

Data Protection for VMware and Application Data in Mission Critical Environments

HP StorageWorks and HP Software Division

Table of contents

Executive summary.....	2
The challenges	2
Backup agent in the virtual machine	4
Backup agent in the console OS/hypervisor.....	5
Backup through proxy server	6
Backup option summary chart	7
Solution technology components.....	8
VMware technology components	8
Virtual Infrastructure 3	8
VMFS.....	9
Raw Device Mapping (RDM)	10
VMware Consolidated Backup (VCB).....	11
HP technology components	11
HP Data Protector software	11
HP StorageWorks EVA Family	12
HP StorageWorks Business Copy EVA (EVA BC)	13
HP StorageWorks Continuous Access EVA (CA EVA).....	13
Implementation overview	14
Implementation configuration	15
Storage and server hardware configuration	16
Implementation details.....	17
VMware Consolidated Backup Integration in HP Data Protector 6.0 software using pre/post execution scripts in a backup specification	18
Performing a Backup	19
VMware Consolidated Backup Integration in HP Data Protector 6.1 software using the VMware online agent integration	22
Data Protector Cell Manager	23
Data Protector VMware Integration component	23
Data Protector Media Agents	23
Configuring the VCB Integration	24
VCBimage backup method	26
VCBfile backup method.....	27
Performing a Zero Downtime Backup.....	27
Conclusion/Summary.....	30
Appendix A	31
For more information.....	32

Executive summary

Enterprise data centers have been rapidly embracing and deploying server virtualization solutions because of the instant and evident benefits server virtualization provides. From fundamental problems such as server sprawls, resource consumption, and server provisioning to more complex challenges such as green computing and disaster recovery, server virtualization provides many different capabilities that help address these common data center intricacies. VMware ESX is the most widely used server virtualization technology in data centers today. In fact, an ESG survey¹ indicates that 70 percent of surveyed virtualization adopters are running VMware in their environment. The rapid adoption of server virtualization is accelerating the deployment of high-utilization virtual machines, which is, in turn, increasing storage capacity usage because physical servers are now running multiple virtual machines, all of which require their own storage capacity usage. This train of events is forcing IT managers to rethink their backup methodologies.

This paper explores and outlines how to apply HP Data Protector 6.0/6.1 software and other key HP StorageWorks storage array technologies to manage and deploy an effective backup strategy in a VMware Virtual Infrastructure, which not only effectively protects virtual machines but also the enterprise applications running inside the virtual machine. HP Data Protector software provides and automates high-performance backup and recovery from disk or tape over unlimited distances. HP Data Protector software is a full-featured and robust backup software that enables enterprise-level capabilities such as continuous backup and instant recovery at a low deployment cost that is 30-70 percent lower than competing products². VMware Infrastructure features and solutions—VMware Consolidated Backup, encapsulation, Raw Device Mapped, Virtual Machine File System (VMFS), and Virtual Machine Disk Format (VMDK)—provide the foundation for building a robust backup solution.

This joint solution demonstrates companies can build and deploy an effective backup solution in their VMware virtualization environment while achieving improved RPO (Recovery Point Objective) and RTO (Recovery Time Objective) of backups, significantly lower backup windows, and efficient data mobility in the process.

The challenges

Performing backups is a mission-critical component of the day-to-day operation of any data center. Though the virtual infrastructure has provided IT managers with effective solutions to address many data center challenges, the rapid deployment of server virtualization has also introduced a new set of backup challenges. IT managers must employ innovative strategies to address business continuity needs of their virtual infrastructure. A survey conducted by Excillio Group Inc (marketing research firm) indicated that 26 percent of surveyed professionals view the implementation of a backup strategy for their Virtual Infrastructure as their top priority for 2008.

¹ ESG Research Report, ESG IT Infrastructure and Service Management Survey, March 2008

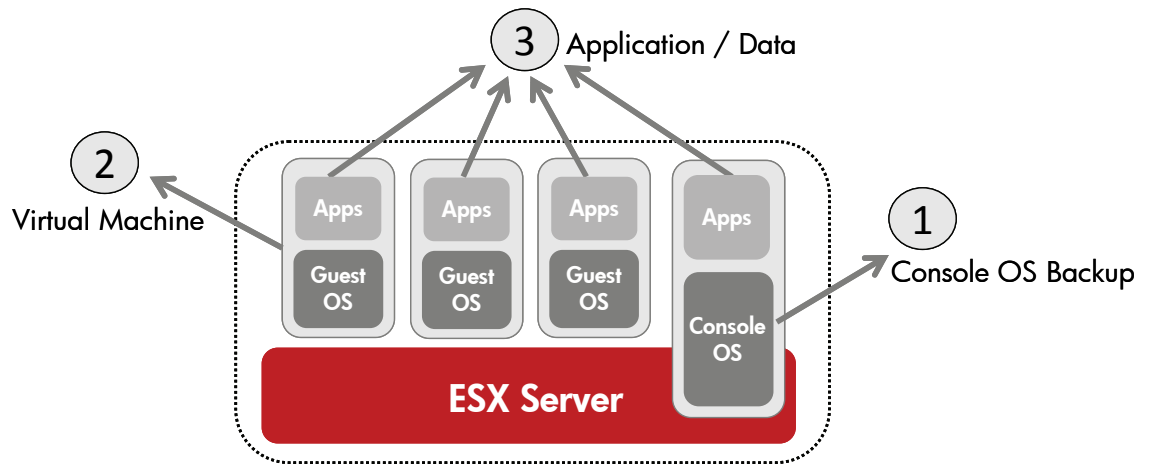
² <http://www.hp.com/go/dataprotector>

So why is the implementation of a backup strategy for a virtual infrastructure not as simple as backing up the set of data and configuration files that encapsulate a virtual machine?

To answer this question effectively, it is important to understand all the components of a virtual infrastructure that need to be backed up and how critical effectively backing up each component is. Figure 1 below depicts the different components of a virtual infrastructure to consider for backup when building a backup strategy:

1. The console OS
2. The virtual machine
3. The application running inside a virtual machine or console

Figure 1: What to backup?



Each of the Virtual Infrastructure components above can be backed up using one or more of the approaches below:

- Backup at the virtual machine level
- Backup at the console OS/hypervisor level
- Backup via proxy server

Each of these backup approaches above presents a different set of challenges and fails to provide a data protection solution that not only is simple to manage and execute but also provides low RTO and RPO. Furthermore, a disaster tolerant backup solution that provides cross-site protection is a critical component of an effective backup solution. To address these challenges, HP Data Protector software provides a backup option at the storage level:

- Backup at the storage level (HP Data Protector software value add with Zero Downtime Backup and Instant Recovery)


Backup agent in the virtual machine

This approach simply is the traditional approach to backups. The backup agents are installed inside each virtual machine exactly as it would be done on a physical server. Backups are performed over the LAN. This method does not require any additional skill set or procedural changes. It allows file-level recovery and the restore process is unchanged. This method is suitable for virtual machine OS and application backups. Through the use of application backup agents, this method can provide application-consistent backups.

Challenges:

1. With multiple backup agents running in each virtual machine, the VMware physical server can quickly be overtaxed for CPU and network bandwidths.
2. This method does not take advantage of the virtual infrastructure encapsulation of virtual machines.
3. Long RTO: Backup windows can be longer as backups must be staggered to avoid overlapping back jobs.

