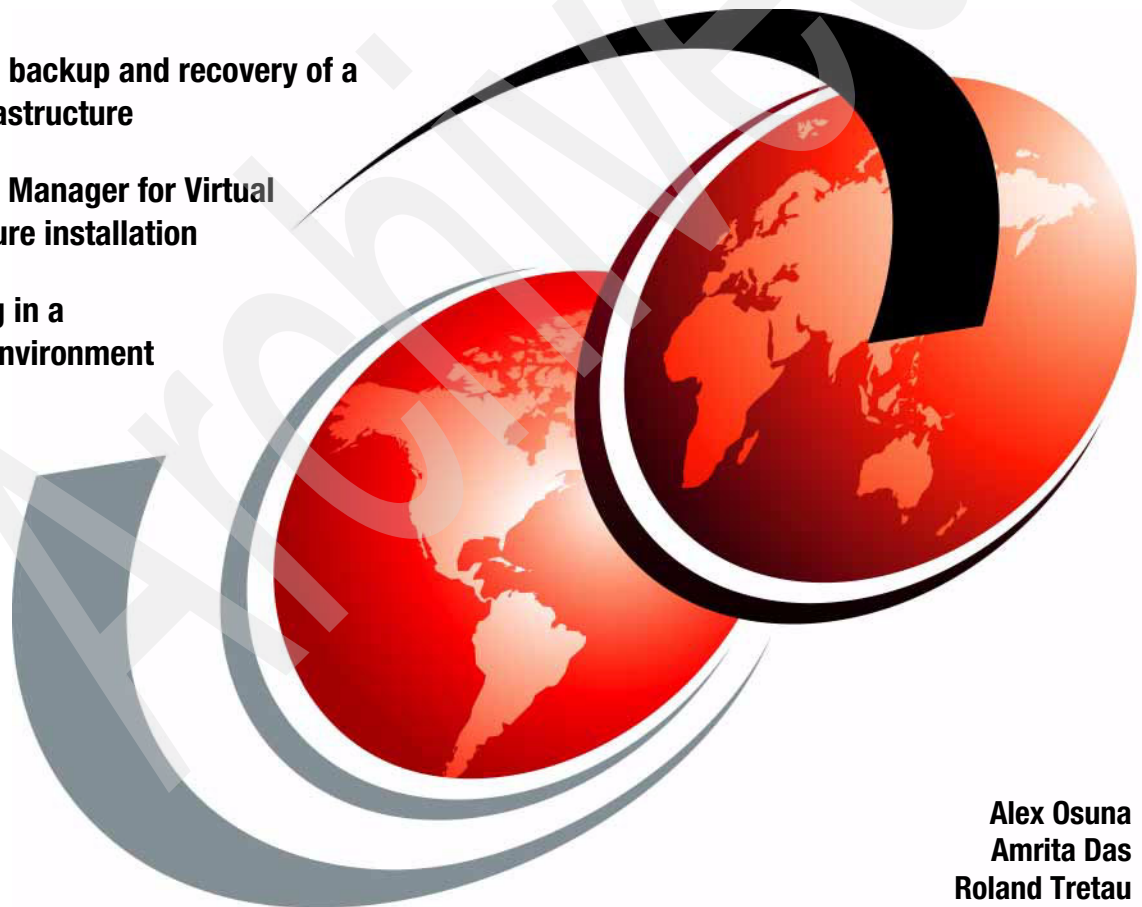


SnapManager 2.0 for Virtual Infrastructure Best Practices

Simplifying backup and recovery of a
Virtual Infrastructure

Using Snap Manager for Virtual
Infrastructure installation

Configuring in a
clustered environment



Alex Osuna
Amrita Das
Roland Tretau



International Technical Support Organization

**SnapManager 2.0 for Virtual Infrastructure
Best Practices**

September 2010

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page vii.

First Edition (September 2010)

This edition applies to IBM System Storage N series SnapManager for VMware® Virtual Infrastructure Version 2.

© Copyright International Business Machines Corporation 2010. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	vii
Trademarks	viii
Preface	ix
The team who wrote this book	ix
Now you can become a published author, too!	x
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Chapter 1. Introduction to SnapManager for Virtual Infrastructure	1
1.1 Overview	2
1.2 Purpose and scope	2
1.3 Installation and configuration of SMVI	3
1.4 User friendly interface	3
1.5 Availability of granular restore options	3
1.6 Single wizard for creating backup jobs	5
1.7 Trigger scripts during backup process	6
1.8 Other features	6
Chapter 2. Concepts	7
2.1 VMware ESX Server architecture	8
2.2 VMware ESX storage options	9
2.2.1 Storage overview: VMFS datastores	9
2.2.2 Storage overview: NAS datastores	10
2.2.3 Storage overview: Raw device mappings	11
2.3 SnapManager backup and recovery	12
2.3.1 Backup retention policy	13
2.3.2 Alert notification	13
2.3.3 Port usage	13
2.4 SnapManager architecture	14
2.5 Components	14
2.5.1 Data ONTAP	14
2.5.2 VMware Virtual Infrastructure	15
2.5.3 SnapManager for Virtual Infrastructure	16
2.5.4 SnapManager for Virtual Infrastructure repository	16
Chapter 3. Planning	17
3.1 Storage configuration	18
3.2 Adding storage systems	19

3.3	SnapManager for Virtual Infrastructure data layout.	20
3.4	vCenter configuration	21
3.5	vCenter user permissions	22
3.6	Distributed resource scheduler implications	23
Chapter 4.	Installation	25
4.1	Installing vCenter and SMVI on a virtual machine.	26
4.1.1	Installing vCenter within a virtual machine	26
4.1.2	Installing SMVI within a virtual machine	26
4.2	Default installation.	27
4.3	Choosing between the GUI and the CLI	28
4.4	Configuring SMVI for shared storage	28
4.5	Configuring SMVI in a clustered environment (MSCS)	29
4.5.1	SMVI configuration	29
4.5.2	Configuring cluster resources for SMVI	30
4.5.3	Confirming SMVI high availability	32
4.6	SMVI in a N series MetroCluster environment	33
Chapter 5.	Process flow	35
5.1	Backup process and implications	36
5.1.1	The backup process	36
5.1.2	Backup process implications.	38
5.2	Scheduled backups and retention policies	39
5.2.1	Backup scheduling	39
5.2.2	Retention policies	39
5.3	Snapshot naming	40
5.4	Scripting	40
5.4.1	Environment variables.	42
5.4.2	Error messages.	43
5.5	Include independent disks and exclude datastores.	45
5.6	Mounting a backup	45
5.7	Single file restore.	46
5.7.1	Pre-requisites	46
5.7.2	Types of file restore sessions	46
5.7.3	Single file restore capability for Linux VMS.	47
5.8	Restore process flow.	48
5.9	Restore enhancements in SMVI 2.0	50
5.9.1	Virtual disk restore	50
5.9.2	Advanced find (more search options)	50
Chapter 6.	Disaster recovery	51
6.1	SnapMirror integration.	52
6.1.1	SnapMirror destinations	52
6.1.2	SnapMirror and deduplication	53

6.2 Configuring the disaster recovery standby site	53
6.2.1 Primary site SMVI server	54
6.2.2 Primary site storage	54
6.2.3 Disaster recovery site	54
Chapter 7. VMware snapshots	55
7.1 VMware snapshots and SMVI	56
7.2 Serialized snapshots	56
7.3 Installation of VMware Tools and VM alignment	57
7.4 Reducing the number of concurrent VMware snapshots	57
7.5 Reducing the amount of disk I/O	57
7.6 Eliminating VMware vCenter snapshots from the backup process	58
Chapter 8. Data consistency in an SMVI environment	59
8.1 Backup	60
8.1.1 Point-in-time consistent backup	60
8.1.2 Coordinated backup process	61
8.2 Recovery	62
8.2.1 Point-in-time consistent recovery	62
8.2.2 Application-consistent recovery	63
8.2.3 Application consistency in combined solution environments	63
8.2.4 Supported configurations in combined solution environments	65
8.2.5 Simplifying supported configurations	66
Chapter 9. Summary of best practices	69
Appendix A. Recovery procedures	73
Loss of SMVI server	74
Recovering from loss of SMVI server when installed on local disk	74
Recovering from loss of SMVI server when installed on a shared device	74
Loss of a VMware ESX host	75
Within a VMware ESX cluster	75
Stand-alone ESX host	76
Primary site loss recovery	76
Appendix B. Sample scripts	79
Sample environment variable	80
Displaying environment variables during backup phases	80
Sample SMVI SnapVault script	80
Appendix C. Troubleshooting single file restore	83
Appendix D. SMVI installation steps	87
Abbreviations and acronyms	93

Related publications 95

IBM Redbooks publications 95

Other publications 95

Online resources 96

How to get Redbooks publications 96

Help from IBM 96

Index 97

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:


This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

IBM®
Redbooks®
Redpaper™

Redbooks (logo) ®
System Storage®
System x®

Tivoli®

The following terms are trademarks of other companies:

Snapshot, RAID-DP, Network Appliance, WAFL, SyncMirror, SnapVault, SnapRestore, SnapMirror, SnapManager, SnapDrive, FlexVol, FlexClone, Data ONTAP, NetApp, and the NetApp logo are trademarks or registered trademarks of NetApp, Inc. in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Preface

This IBM® Redbooks® publication provides best practices for SnapManager® for Virtual Infrastructure 2.0. We address the resource utilization issue typically found within virtual environments by leveraging the underlying Snapshot™ technology, which enables you to create point-in-time copies of your virtual machines or entire data stores and then restore from these backup copies at any level of granularity, datastore, VM, disk (VMDK), or guest file, simply and quickly when required. In addition, we provide best practices for protecting the SMVI server and recovering in case of a disaster.

The reader will gain a deep understanding of how to implement SnapManager for Virtual Infrastructure in VMware vSphere environments. Furthermore, we explain the seamless integration of N series storage solution, including MetroCluster, so customers can use storage and virtualization technologies to create dynamic infrastructures that can create tremendous business value.

The team who wrote this book

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Alex Osuna is a Project Leader at the International Technical Support Organization, Tucson Center. He writes extensively on all areas of data storage. Before joining the ITSO five years ago, Alex worked in the IBM Tivoli® Western Region as a Principal Systems Engineer supporting Tivoli Storage products. Alex has over 32 years in the I/T industry and with the majority of them focused on data storage. Alex holds certifications from IBM, Red Hat, Microsoft®, and the Open Group.

Amrita Das is a Technical Marketing Engineer at Network Appliance™ Inc. Their engineering perspective has made her a valuable contributor toward the writing of this book.

Roland Tretau is an Information Systems professional with over 15 years experience in the IT industry. He holds Engineering and Business master's degrees, and is the author of many storage related Redbooks publications. Roland has a solid background in project management, consulting, operating systems, storage solutions, enterprise search technologies, and data management.

Thanks to the following people for their contributions to this project:

Kyle Burell
Sandro De Santis
Rucel F. Javier
IBM

Emma Jacobs
International Technical Support Organization, San Jose Center

Keith Aasen
Josh Bonczkowski
John Ferry
Lisa Haut-Mikkelsen
Antony Jayaraj
Yateendra Kulkarni
Gabriel Lowe
Niels Reker
Leo Yaroslavsky
NetApp®

Jürgen Mutzberg
VMware

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks publications form found at:

ibm.com/redbooks

- ▶ Send your comments in an email to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks publications weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Introduction to SnapManager for Virtual Infrastructure

With the adoption of virtualization technologies, data centers have been transformed and the number of physical servers drastically reduced. Virtualization has had many positive effects, not only in reducing the number of physical systems, but also in reducing network, power, and administrative overhead.

In contrast to physical environments, where server resources are under utilized, fewer resources are available in virtualized environments. Where each physical server had dedicated network and CPU resources, virtual machines (VMs) must now share those same resources. This can result in performance issues, especially while backing up the virtual environment, because many virtual machines impact host network and CPU resources concurrently. As a result, after completion during non-business hours, backups have seen their backup window grow.

In this chapter we introduce you to IBM System Storage® N series SnapManager for VMware Virtual Infrastructure (SMV) and provide details about its features.

1.1 Overview

SnapManager for Virtual Infrastructure addresses the resource utilization issue typically found within virtual environments by making use of the underlying Snapshot technology, which enables you to create point-in-time copies of your virtual machines or entire data stores and then restore from these backup copies at any level of granularity, datastore, VM, disk (VMDK), or guest file, simply and quickly when required (Figure 1-1). This is all done on our storage systems, freeing your servers to run applications, not backups.

