

AVAMAR INTEGRATION WITH VMWARE

Nadine Francis

Table of Contents

Virtualization Leading IT Transformation	3
Why the IT Transformation?	4
How can Backup & Recovery accelerate this journey?	4
Why Avamar?	6
Avamar Virtual Edition	6
Guest-Level Backup & Recovery	7
VMware guest backup capabilities	7
VMware guest backup limitations	8
Image-Level Backup & Recovery	8
VMware Image backup capabilities	9
VMware Image backup limitations	9
Why CBT?	9
Backup & Recovery with CBT	10
The requirements	12
VADP Snapshot Consideration	12
Avamar Integration with Data Domain	12
AVE with Data Domain	12
Licensing Guide	14
Appendix	17

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Virtualization Leading IT Transformation

The IT Transformation Storymap graphically depicts the aspects of transformation to deliver IT as a service across infrastructure, application, and operating model, using an analogy from automobile manufacturing.

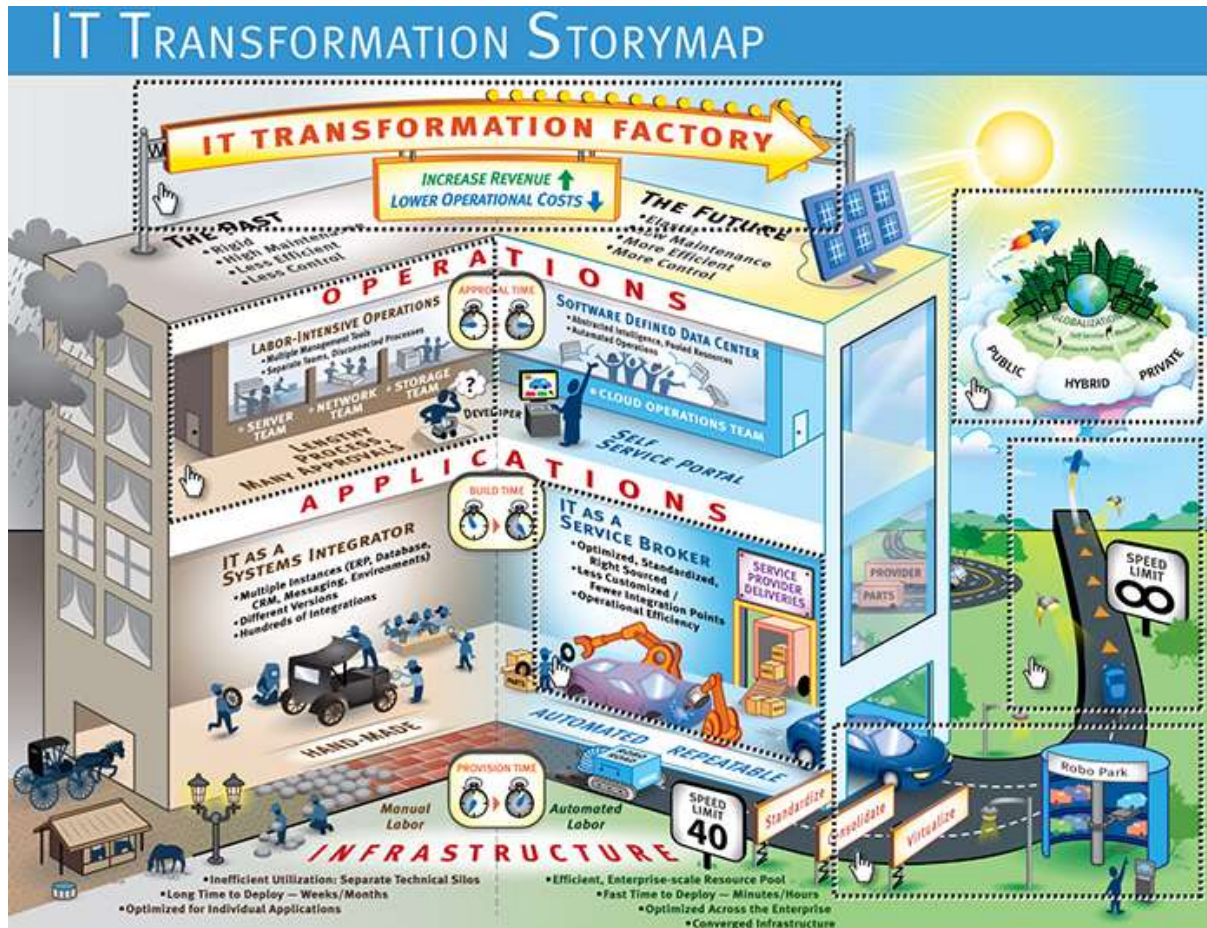


Figure 1: IT Transformation Storymap

Generally, EMC believes there are three steps for the IT transformation.

Step 1: Transform existing infrastructure to a cloud infrastructure. This is the move from rigid application-based silos to a truly dynamic infrastructure that can address the needs being placed on IT. In this step, virtualization is relied on to reach what we call the “software-defined data center”.

Step 2: Transform the operating model to deliver IT as a service. IT must act and operate like a Service Provider.

Step 3: Build applications for the cloud. Leverage Big Data to build mobile and predictive analytic applications.



Figure 2: The 3 steps of IT Transformation

While we still have work to do, improvements to make and kinks to work out. This will always be the case. Transformation is eternal. The need to adapt will be increasingly important as the pace of business continues to accelerate and complexity manages to creep in as scale increases. Transformation is not a simple process. It is a journey that takes patience and a collective effort from all parts of the organization, but the results are worth the trip.

Why the IT Transformation?

These transformations are happening across infrastructure, applications, and operations. They increase revenue and lower operational costs for the business while providing a multitude of services are the key goals for many IT organizations. This requires services that are both compelling and meaningful to the business and may be delivered by both internal and external providers to meet these goals.

How can Backup & Recovery accelerate this journey?

By enhancing backup and recovery, EMC and VMware together can help its customers unlock additional business value and accelerate their journey to broader, potentially 100 percent, virtualization.

Enhanced backup and recovery leads to extended business value that fuels the journey to 100 percent virtualization.

Avamar's enhanced backup and recovery enables you to:

- Increase consolidation ratios
- Perform 90 percent faster backups
- Experience 40- to 60-times data reduction and a 99 percent reduction in replication bandwidth
- Have file, image, and site recovery

Avamar's extended business value translates into:

- Lower backup infrastructure costs
- Simplified management
- Improved operations
- Investment protection
- Reduced risk