Figure 2: Backup agent in Virtual Machine



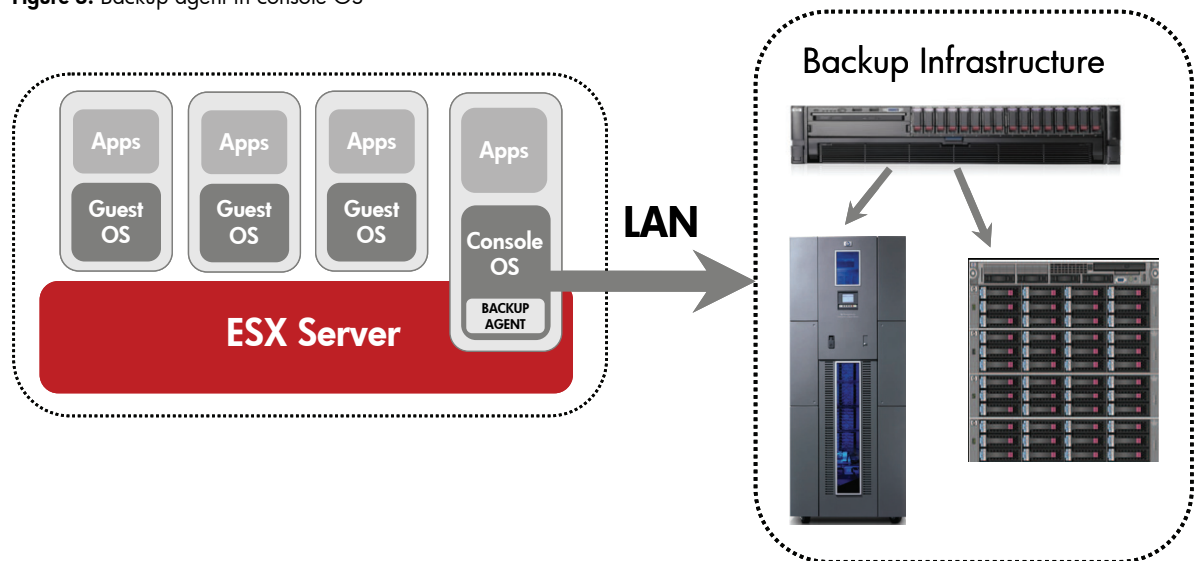
Backup agent in the console OS/hypervisor

This method involves installing a backup agent within the console OS and then backing up the set of VMDK files and configuration files that encapsulate each virtual machine. A single backup agent is needed and installed in the console OS and backups are performed through the LAN. This method provides fast image level recovery of virtual machines.

Challenges:

1. In order to ensure application consistency during the backup, scripting must be employed to shutdown and start up the virtual machine.
2. Virtual machine snapshots can eliminate downtime; however, they will not provide an application data consistency as the backups will be crash consistent.
3. This method does not provide file level backups or incremental backups.
4. The availability of a service console is a must.
5. Long RTO: No file level restores available making restore windows larger than sometimes necessary.

Figure 3: Backup agent in console OS




Backup through proxy server

By integrating a third-party backup application with VMware Consolidated Backup (VCB), the backup through proxy allows LAN free backups of virtual machines at the file and image level, and requires deployment of the VMware VCB framework. This method allows taking the backup load off the ESX server by having a third party host perform the backup. The third-party host has access to the same SAN volumes as the ESX server. The VCB framework allows the ESX server to flush the file system before creating snapshots of VMDKs to be backed up. A vLUN driver within the framework allows the VMDKs to be presented to the third party backup host for access. Finally, scripting utilities are provided to assist with automation.

Challenges:

1. Simplicity depends on integration with third-party backup application.
2. File level recovery requires that a backup agent be installed in the guest
3. Without VSS, windows image level backup are crash consistent
4. VCB (VMware Consolidated Backup) does not provide support for enterprise applications like Exchange, Oracle, SAP, and SQL.

Figure 4: Backup via proxy server



Backup option summary chart

CHART 1: Backup option chart

• 100% Protection	• Agent in VM	• Agent in ESX console	• VCB	• HP DP ZDB/IR + VCB
• Virtual Machine	<ul style="list-style-type: none"> • Yes • Severe performance penalty 	<ul style="list-style-type: none"> • Yes 	<ul style="list-style-type: none"> • Yes • Windows only 	<ul style="list-style-type: none"> • Yes • Zero performance impact • Windows only
• Databases & Applications	<ul style="list-style-type: none"> • Yes • Severe performance penalty 	<ul style="list-style-type: none"> • No 	<ul style="list-style-type: none"> • No 	<ul style="list-style-type: none"> • Yes • Zero performance impact

(1) See Appendix A for more details.

Seventy four percent of surveyed IT professionals to a Data Protection indicated that the recovery SLAs for mission critical applications in their environment was within four hours time.

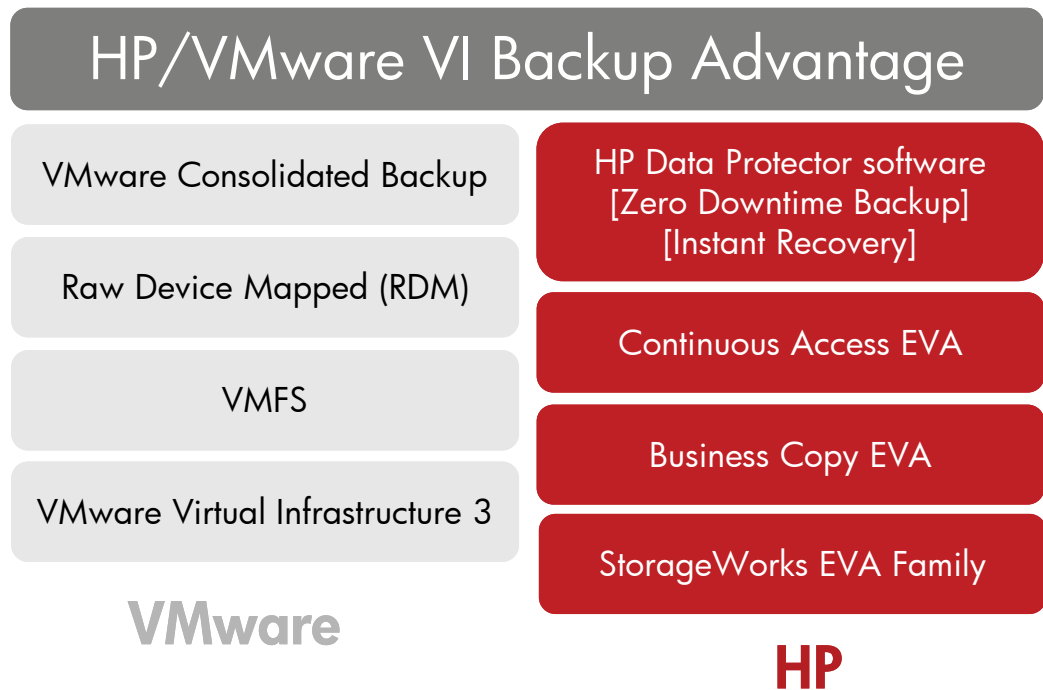
Source: Excillio Group Inc, Data Protection Report – January 2008.

To address these difficult challenges, this paper examines a solution based on key HP Data Protector software and HP StorageWorks Enterprise Virtual Array (EVA) technologies, which considerably reduce the complexity of deploying a backup strategy in a VMware environment, reduce recovery times while providing continuous backup protection, and a disaster-tolerant backup approach:

- **Reduce RPO and RTO:** Zero Downtime Backup and Instance Recovery shrink RPO and RTO.
- **Reduce complexity of backup strategy:** Backups managed and performed through simple interface.
- **Provide disaster tolerance:** HP Continuous Access EVA protects against complete site failure.

The solution technology components explored in the paper are as follows:

Figure 5: Solution technology components



Solution technology components

VMware technology components

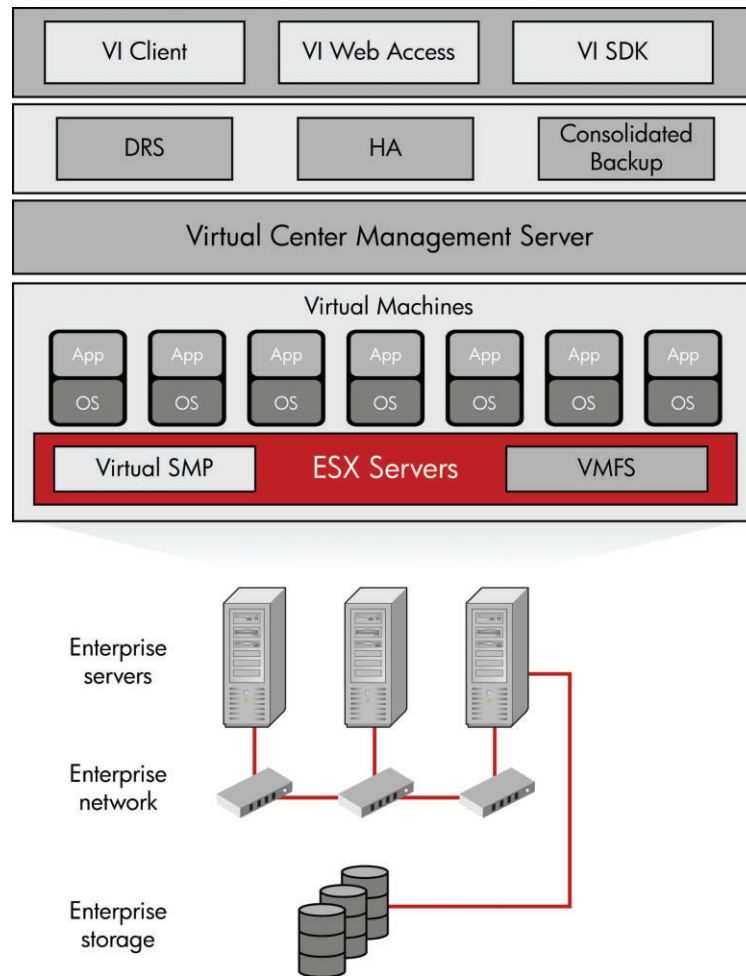
Virtual Infrastructure 3

VI 3 allows enterprises and small businesses alike to transform, manage, and optimize their IT environments through virtualization.