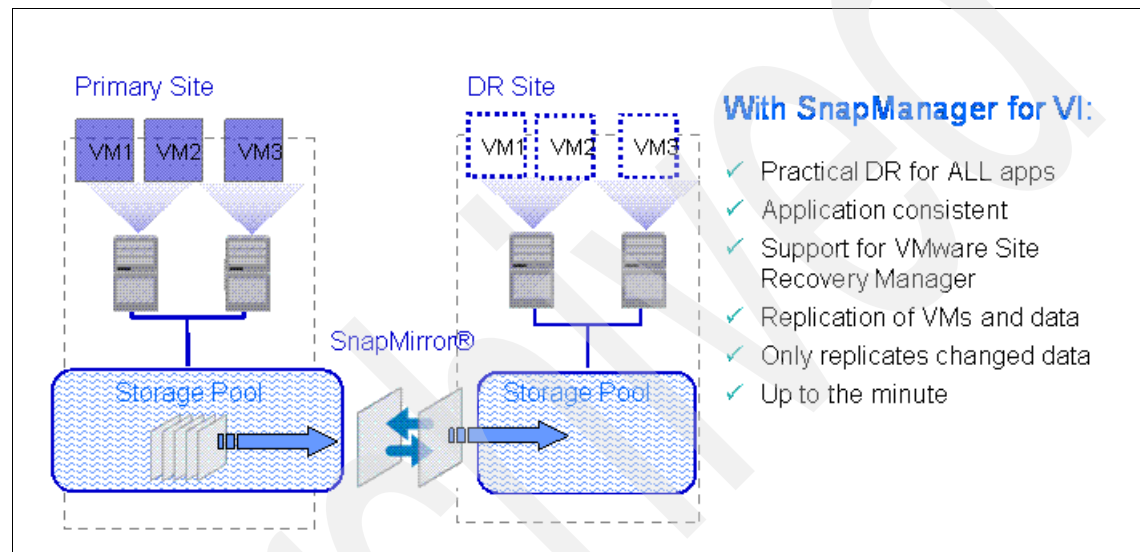


Figure 1-1 Datastore backups

1.2 Purpose and scope

The purpose of this book is to provide best practices for deploying SMVI to back up and recover VMware ESX virtual machines and their associated datastores. In addition, this book provides best practices for protecting the SMVI server and recovering in case of a disaster. Intended audience.

Readers can benefit from reviewing the following publications:

- ▶ *IBM System Storage N series Data ONTAP 7.3 System Administration Guide*, GC52-1279-04
- ▶ *IBM System Storage N series SnapManager 2.0 for Virtual Infrastructure Installation and Administration Guide*, GC53-1145-02

1.3 Installation and configuration of SMVI

Snap Manager for Virtual Infrastructure (SMVI) can be quickly installed and configured for use in a new or existing VMware vSphere environment, saving valuable time during backups and allowing quick and efficient restorations, thus reducing administrative overhead.

1.4 User friendly interface

The user interface (UI) has been greatly improved to provide a more user friendly experience, as shown in Figure 1-2.

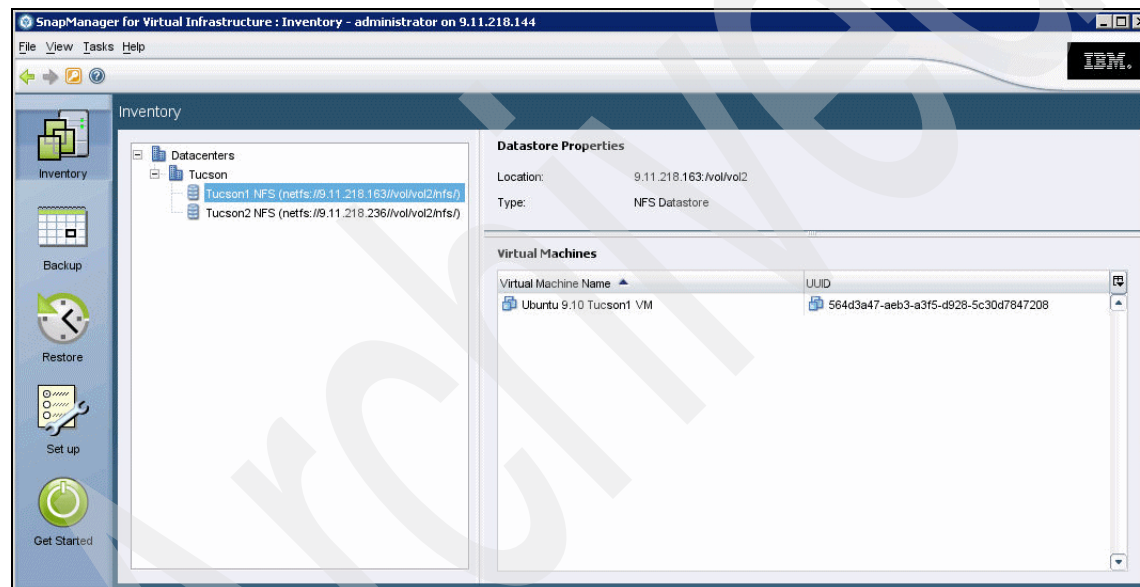


Figure 1-2 SMVI user interface

1.5 Availability of granular restore options

Granular restore options are available as displayed in Figure 1-3:

- ▶ Single File Restore (SFR)
- ▶ VMDK restore

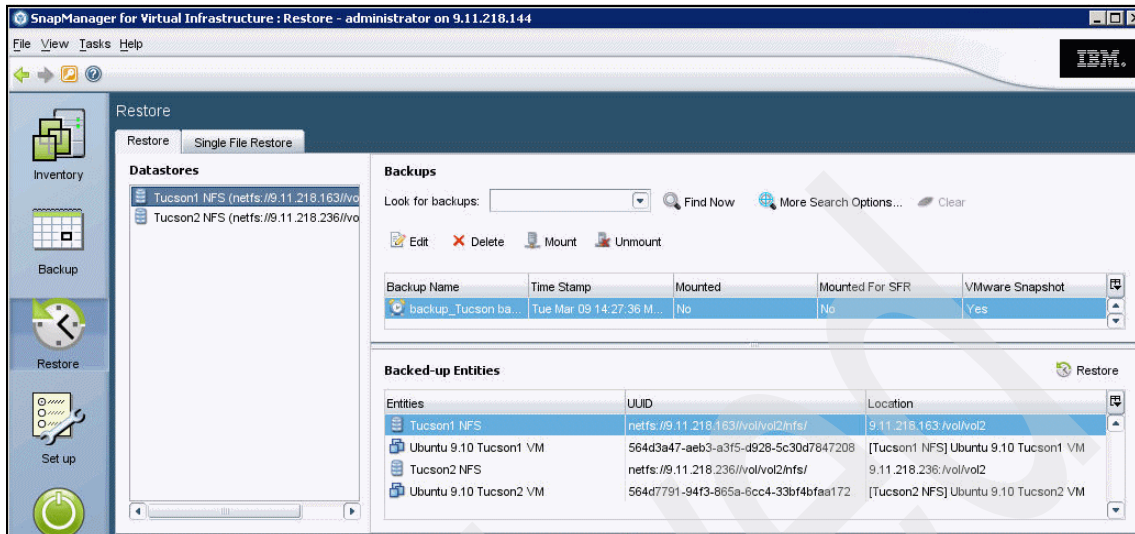


Figure 1-3 Restore options

Figure 1-4 shows the possibility to select whether you want to restore the entire disk or the virtual disk. Furthermore, you can select the original or a new location.




Figure 1-4 SMVI 2.0 enhanced restore options

1.6 Single wizard for creating backup jobs

The wizard enables you to create manual and scheduled backup jobs. Simply select **Backup** on the right pane, name your new backup job, and select the per-backup job options as displayed in Figure 1-5:

- ▶ Exclude datastores from backup
- ▶ Include datastores with independent disks in backups
- ▶ Perform VMware vSphere consistent backups.

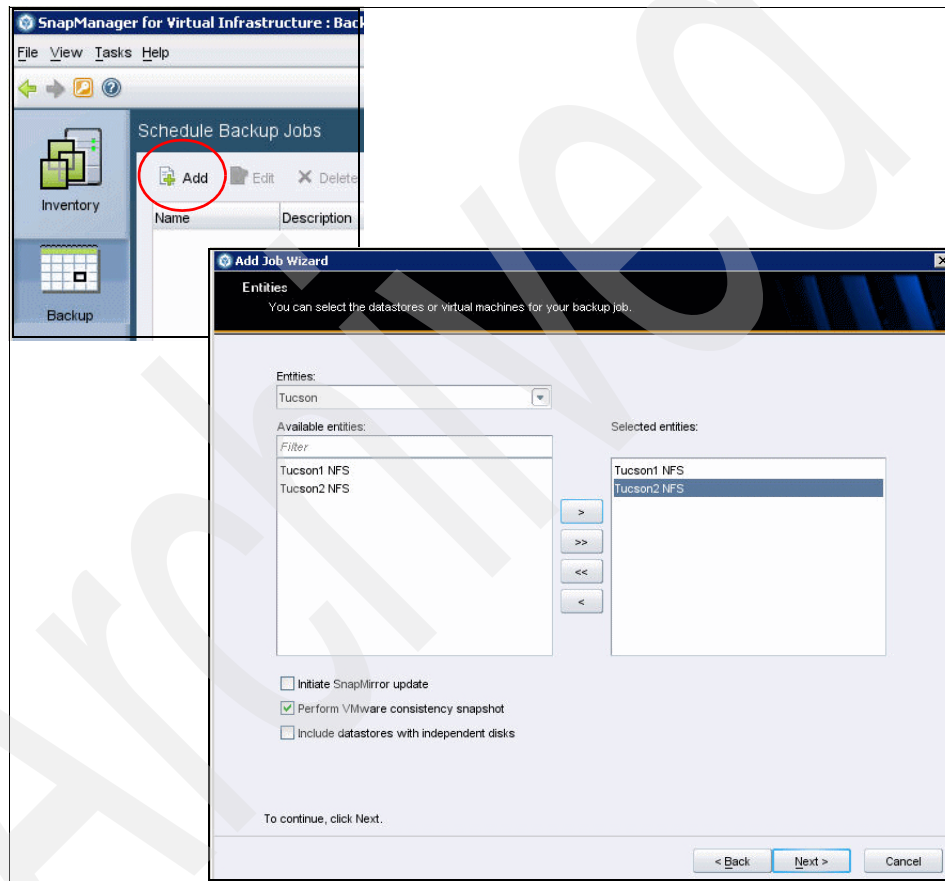


Figure 1-5 Backup wizard

1.7 Trigger scripts during backup process

Figure 1-6 shows that administrators are able to run designated pre, post, and failure scripts during the backup process. Selecting the scripts is part of the backup wizard.

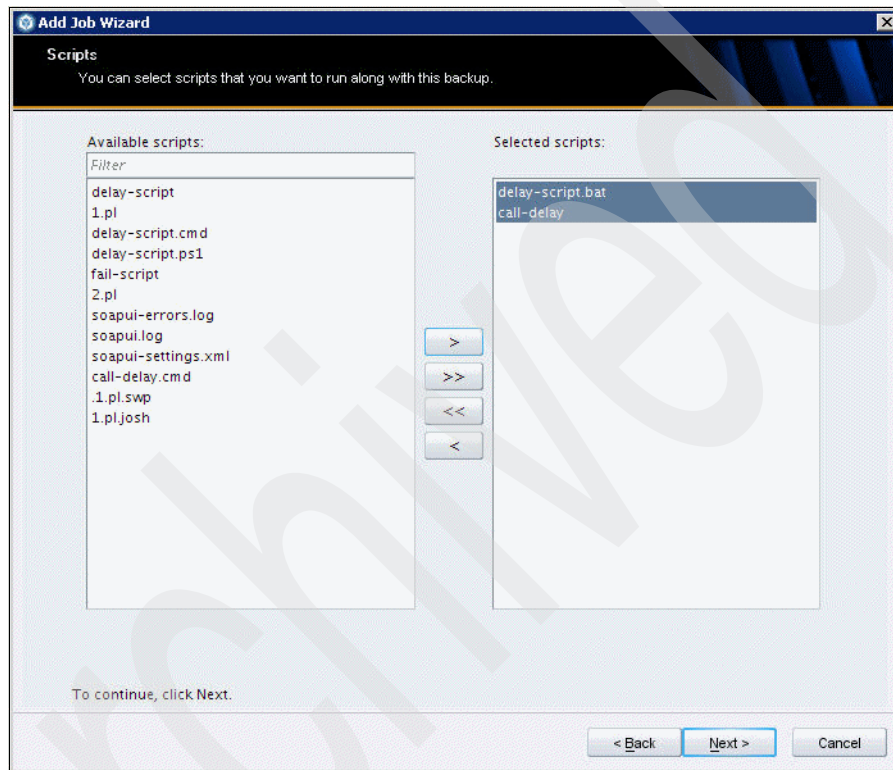


Figure 1-6 Run individual scripts

1.8 Other features

Here we list additional SMVI 2.0 features:

- ▶ Consistent backup naming
- ▶ Serialization of VMware vSphere snapshots
- ▶ ASUP (AutoSupport) logging
- ▶ vFiler unit support for multiple IP addresses
- ▶ Advanced Find option to find specific backups

Concepts

In this chapter we provide an introduction to the architecture of VMware vSphere and the VMware ESX storage options. We also describe SnapManager for Virtual Infrastructure, including backup and recovery.

2.1 VMware ESX Server architecture

VMware ESX Server is Virtual Infrastructure partitioning software that is designed for server consolidation, rapid deployment of new servers, increased availability, and simplified management. It helps to improve hardware utilization, save space, IT staffing, and hardware costs.

Many people might have had earlier experience with VMware's virtualization products in the form of VMware Workstation or VMware GSX Server. VMware ESX Server is quite unique compared to other VMware products in that it runs directly on the hardware, offering a mainframe-class virtualization software platform that enables the deployment of multiple, secure, independent virtual machines on a single physical server.

VMware ESX Server allows several instances of operating systems, such as Windows® Server 2003, Windows Server 2008, Red Hat and (Novell) SuSE Linux®, and more, to run in partitions that are independent of one another; therefore, this technology is a key software enabler for server consolidation that provides the ability to move existing, unmodified applications and operating system environments from a large number of older systems onto a smaller number of new high-performance IBM System x® platforms.