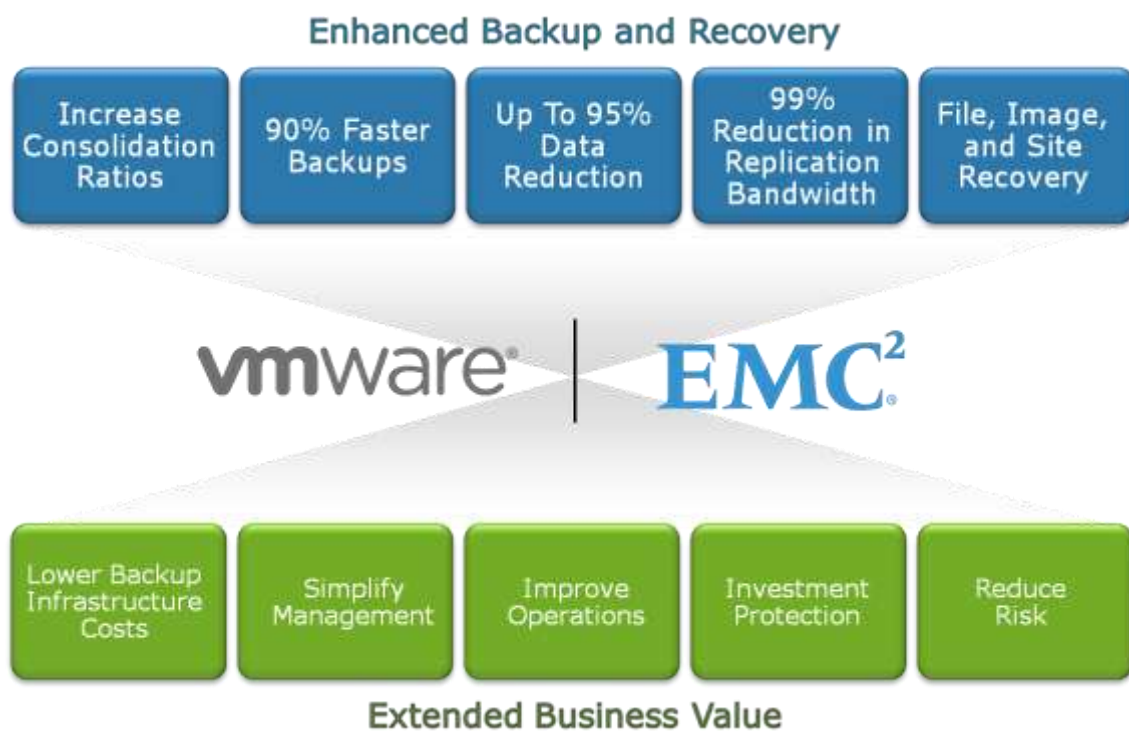


Figure 3: Benefits of IT Transformation on Business Value & Backup and Recovery Systems

Why Avamar?

EMC Avamar® provides variable-length client-side deduplication to accelerate the virtualization journey by providing extremely fast and efficient backup and recovery for the VMware environment. Avamar protects virtual machines (VMs) by deduplicating data at the client level so that only new, unique, sub-file, variable-length data segments are sent during daily full backups. While traditional backup software moves upward of 200 percent of the primary backup data on a weekly basis, Avamar moves as little as 2 percent over the same seven-day period, removing backup bottlenecks and enabling even greater levels of virtualization.

In many ways, Avamar was really built for VMware backup. Additionally, Avamar and VMware is a great match as it supports all major modes of VMware backup, either guest or image. We will discuss them in depth in the next sections.

Avamar Virtual Edition

Avamar Virtual Edition (AVE) for VMware is Avamar server software deployed in a VMware virtual machine. AVE for VMware is an ideal product for small, medium, or remote offices that have standardized on VMware infrastructure. AVE for VMware provides:

- Rapid, cost-effective deployment.
- Simplified management by virtualizing all aspects of the backup and recovery offering.
- Lower cost by leveraging shared server and storage infrastructure.
- Replication in virtual environments as well as between virtual and physical environments for disaster recovery.
- Rapid return on investment.
- VMware vMotion support.
- Scalability with support for up to two Avamar Virtual Edition for VMware virtual appliances per ESX server.
- Capability to work on any server and iSCSI, SAN, or DAS storage that meets the combined Avamar and VMware specifications.

Starting with Avamar 7.0, each Avamar Virtual Edition for VMware virtual appliance supports up to 4TB of licensed disk storage capacity provisioned to a virtual machine on a VMware ESX server. Other available capacities: 500GB, 1TB, and 2TB.

Guest-Level Backup & Recovery

With the guest backup option, the Avamar client software is installed on the individual virtual machines. Backup configuration for this method is identical to that of a physical machine. Backup agents are installed on the virtual machines and send their backup data to the backup application for storage.




Figure 4: Guest-Level Backup & Recovery

VMware guest backup capabilities

- Highest level of data deduplication.
- Support for backup of applications inside the virtual machines.
- Application-consistent backups.
- Support for partial or file-level restores.
- Identical backup methods for physical and virtual machines.
- No requirement for advanced scripting or VMware software knowledge.
- Unchanged day-to-day procedures for backing up.

VMware guest backup limitations

- Since each virtual machine has a separate backup client installed, ESX servers with a large number of virtual machines may experience a strain on resources, especially memory, if all machines are backed up at the same time.
- No Image-Level restore option.
- File-level restore is done in one step; however, full system restore is done in two steps.