Virtualization creates an abstraction layer that decouples the physical hardware from the operating system to deliver greater IT resource utilization and flexibility. Multiple virtual machines, running a range of operating systems (such as Microsoft® Windows® Server 2003, or Linux) and applications run in isolation and also side-by-side on the same physical server.

VI 3 delivers comprehensive virtualization, management, resource optimization, application availability, and operational automation capabilities in an integrated offering. Figure 6 provides a logical view of the components of a VI 3 implementation.


Figure 6: VI 3 implementation



VMFS

VMFS is a high-performance clustered file system that leverages shared storage to allow multiple instances of VMware ESX to read and write to the same storage, concurrently. A Virtual Machine File System serves as a repository for virtual machines and virtual machines state. Each virtual machine's files are encapsulated in its own directory. VMFS also house other files such as templates, ISO images for fast deployment, or virtual machines. VMFS volumes can be accessed through the service console through the mount point `/vmfs/volumes`.


Figure 7: VMFS



Raw Device Mapping (RDM)

A Raw Device Mapping (RDM) is a special file in a VMFS volume that acts as a proxy for a raw device. RDM LUN supports direct access³ between the virtual machine and physical storage subsystem, and is useful with SAN snapshots and other layered applications running in a virtual machine. There are two types of RDM devices—physical compatibility RDM, also called pass-through RDM, and virtual compatibility RDM, also known as non-pass-through RDM. Pass-through RDM volumes allow SCSI commands to pass through directly from the guest operating system to the SAN. Thus, they enable scalable backups better using native SAN features. Non-pass-through RDM volumes preserve some features of VMware disk such as snapshots, however, a larger set of SCSI commands are filtered when using non-pass-through RDM.

Figure 8: Raw Device Mapping (RDM)



³ Fibre Channel or iSCSI only


VMware Consolidated Backup (VCB)

VMware Consolidated Backup (VCB) takes the backup off the ESX server host, eliminates the backup window, removes backup traffic from the LAN, and eliminates the need to run backup agents inside the virtual machines to perform file-level backups of virtual machine data.

VCB leverages new capabilities in VMware tools to quiesce the file system inside the virtual machine at the time a snapshot is taken. This process ensures that a consistent snapshot is taken as all pending changes are flushed to disk before the snapshot is taken.

With the use of an agent running in another physical machine, this physical machine is able to mount the snapshot as if it was a disk physically attached to this machine. The backup agents running in this physical machine can then backup the content of the mounted snapshot.

Figure 9: VMware Consolidated Backup



HP technology components

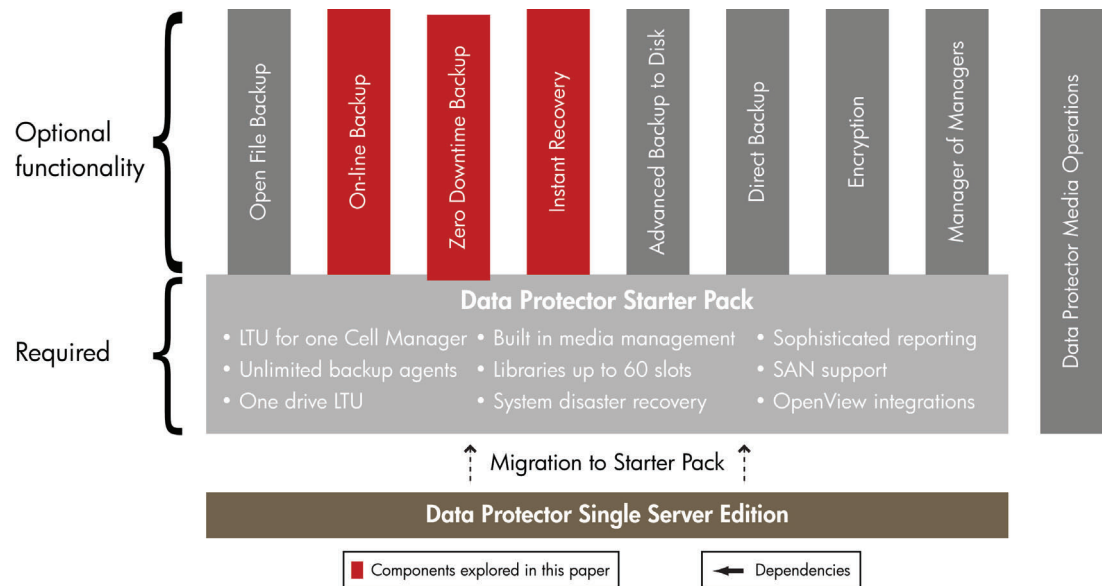
HP Data Protector software

HP Data Protector software automates high performance backup and recovery, from disk or tape, over unlimited distances, to enable 24x7 business continuity, and improve IT resource utilization. As an integral component of the fast-growing HP Software portfolio, which includes storage resource management, archiving, replication, and device management software, HP Data Protector software also fully integrates with the HP OpenView management solutions.

HP Data Protector software simplifies the use of complex backup and recovery procedures with the fastest installation, automated routine tasks, and easy-to-use features. HP Data Protector software is the ideal solution to reduce IT costs and complexity while remaining reliable and scalable to grow from single server environments to the largest distributed enterprise infrastructures.

HP Data Protector software has the following structure:

Figure 10: HP Data Protector Software Structure



Zero Downtime Backup

Zero Downtime Backup (ZDB) provides no impact backup by performing backups on a copy of the production data. ZDB provides the option to copy or move data to tape. As of Data Protector 6.0 and newer, ZDB also allows incremental backups. Zero Downtime Backup is suitable for application-aware data backups and non-disruptive protection. ZDB integrates with array-based replication and provides backup of the replica of the production data locally or at a remote site. ZDB is easily configured through a step-by-step user interface and provides administrators complete control on the protection automation, the replica specifications, and the backup schedule.

Instant Recovery

Instant Recovery allows HP Data Protector software to recover data directly from the replicas on disk instantaneously. With Instant Recovery, ZDB can now keep track of multiple rotating replicas on disk. Instant Recovery is suitable for application/environments that require a very fast RTO (minutes instead of hours). Instant recovery provides administrators with great management and automation flexibility.

HP Data Protector software acquisition and deployment costs are 30-70 percent less than competition

HP StorageWorks EVA Family

EVA is HP mid-range storage system. The EVA supports various operating systems, boot from SAN, asymmetric active-active controllers, 2-Gb and 4-Gb Fibre Channel host connections, multi-pathing, SCSI, and FATA drives. Maximum LUN size as of this writing is 2 TB and the maximum number of LUN per storage system is 1,024. Designed for the data center where there is a critical need for improved storage utilization, and scalability, the EVA meets application-specific demands for consistent high transaction I/O for the customer, and provides easy capacity expansion, instantaneous local replication, and simplified storage administration. HP also brings integrated iSCSI connectivity to the EVA and Remote Replication, SAN over WAN, with EVA Continuous Access EVA.

HP StorageWorks Business Copy EVA (EVA BC)

Business copy EVA is a feature of the EVA that allows the creation of point-in-time copies of logical units local to the array. These copies can be rapidly created and deployed for various business critical reasons such as:

- Back up of data with very little to no impact on applications
- Running tests on applications against real data
- Quick restore of a logical units data in case of data loss
- Data mining for marketing or business improvement purposes

In order to meet different usage requirements, BC EVA comes in three flavors:


- Snapshots
- Snapclone
- Mirrorclones

HP StorageWorks Continuous Access EVA (CA EVA)

Continuous Access (CA) EVA is another feature that allows data replication between two or more EVAs. Data replication can be done synchronously or asynchronously. CA EVA supports various interconnection technologies such as FCIP and Fibre Channel. Additionally, the EVA also supports bi-directional replication. Data replication between sites is most widely used when creating a true disaster-tolerant data center.