The architecture of VMware ESX Server is shown in Figure 2-1.




Figure 2-1 VMware ESX Server architecture

Instead of deploying multiple servers that are scattered around a company and running a single application on each server, they can be consolidated together physically, as they simultaneously enhance system availability. VMware ESX Server allows each server to run multiple operating systems and applications in virtual machines, providing centralized IT management. Because these virtual machines are completely isolated from one another, if one were to go down, it does not affect the others, which means that not only is VMware ESX Server software great for optimizing hardware usage, it can also give the added benefits of higher availability and scalability.

2.2 VMware ESX storage options

VMware ESX supports three types of storage configurations when connecting to shared storage arrays: VMFS datastores, NAS datastores, and raw device mappings. We assume that customers understand that shared storage is required to enable high-value VMware vSphere features such as HA, DRS, vMotion, and Fault Tolerance. The goal of the following sections is to provide customers with information to consider when designing their virtual data centers.

VMware virtualization technology makes it easy for customers to use all of these storage designs at any time or simultaneously. The following section reviews these storage options and summarizes the unique characteristics of each architecture. For information regarding deploying with VMFS, NFS, and RDMs, see the *VMware ESX and ESXi Server Configuration Guide*, at the website:

http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxi_i_vc_setup_guide.pdf

2.2.1 Storage overview: VMFS datastores

The VMware Virtual Machine File System (VMFS) is a high-performance clustered file system that provides datastores, which are shared storage pools. VMFS datastores can be configured with LUNs accessed by Fibre Channel, iSCSI, or Fibre Channel over Ethernet. VMFS allows traditional LUNs to be accessed simultaneously by every ESX Server in a cluster.

VMFS provides the VMware administrator with a fair amount of independence from the storage administrator. By deploying shared datastores, the VMware ESX administrator is free to provision storage to virtual machines as needed. In this design, most data management operations are performed exclusively through the VMware vCenter Server.

For applications that traditionally require storage considerations, in order to make sure their performance can be virtualized and served by VMFS, it is best to deploy the virtual disks on a datastore that is connected to all nodes in a cluster but is only accessed by a single VM.

This storage design can be challenging in the area of performance monitoring and scaling. Because shared datastores serve the aggregated I/O demands of multiple VMs (Figure 2-2), this architecture does not natively allow a storage array to identify the I/O load generated by an individual VM. This issue can be exacerbated by spanning VMFS volumes across multiple LUNs.

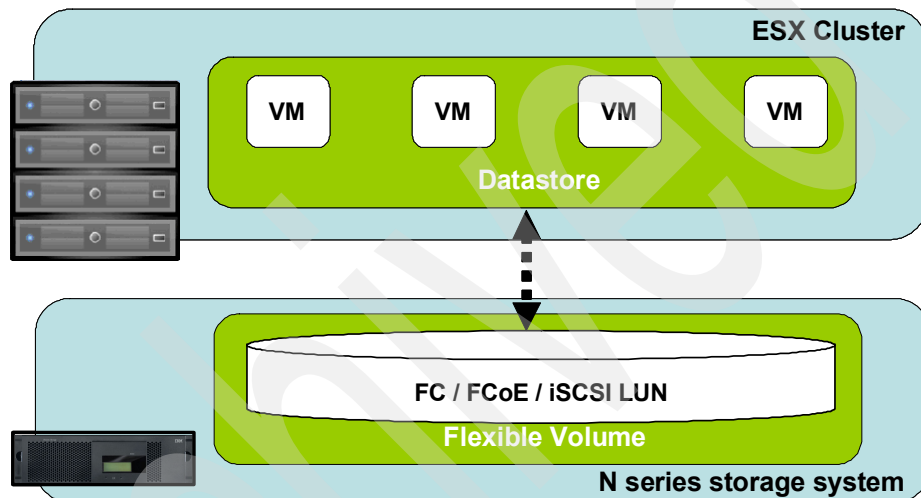


Figure 2-2 ESX server connected to a VMFS datastore using FC or iSCSI

2.2.2 Storage overview: NAS datastores

In addition to VMFS, VMware vSphere allows a customer to use enterprise-class NFS servers in order to provide datastores with concurrent access by all of the nodes in an VMware ESX cluster. This method of access is very similar to that with VMFS. NFS provides high performance, the lowest per-port storage costs (as compared to Fibre Channel solutions), and certain advanced data management capabilities.

Deploying VMware vSphere with N series NFS datastores is the easiest means to integrate VMware virtualization technologies directly with WAFL®.

Examples of this transparent integration include production-use data deduplication, immediate zero-cost VM and datastore clones, array-based thin provisioning, and direct access to array-based Snapshot copies.

N series provides additional VMware vSphere integrated tools for NFS such as SnapManager for Virtual Infrastructure.

Figure 2-3 depicts an example of this configuration. Note that the storage layout appears much like that of a VMFS datastore, yet each virtual disk file has its own I/O queue directly managed by the N series system. Combining N series advanced NFS servers with VMware's high-performance NFS implementation can provide I/O to shared datastores that is on par with that of other storage protocols such as Fibre Channel.




Figure 2-3 ESX server connected to an NFS datastore

2.2.3 Storage overview: Raw device mappings

VMware ESX allows for virtual machines to have direct access to LUNs for specific use cases such as physical to virtual (P2V) clustering or storage vendor management tools. This type of access is referred to as a Raw Device Mapping and can be configured with Fibre Channel, iSCSI, and Fibre Channel over Ethernet. In this design, VMware ESX acts as a connection proxy between the VM and the storage array.

Unlike VMFS and NFS, RDMs are not used to provide shared datastores. RDMs are an enabling technology for solutions such as virtual machine and physical to virtual machine host-based clustering, such as with Microsoft Cluster Server (MSCS). RDMs provide traditional LUN access to a host, so they can achieve high individual disk I/O performance, and they can be easily monitored for disk performance by a storage array.

The N series can enhance the use of RDMs by providing array-based LUN level thin provisioning, production use data deduplication, advanced integration components such as SnapDrive®, VM granular Snapshot copies, and FlexClone® zero-cost cloning of RDM-based data sets.

The challenges of this solution are that VMware ESX clusters might have to be limited in size, and this design requires ongoing interaction between storage and VMware vSphere administration teams. Figure 2-4 shows an example of this configuration.

RDMs are available in two modes: physical and virtual. Both modes support key VMware vSphere features such as VMotion and can be used in both HA and DRS clusters. The key difference between the two technologies is the amount of SCSI virtualization that occurs at the VM level. This difference results in certain limitations around MSCS and VMware vSphere Snapshot use case scenarios.

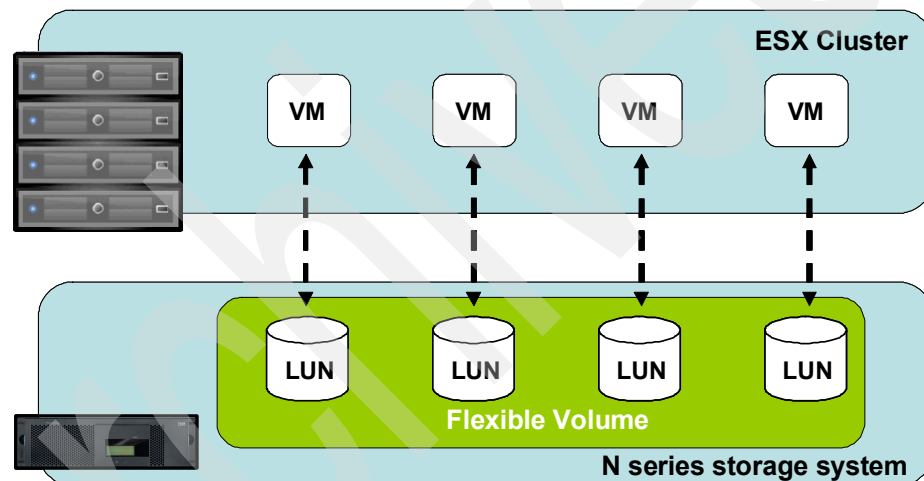


Figure 2-4 ESX Cluster connected to a VMFS datastore using FC or iSCSI

2.3 SnapManager backup and recovery

SnapManager for Virtual Infrastructure provides local backup and recovery capability with the option to replicate backups to a remote storage system by SnapMirror® relationships.

Backups can be performed on individual virtual machines or on datastores with the option of updating the SnapMirror relationship as part of the backup on a per-job basis. Similarly, restores can be performed at a level of datastore, individual virtual machine, VMDK, or individual file within the guest OS.

With the exception of identifying the relevant vCenter server and N series storage, very little configuration has to be completed before backups can be scheduled within SMVI. No profiles or databases are required because SMVI uses an XML catalog file to record information such as when a backup was created, which virtual machines or datastores were backed up, and how long each backup has to be retained.

2.3.1 Backup retention policy

Policies can be created specifying the retention period on a per-scheduled-backup job basis, allowing the administrator flexibility to meet varying service-level agreement levels within their environment. Retention can be specified by number of days or number of backups, or maintained indefinitely until manually deleted.

2.3.2 Alert notification

Alert notifications are created on a per-scheduled-backup job basis and are sent by e-mail to administrator defined accounts. Alert notification can be configured so that the specified account is e-mailed after every backup, although this is not desirable because the number of e-mails can become unmanageable. Configuring alerts to notify administrators after an error or warning within a backup offers a more useful and practical alert level.

2.3.3 Port usage

For SMVI, make sure that the following ports are kept open:

- ▶ The SMVI client uses port 8043 to communicate with the SMVI server.
- ▶ By default, the SMVI server tries to communicate with the N series controller on port 443 using HTTPS. If HTTPS is not enabled, it will fall back to HTTP using port 80.

2.4 SnapManager architecture

Figure 2-5 illustrates the SnapManager for Virtual Infrastructure architecture and the components that work together to provide a comprehensive and powerful backup and recovery solution for VMware vSphere environments.

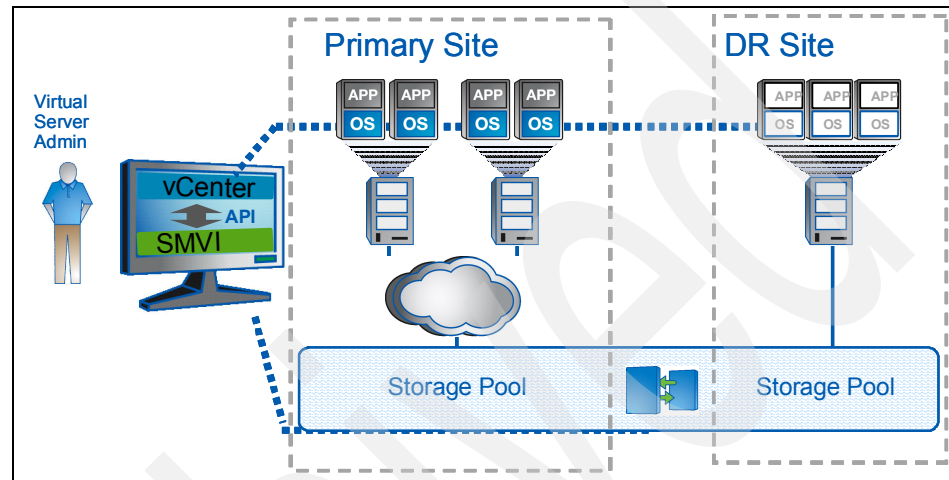


Figure 2-5 Architecture overview

2.5 Components

This section outlines the components used to create a SnapManager for Virtual Infrastructure solution.

2.5.1 Data ONTAP

SnapManager for Virtual Infrastructure will only function within an N series storage environment. SMVI requires that the primary storage, where the virtual machines actually reside, and the secondary storage used as the SnapMirror destination, are run Data ONTAP® storage software. These storage systems can be either Data ONTAP-specific physical storage systems or vFile units. SMVI supports Data ONTAP 7.2.x and 7.3.x on physical N series storage systems and 7.2.4 and higher or 7.3.x on vFile units. In addition, the following licenses are required:

- ▶ SnapRestore®
- ▶ The required protocol license (NFS, FCP, iSCSI)
- ▶ SnapMirror (if required)

- ▶ FlexClone:
 - FlexClone is required for mount operations of NFS datastores. FlexClone is not required for NFS VM in-place VMDK restores (SMVI uses ZAPIs). FlexClone is required for out-of-place NFS VMDK restores. Also, FlexClone is optional for mounting VMFS datastores. SMVI uses LUNClone when FlexClone is not available.
 - For SMVI versions 1.2x, FlexClone is needed only if you want to perform NFS datastore mounts. For 2.0, FlexClone is required for NFS datastore mounts and NFS Single File Restore operations, because the SFR workflow internally mounts NFS datastores.

2.5.2 VMware Virtual Infrastructure

SMVI supports specific VMware versions:

- ▶ VMware Virtual Infrastructure 3:
 - VMware ESX 3.5 Update 4 and later
 - VMware ESXi 3.5 Update 4 and later
 - VMware VirtualCenter 2.5 Update 4 and later
- ▶ VMware vSphere 4:
 - VMware ESX 4 and later
 - VMware ESXi 4 and later
 - vCenter Server 4 and later

For the most current information, see the N series interoperability matrix at:

<http://www-03.ibm.com/systems/storage/network/interophome.html>

Because SMVI communicates with vCenter rather than individual systems during backups, the clustering of VMware ESX hosts with the high-availability (HA) and Distributed Resource Scheduling (DRS) features enabled is supported. SMVI supports both VMFS datastores over iSCSI and FCP protocols and NFS datastores. Any combination of these datastores and protocol types can be supported by a single SMVI server.