Image-Level Backup & Recovery

Avamar is tightly integrated to vStorage APIs for Data Protection (VADP) for agentless backups. VADP enables backup products to perform centralized, efficient, off-host backups of vSphere virtual machines. Introduced in vSphere 4.0, VADP replaces the VMware Consolidated Backup Framework (VCB) for virtual machine backups. Since VADP integrates directly with Avamar, no additional software is required to be downloaded and installed. Incremental backups and restores are supported through the use of change block tracking (CBT) which will be explained in detail later in this article.

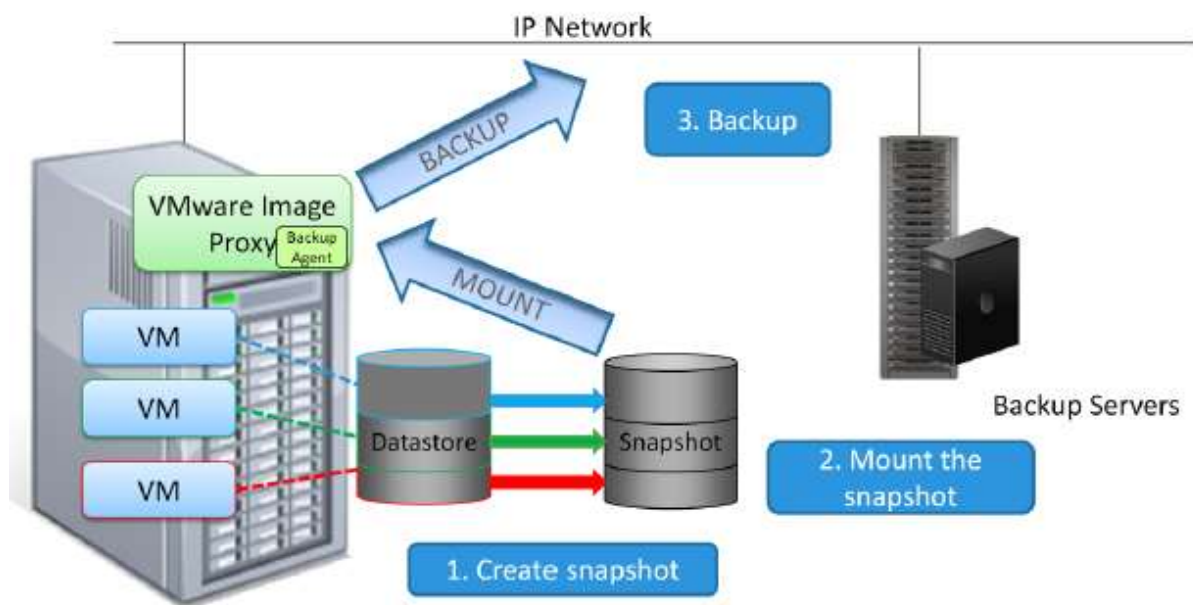


Figure 5: Image-Level Backup & Recovery

VMware Image Backup uses a proxy server to handle backup processing. Using VADP, Avamar can back up VMware virtual machines without using backup agents inside each virtual machine. This proxy server can communicate with the vCenter server to mount a

snapshot of a particular VM's vmdk to perform image-level backup of that virtual machine. With this method, deduplication is provided on the file level as well as on the .vmdk level.

Performing an image-level backup of a VMDK, the ability to execute a bare metal restore replacing a VM instance or creating a separate instance is also possible and supported. This is accomplished by creating a snapshot of the virtual machine, then backing up that snapshot. No downtime for the virtual machine is required. Backup processing can be offloaded from the ESX hosts, depending upon the method of deploying the proxy. File-level backups and restores are supported with image level backup.

VMware Image backup capabilities

- Perform backup and restores without installing agents on individual virtual machines.
- Perform backups from a centralized location, thus offloading backup processing from the ESX hosts.
- Perform backups any time because backups are non-disruptive to the virtual machines being backed up.
- The ability to perform backups at any time provides greater flexibility in scheduling backups and increased backup window times.