A copy set is a replicated vdisk and a Data Replication (DR) group is a data replication group that is comprised of replicated vdisks (Copy Sets). Each DR group acts as a consistency group and all copy sets within that group share a single write history log. Thus, a DR group is the primary level of CA management. All CA management actions such as Write Mode, Failsafe Mode, Suspend Mode, and Failover are performed on a DR group and not on each copy set. Replication Solution Manager is the preferred tool to manage BC and CA on EVA.

Figure 11: EVA Continuous Access



The number of replication group created should be tailored to the specific user environment. For instance, the creation of replication groups can be based on a number of decision factors, for example:

- The disk resources for each application running in a virtual machine may require their own DR group. This allows failing over just one or multiple applications individually to a different virtual machine without having to failover the virtual machine they are running on.
- A virtual machine or group of virtual machines may need to have all of their disk resources in a single DR group. Virtual machines booting off of VMDK files on the same VMFS volume would need to all be failed over together when the LUN where the VMFS resides is failed over.

The EVA provides a limited number of DR groups so understanding your environment and its replication granularity requirements will help you reduce the number of DR groups required for your environment and provide improved efficiency. Consult with your HP field representative for data replication groups' implementation strategies.

Implementation overview

As discussed above, traditional backup approaches present many challenges to the deployment of an effective backup strategy in a virtual environment. Surveyed data collected by VMware indicated that 80-85 percent of VMware virtual machines are created on VMFS volumes. This means that any effective backup solution has to tightly integrate with VMware Consolidated Backup framework in order to leverage the many benefits it provides to backup the VM operating system, OS, and application configuration, which are encapsulated and saved on VMFS volumes. This paper demonstrates how to deploy a complete and effective backup strategy that leverages the benefits of VMware Consolidated Backup and complements it with proven HP backup technologies, Zero Downtime Backup and Instant Recovery, to meet data center SLAs for RPO and RTO, while providing a robust disaster tolerant backup solution. The backup and restore processes are described below:

Backup process

1. Virtual machine level backup (OS, OS and application configuration) provided through Data Protector Integration with VMware Consolidated Backup
2. Application data-level backup provided through Data Protector using the Data Protector online agents
3. Data Protector Zero Downtime Backup provides continuous backup of the application data of the application running inside a VM, while providing local and remote disaster tolerance by leveraging EVA Business Copy and Continuous Access respectively.

Restore process

1. The full virtual machine image is restored through VMware Consolidated Backup proxy host
2. For local replication, if the data is on tape then it can be restored directly to the VM. If the backup data still exist as a replica (EVA snapclone or XP mirror), then it can be restored through Instant Recovery
3. For remote replication, whether the data is on disk or tape, recovery is done directly into the virtual machine

Figure 12a: ZDB/IR





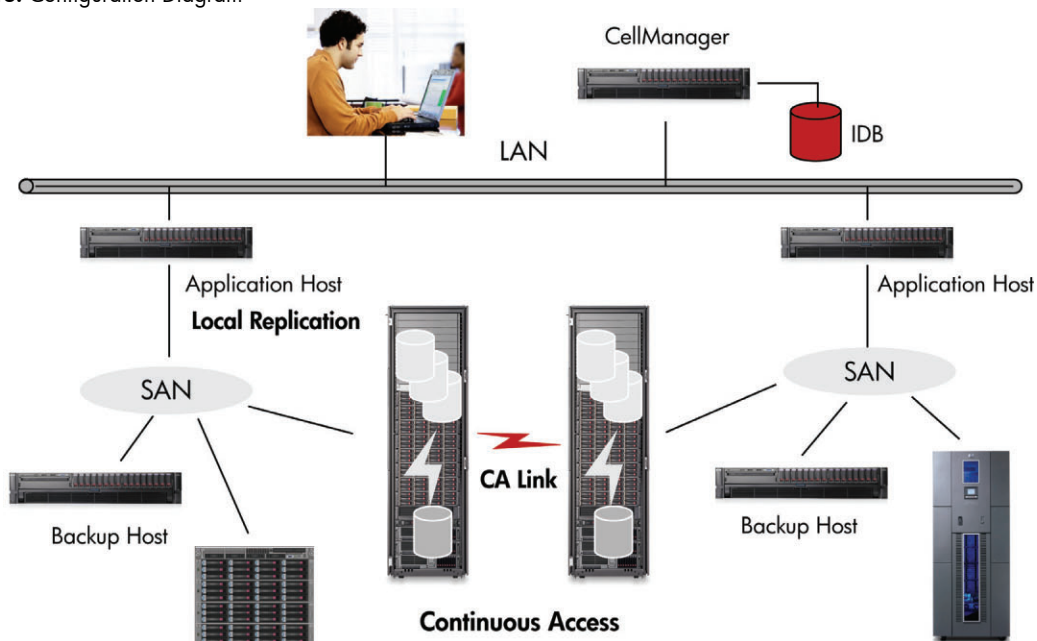
Figure 12b: VMware Consolidated Backup



Implementation configuration

The environment used for this case study consisted of two EVAs linked together through a CA link. Each EVA was accessed by an ESX server and a backup server. Each backup server was also connected to a tape device or virtual library. Additionally, a data protector cell manager was configured in the environment along with two Command View EVA management stations that are accessed for array management and configuration.

Figure 13: Configuration Diagram



Storage and server hardware configuration

Table 1 and 2 respectively detail the storage and servers' hardware configuration:

Table 1: EVA8100 configuration (Production system Local site)

1 set (2 pairs) of EVA8100 controllers
2 disk shelves
4GB cache per controller
28 * 300 GB 10K rpm disk drives
EVA firmware revision: XCS 6110

Table 2: EVA8100 configuration (Mirror site system Remote site)

1 set (2 pairs) of EVA8100 controllers
2 disk shelves
4GB cache per controller
28 * 300 GB 10K rpm disk drives
EVA firmware revision: XCS 6110

Table 3: EVA4400 configuration (Local site virtual file library)

1 set (2 pairs) of EVA4400 controllers
1 disk shelf
4GB cache per controller
12 * 146 GB 10K rpm disk drives
EVA firmware revision: 09003000

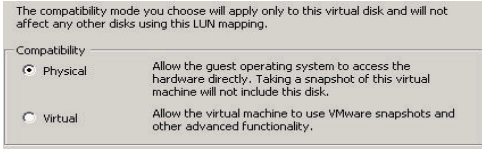
Table 4 ESX Server configuration

ESX Server 3.5 build 64607	ESX Configuration parameters
DL365 G1 (4CPU x 2.2GHz) dual-core AMD	Disk.UseLunReset = 1
12 GB RAM	Disk.UseDeviceReset = 0
1 Dual Channel 4 GB QLogic HBA	Disk.MaxLUN = 256
2 Gig-E Broadcom network cards	LVM.EnableResignature = 0
	LVM.DisallowSnapshotLUN = 1

Table 5: Virtual machine configuration (VMware Consolidated Backup Test)

Windows Server 2003 R2	
VM boot disk: 30GB vmdk	
1 LSI Logic virtual SCSI adapter	
2 Virtual Ethernet adapters	
2 GB RAM	
2 virtual CPU	

Table 6: Virtual machine configuration (Oracle ZDB/IR Backup Test)

Windows Server 2003 SP1	
VM boot disk: 30GB vmdk	
RDM 1: Oracle control files (c:\mnt\)\ – 2GB	RDM configuration All RDM LUNs must be set to physical compatibility mode. 
RDM 2: Oracle Data Files (E:\) – 25GB	
RDM 3: Oracle Redo Logs (D:\) – 60GB	
RDM 4: Oracle Archive Logs (E:\) – 120GB	
2 LSI Logic Virtual SCSI adapters. (1 for boot LUN vmdk, 1 for RDM LUN)	
2 virtual CPUs	
2 virtual NICs	
2 GB RAM	
Database configuration	
Single instance Oracle Database 10g	
NTFS file system for control, redo and archive logs	
Oracle database SID: orcl	
Oracle base: c:\oracle	
Oracle home directory: c:\oracle\product\10.2.0\db_1	
Oracle user: Domain Administrator	

Implementation details

As discussed previously in this paper, implementing a backup strategy for a virtual infrastructure is a very challenging task. Figure 1 showed that depending on the data center need, an administrator would need to backup the virtual machine, the application and/or its data, and the ESX server itself. Without advanced technologies such as data protector ZDB/IR and VMware consolidated backup, the administrator would not only have to perform backups for each of these manually but the virtual infrastructure RTO and RPO would be much too high. For instance, to backup a running database the administrator would have to leave this database in backup mode for long periods of time. Furthermore, to restore application data, the application would have to be down for the entire length of the restore.