2.5.3 SnapManager for Virtual Infrastructure

SnapManager for Virtual Infrastructure has two components:

- ▶ SnapManager Server
- ▶ SnapManager Client (GUI/CLI)

SMVI can be installed on any Windows platform (XP, 2003, Vista, 2008) that has connectivity with the vCenter server. When possible, install SMVI on the vCenter server to reduce the impact of network outages between the two components (no impact on the vCenter service when SMVI is installed on the same server).

2.5.4 SnapManager for Virtual Infrastructure repository

The SMVI repository consists of several .XML files that are placed on local storage if the default locations are accepted during installation. These files are critical to recovering SMVI in case a failure occurs at any number of levels. Place the repository files on shared storage, ideally within a Microsoft Cluster Server (MSCS) if possible, otherwise be sure to back them up regularly. Further details on configuring SMVI for shared storage and cluster configurations can be found in Chapter 4., “Installation” on page 25.



Planning

In this chapter we discuss important considerations when planning for your SnapManager for Virtual Infrastructure environment, including storage systems, data layout, configuration, and user permissions.

3.1 Storage configuration

SnapManager for Virtual Infrastructure seamlessly integrates into the VMware virtualized environment (see 2.5, “Components” on page 14 for supported versions of VMware ESX and vCenter).

NFS datastores are supported as well as VMFS datastores accessed by FCP and iSCSI protocols. When deploying the VMware Virtual Infrastructure on N series storage systems that will be supported by SMVI.

Figure 3-1 shows a sample of how NFS datastores are displayed.

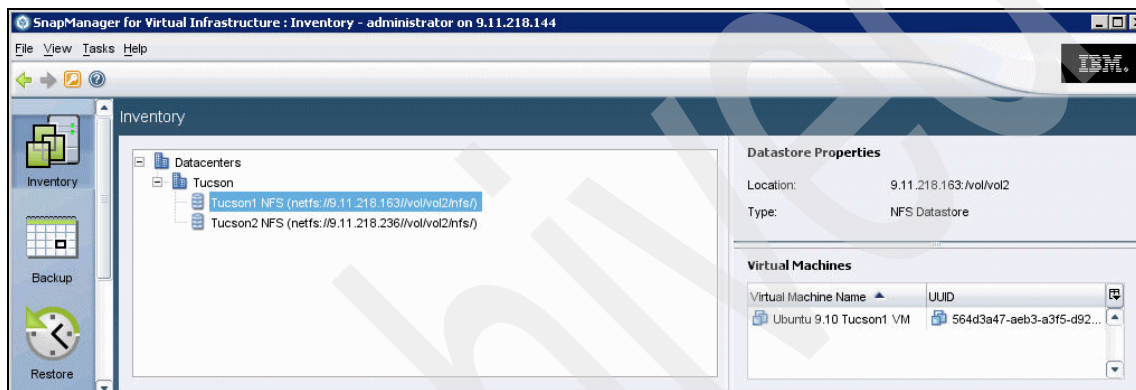


Figure 3-1 Datastores

3.2 Adding storage systems

During the configuration of SnapManager for Virtual Infrastructure, all primary and secondary storage must be identified within the Setup window of the GUI or Create a custom storage account other than root as a security precaution. To add an N series storage system, go to **Setup** and select **Add** as shown in Figure 3-2.




Figure 3-2 Adding a storage system

A further step must be taken as a security precaution when configuring SnapManager for Virtual Infrastructure and the associated storage systems. It is best to enable Secure Sockets Layer (SSL) on all storage systems identified to SMVI. Enabling SSL ensures account passwords are encrypted when transmitted to the storage system.

3.3 SnapManager for Virtual Infrastructure data layout

As indicated by N series best practices for VMware vSphere, any transient and temporary data such as the guest operating system swap file, temp files, and pagefiles, needs to be moved to a separate virtual disk on another datastore. The reason is that snapshots of this type of data can consume a large amount of storage in a very short period of time due to the high rate of change.

When a backup is created for a virtual machine with SnapManager for Virtual Infrastructure, SMVI is aware of all VMDKs associated with the virtual machine and will initiate a Snapshot copy on all datastores upon which the VMDKs reside. For example, if a virtual machine running Windows as the guest operating system has its C drive on datastore ds1, data on datastore ds2 and transient data on datastore td1, SnapManager for Virtual Infrastructure will create a Snapshot copy against all three datastores at the underlying volume level, thus defeating the purpose of separating temporary and transient data.

In order to exclude the datastore containing the transient and temporary data from the SnapManager for Virtual Infrastructure backup, configure the VMDKs residing in the datastore as “Independent Persistent” disks within vCenter. After being configured, the transient and temporary data VMDKs will be excluded from both the VMware vCenter snapshot and the N series Snapshot copy initiated by SnapManager for Virtual Infrastructure.

Also, create a datastore dedicated to transient and temporary data for all virtual machines with no other data types or VMDKs residing upon it. This will avoid a Snapshot copy being taken against the underlying volume as part of the backup of another virtual machine. Do not deduplicate the data on this datastore.

SnapManager 2.0 for Virtual Infrastructure has the option to include independent disks and exclude datastores from backup as displayed in Figure 3-3. We cover this topic in more detail in 5.5, “Include independent disks and exclude datastores” on page 45.

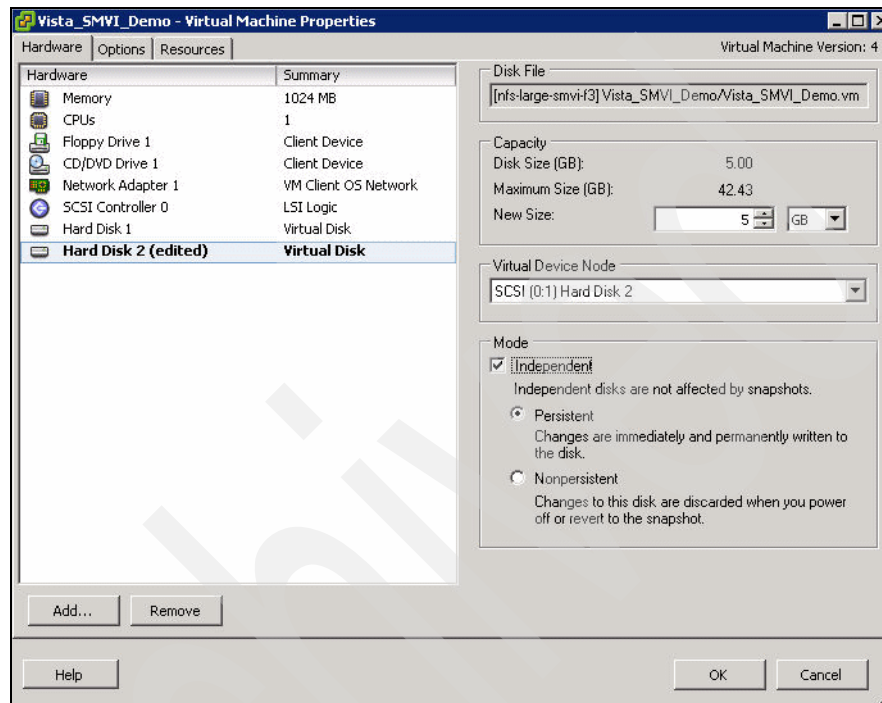


Figure 3-3 vCenter configuration of an independent disk

3.4 vCenter configuration

After installing SnapManager for Virtual Infrastructure, you must configure it for the specific environment. As part of this configuration, a vCenter server must be identified. SMVI can only communicate with one vCenter server at a time.

In many environments, more than one vCenter server is in use. Because SMVI can communicate only with one vCenter server at a time, it is best to install one SMVI server per vCenter server, thus enabling all scheduled backups to run successfully while ease of administration is maintained. Be aware that this configuration will not increase the cost of the solution, because SMVI is licensed on a per- VMware ESX host or storage-tier basis, rather than on a per-SMVI server basis.

Although the specified vCenter server within SMVI can be changed from both the GUI and the CLI, doing this is not desirable in an environment with multiple vCenter servers, because all jobs scheduled to run against a vCenter server that is not currently specified within the SMVI setup will fail.

3.5 vCenter user permissions

The SMVI vCenter connection account requires the following vCenter permissions for VMware ESX 4.0:

- ▶ Datastore:
 - Remove Datastore
 - Browse Datastore
 - Remove File
 - Rename Datastore
- ▶ Host configuration:
 - Storage Partition Configuration
 - Advanced Settings
 - Change Settings
- ▶ Virtual machine:
 - Interaction:
 - Power off
 - Power on
 - Configuration:
 - Add Existing Disk
 - Add New Disk
 - Remove Disk
 - Add or Remove Device
 - Advanced
 - State:
 - Create Snapshot
 - Remove Snapshot
 - Revert to snapshot

3.6 Distributed resource scheduler implications

VMware's Distributed Resource Scheduler, or DRS, pools the resources of ESX hosts within a cluster, dynamically migrating virtual machines by VMotion to ensure optimal resource availability and virtual machine performance.

Within a DRS configuration, there are three possible levels of automation:

- ▶ **Manual:** vCenter indicates which virtual machines have to be migrated; however, migrations are manually implemented by an administrator.
- ▶ **Partially Automated:** Virtual machines are automatically placed onto a VMware ESX host within the cluster by vCenter during power on. However, migrations are suggestions only and must be manually implemented by an administrator.
- ▶ **Fully Automated:** VMware ESX Virtual machines are automatically placed onto VMware ESX hosts at power on and are automatically migrated to attain the best use of resources.

Within a Fully Automated configuration, there are five possible levels of aggression that can be selected to determine the threshold for moving virtual machines. The sliding scale ranges in aggression from most conservative, where vCenter only migrates virtual machines to satisfy cluster constraints, to most aggressive, where vCenter automatically applies migrations that promise even a slight improvement to the cluster's load balance.

Depending on the environment, the level of aggression has a direct effect on the number of migrations that will occur. In certain cases, when the highest level of aggression has been selected, virtual machines are constantly undergoing migration. Such constant migration can cause performance problems within the VMware vSphere environment due to the level of overhead involved in migrating a virtual machine. It can also cause errors to occur during an SMVI backup job.

SnapManager for Virtual Infrastructure is VMotion aware; as long as the storage system a datastore resides upon is known to SMVI, it does not matter which host a virtual machine resides upon. However, SMVI cannot back up a virtual machine that is actively undergoing migration. In case a backup is run against a datastore that has virtual machines actively being migrated, an error will be generated and those particular virtual machines will not be backed up. As a result, the level of aggression and the number of migrations occurring must be carefully monitored.

In case the backups begin to experience errors when DRS is configured in Fully Automated mode, it is best to scale back the level of aggression so that virtual machines are migrated only when a significant gain in performance can be achieved. This will improve not only the success rate of the backups, but the overall virtual machine performance as well.

Installation

There are a number of options available when installing SnapManager for Virtual Infrastructure, ranging from the default installation on local disk to a clustered configuration with a Microsoft Cluster Server.

In this chapter we describe the supported installations. However, the preferred way of installing SMVI is with the catalog files residing upon shared storage, which provides quick and easy recovery of the SMVI server and its associated backup files in case of a failure.

4.1 Installing vCenter and SMVI on a virtual machine

In the following sections we outline how to install SMVI and vCenter on a virtual machine.

4.1.1 Installing vCenter within a virtual machine

Running VMware vCenter (Virtual Center) within a virtual machine is fully supported by VMware to the same degree as though it were installed on a physical server. However, there can be ramifications in a SnapManager for Virtual Infrastructure environment if you do not strictly adhere to VMware's best practices for installing vCenter on a virtual machine.

Specifically, the database associated with vCenter must not be installed on a VMware virtual machine protected by SMVI, but rather must be installed on a physical system. In case the database is installed in a virtual machine backed up by SMVI, vCenter will fail due to vCenter time-outs caused by the VMware snapshot process.

Further information about running VirtualCenter within a virtual machine can be found in the VMware Technical Note titled "Running VirtualCenter in a Virtual Machine", at the following website:

http://www.vmware.com/pdf/vi3_vc_in_vm.pdf

4.1.2 Installing SMVI within a virtual machine

Installing SnapManager for Virtual Infrastructure within a virtual machine is fully supported, and the installation process is identical to that of a physical machine. Just as it is best to install both SMVI and vCenter on the same physical server, it is also best to install both SMVI and vCenter on the same VMware virtual machine.

Although SnapManager for Virtual Infrastructure is supported when running within a virtual machine, you must take additional steps if SMVI will be "backing up itself" to avoid potential issues after a restore of the SMVI virtual machine.

Backups: SMVI "backing up itself" works for crash consistent snapshots. It is best to use the Windows native backup utility within the guest or another backup application.

When the virtual machine running SnapManager for Virtual Infrastructure is backed up by SMVI, the current state of the SMVI backup workflow is captured within the Snapshot copies. In case the virtual machine is restored at a later date, SMVI within the virtual machine assumes that the host has failed in mid-workflow and attempts to resume backups at the point at which the Snapshot copy was taken, which can cause backups to run multiple times even if they were originally successful.