VMware Image backup limitations

- Requires at least ESX 4.0.
- Requires a moderate amount of VMware knowledge to set up and configure backups and restores.
- Backups consume ESX server resources, including CPU, RAM, and disk.
- Backup requires a temporary virtual machine snapshot.
- Backup is crash-consistent snapshot of full VM image and may not support a full system restore without data loss.
- File level restore from an image backup for Windows and for Linux, starting with Avamar 6.1.

Why CBT?

Starting with Avamar 6.0, support for vSphere Changed Block Tracking (CBT) enables faster backup and recovery by tracking incremental changes on virtual machines, introducing incremental protection instead of requiring full backups and full restores.

Another enhancement with Avamar 6.0 is proxy VM load balancing. With proxy VM load balancing, VM backups utilize a pool of backup server resources balanced for higher throughput and backup speed. Provisioning of two proxy servers enables the VADP process to be performed on two VMs simultaneously for backup by Avamar. Previously, VMs were assigned manually by the administrator to specific proxy servers to assure proper load balancing. However, as of Avamar 6.0, this load balancing occurs natively. Avamar will now utilize any idle proxy server to automatically load balance workload and maximize parallelism through the backup process. Flexible image recovery options enable administrators to restore any backup image to a virtual machine, not just the most recent full backup of that specific VM. This vCenter integration feature enables Avamar to create a VM in a vSphere environment as part of the recovery process.

Backup & Recovery with CBT

The Avamar agent is installed on a Proxy VM, off-loading the backup process from any of the VMs where the applications are running. Through vSphere, each VM is dynamically mounted to the proxy without physically moving data across the network, enabling Avamar to back up numerous virtual machines in just minutes. With CBT, VMware presents only changed blocks to the Avamar agent, where each block is broken into variable length segments and further evaluated for uniqueness. Only the unique segments are sent for backup, achieving the fastest backup possible. Conversely, the restore process also leverages CBT for faster recovery. Avamar enables full VMDK or file-level restore to the original VM, an existing VM, or a new VM directly from the Avamar user interface as shown in Figure 6-1, 6-2, and 6.3.




Figure 6-1: CBT Backup

The VMware vStorage API for Data Protection determines the changed blocks. Then, the proxy reads only these changed blocks, deduplicates them, and sends the unique blocks to Avamar as depicted in Figure 6-1.




Figure 6-2: Standard Restore

Prior to Avamar 6.0, Restore operations were an "all-or-nothing" implementation as depicted in Figure 6-2. When restoring a particular VM image from a backup within Avamar, the entire VMDK would need to be sent over the network and restored. Recovery of only the individual changed blocks since the last backup within that VM image was not possible.

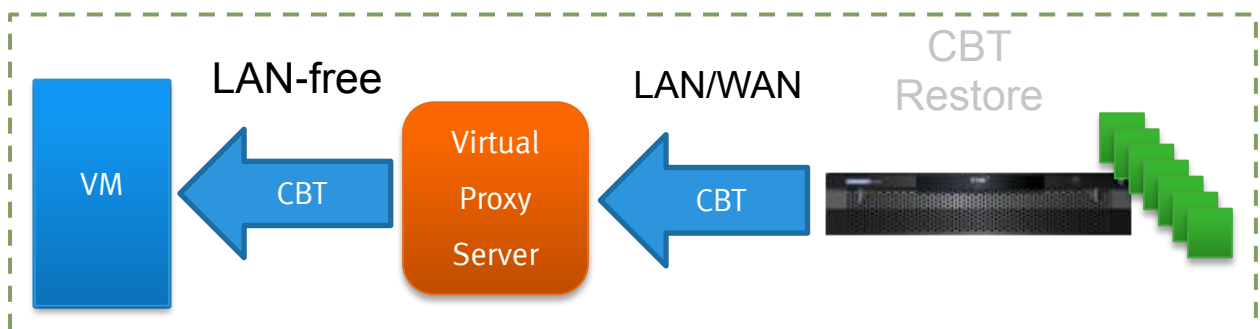


Figure 6-3: CBT Restore

CBT Restore (Figure 6-3) is similar to CBT Backup as it creates a snapshot and uses the VMware vStorage API Data Protection to determine the changed blocks. Then, only the changed blocks are read and recovered to the original VM.