In its simplest implementation, HP Data Protector software ZDB/IR is fully functional with various file systems and in more complex implementations; ZDB/IR is fully integrated with many enterprise class applications such as SAP, SQL, and Oracle. The case study presented in this paper will focus on backups of a virtual machine using VMware Consolidated Backup framework while the enterprise application in this case Oracle Database 10g is backed up and restored using HP Data Protector software Zero Downtime Backup and Instant Recovery.

This paper assumes that its readers are familiar with installing, configuring, and using, VMware VI3 components and HP Data Protector software. Else, please refer to HP and VMware documentation that highlights in ample details installation and configuration steps.

VMware Consolidated Backup Integration in HP Data Protector 6.0 software using pre/post execution scripts in a backup specification

Out of the various virtual infrastructure backup options previously explored VMware Consolidated Backup is the solution of choice when backing up virtual machines because it provides the flexibility of VM-level incremental or full backups, removes the load of performing backups from the ESX server and essentially eliminates backup windows.

VMware consolidated Backup consists of two use cases:

1. Perform Backup
 - Full image backup
 - File Level backup (Windows only)

2. Perform Restore

Though this paper will not cover configuration steps in details ensure the following configuration steps are properly performed. (Detailed instructions are available in the HP OpenView Storage Data Protector Installation and Licensing Guide):

1. Install HP Data Protector software
 - Install Cell Manager
 - Install Backup proxy Server
 - o Ensure disk agents are installed on backup proxy host
 - o If backup proxy host is connected directly to a backup media then the media agents must also be installed on the backup proxy host
2. Install VMware Consolidated Framework on backup proxy host
3. VMware Consolidated Backup HP Data Protector software integration scripts
 - vmwarepreexec.cmd
 - vmwarepostexec.cmd
 - vcbmount.js

NOTE:

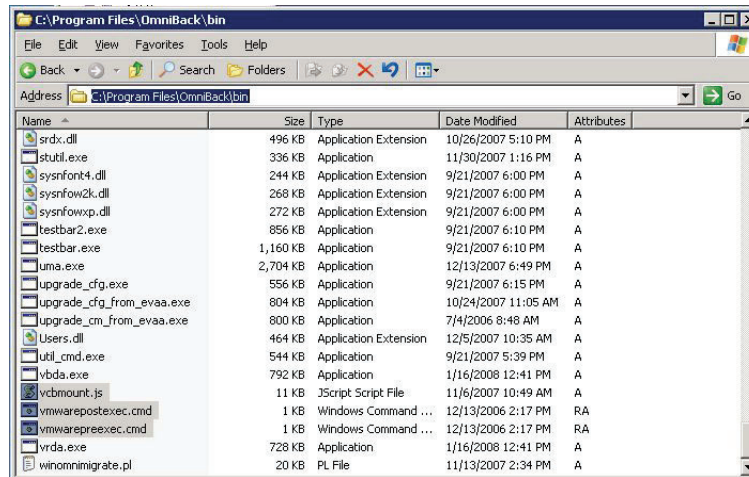
You can download the Data Protector 5.5/6.0 - VMware Consolidated Backup (VCB) integration packet from :

<http://h20000.www2.hp.com/bizsupport/TechSupport/SoftwareDescription.jsp?lang=en&cc=us&prodTypeId=18964&prodSeriesId=3241176&prodNameId=3241177&swEnvOID=54&swLang=13&mode=2&taskId=135&swItem=co-47153-5>

This package contains the scripts and the document describing the integration of VMware Consolidated Backup 3.0.1 snapshot capabilities with HP Data Protector software 5.5 and 6.0.

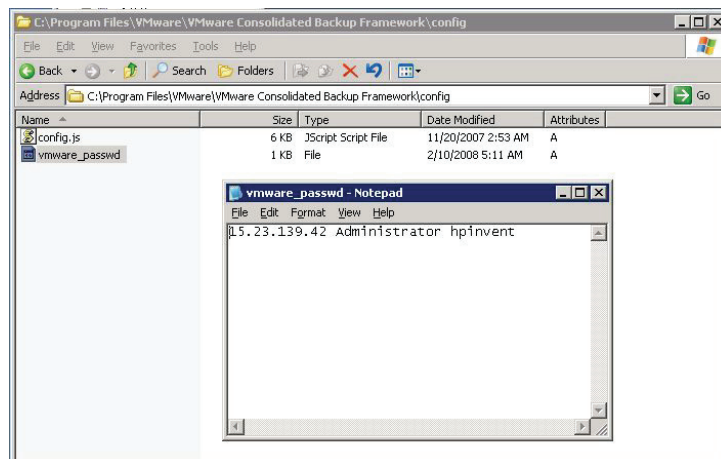
The scripts above must be copied to the bin directory of the HP Data Protector software installation of the backup proxy host as shown in figure 14:

Figure 14: Integration scripts



4. Create the "vmware_passwd" file in the config directory of the VMware Consolidated backup framework installation directory on the backup proxy host. This file should contain the IP or server name of the virtual center server, its user name and password in the format shown below:

Figure 15: vmware-passwd file



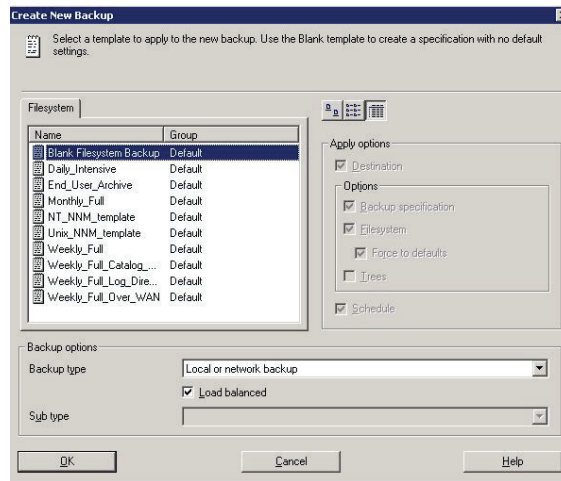
Performing a Backup

(For an example for full image backup refer to HP Data Protector software VCB integration whitepaper for more details and file level backup description)

Before creating a new backup specification for virtual machines using Data Protector and VCB, it is critical to ensure that the backup host mount point for the virtual machine snapshots will have sufficient space to house all virtual machines being backed up.

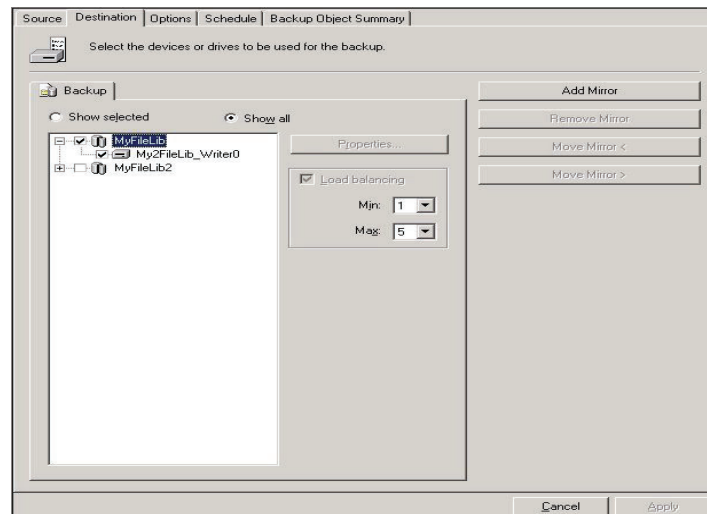
To begin a new backup specification, run the Data Protector software manager and in the backup context window create a file system backup specification by selecting the “Blank Filesystem Backup.”

Figure 16: Blank Filesystem Backup



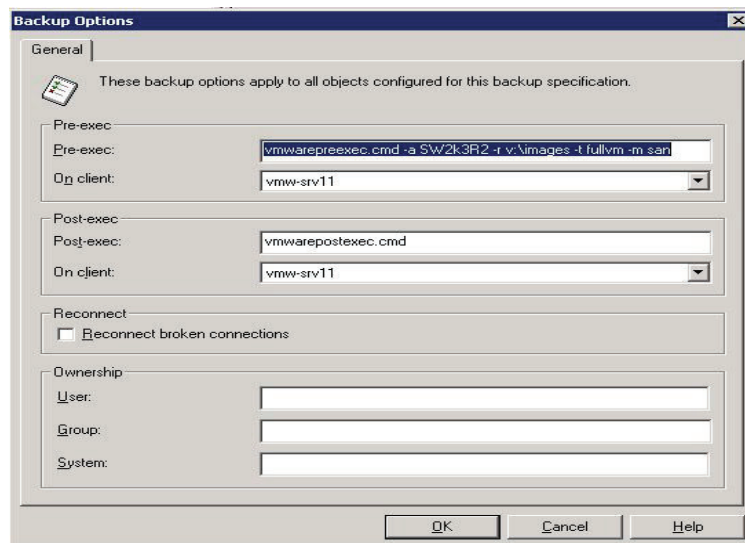
- Next select the directory and mount point where the full image backup snapshot will be mounted for backup.
- Next select the backup device where the backups will be saved. In this experiment, a file library was created from LUNs on an EVA4400.