To prevent this issue, take the following steps when SnapManager for Virtual Infrastructure is installed on a virtual machine:

1. Set the Startup Type of the SnapManager VI Windows Service to Manual. This will prevent SnapManager for Virtual Infrastructure from restarting automatically when the virtual machine it's installed on is powered up after it has been restored.
2. After restoring the SMVI virtual machine, power the VM on and remove the contents of the %PROGRAMFILES%\IBM\SMVI\server\crash directory. SMVI uses this directory after a crash to resume failed backups.
3. Start the SnapManager VI Windows Service.

Also monitor the backup of the virtual machine containing SnapManager for Virtual Infrastructure. In case time-outs occur during the VMware snapshot process, configure the backups to run with the VMware snapshots disabled on the individual virtual machine.

4.2 Default installation

SnapManager 2.0 for Virtual Infrastructure is easy to install and configure. Here we provide a brief overview of the steps involved in the installation of SMVI:

1. Verify that all the prerequisites mentioned in the *Installation and Administration Guide* have been met. Make sure that:
 - The vCenter and ESX versions are supported by SMVI.
 - Data ONTAP is installed on the N series storage systems and the following licenses are enabled:
 - The correct protocol (FCP, iSCSI, or NFS)
 - SnapRestore
 - SnapMirror (as required)
 - FlexClone
2. Download the appropriate SnapManager 2.0 for Virtual Infrastructure software from the IBM support website. Follow the instructions in the *Installation and Administration Guide* to install SMVI. A sample installation is shown in Appendix D, “SMVI installation steps” on page 87.

4.3 Choosing between the GUI and the CLI

All SnapManager for Virtual Infrastructure commands can be performed using either the GUI or the CLI, with certain exceptions. The creation of scheduled jobs and their associated retention policies and Single File Restore can only be performed through the GUI.

4.4 Configuring SMVI for shared storage

If shared storage is available, it is best that the SMVI server be installed and configured with the configuration files residing on the shared storage. These steps are for disks that have been formatted by the OS and will not be supported for mapped drives or mount points. In the steps provided here, the shared storage device is depicted as the H drive:

1. Install SMVI.
2. Stop the SnapManager VI Windows service.
3. Within a directory residing upon shared storage, create the following directories:
 - H:\IBM\SMVI\server
 - H:\IBM\SMVI\server\etc (stores credentials for vCenter)
 - H:\IBM\SMVI\server\repository (stores backup catalogs)
 - H:\IBM\SMVI\server\crash (stores required files to resume operation)
4. Update the SMVI configuration files to identify the new location of the repository and crash folders on the shared storage. The following two files must be edited:
 - %PROGRAMFILES%\IBM\SMVI\server\etc\smvi.config
 - %PROGRAMFILES%\IBM\SMVI\server\etc\smvi.override
5. Two parameters must be changed within the smvi.config file to reflect the path to the shared storage (this parameter will have C: entered after a default installation):
 - smvi.repository.path=H:\\IBM\\SMVI\\server\\repository
 - flow.persistence.embedded.storageLocation=H:\\IBM\\SMVI\\server\\crash (must match the changes made to the smvi.config and smvi.override files).
6. One parameter must be changed within the smvi.override file to reflect the path to the shared storage (this parameter will have the C: entered after a default installation):
credential.persistence.file.path=H:\\IBM\\SMVI\\server\\etc\\cred

7. Copy the credential file from c:\program files\IBM\smvi\server\etc to the \etc directory in the new path structure. If this is not done, then use the CLI to set the vcserver IP address and credentials.
8. Start the SnapManager VI Windows service.
9. SMVI will now be configured to use the shared storage. Recovering from the loss of the SMVI Server when using this configuration is covered in Appendix A., “Recovery procedures” on page 73.

4.5 Configuring SMVI in a clustered environment (MSCS)

SnapManager for Virtual Infrastructure can be installed within a Microsoft Clustered Solution, if desired. The following section details the configuration of the SnapManager for Virtual Infrastructure server in a Microsoft clustered environment. In this example, the following assumptions are made:

- ▶ Windows 2003 Server SP2. In this example, a two node cluster is assumed.
- ▶ You have access to shared storage, specifically a LUN created on shared storage mounted as drive letter Q to be used as the cluster’s quorum disk, and a second device on which to place SMVI’s configuration files.
- ▶ The Cluster Administration Utility has been used to configure basic MSCS resources.
- ▶ A domain account has been created with access to the vCenter server.

After the cluster has been configured and tested, use the following steps to configure and test SnapManager for Virtual Infrastructure. The H drive represents the shared storage device available for the SMVI configuration files.

4.5.1 SMVI configuration

This section explains the steps to configure SMVI:

1. Perform a full installation of SnapManager for Virtual Infrastructure on both nodes within the cluster according to the *SnapManager for Virtual Infrastructure Installation and Administration Guide*.
2. Follow the steps in 4.4, “Configuring SMVI for shared storage” on page 28, on both cluster nodes.
3. Restart the SnapManager VI Windows service on node 1; this will reconfigure SMVI to use the folders created on the shared drive.
4. Stop the SnapManager VI Windows service on node 2.

4.5.2 Configuring cluster resources for SMVI

The following steps outline the configuration of the MSCS cluster resources for SMVI:

1. Start the Cluster Administrator and connect to the appropriate cluster.
2. Expand **Groups** in the left window pane of Cluster Administrator.
3. Right-click the Cluster Group and select **New Resource**.
4. Enter the following information in the New Resource window as shown in Figure 4-1.

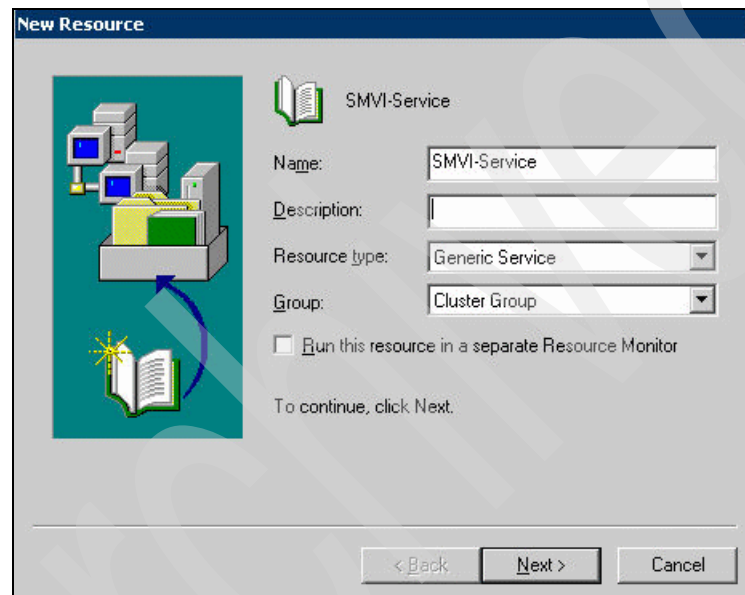


Figure 4-1 MSCS cluster new resource

5. Ensure that all the cluster nodes appear under the “Possible Owners” list. Move cluster nodes from “Available Nodes” to “Possible Owners” as necessary.

6. Within the Dependencies window, select the shared H: drive containing SnapManager for Virtual Infrastructure's configuration files. In Figure 4-2, the H: drive has been named SMVI-Data- Disk within the cluster.

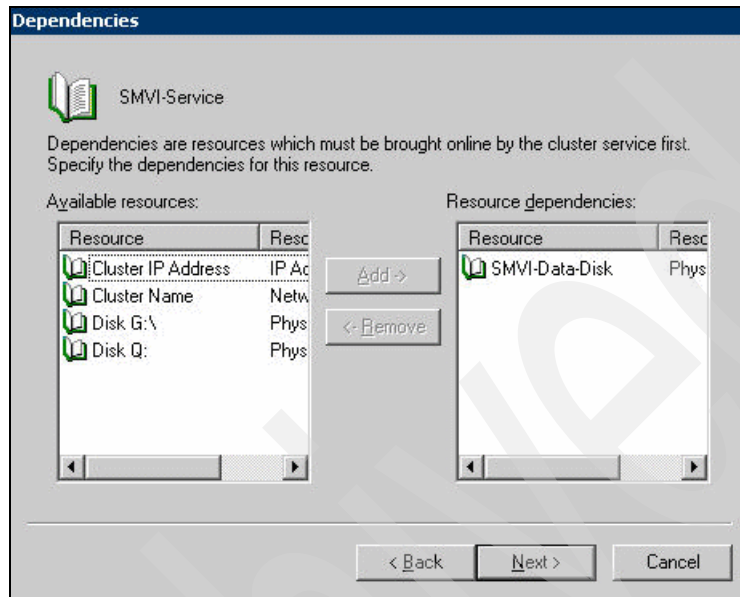
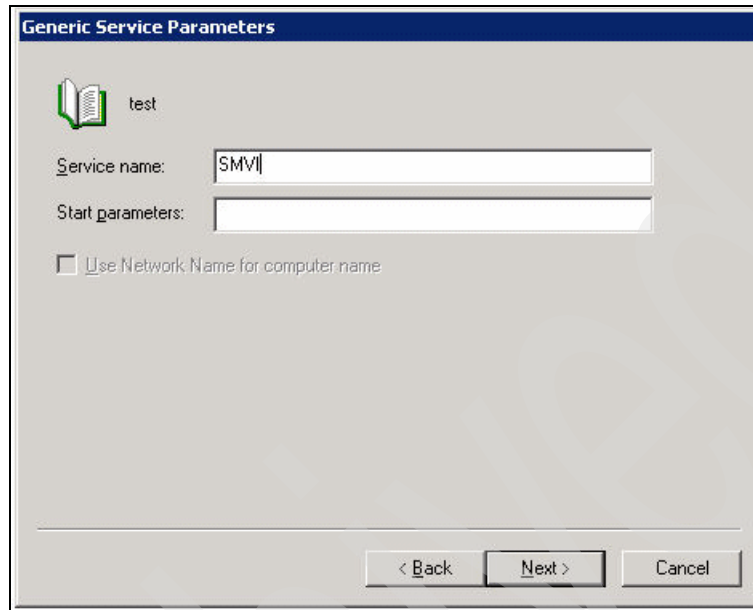


Figure 4-2 SMVI service dependencies

7. Enter the Service Name as SMVI and leave the Service Parameters field blank (see Figure 4-3).



Generic Service Parameters

test

Service name: SMVI

Start parameters:

☐ Use Network Name for computer name

< Back Next > Cancel

Figure 4-3 SMVI service parameter

8. Leave the Registry Replication window blank and click **Finish**.
9. Bring the resource online by right-clicking **SMVI-Service**.
10. Bring online.

4.5.3 Confirming SMVI high availability

After the SnapManager for Virtual Infrastructure cluster resource has been configured as described above, follow the steps listed below to test the configuration and determine that SMVI is running in a high availability mode:

1. Start the SnapManager for Virtual Infrastructure GUI on either node or on a client system with network connectivity to the cluster. Then connect to the SMVI server specifying the cluster's virtual IP address.
2. Within the Setup window of SMVI, enter the appropriate vCenter and storage information. Instructions on how to configure SMVI can be found in the *SnapManager for Virtual Infrastructure Installation and Administration Guide*.
3. Perform a backup using SnapManager for Virtual Infrastructure.

4. Failover the cluster group to the second node. When the failover is complete, connect to the cluster's virtual IP address and ensure that the backup taken in step 4 is available for restore.

4.6 SMVI in a N series MetroCluster environment

N series MetroCluster configurations consist of a pair of active-active storage controllers configured with mirrored aggregates and extended distance capabilities to create a high-availability solution. The primary benefits include:

- ▶ Higher availability with geographic protection
- ▶ Minimal risk of lost data, easier management and recovery, and reduced system downtime
- ▶ Quicker recovery when a disaster occurs
- ▶ Minimal disruption to users and client applications

A MetroCluster (either Stretch or Fabric) behaves in most ways just like an active-active configuration. All of the protection provided by core N series technology (RAID-DP®, Snapshot™ copies, automatic controller failover) also exists in a MetroCluster configuration. However, MetroCluster adds complete synchronous mirroring along with the ability to perform a complete site failover from a storage perspective with a single command.

The following N series MetroCluster types exist and work seamlessly with the complete VMware vSphere and ESX server portfolio:

- ▶ *Stretch MetroCluster* (sometimes referred to as nonswitched) is simply an active-active configuration that can extend up to 500m depending on speed and cable type. It also includes synchronous mirroring (SyncMirror®) and the ability to do a site failover with a single command.
- ▶ *Fabric MetroCluster* (also referred to as switched) uses four Fibre Channel switches in a dual-fabric configuration and a separate cluster interconnect card to achieve an even greater distance (up to 100 km depending on speed and cable type) between primary and secondary locations.

The integration of the MetroCluster and VMware vSphere is seamless and provides storage and application redundancy. In addition to connecting to the vSphere environment using FCP, iSCSI, or NFS, this solution is able to serve other network clients with CIFS, HTTP, and FTP at the same time. The solution shown in Figure 4-4 provides redundant VMware server, redundant N series heads, and redundant storage.