The Avamar software will automatically evaluate the workload between both restore methods (Full Image Restore, or a recovery leveraging CBT) and perform the method that will result in the fastest restore times for your particular scenario or environment.

The requirements

Virtual machines with virtual hardware version 7 and later running on ESX/ESXi 4.0 and later hosts support CBT. Disks cannot be in physical compatibility RDM mode and cannot be mounted by multiple virtual machines. This feature is enabled and disabled for each virtual machine. Because there is a slight VMware performance impact, CBT is disabled by default.

VADP Snapshot Consideration

By default, VMware snapshots are created on the datastore where the virtual machine's configuration files are located. To minimize the time required for taking and removing a snapshot, schedule backups for a virtual machine when there is less I/O activity occurring on the virtual machine datastore. It is recommended to keep at least 20% free space on all datastores for snapshot management. It is important to ensure that the snapshot working directory supports the size of all the disks attached to the virtual machine. In cases where the virtual machine has a high change rate during backups, the snapshot working directory must be able to support any snapshot size while the backup is running. Special consideration is given as well for support of virtual compatibility RDM disks.

Avamar Integration with Data Domain

Data Domain[®] system integration enables deployment of enhanced target-based deduplication in an Avamar environment. It creates the ability to align backups, based on their attributes, with the most effective deduplication approach (e.g. source or target). For large enterprise application backups such as Oracle, SQL, SharePoint, Exchange, DB2, SAP, Sybase, Hyper-V, and VM Images, Avamar can be configured with Data Domain Boost (DD Boost) software to send backup data directly to a Data Domain system. DD Boost is a software API that is automatically installed with the Avamar agent on supported clients. It provides a backup connection to the Data Domain appliance and manages the initial block segmentation portion of the deduplication process. With DD Boost, Avamar sends the backup stream directly to the Data Domain target and the related metadata to the Avamar Server. This separation process allows Data Domain targets to be centrally managed from the Avamar Server management console.

AVE with Data Domain

With Avamar 6.0 and above, VMware image backup data can be sent to a Data Domain device. AVE can be used to store Avamar's checkpoints, which includes a backup catalog (metadata), clients, and more. As shown in Figure 7, all data center workloads are sent to the Data Domain system through Avamar's integration to DD Boost software for backup

storage, while the Avamar checkpoint is stored on the AVE virtual appliance. AVE will then copy the checkpoint to the Data Domain system.

As discussed earlier, Avamar provides a single pane of glass to manage the entire backup process.

The use of AVE reduces the hardware required as no Avamar Data Store hardware is required, which in turn streamlines deployment and lowers costs.

