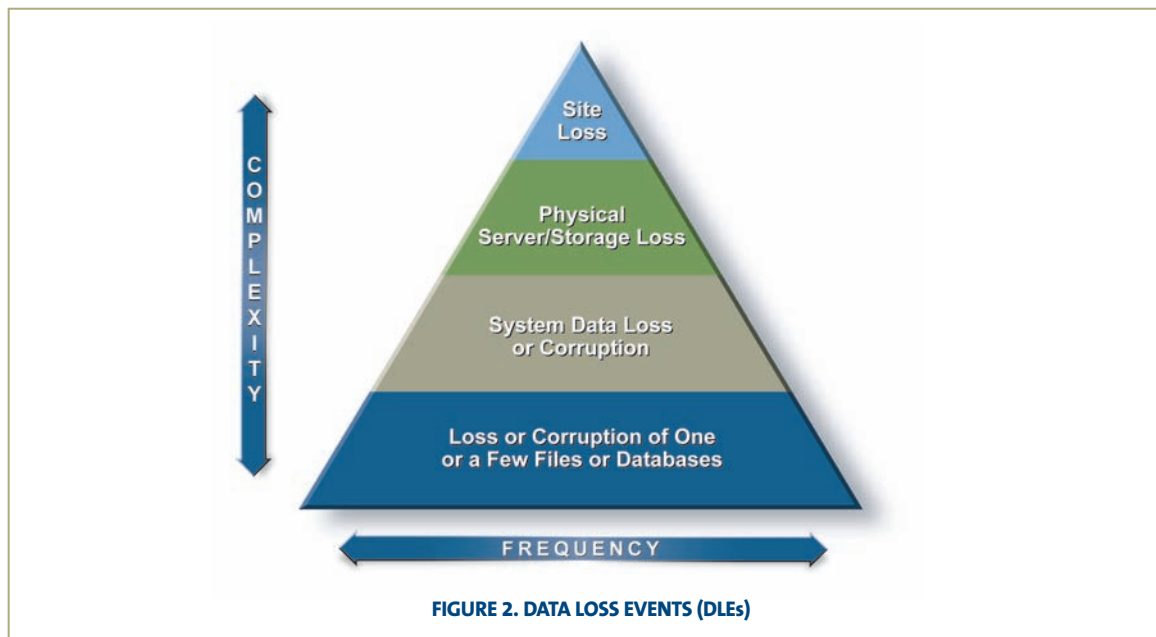
The beneficial synergy of using LiveVault to protect ESX virtual machines is that:

- LiveVault provides excellent protection for data corruption events and file and database losses (i.e., losses that do not involve a whole system). Such losses are not within the purview of ESX which is concerned with the flexible operation, management and migration of systems – virtual systems.
- ESX provides excellent support for operating system level disaster recovery, i.e., "whole system" recovery. With ESX snapshots you can recover whole systems quickly and easily. LiveVault also provides recovery of whole systems. The combination of ESX's and LiveVault's capabilities can simplify and streamline system-level recoveries and/or provide alternate mechanisms to be used depending on the nature of a disaster and the level of preparedness.
- LiveVault provides immediate low-cost off-site protection. The ESX high availability option used between geographic sites at a distance is expensive (especially in bandwidth cost), or effectively impossible (if bandwidth latency is high). On the other hand, LiveVault has been designed to provide server backup to off-site vaults over a wide area network. For protection against total site outages, LiveVault provides effective protection using the minimum resources.

Let us examine these points in more detail. First, let us categorize the types of data loss events.

DATA LOSS EVENTS (“DLE”S)

While disaster recovery, “DR”, is often “top of mind”, there are actually many types of data loss events, some of which are indeed disasters if an organization is not prepared, but all of which a competent IT organization should be able to respond to expeditiously. Figure 2. shows conceptually different types of data losses and their frequency. The bottom level reflects the fact that most data loss events are not total disasters. Empirically Iron Mountain Digital knows that over 80% of restore requests are to retrieve just one or two files or databases. The data may be critical, but from an IT perspective, the response should not require much effort. The top triangle reflects site disasters which only account for about 1% of restores but which require planning and investment in order to handle within a reasonable time frame. This paper will review the use of LiveVault on ESX from the bottom up, starting with the common cases of selected file and directory restores and then discussing major catastrophes.