Figure 17: File Library Backup repository



- Next enter a description for the backup specification. This name should be specific to the type of backup performed and also indicative of what is being backed up.
- Click the advanced button under the backup specification options to enter the pre-exec and post-exec options for VCB backup.

Figure 18: VCB pre/post exec options

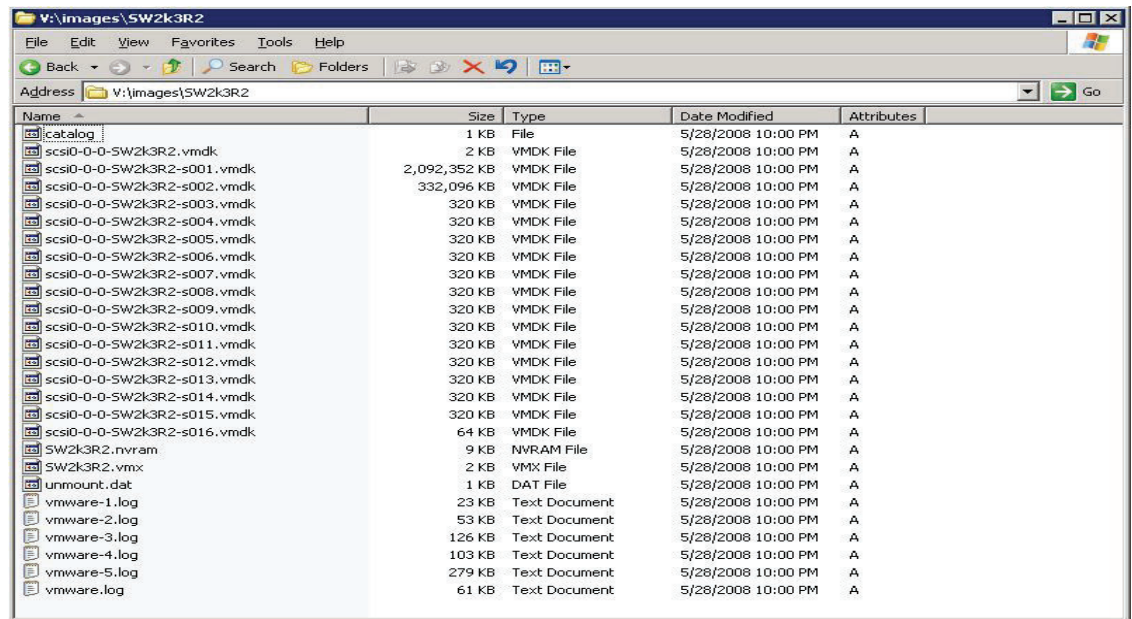


The pre-exec script option above will backup the virtual machine named "SW2K3R2". The mount point for the backup files will be "v:\images" and a full backup will be performed.

- Apply these changes and under the schedule tab, if desired configure a backup schedule.
- Finally, run and test this backup specification. During the backup, the Data Protector summary pane will display the backup progress as shown below and also report any warnings, failures, and completion events.

During the backup operation the backup mount point “v:\images” is populated with all the files that encapsulate the virtual machine in question.

Figure 19: VCB Full Backup mount point



This process provides a crash consistent backup copy of the virtual machine, the OS running inside of it, and all virtual machine specific configuration files. Since most virtual machines are installed on virtual disks, VCB provides better backup granularity, and much more space efficiency of backups than SAN backups of an entire VMFS volume that may house many more VMs.

VMware Consolidated Backup Integration in HP Data Protector 6.1 software using the VMware online agent integration

Even here VMware Consolidated Backup is the solution of choice when backing up virtual machines because it provides the flexibility of VM-file level incremental/differential backups (VCB file), full backups (VCB file) and virtual machine image (VCB image) backups. VCB removes the load of performing backups from the ESX server and essentially eliminates backup windows.

Data Protector integrates with VMware virtual infrastructure through the Data Protector VMware integration agent, which channels communication between the Data Protector Session Manager and the clients in the VMware environment. The Data Protector VMware integration agent communicates with the virtual infrastructure through VI SDK, a web-service API.

VCB consists of two use cases:

1. Perform Backup
 - Full image backup
 - File Level backup (Windows only)
2. Perform Restore

Data Protector supports environments where ESX Server systems are managed through a VirtualCenter Server system (VirtualCenter environments) as well as environments with standalone ESX Server systems (standalone ESX Server environments). Mixed environments, in which some of the ESX Server systems are managed through a VirtualCenter Server system and some are standalone, are also supported. You can even have multiple VirtualCenter Server systems in your environment, each managing its own set of ESX Server systems.

Though this paper will not cover configuration steps in details, please make sure the following configuration steps are properly performed (Detailed instructions are available in the HP Data Protector A.06.10 Integration guide for VMware Virtual Infrastructure, Sybase, Network Node Manager, and Network Data Management Protocol Server):

Data Protector Cell Manager

The Data Protector Cell Manager can be installed on a virtual machine, VirtualCenter Server system, backup proxy system, or a separate system outside the VMware Virtual Infrastructure environment.

Data Protector VMware Integration component

The Data Protector VMware Integration component must be installed on the following clients:

- All ESX Server systems from which you plan to back up virtual machines
- VirtualCenter Server systems (if they exist)
- Backup proxy systems (if you plan to use the VCBfile and VCBimage backup methods)
- Windows systems (physical or virtual) to which you plan to restore file systems of virtual machines

The component consists of the following parts:

- **vmware_bar.exe** is activated during backup and restore
- **util_vmware.exe** is activated during configuration and mounting on backup proxy systems

Data Protector Media Agents

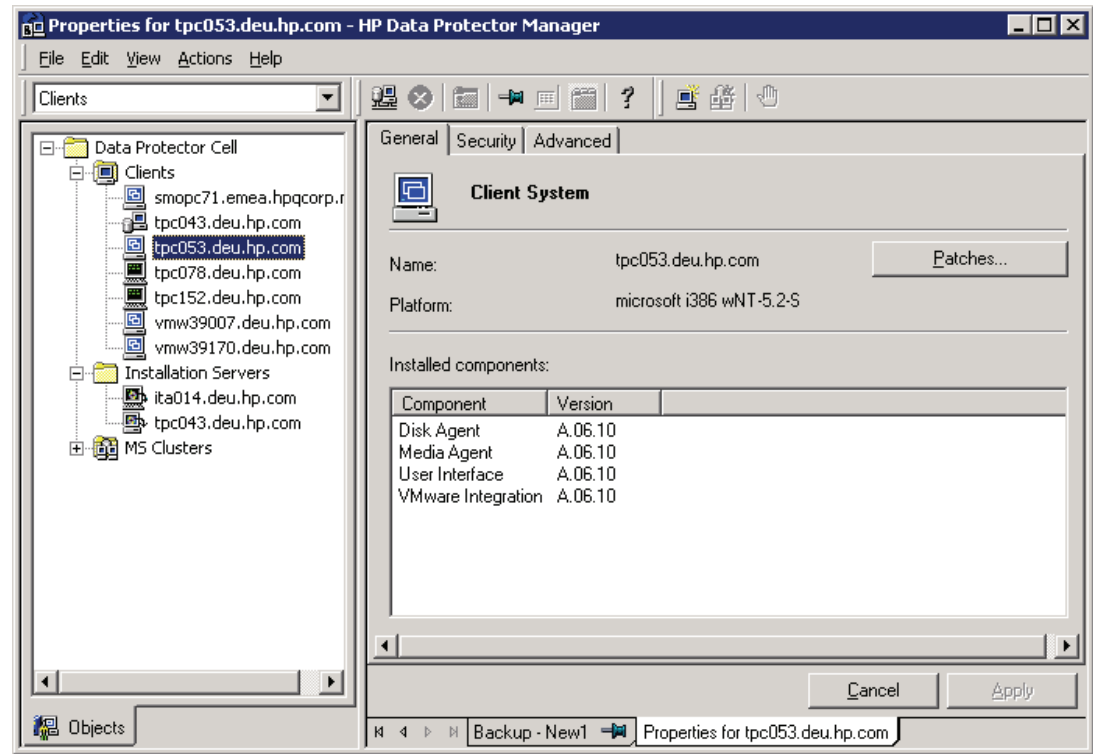
Data Protector Media Agents can be installed on ESX Server systems, VirtualCenter Server systems, backup proxy systems, or separate systems outside the VMware Virtual Infrastructure environment.