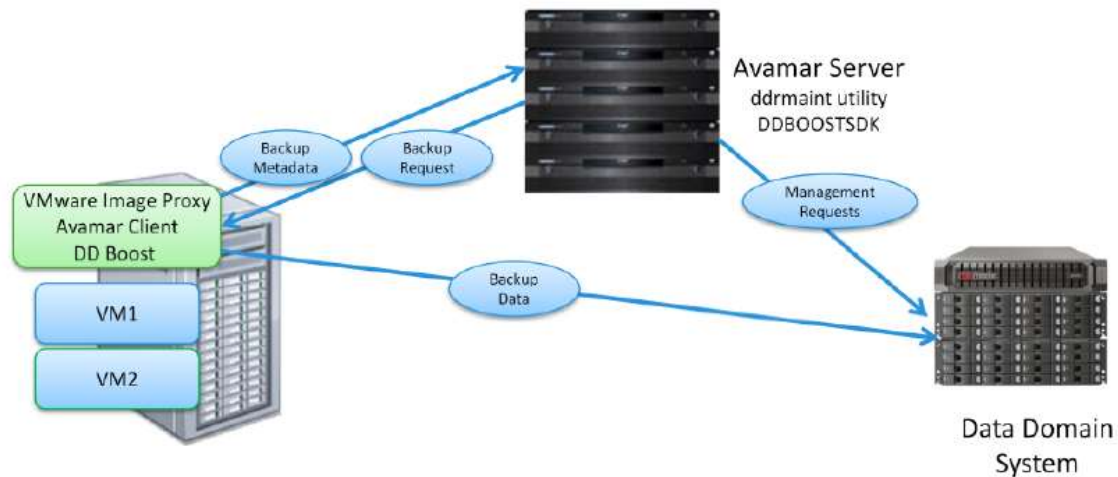


Figure 7: Avamar integration With Data Domain

Storage and retrieval of backups on a Data Domain system may be faster than on an Avamar server in environments with large, active databases. Beginning with Release 6, Avamar provides the ability to selectively perform backups and restores for VMware image backups to a Data Domain device. Backups, restores, and replication between Data Domain devices of Avamar backup data stored on Data Domain systems are configured, managed, and reported in Avamar. DD Boost software is installed with the Avamar client software.

When using Data Domain as a storage device for Avamar backups, all backups for a client must be stored on the same device, either the Avamar server or a specific Data Domain system. When moving an Avamar client to backup to another device, a full backup must be performed on the new device.

The Linux and Windows VMware Image plug-ins for backup allow the user to specify whether to use CBT and to send the backup data to a Data Domain system. To use CBT, the functionality must also be enabled at the virtual machine to write the backup data to Data Domain.

CBT is also a restore option for VMware Image Linux and Windows plug-ins. There are no user-configurable options available for the Linux and Windows VMware file-level restore plug-ins.

Now, with Avamar 7.0, VMware VM images backed up with Avamar to a Data Domain system can be accessed instantly and run from the Data Domain system. This provides the end user access during a critical time period when the end user cannot wait (even a few minutes) for a restore back to the production environment.

Licensing Guide

Deciding which backup method to employ for backing up virtual machines depends upon many factors. These include ease of use, efficiency and impact of backup processing on resources, and backup and restore capabilities.

The guest backup method is not dependent upon the version of ESX server. Rather, it is dependent upon whether the operating system and applications running on the virtual machine are supported by the Avamar. VMware Image Backup and Restore with VADP are supported with ESX server version 4.0 and above.

As mentioned earlier, guest backups and restores are performed using the same procedures as those for physical machines. Thus, there are the same advantages and considerations for guest backups as for physical backups. Installing and managing a backup agent on every machine are required. Because a backup agent is running on the virtual machine, resources of the virtual machine, including CPU and memory, will be consumed during backups. However, unless all backups for all virtual machines running on an ESX server are running at the same time, there is no impact on the ESX server itself.

Employing VMware image backups and restores require learning new configuration, backup, and restore procedures. Also, the backup administrator should be familiar with VMware technology and the topology of the individual site, such as which ESX servers host, each datastore, and which datastores contain data for which virtual machines.

With guest backups, it is possible to employ VSS and online backup technologies to ensure backup consistency. Backups can be customized, advanced features can be employed, and pre- and post-scripts can be run. Backups are highly optimized, providing the ability to exclude temp files, swap files, operating system files, and other files that are not required for the backup. Database backups provide the ability to truncate transaction logs. However,

virtual machines must be powered on in order to be backed up. There must be a network connection between each virtual machine and Avamar. Individual files and directories can be restored for file system data. This ability is dependent upon the capabilities of database applications and agents.

Image backups are supported for all virtual machines and the virtual machines need not be running or connected to the backup application for backups to occur. Because database applications are not involved in the backup, database consistency in backups cannot be assured. Backups are crash-consistent snapshots of the VM image, but may or may not support full system restore without data loss. Also, backups cannot be optimized for specific applications or customization. For example, exclude lists are not supported. In addition to restoring an entire image from backup, backup software may provide the ability to restore individual files and directories. To summarize it, let's look at the following table.