LOSS OR CORRUPTION OF ONE OR A FEW FILES OR DATABASES

“Leave it to LiveVault”. Corruption is the most common reason for doing restores, and the most common cause of corruption is human error. Protection from site catastrophes may be the most important motivation for buying LiveVault, but recovering from less severe losses is the most common actual usage. Such restores should be easy and simple to carry out. LiveVault is aware of the file and directory structures and its design and function are targeted at managing the backup and restore of files and directories. The LiveVault web interface provides simple and immediate means to backup exactly what needs to be backed up, to specify how long the backups must be retained, to provide continuous backup and multiple recovery points, and a simple means to restore specific items. Furthermore, these functions are designed to operate highly efficiently to provide immediate off-site backup and fast restores via features such as delta backups, DeltaRestore™ and TurboRestore devices. For these reasons, LiveVault is the best choice for meeting the needs for file level backups and restores, off-site storage and for archiving backups.

By contrast, ESX's features are targeted at managing virtual machine contexts, not specific databases or files within a virtual machine. ESX does have the ability to snapshot a virtual machine, either a full snapshot or an incremental snapshot based on prior snapshots. (To be useful an incremental snapshot requires a complete chain of incremental snapshots going back to a full image.) However, using this facility as the basis to be able to retrieve a file is a poor choice because there is no enforced retention and no off-site storage. Using ESX snapshots would require careful procedural design, and careful, consistent execution. History shows that such "home grown" procedures, often found with tape backup, are a primary source of errors that lead to recovery failures. To restore a specific file or directory, an administrator would instantiate a snapshot (or a chain of snapshots) in a separate virtual machine, start the machine, log in and then copy the desired files to the virtual system where they are needed.

SYSTEM DATA LOSS OR CORRUPTION

"Leave it to VMware." In the context of virtual machines, the concern here is the recovery of a virtual machine that "goes bad". In the case of Windows, "system data loss" refers to the operating system files and the system state. The assumption is that the underlying physical machine is operating properly. VMware provides the best means for handling this type of event. The administrator would want to have VMware snapshots made from time to time. In the event of a failed or corrupted virtual machine, the machine can be easily reverted to a prior snapshot state. For example, an administrator could enable snapshots in VMware before installing a Windows patch, test the patch, and if the patch breaks something the administrator can restore using the VMware snapshot. An administrator would not restore from a LiveVault backup version.

When there is a need to recover a system, if the VMware snapshot is older than the most recent LiveVault backup (before the problem occurred), an administrator could take advantage of both backups. The VMware snapshot allows a quick and easy restore of the system. Then a restore from a more recent LiveVault backup will recover the most recent changes (prior to the problem). This is efficient because LiveVault will do a DeltaRestore. LiveVault's DeltaRestore feature ensures that only blocks of data that are actually newer and different will be transferred. Thus if the full system is backed up with LiveVault, LiveVault provides an alternate recovery mechanism, and may have more current data from which a re-instantiated VMware snapshot can be refreshed.

TOTAL SITE OR PHYSICAL SERVER LOSS

"Use VMware and LiveVault together." An organization may have one or more physical ESX host systems. Several physical hosts typically share a SAN and are managed as an ESX cluster. In normal operation there are many benefits to this including the possibility of automatic load balancing and support for high availability that will automatically move virtual machines off of a physical host on which maintenance must be performed or errors are being detected, and is thus likely to fail in the near future.

Single ESX Host

If there is only one ESX host at a site, the typical approach to providing recovery from site and physical server loss is to execute backup software in each virtual machine. This software must be capable of backing up the full system. ESX's capabilities are not important at backup time. Off-site protection requires that the backup data be moved off-site either physically (on tape) or electronically.

LiveVault provides excellent system-level protection, and the ESX virtual machine environment provides excellent assistance and increased flexibility when a system recovery is needed. Recovery is the time that there is valuable synergy between LiveVault and ESX. First, the full system image backed up by LiveVault contains a generic virtual machine image. This means that the physical ESX environment set up after a catastrophe does not have to match the physical environment that was lost or failed. The user has great flexibility as to the hardware used as the recovered environment. (While it may not be advisable to migrate “hot” or running virtual machines between different hardware types, for “cold” recoveries, as is the case after a disaster, virtual machine backups can be instantiated on different hardware without issue.)