NOTE:

You can install Data Protector VMware clients and Data Protector Media Agents remotely, by distributing the software using the Data Protector 6.1 Installation Server.

Figure 20 depicts the VMware integration component deployed on a VMware VirtualCenter server.

Figure 20: VMware Integration Component on VirtualCenter Server



Configuring the VCB Integration

Configure the integration as follows:

NOTE:

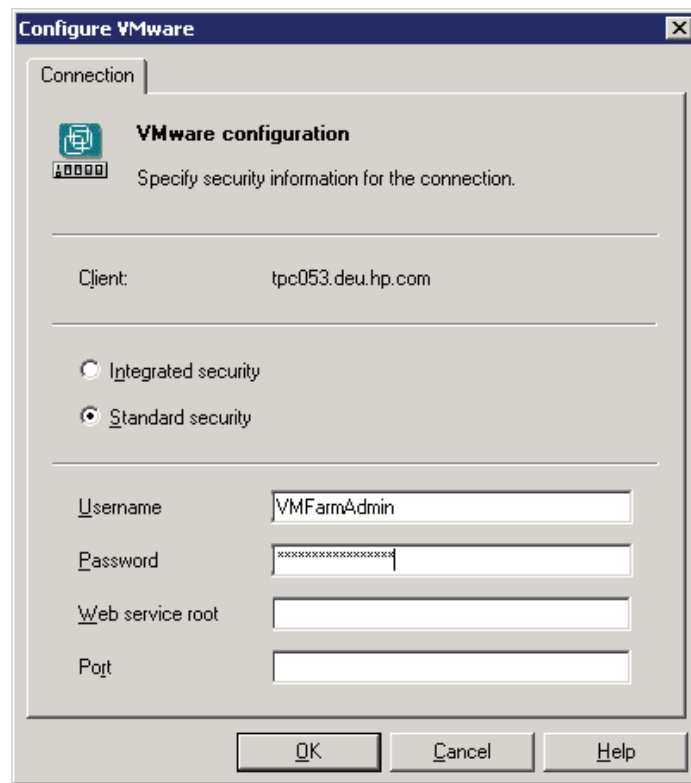
Detailed VCB integration configuration details can be found in the guide: "HP Data Protector A.06.10 Integration guide for VMware Virtual Infrastructure, Sybase, Network Node Manager, and Network Data Management Protocol Server"

The VMware integration component requires certain users to be setup properly.

VirtualCenter users (VirtualCenter environment)	For each VirtualCenter Server system, identify the Windows operating system user who administrates the VirtualCenter Server.
ESX Server users (standalone ESX Server environment)	For each standalone ESX Server system, identify an operating system user who has read, write, and execute permissions on the related datastores.

- Configure users as described in HP Data Protector A.06.10 Integration guide for VMware Virtual Infrastructure, Sybase, Network Node Manager, and Network Data Management Protocol Server: "Configuring VMware users".
- If your ESX Server systems are configured in a cluster, check the cluster settings as described in: "Configuring clusters".
- If your virtual machines reside on iSCSI datastores, restart the Data Protector Inet service on the related backup proxy system under a network domain user account that has read-write permissions for the directories described in: "Configuring backup proxy systems".
- Provide Data Protector with login information to VMware management clients as described in "Configuring VMware management clients".

Figure 21: Configuring a VirtualCenter Server system



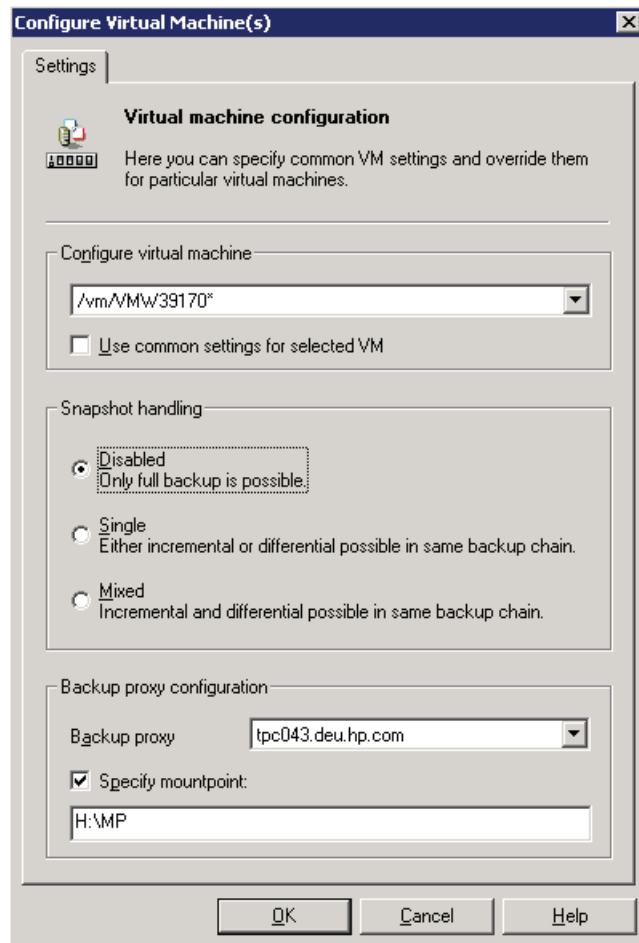
- For each virtual machine, specify details on how to perform various backup methods as described in "Configuring virtual machines".

For the VCBfile and VCBimage backup methods specify, which backup proxy system and mount points should be used to back up virtual machines or filesystems.

You can configure each virtual machine separately or all together. Configuration settings for virtual machines of the same data center are saved in a separate configuration file on the Cell Manager. The file is named `VMwareManagementClient%DatacenterPath`. It is used for all backup sessions involving this particular datacenter.

To configure virtual machines, use the Data Protector GUI or CLI.

Figure 22: Configuring virtual machines and their backup proxy settings



VCBimage backup method

For the VCBimage backup method, you need to have at least one backup proxy system configured in your environment. A backup proxy system is a Windows system that has the VCB software installed. For details on the VCB software, see the VMware documentation. During a VCBimage backup, Data Protector invokes VCB to mount virtual machines on a backup proxy host. Before a virtual machine is mounted, VCB creates a virtual machine snapshot to put the virtual machine into a consistent state. Once mounted, the virtual machine is copied (exported) to the backup proxy local disk.

NOTE:

Depending on virtual machine disk sizes, the copy operation can be very time-consuming. It may take longer than the default Data Protector Session Manager timeout, which is 10 minutes. If the timeout is reached, the session is automatically aborted. To solve the problem, extend the timeout by resetting the Data Protector *SmWaitForFirstBackupClient* global options variable. For details, on how to set the variable, see the Data Protector software online help index: "global options".

After the virtual machine copy is created, it is transferred to Data Protector media. At the end, the virtual machine is unmounted. Consequently, the virtual machine copy is removed from the backup proxy. The virtual machine snapshot is removed as well. The VCBimage backup method backs up only the current state of a virtual machine. Information about the snapshot tree and the changes made on non-active snapshot branches are not included in the backup. As a consequence only full backups are supported.

VCBfile backup method

During a VCBfile backup, Data Protector invokes VCB to mount NTFS filesystems of Windows virtual machines on a backup proxy host. Before a filesystem is mounted, VCB creates a virtual machine snapshot to put the files into a consistent state. Once the filesystem is mounted, the files are transferred directly to Data Protector media while the backup proxy is only referencing them. At the end, the filesystem is unmounted and the virtual machine snapshot is removed. The VCBfile backup method enables you to back up NTFS filesystems of virtual machines running Windows. Filesystems of other guest operating systems cannot be backed up. VCBfile backup supports either full or incremental or differential backups.

NOTE:

Disk space: Virtual machine operations that are performed during backup require additional disk space on the data stores. Data Protector checks for each virtual machine or filesystem separately whether the required virtual machine operation can be safely performed (whether enough disk space is available). If not, the backup of that particular virtual machine or filesystem is skipped.

For the VCBimage and VCBfile backup methods, disk space is needed also on the backup proxy system for mounting virtual machines and filesystems. VCB checks whether enough disks space is available and informs Data Protector of it. If not, Data Protector skips the backup of that particular virtual machine or filesystem.