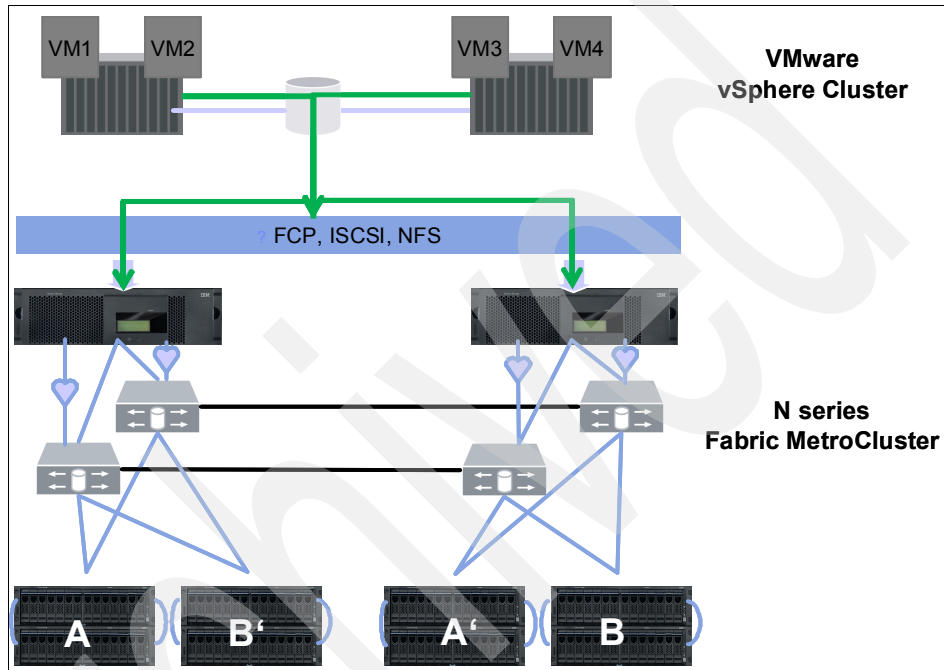


Figure 4-4 MetroCluster and VMware vSphere integrated solution

For more information about N series MetroCluster, refer to the “MetroCluster” chapter in the Redbooks publication, *IBM System Storage N series*, SG24-7129, at the following website:

<http://www.redbooks.ibm.com/abstracts/sg247129.html?Open>

Process flow

In this chapter we provide detailed information about the typical SnapManager for Virtual Infrastructure architecture on the primary site storage, and explain how it will be used in the backup process flow.

5.1 Backup process and implications

This section outlines backup process and backup implications.

5.1.1 The backup process

SnapManager for Virtual Infrastructure uses N series Snapshot technology to create fast and space efficient backups of VMware datastores and their associated virtual machines. These backups offer point-in-time images, or copies, of the virtual machines and are stored locally on the same storage platform on which the datastores physically reside.

In addition to the Snapshot copy stored locally, SnapManager for Virtual Infrastructure also provides an option to update an existing SnapMirror relationship upon the completion of a backup. This can be selected on a per-backup-job basis as required by the administrator.

The backup process flow is identical for both a manual and a scheduled backup job, as shown in Figure 5-1.

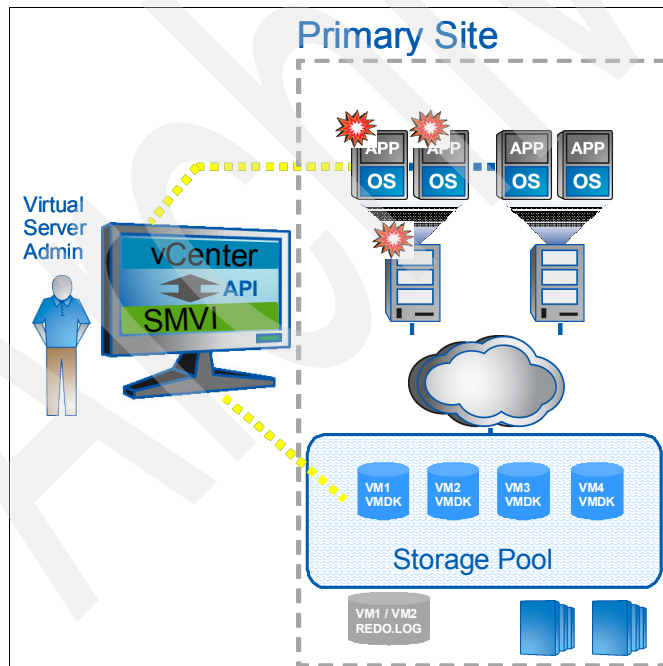


Figure 5-1 Virtual infrastructure and associated storage during an SMVI backup

Figure 5-1 on page 36 represents a high-level overview of the typical SnapManager for Virtual Infrastructure architecture on the primary site storage and is used in detailing the backup process flow:

1. A backup is initiated within SMVI:
 - Individual VMs backup: A VMware snapshot is created for each virtual machine selected for backup that is powered on at the time of backup.
 - Datastore backup: A VMware datastore snapshot is created for every virtual machine that is powered on within the datastore that has been selected for backup.
 - Virtual machines: Regardless of backup type, virtual machines that are powered off during a backup are backed up; however, no VMware vSphere snapshot is required.
2. The VMware snapshot preserves the state of the virtual machine and is used by SMVI during restores to revert the virtual machine back to the backup point-in-time state. VMware snapshots initiated by SMVI capture the entire state of the individual virtual machines, including disk and settings state; however additional steps must be taken to quiesce an application within a virtual machine. Information about application consistency is provided in Chapter 8., “Data consistency in an SMVI environment” on page 59. SMVI also gives users the option to disable VMware snapshots and just take an N series Snapshot copy on the volume underlying the datastore.
3. After all the VMware snapshots have completed for a datastore, SMVI initiates an N series Snapshot copy on the volume underlying the datastore. The N series Snapshot copy is at the volume level regardless of the type of backup selected: individual virtual machines or datastore. During the backup of a virtual machine with VMDKs residing upon multiple datastores, SnapManager for Virtual Infrastructure initiates an N series Snapshot on all the underlying volumes. As a result, multiple datastores are displayed within SMVI's restore window regardless of how many datastores were selected for backup.
4. Upon completion of the N series Snapshot copy, SMVI removes the VMware vCenter snapshot to reduce space and performance overhead. Although the vCenter snapshot is removed within vCenter, one VMware snapshot is maintained within the backup for each virtual machine that was in a powered-on state. This snapshot is maintained so it can be used by the restore process to revert the virtual machine to its point-in-time state.
5. Upon completion of the local backup, SnapManager for Virtual Infrastructure updates an existing SnapMirror relationship on the volume underlying the datastore if the SnapMirror option was selected. SnapMirror is discussed in further detail in a later section of this document.

5.1.2 Backup process implications

As documented before, the process flow is similar regardless of the type of backup performed, individual virtual machines or datastore. However, the number of virtual machines that will undergo a VMware vCenter snapshot will vary depending on the type of backup selected, as will the number of N series Snapshot copies per volume.

When an individual virtual machine is selected for backup, only that virtual machine has a VMware vCenter snapshot created. As a result, only the selected virtual machine is available for restoration, even though the entire underlying datastore volume was protected by an N series Snapshot copy.

Conversely, when a datastore is selected for backup, every virtual machine in the datastore in a powered-on state has a VMware vCenter snapshot created, before a single N series Snapshot copy of the volume is performed.

Unlike the individual virtual machine backup, any virtual machine that resided upon the specified datastore can be selected for restoration. The entire datastore need not be restored unless required; individual virtual machines can be selected from the backup as necessary.

Configure datastore-level backups whenever possible, especially for scheduled backups. This not only reduces administrative overhead by lessening the number of backups that need to be configured and tracked, but also reduces the number of N series Snapshot copies per volume.

Multiple N series Snapshot copies of the volume are required if each virtual machine is selected for backup individually, rather than the single N series Snapshot copy required per datastore backup.

Reducing the number of N series Snapshot copies per volume increases the number of backups that can be retained.

To reduce both storage and administrative overhead as well as increase the restoration options available to the administrator, it is best to configure datastore-level backups as much as practically possible.

Datastore: For SMVI, the first 24 characters of a datastore must be unique.

5.2 Scheduled backups and retention policies

The limit of 255 Snapshots per volume must be taken into consideration when scheduling backups and configuring the associated retention policies. The number of Snapshot copies per volume can be managed with the proper scheduling and retention policies on a per-scheduled-backup basis while still meeting Service level agreements (SLAs) on the virtual machines.

5.2.1 Backup scheduling

Five scheduling options are available when creating a backup within SMVI:

- ▶ Hourly
- ▶ Daily
- ▶ Weekly
- ▶ Monthly
- ▶ None

The frequency of the backup has a direct bearing on the number of N series Snapshot copies taken on the underlying volume, regardless of the type of backup performed: individual VMs or datastore.

The backup frequency, as well as the number of various backups performed against a datastore (for example, one backup running against datastore ds_1 weekly and another monthly) must be taken into account when specifying the retention policy so as not to exceed the maximum number of Snapshots per volume. In case the number of Snapshot copies exceeds 255 on any given volume, future backups against that volume will fail. Selecting the **None** option and choosing to delete the job are equivalent to the manual backup-job creation option.

5.2.2 Retention policies

Three options are available to the administrator when configuring retention policies for a scheduled backup:

- ▶ Maximum number of days
- ▶ Maximum number of backups
- ▶ Indefinitely

Preferably, use the policies not only to meet specific SLAs, but also to maintain a supported number of N series Snapshots on the underlying volumes, which can be achieved by using either of the first two of these options.

For example, by setting a retention policy of 30 backups on an hourly backup, we can limit the maximum number of Snapshots associated with the backup to 30. However, if the retention policy is configured as 30 days, the Snapshot limit per volume is reached after 10 days and the backups begin to fail from that point on (24 backups per day, reaching the 255 Snapshot limit on the 11th day).

The third option, *Indefinitely*, has to be used with caution. When selecting this option, backups and the associated N series Snapshot copies are maintained until manually deleted by the administrator. These Snapshot copies are included in the maximum number supported on a volume. Of further note, the N series Snapshot copies associated with manual backups must also be considered when determining the number of Snapshot copies maintained against a volume.

5.3 Snapshot naming

SMVI 2.0 includes the following changes to the snapshot naming convention:

- ▶ Jobs scheduled by GUI have the snapshot naming pattern “**smvi_{jobName}_{novmsnap}_{date}{time}**”. A job has multiple backups associated with it, taken at various times depending on the schedule, only the most recent backup of the job has snapshots with the name “**smvi_{jobName}_{novmsnap}_recent**”.

When the job runs the next time, the current recent Snapshot copy of the job is renamed to **smvi_{backupName}_{novmsnap}_{date}{time}**, where {date}{time} is the start time of the backup. The [novmsnap] String is inserted depending on whether VMware snapshots were taken for that particular backup.

- ▶ Manual backups done through the CLI have the snapshot naming pattern “**smvi_{backupName}_{date}{time}**”. Manual backups never have **recent** in their snapshot names.

5.4 Scripting

SnapManager for Virtual Infrastructure provides users the ability to run pre-, post- and failure backup phase scripts. These scripts are any executable process on the operating system in which the SMVI server is running. When defining the backup to run, the pre-, post-, and failure backup scripts can be chosen using either the SMVI GUI or CLI. The scripts must be saved in the <SMVI Installation>/server/scripts/ directory. Each chosen script runs as a pre-, post-, and failure backup script.

From the GUI, the user can select multiple scripts using the backup creation wizard or when editing an existing backup job as displayed in Figure 5-2. The UI will list all files found in the server/scripts/ directory. SMVI runs the scripts before creating the VMware snapshots and after the cleanup of VMware snapshots.

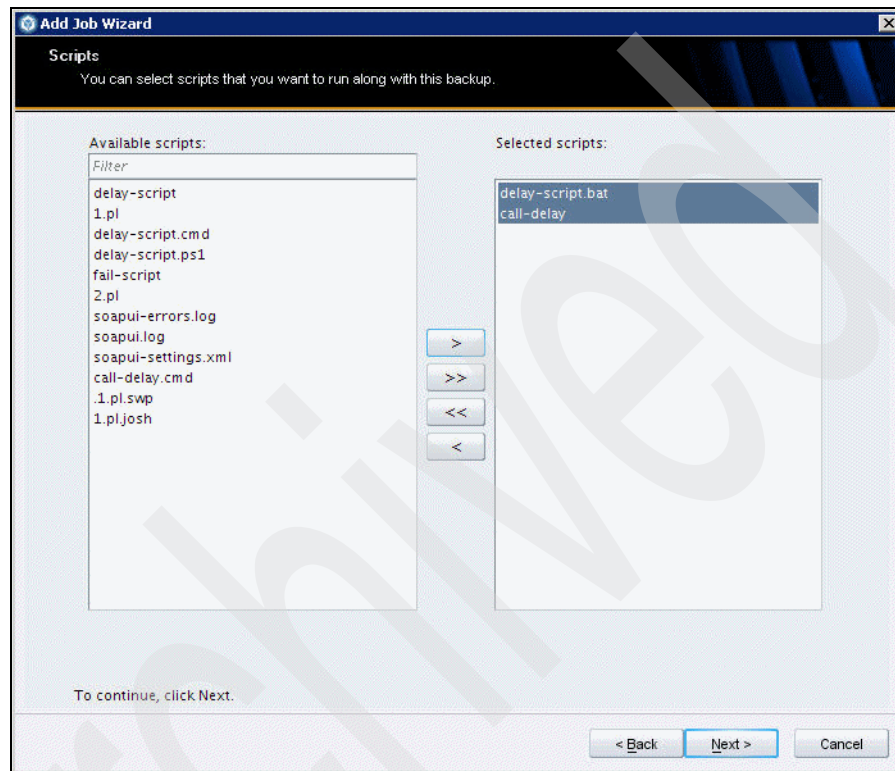


Figure 5-2 Select backup scripts from the GUI

When SMVI starts each script, a progress message is logged indicating the start of the script. When the script completes, or is terminated by SMVI because it was running too long, a progress message is logged to indicate the completion of the script and if the script was successful or failed. If a script is defined for a backup but is not found in the scripts directory, a message is logged stating that the script cannot be found.