Second, LiveVault full system recovery needs a base level operating system so that the LiveVault agent can be loaded to orchestrate the system recovery (operating system plus applications plus data). With ESX, the creation of new virtual machines, pre-loaded with a base-level operating system, is straightforward, and more importantly can usually be done by cloning from a previously created virtual machine template. Thus ESX eliminates the concern of full system recovery onto physical servers where the hardware environment may create issues and the ESX templates can simplify and speed the recovery of multiple virtual machines.

Multiple ESX Hosts

Multiple ESX hosts can share one SAN as an ESX cluster, and/or there can be multiple ESX clusters or hosts with separate disk systems. In multi-host environments it is useful to recognize site disaster, disk storage disaster, and host failures as different data loss events. Some DR protection strategies are suited to some but not all of these DLEs.

Site Disaster

For protection against a *site disaster*, an organization will want to have recent and frequent LiveVault off-site backups of at least the data sets and databases (the “data”), and possibly the full system images. LiveVault is a low cost solution that provides the least complex, least error-prone means of ensuring off-site protection. If only data is protected, then depending on the complexity of the applications and systems, an organization may also wish to create and backup templates from which new virtual machines could be cloned after a disaster prior to restoring the data. As noted above, the generic nature of ESX virtual machines provides peace of mind about the ability to recover onto differing hardware environments that might be available after the catastrophe. LiveVault’s continuous protection and emphasis on reliability provides peace of mind that current data will be available for recovery.

Disk Storage Disaster

ESX’s ability to capture snapshots of virtual machines can be used to achieve local protection against a *disk storage disaster*. These snapshots can be moved to a different storage system and thus be available to support a local recovery. (This does require additional disk space on the second storage system.)

An organization may elect to reduce the frequency with which such snapshots are made and moved and rely on LiveVault to provide the most recent data. After a virtual machine is restored from an ESX snapshot, running a LiveVault restore from LiveVault's most recent backup data will cause LiveVault to only restore the blocks of data that are different or missing in the restored snapshot. This is LiveVault's DeltaRestore feature.

Host Failure

ESX snapshots can also be used to provide local protection against *host failures* within an ESX cluster. For this protection there is no need to move the snapshots off the SAN and thus it may be more acceptable to have frequent full and incremental snapshots.

To speed recoveries for all types of DLEs other than site failures, LiveVault offers a TurboRestore Appliance option. This is a local system that can hold the most recent backups so that LiveVault restores can be done over the LAN without data moving over the Internet.

CONCLUSION

With ESX virtual machines, there is little need to worry about hardware differences when planning or executing a disaster recovery, and VMware provides a range of capabilities, that are useful for managing and recovering virtual systems within a site. LiveVault provides unobtrusive off-site protection for virtual machines with backup archive history up to seven years. LiveVault also provides a direct, simple method of specifying what data needs to be backed up and easy restores of specific files and databases when needed – a common restore scenario.

The combination of VMware and LiveVault strengths provides an organization with excellent flexibility and reliability in protecting and recovering the data assets of an organization from all types of data loss events. The combination allows an IT group to efficiently meet the full range of backup, restore and retention requirements.

This page intentionally left blank.

© 2008 Iron Mountain Incorporated. All rights reserved. Iron Mountain, the design of the mountain and LiveVault are registered trademarks, and Iron Mountain Digital, DeltaRestore and TurboRestore are trademarks of Iron Mountain Incorporated. All other trademarks and registered trademarks are the property of their respective owners.

 **IRON MOUNTAIN** DIGITAL™
120 Turnpike Road
Southborough, Massachusetts 01772
(800) 899-4766

Iron Mountain Digital, the world's leading provider of data backup/recovery and archiving software and storage as a service (SaaS), offers a comprehensive suite of data protection and e-records management software and services to thousands of companies around the world. For more information, visit our Web site at www.ironmountain.com/digital.