Concurrent sessions: Backup sessions that use the same devices or back up the same data center cannot run concurrently. If multiple sessions are started, one session waits for the other to complete.

The remaining configuration steps are standard configuration settings such as:

- Define a target backup device
- Define backup options such as protection time
- Define a backup schedule

Performing a Zero Downtime Backup

Despite providing an effective way to back the virtual machine and its content, VCB is not application aware, thus it does not provide consistent backups of application data. Application level backup can be accomplished using HP Data Protector software Zero Downtime Backup and Instant Recovery. This paper will demonstrate how backup of an Oracle Database is performed using Zero Downtime Backup. Zero Downtime Backup, allows two types of backups: Local and remote. Local copies can be instantaneously played back using Instant Recovery and remote copies protect against disasters.

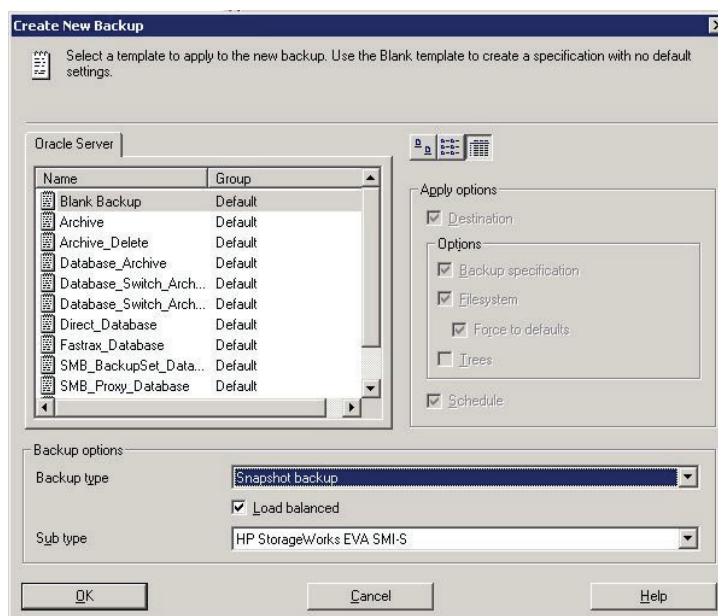
Before being able to perform a Zero Downtime Backup and Instant Recovery, a series of configuration settings are required. Though these configuration steps will not be discussed in detailed in this paper, an overview of the required steps is shown below:

- Configure supported Command View EVA management stations
- Configure SMI-S on the command view EVA management stations
- Obtain and install, BC licenses for each EVA in the configuration
- Obtain and install, CA EVA licenses for each EVA in the configuration
- Ensure SAN is properly configured to enable EVA Continuous Access
- Install and configure the virtual machine
 - The virtual machine configuration used for the Oracle Database is shown in table 4 above.
- Install and configure the Oracle Database
 - For DB/IR, it is required to configure the database such way were the data files and control files are created on separate LUNs on the EVA array.
- Install the following data protector components in the virtual machine:
 - Disk agent
 - User interface
 - HP StorageWorks EVA SMI-S agent
 - Oracle Integration package

Upon successfully completing these steps, a new backup specification for Zero Downtime Backup can be created. From the HP Data Protector software manager GUI perform the following steps to create a new Oracle backup specification:

- Select empty template
- Select as backup type: Snapshot backup with the sub-type HP StorageWorks EVA SMI-S

Figure 23: Backup type using EVA SMI-S Agent



- Specify the application (virtual machine running Oracle database) and backup server.
- Select a replica option
 - BC for local replication (This option allow Instant Recovery)
 - CA + BC for remote replication (recovery is done through LAN)
- Select “Track the replica for instant recovery” and specify a number of replicas to rotate through.
 NOTE: Selecting this option will automatically grey out the snapshot type selection dropdown. This is because replicas tracked for Instant Recovery require snapclones (for EVA) or mirrors (for XP).
 Tracking replicas when creating a Zero Downtime Backup specification is a critical component for being able to utilize Instant Recovery
- Next, select the application database

Figure 24: Selection of the application database to be backed up

The screenshot shows a window titled "Specify the application that you want to back up." with a folder icon. It contains two main sections: "Application" and "User and group".

Application Section:

- Client:** A dropdown menu showing "15.23.138.96".
- Application database:** A dropdown menu showing "orcl".

User and group Section:

- Username:** An empty text input field.
- Group name:** An empty text input field.

- Select control files and database
- Select the destination library device
- Create a backup schedule if desired
- Finally save the backup specification

Conclusion/Summary

By combining HP Data Protector software Zero Downtime Backup and Instant Recovery with VMware Consolidated Backup virtual infrastructure IT managers now have an effective backup solution that protects their VM and application data while significantly reducing RPO and RTO requirements:

- **Reduce RPO and RTO**
 - Application data can be restored in matter of minutes
 - Virtual machine backups provide more granularity than regular SAN backups.
- **Reduce complexity of Backup strategy**
 - Backups managed and performed through simple interface.
 - Once configured all backup can be carried out on an automated schedule.
 - With ZDB/IR an administrator has the flexibility of leaving backup on disk (as replica) for much faster restore and recovery granularity from a point in time perspective.
- **Provide disaster tolerance**
 - ZDB/IR through continuous access provide across site replication ensuring that backup are available for restore when disaster strikes.

Appendix A

Backup Option	Pros	Cons
Backup agent in virtual machine	<ul style="list-style-type: none">• Traditional backup/restore method• No additional skill set required• File level backup available for all OS• Provides ability for application consistent backups	<ul style="list-style-type: none">• Puts significant load on system resources (CPU, network bandwidth)• Does not take advantage of virtualization encapsulation• Backups must be staggered and LAN backups take long, both yielding longer RTO
Backup agent in the ESX console	<ul style="list-style-type: none">• Fast image level VM backup by leveraging virtualization encapsulation• Only a single backup agent is needed	<ul style="list-style-type: none">• Requires scripting to enable application backup consistency• File-level or incremental backups are not possible• Restore window are large because full images restores are the only option• Puts load on the console OS
Backup via proxy server	<ul style="list-style-type: none">• File and image level backups are both possible• Backup load is offloaded to a third-party server	<ul style="list-style-type: none">• Integration with third-party backup software may be complex• File-level recovery still require a backup agent in the virtual machine• No integration with third party enterprise application
HP/VMware backup advantage [Solution described in this paper]	<ul style="list-style-type: none">• File and image level backups are both possible• Backup load is offloaded to a third party server• HP Data Protector Software is fully integrated with enterprise applications like (Oracle, SAP)• Quick and easy integration with VMware VCB• Instant recovery capability provides the ability for very small RTO.	<ul style="list-style-type: none">• Virtual machine OS file level recovery may require a backup agent to be installed.

For more information

- HP Data Protector software 6.0 / 6.1 User Guide
- HP OpenView Storage Data Protector Zero Downtime Backup Concepts Guide
- HP OpenView Storage Data Protector Zero Downtime Backup Administrator's Guide
- HP OpenView Storage Data Protector Zero Downtime Backup Integration Guide
- Instant Recovery for HP StorageWorks Enterprise Virtual Array in CA+BC Configuration
[\[http://www.hp.com/go/dataprotector\]](http://www.hp.com/go/dataprotector)
- Data Protector 5.5/6.0 - VMware Consolidated Backup (VCB) integration packet
[\[http://h20000.www2.hp.com/bizsupport/TechSupport/SoftwareDescription.jsp?lang=en&cc=us&prodTypeId=18964&prodSeriesId=3241176&prodNameId=3241177&swEnvOID=54&swLang=13&mode=2&taskId=135&swItem=co-47153-5\]](http://h20000.www2.hp.com/bizsupport/TechSupport/SoftwareDescription.jsp?lang=en&cc=us&prodTypeId=18964&prodSeriesId=3241176&prodNameId=3241177&swEnvOID=54&swLang=13&mode=2&taskId=135&swItem=co-47153-5)
- HP StorageWorks EBS Solutions guide for VMware Consolidated Backup with HP Data Protector
[\[http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01116446/c01116446.pdf?jumpid=reg_R1002_USEN\]](http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01116446/c01116446.pdf?jumpid=reg_R1002_USEN)
- Using VMware ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery
Using VMware Infrastructure for Backup and Restore
[\[http://www.vmware.com/\]](http://www.vmware.com/)
- SEPATON Addresses Backup of Virtualized Servers
[The Impact of Server Virtualization on Storage, December 2007 -
[http://www.enterprisestrategygroup.com/\]](http://www.enterprisestrategygroup.com/)