	Guest Backup	VMware Image Backup
ESX Server Version	<ul style="list-style-type: none"> • Not dependent upon ESX version. 	<ul style="list-style-type: none"> • 4.0 and above.
Ease of use	<ul style="list-style-type: none"> • Fits in with traditional backup environment procedures. • Requires backup agent on each virtual machine (VM). • No VMware knowledge is required. 	<ul style="list-style-type: none"> • Requires initial setup and configuration that is different from traditional methods. • Backup agent is not installed on every VM. • Moderate VMware knowledge is required. • Knowledge of site's vCenter topology is required.
Efficiency regarding resources	<ul style="list-style-type: none"> • Backup doesn't consume ESX server resources. • Backups consume guest VM resources. 	<ul style="list-style-type: none"> • Backups consume ESX server resources.

Backup support	<ul style="list-style-type: none"> • Supports VMs running OS and applications for which backup software provides support. 	<ul style="list-style-type: none"> • Image backups are supported for all VMs.
Backup	<ul style="list-style-type: none"> • Ensures consistency by quiescing VM before backup. • VM must be online for backup to occur. • Each VM must have a network connection to the backup application. 	<ul style="list-style-type: none"> • Backups are crash-consistent snapshots of VM image; may or may not support full system restore. • VMs don't need to be running. • VMs don't need to have a network connection to the backup application.
Restore	<ul style="list-style-type: none"> • Individual file and folder restores are supported. 	<ul style="list-style-type: none"> • Individual file restores are supported for specific OSs.

Appendix

- EMC Software Price Guide: Backup and Recovery.
http://powerlink.emc.com/km/live1/en_US/Sales_Support/Sales_Guide/EMC_SoftwarePriceGuide_BackupRecovery.pdf?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tTZWFyY2hSZXN1bHRzRXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwNzkwMzY1LGRhdGFTb3VyY2U9RENUTV9lbnVU18w
- Data Sheet: EMC Avamar for VMware.
http://powerlink.emc.com/km/live1/en_US/Offering_Basics/Offering_Collateral/h2823-avamar-vmware-ds.pdf?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tTZWFyY2hSZXN1bHRzRXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwNzUwMmRkLGRhdGFTb3VyY2U9RENUTV9lbnVU18w
- Customer Presentation: EMC Avamar for VMware Environments.
http://powerlink.emc.com/km/live1/en_US/Offering_Basics/Presentation/avamar-vmware-environments.pptx?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tTZWFyY2hSZXN1bHRzRXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwNzRmY2M3LGRhdGFTb3VyY2U9RENUTV9lbnVU18w
- Customer Presentation: What's New in EMC Avamar 7
http://powerlink.emc.com/km/live1/en_US/Offering_Basics/Presentation/Avamar-whats-new.pptx?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tDb250ZW50RXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwNzU2MGNILGRvY3VtZW50VHlwZT1wCHR4LG5hdmVOb2RIPtBiMDE0MDY2ODAxZmEwYmRfR3JpZA
- Customer Presentation: What's New in EMC Avamar 6.1
- Customer Presentation: What's New in EMC Avamar 6.0
- White Paper: Backup & Recovery for VMware Environments with Avamar 7
- http://powerlink.emc.com/km/live1/en_US/Offering_Technical/White_Paper/h8945-backup-for-vmware_environments-with-avamar-v6-wp.pdf?mtcs=ZXZlbnRUeXBIPUttQ2xpY2tDb250ZW50RXZlbnQsZG9jdW1lbnRJZD0wOTAxNDA2NjgwNzUwYmM0LGRvY3VtZW50VHlwZT1wZGYsbmF2ZU5vZGU9MGllwMTQwN
- Backup Recovery Solution: VMware Integration (Expert Material).
- Data Protection Strategy and Vision presentation

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.