The SMVI server will maintain a global configuration value to indicate the amount of time that a script can execute. After a script has run for this length of time, the script is terminated by the SMVI server in order to prevent run away processing by scripts. If the SMVI server has to terminate a script, it will implicitly be recognized as a failed script and might force termination of the SMVI backup in the pre-backup phase.

With the default settings, SMVI will wait for up to 30 minutes for each script to complete in each phase. This default setting can be configured using the following entry in the server/etc/smvi.override file:

```
smvi.script.timeout.seconds=1800
```

SMVI backup scripts will receive inputs from environment variables. This will allow sending the inputs in a way that avoids CLI line length limits. The set of variables will vary based on the backup phase.

5.4.1 Environment variables

The scripts can expect the environment variables listed in Table 5-1 using the appropriate phases.

Table 5-1 List of environment variables

Variable	Content	Format	Phase
BACKUP_NAME	Name of the backup		All phases
BACKUP_DATE	Date of the backup	yyyymmdd	All phases
BACKUP_TIME	Time of the backup	HHMMss	All phases
BACKUP_PHASE	Phase of the backup	PRE_BACKUP, POST_BACKUP, or FAILED_BACKUP	All phases
VIRTUAL_MACHINES	The number of VMs in the backup		
VIRTUAL_MACHINES.#	One of the defined virtual machines	Uses the fixed format <VM name> <VM UUID> <power state - POWERED_ON POWERED_OFF SUSPENDED> <VM snapshot is taken - true false> <ip addresses>.	All phases
STORAGE_SNAPSHOTS	The number of storage snapshots in the backup		POST_BACKUP
STORAGE_SNAPSHOT.#	One of the defined storage snapshots	Uses the standard convention of <filer>:/vol/<volume>: <ONTAP snapshot name>.	POST_BACKUP

Sample environment variables and sample scripts are available in Appendix B, “Sample scripts” on page 79.

Upon completion of the script, SMVI expects an exit value of zero (0) to indicate success. Any exit value other than zero (0) indicates a failure by the script. If a script fails in the PRE_BACKUP phase, the SMVI backup will fail. Script failures in the POST_BACKUP phase will result in a WARN message but will not fail the backup. Scripts can write output to stdout/stderr that is read by the SMVI server and collected for each script that is run. The entire contents of stdout/stderr are logged in the SMVI server log file along with the exit value of the script.

The most common problem is that the script is written in a particular language, but the required binaries for that language are not on the system PATH. In this case, do a quick check to open a command prompt and type `echo %PATH'%`. If the path to the language is not on the path, it must be added. Another check is to create a `.cmd` script which echoes the input variables, including PATH. Then a backup can be run with this script and the output can be viewed from the SMVI server log file. If a script is developed using perl, do not use backticks (```) to run another process: instead, use the `system()` call.

When running a backup from the CLI, a new option is available to define which scripts are run as both pre- and post-backup scripts:

```
smvi backup create  
    [-scripts {script name} [script name ...]
```

: (optional) name of the scripts to run with this backup, multiple script names can be specified. The script name must match the name of the script as found on the SMVI server. Only include the script name, not the path to the script. Scripts must reside in the scripts directory on the SMVI server and cannot be in a sub-folder. The order in which scripts are run is not guaranteed.

5.4.2 Error messages

All output messages from the scripts are stored in the SMVI log. If a script fails in the pre-backup phase, the backup will fail with a message similar to the one shown in Example 5-1.

Example 5-1 Error message

```
smvi backup create -id Empty-Test-VM-1 -scripts fail-script  
[08:15] Starting backup request  
[08:15] Excluding datastore Development (Backup  
backup_josh-test-schedules_20090729144239) from the backup for virtual  
machine Empty-Test -VM-1. Empty-Test-VM-1 has only independent disks on  
Development (Backup backup_josh-test-schedules_20090729144239).
```

```

[08:15] Backing up datastore(s) ([Development
(netfs://10.60.231.19//vol/jb_smvi_datastore2/)])
[08:15] Backing up the following virtual machine(s) ([Empty-Test-VM-1])
[08:15] Script fail-script is starting in phase PRE_BACKUP
[08:15] [ERROR] Script fail-script failed with exit code 1
[08:15] Script fail-script is starting in phase FAILED_BACKUP
[08:15] Script fail-script completed successfully in phase
FAILED_BACKUP
[08:15] Storing logs for backup_5c03867780fc40811168da12e907159d in
file
./repository/logs/unscheduled/backup_backup_5c03867780fc40811168da12e90
7159d.xml
SMVICLI-0101: Command failed

```

Failures in the post-backup phase will not fail the backup. Example 5-2 shows an error of a post backup phase failure output.

Example 5-2 Error in post backup

```

smvi backup create -id Empty-Test-VM-1 -scripts fail-script
[08:20] Starting backup request
[08:20] Excluding datastore Development (Backup
backup_josh-test-schedules_20090729144239) from the
backup for virtual machine Empty-Test-VM-1. Empty-Test-VM-1 has only
independent disks on
Development (Backup backup_josh-test-schedules_20090729144239).
[08:20] Backing up datastore(s) ([Development
(netfs://10.60.231.19//vol/jb_smvi_datastore2/)])
[08:20] Backing up the following virtual machine(s) ([Empty-Test-VM-1])
[08:20] Script fail-script is starting in phase PRE_BACKUP
[08:20] Script fail-script completed successfully in phase PRE_BACKUP
[08:20] Creating storage snapshots for all datastores/virtual machines
that are being backed up.
[08:20] Script fail-script is starting in phase POST_BACKUP
[08:20] Script fail-script failed with code 1 in phase POST_BACKUP
[08:20] Storing logs for backup_449142621bee95e72b598e8e45bad68c in
file
./repository/logs/unscheduled/backup_backup_449142621bee95e72b598e8e45b
ad68c.xml
[08:20] Backup backup_449142621bee95e72b598e8e45bad68c of
datastores/virtual machines is complete.
SMVICLI-0100: Command completed successfully

```

5.5 Include independent disks and exclude datastores

Preferably, create a datastore dedicated to transient and temporary data for all virtual machines with no other data types or VMDKs residing upon it, in order to avoid a Snapshot copy being performed on the underlying volume as part of the backup of another virtual machine. It is best to exclude datastores that contain transient and temporary data from the backup, this ensures that snapshot space is not wasted on transient data with a high rate of change. In SMVI 2.0, when selected entities in the backup span multiple datastores, one or more of the spanning datastores might be excluded from the backup.

After being configured, the transient and temporary data VMDKs are excluded from both the VMware vCenter snapshot and the N series Snapshot copy initiated by SnapManager for Virtual Infrastructure. In SMVI 1.0 datastores with only independent disks were excluded from the backup. In SMVI 2.0 there is an option to include them in the backup. Datastores with a mix of independent disks and normal disks or configuration files for a VM are included in the backup irrespective of this option.

In case you have a normal disk and an independent disk for backup on the same datastore, this datastore is always included for backup irrespective of the “include datastore with independent disk” option. It is best that a separate datastore be designated exclusively for swap data or system data.

Restore from backup: If you exclude non-independent disks from the backup of a VM, that VM cannot be completely restored. Only virtual disk restore and single file restore can be performed from such a backup.

5.6 Mounting a backup

SnapManager for Virtual Infrastructure provides the option of mounting a backup. The mounted backup is a clone of the protected datastore. After being mounted, the backup is displayed within vCenter and can be browsed in the same manner as a standard datastore. The mount operation can be used to verify the contents of a backup.

To mount a backup, complete the following steps:

1. In the Datastores pane of the Restore window, select a datastore to display a list of its backups in the Backups pane.
2. In the Backups pane, select the backup that has to be mounted.
3. Click **Mount**.

4. When prompted, enter the name of the VMware ESX server on which you want to mount the datastore.
5. Browse the mounted backup through vCenter in the same manner as you might browse a standard datastore.

5.7 Single file restore

SMVI 2.0 supports restoring one or more files from the virtual machine without having to restore the entire virtual machine.

5.7.1 Pre-requisites

Following are the prerequisites for single file restore (SFR):

- ▶ There is an existing process (for example, help desk/ticketing system) to be used by users to initiate file restore requests.
- ▶ The authentication and authorization of the user to be able to restore files belonging to a particular VM is done by this external tool / process, which can be unique for every enterprise and is outside the scope of SMVI.
- ▶ When the request to restore files on the guest OS comes to the SMVI administrator through this tool, it is assumed that the user asking for the restore of files of the particular guest OS is authorized to do so.

5.7.2 Types of file restore sessions

SnapManager for Virtual Infrastructure automates the process of restoring single files based on the relationship between the source virtual machine (which was backed up) and the destination virtual machine (the VM to which files will be restored). The source and destination VM can be the same or other virtual machines.

SnapManager for Virtual Infrastructure supports three types of restore sessions.

- ▶ **Self Service:** The SnapManager for Virtual Infrastructure administrator creates a restore session using SMVI. Users can then install the Restore Agent (RA) on the destination virtual machine, browse the mounted backups on a guest virtual machine, and restore the individual disk file.
- ▶ **Administrator Assisted:** This type of file restoration is basically the same as self-service, except that the SnapManager for Virtual Infrastructure administrator runs the Restore Agent and copies the restored files to a shared location that the user has access to.

- ▶ **Limited Self-Service:** The SnapManager for Virtual Infrastructure administrator finds the backup copy within a user-specified range of backups and attaches the backed-up disks to the destination virtual machine. The users can then run Restore Agent on a destination virtual machine, browse the mounted backups, and restore the individual disk file.

For SFR, it is best that the source and destination VMs have the latest VMware Tools running.

You must also configure global notification settings before you can successfully use the SFR feature. Because the workflow uses e-mail to send out details like the Restore Agent installer link and a configuration file. A VM that has a network assigned on a distributed vSwitch cannot be selected from the SFR wizard; then you have to manually enter the VM name while creating the SFR session. It is best to have port 8043 open on the SMVI server if there are firewalls on vCenter or the guests involved in an SFR if it is a Self Service SFR. For more detailed steps on troubleshooting an SFR, refer to Appendix C., “Troubleshooting single file restore” on page 83.

5.7.3 Single file restore capability for Linux VMS

In order to perform a manual SFR on Linux, you must use the administrator assisted SFR mode. This mode will pre-attach the disks from the backups to the VM. Restore Agent is not available for Linux VMs.

The steps to perform SFR for Linux VMs are as follows:

1. The SMVI administrator creates a new SFR admin-assisted restore session and pre-mounts the disks from the backups to the virtual machine.
2. A user on the VM must perform the following actions.

Privileges: These actions require root privileges (either directly as root or through sudo).

- a. List existing disks:

```
# fdisk -l
```

- b. Rescan SCSI bus for new devices:

If your distribution ships it, there is a rescan-scsi-bus.sh script that can be run to perform the rescan. It is best to use this script for that purpose.

If you do not have access to the script, you can issue the following command:

```
# echo “- -” > /sys/class/scsi_host/host0/scan
```

Your system might contain several directories of the form /sys/class/scsi_host/host[0-9]/. If it does, you can issue the echo command for each directory.

c. List the new disks:

```
# fdisk -l
```

Compare this output against the first command to locate the new disks. This will also give you the partitions on the new disks along with the file system types.

d. Mount the new partitions:

Use the 'mount' command to mount your Linux partitions. You have to know what file system type your partitions are. Likely partition types are ext3, ext4 and reiserfs.

```
# mount -t ext3 /dev/sdb1 /mnt/disk1_partition1/
```

If your disks were constructed using LVM, you need to reconstruct your LVM disk group manually in order to view your disk data. That process is beyond this list of commands.

e. Copy data from your newly mounted backup disks:

f. Unmount the partitions when complete. Use the 'unmount' command to unmount your Linux partitions when you are finished copying the data. Because we are using the admin assisted SFR session, the virtual disks that were attached to the VM are removed automatically by the server. Unmounting the disks first will just aid in the cleanup of the system:

```
# unmount /mnt/disk1_partition1/
```

5.8 Restore process flow

The methodology used to restore a virtual machine or datastore is based on both the type of restore selected, as well as the environment on which the restore is being performed. Table 5-2 summarizes how restores are performed using SnapManager for Virtual Infrastructure.

Table 5-2 SMVI restore types

Datastore type	Restore type	Restore method	VMware ESX Credentials	VMware ESX Platforms supported
VMFS	Single file	Clone, mount, RA	No	3.5 +

Datastore type	Restore type	Restore method	VMware ESX Credentials	VMware ESX Platforms supported
VMFS	In place or out-of-place disk restore	Clone, mount, and copy	No	3.5 +
VMFS	VM	Clone, mount, and copy	No	3.5 +
VMFS	Datastore	Single file SnapRestore	No	3.5 +
NFS	Single file	FlexClone, mount, RA	No	3.5 +
NFS	Out-of-place disk restore	FlexClone, mount, and copy	No	3.5 +
NFS	VM, datastore, or in-place disk restore	Single File SnapRestore	No	3.5 +

With these differences in restore types aside, the process flow used by SnapManager for Virtual Infrastructure during a restore is as follows:

1. SMVI powers off any virtual machine that is being restored, if it is in a powered on state, unless this is for a Single File Restore.
2. Files are restored as described previously based on restore and datastore type.
3. The virtual machines are reloaded.
4. The virtual machines are reverted to the VMware snapshot that was taken at the time of backup, if the virtual machine was running during backup, thereby reverting the virtual machines to the specified point in time.
5. SMVI removes the VMware snapshot to avoid performance overhead.

At this point, the restore is complete and virtual machines can be powered on. Be aware that in SnapManager 2.0 for Virtual Infrastructure virtual machines must be registered with vCenter if they did not exist in inventory at the time of the restore. This can be accomplished by browsing the datastore on which the virtual machine resides, selecting the virtual machine folder; right-clicking the restored virtual machine's .vmx file, and selecting "Add to Inventory" from the drop-down menu.

5.9 Restore enhancements in SMVI 2.0

In the following topics, we describe the restore enhancements in SMVI 2.0.

5.9.1 Virtual disk restore

SMVI 1.x allows restore at datastore and virtual machine granularity. Restore enhancements in SMVI 2.0 extend restore granularity to one more level adding the capability to restore individual virtual disks (VMDKs) of a VM.

The virtual disk restore feature also allows disks to be restored to another datastore, referred to as out-of-place restore, which allows VM administrators to test the consistency of the disks before attaching the disk to the VM. SMVI does not attach the restored disks automatically if the VMDK is removed from the current configuration of the VM or restored to another datastore. Out-of-place restore copies the virtual disk files to the destination datastore after which users have to attach the disks manually.

For out-of-place virtual disk restores, if VMware snapshots already existed before the backup, use vmkfstools to consolidate snapshots before attaching restored out of place VMDKs:

```
vmkfstools -i "tail-end-of-disk-00000N.vmdk" "newDiskToAttach.vmdk"
*N - Notifies the number of snapshots
```

If five VMware snapshots are taken, then there will be five delta disks; hence, this operation has to be performed on disk-000005.vmdk. To learn more about consolidating VMware snapshot disks, refer to the VMware KB article “Consolidating Snapshots”, at the following website:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1007849

Attention: Consolidating snapshots can take time to complete.

5.9.2 Advanced find (more search options)

Using the Advanced Find feature, the VM administrator can search for these types of backups:

- ▶ Recent backup
- ▶ Backups taken within a specific time range
- ▶ Backups with VMware snapshot
- ▶ Backups mounted for SFR



Disaster recovery

In this chapter we provide disaster recovery information for SnapManager for Virtual Infrastructure.

6.1 SnapMirror integration

SnapMirror is a data protection feature of Data ONTAP. It mirrors a local Snapshot copy of data from the primary storage system to a secondary storage system; typically in a remote location.

SnapMirror relationships cannot be configured through SnapManager for Virtual Infrastructure, but SMVI can update an existing SnapMirror relationship on the volume underlying the datastore or virtual machine. Preferably, test the SnapMirror relationship from the storage system command line before updating through SMVI. This will aid in identifying where any potential issues might occur. In case the SnapMirror update is successful from the command line interface but fails from within SMVI, the administrator will have a better understanding of where to concentrate troubleshooting efforts.

Of further note, the destination storage must be identified within SMVI in the same manner as the relationship is configured on the storage system. For example, if a SnapMirror relationship is configured on the storage system using IP addresses rather than a DNS name, the secondary storage must be identified to SMVI by IP address also. Conversely, if identified by system name on the storage, the same name must be entered within the SMVI setup.

Because SnapManager for Virtual Infrastructure provides support for volume SnapMirror only, it is best to map one volume per datastore.

During backup creation, SnapManager for Virtual Infrastructure provides the option of updating an existing SnapMirror relationship so that every time a Snapshot is created the data is transferred to a remote storage system. Whenever the backup of a virtual machine or datastore is initiated with the SnapMirror option, the update starts as soon as the backup completes, outside of the current SnapMirror schedule. Customers can schedule SnapMirror updates on a volume outside of SnapManager for Virtual Infrastructure. For example by configuring regular SnapMirror updates on a filer outside of the SMVI schedule, you can cut down on the time required to update the mirror when SMVI runs (because the mirror is updated in the interim). However, keep in mind that the updates must be scheduled in such a way that they do not conflict with the SMVI backup.

6.1.1 SnapMirror destinations

SnapManager for Virtual Infrastructure supports one SnapMirror destination per volume. In case a SnapMirror update is selected as part of an SMVI backup on a volume with multiple destinations; the backup will fail.

If multiple SnapMirror destinations are required, use a tiered approach when configuring the SnapMirror relationships. For example, if the data must be transferred to four destinations, configure one destination from the primary storage system supported by SMVI to one destination and three additional destinations from the secondary storage through the storage system command line interface.

6.1.2 SnapMirror and deduplication

Preferably, do not use deduplication with Sync SnapMirror. Although technically it will work, the integration and scheduling of deduplication with Sync SnapMirror are complicated to implement in the type of rigorous real-world scenarios that demand synchronous replication.

When configuring volume SnapMirror and deduplication, it is important to consider the deduplication schedule and the volume SnapMirror schedule. As a best practice, start volume SnapMirror transfers of a deduplicated volume after deduplication has completed (that is, not in the middle of the deduplication process). This is to avoid sending un-deduplicated data and additional temporary metadata files over the network. If the temporary metadata files in the source volume are locked in Snapshot copies, they also consume extra space in the source and destination volumes. Volume SnapMirror performance degradation can increase with deduplicated volumes.

The scenario described previously has a direct impact on backups configured within SnapManager for Virtual Infrastructure when the SnapMirror update option has been selected. Avoid scheduling an SMVI backup with the SnapMirror update option until a time when volume deduplication is known to be complete. Although a few hours must be scheduled to ensure that this issue is avoided, the actual scheduling configuration will be data and customer dependent.

6.2 Configuring the disaster recovery standby site

In this section we provide an example of how to prepare and configure a true disaster recovery standby site in a SnapManager for Virtual Infrastructure environment. Our example assumes that the destination equipment is dedicated to a disaster recovery plan in case of primary site loss.

Detailed steps for recovery from a disaster in this type of environment are provided in Appendix A., “Recovery procedures” on page 73.

6.2.1 Primary site SMVI server

The SMVI server must be configured with the repository residing upon shared storage as discussed earlier. Preferably, take Snapshot copies of the repository's underlying volume, and establish a SnapMirror relationship on that volume to the secondary storage. SnapMirror updates on the repository volume can be scheduled as required on an individual basis based on changes within the SMVI environment. This will allow quick recovery of the repository on the secondary site in case of a failure.

SMVI server: In this scenario, configure the destination-side SMVI server to use shared storage before establishing the relationship as described in step 3 in the following topics.

6.2.2 Primary site storage

Configure the primary site storage per a normal SMVI installation with SnapMirror relationships established with the secondary storage. This scenario assumes that backups have been running within SMVI and SnapMirror relationships have been updated as part of those backups.

6.2.3 Disaster recovery site

Although there are a number of options when it comes to configuring the secondary site, here is a quick and efficient way to configure the environment for rapid recovery if the secondary site is a true "standby" site:

1. Install SMVI on the vCenter server (or other Windows server as the customer's environment dictates).
2. Configure SMVI to use the volumes on the destination side (secondary site) storage systems.
3. Enter the vCenter server and storage system IP addresses or names within the SMVI Setup window.
4. Run the `smvi servercredential set` command from the CLI if necessary.
5. Stop the SMVI service within Windows.
6. Establish the SnapMirror relationship on the underlying volume from the primary site to secondary site. (Volumes used for SMVI on the destination side storage must be used as the SnapMirror destination volumes.)

VMware snapshots

In this chapter we discuss the integration of VMware snapshots and SnapManager for Virtual Infrastructure (SMVI).

7.1 VMware snapshots and SMVI

As previously discussed, the VMware snapshot preserves the state of the virtual machine and is used by SnapManager for Virtual Infrastructure during restores to revert the virtual machine back to the backup point-in-time state. VMware snapshots can capture the entire state of the individual virtual machines, including memory state, settings state and disk state. However, the memory state is not captured during a backup initiated by SnapManager for Virtual Infrastructure by design to aid in backup performance.

By default, SnapManager for Virtual Infrastructure initiates a VMware snapshot as part of the backup process, creating a snapshot on each virtual machine undergoing backup that is in a powered-on state.

VMware snapshots can cause issues in certain environments. In case the datastore that the virtual machines reside upon experiences heavy disk I/O, VMware snapshots can take a long time to create and might eventually time out and fail. Although this occurs at a VMware level, SMVI depends upon the VMware snapshots for a backup to complete successfully.

In case this issue is encountered during VMware snapshot creation, whether initiated manually through vCenter or through SMVI, the administrator must reduce the number of concurrent VMware snaps, reduce the amount of disk I/O, or eliminate the VMware snapshots from the SMVI backup process.

7.2 Serialized snapshots

VMware advises creating and deleting VMware snapshot operations in a serial manner to reduce snapshot errors. Snapshot operations consume resources from the VMware ESX server and taking too many snapshots at a time might adversely affect the quality of service of the VMs. If various VMware ESX servers attempt snapshots on the same datastore, lock contentions might happen, which might also affect performance. In SMVI 2.0, the number of snapshots created or deleted at a time during backup is controlled by the property:

```
vmware.max.concurrent.snapshots.
```

The default value of this property is 3 and can be changed by setting it in “smvi.override” and restarting the SMVI service, meaning that by default, three snapshots will be created or deleted per datastore during backup. Setting this value higher might help with faster backups, but can increase snapshot errors as mentioned previously. Setting this value too low can cause slower backups.

7.3 Installation of VMware Tools and VM alignment

Preferably, install the latest VMware Tools on virtual machines to enable successful backups; Also, it is best to align VMs.

7.4 Reducing the number of concurrent VMware snapshots

In case a virtual machine consistency is required, it is possible to limit the number of concurrent VMware snapshots by making backup configuration changes within SnapManager for Virtual Infrastructure. This can be achieved by creating multiple jobs per datastore, one per individual virtual machine or a few virtual machines, rather than at the datastore level as a whole.

There are, however, implications when using this approach to address VMware snapshot time-out issues. First, not only does this approach create additional administrative overhead, but it also might not correct the issue (snapshot failure) if the datastore the virtual machines reside upon experiences heavy disk I/O.

Second, the number of N series Snapshot copies per volume will increase as more backup jobs are created on a given datastore, thus increasing storage overhead and reducing the amount of time a backup can be retained before reaching the maximum number of N series Snapshot copies per volume.

7.5 Reducing the amount of disk I/O

Although reducing the amount of disk I/O on a datastore can potentially alleviate VMware snapshot issues, this can be hard to achieve after datastores and the underlying storage platforms have been configured and are in use. Reducing the number of virtual machines per datastore, and therefore per volume, is not always practical or possible depending on the amount of available, unused storage.

7.6 Eliminating VMware vCenter snapshots from the backup process

SnapManager for Virtual Infrastructure provides an option to disable the taking of VMware snapshots, thereby eliminating this potentially problematic step from the backup process.

In an environment experiencing heavy disk I/O, this will greatly increase both the speed and success rate of SnapManager for Virtual Infrastructure backups. Although disabling VMware snapshots is the preferred best practice in case of backup failures because of heavy disk I/O, be aware that this results in virtual machines being crash-consistent only, because an N series Snapshot copy is taken without first quiescing the guest operating systems. Whereas most virtual machines will not suffer any adverse effects after a restore from a crash-consistent backup, take care before disabling VMware snapshots on virtual machines that contain critical applications, such as databases or e-mail servers.



Data consistency in an SMVI environment

In this chapter we address data consistency in context of SMVI backup operations.

8.1 Backup

A backup can be point-in-time consistent (often called crash-consistent) or application-consistent, depending upon the type of technologies used to perform the backup.

8.1.1 Point-in-time consistent backup

A point-in-time consistent backup is the simpler of the two types of backup because it requires minimal coordination among data components. There are two levels of point-in-time consistent backups possible, and SMVI supports both.

- ▶ Point-in-time consistent backup:

For a VMware virtual machine running on VMware ESX or on a hosted hypervisor using unbuffered host I/O, a non-quiesced VMware snapshot provides the same level of point-in-time consistency to the software running inside a guest as a Data ONTAP Snapshot copy of the backing storage for this virtual machine. However, because the VMware snapshot delta files created during the non quiesced VMware snapshot are redundant in an SMVI scenario in which VMware's guest file system consistency is not being used, SMVI uses the second method, a Data ONTAP Snapshot copy of the backing storage (without first triggering VMware snapshots of the virtual machines) to provide a point-in-time consistent backup with no guest file system consistency.

- ▶ Point-in-time consistent backup with guest file system consistency:

VMware Tools can also create quiesced snapshots of virtual machines using file system sync. The sync command comes natively with UNIX® and Linux systems, and the Windows version of VMware Tools includes an implementation of the sync driver for Windows. The sync driver provides file system consistency for the virtual machine snapshot by flushing the file system buffers and freezing I/Os while the virtual machine snapshot is being taken. By default (although an option exists to turn off taking VMware snapshots), SMVI triggers quiesced snapshots of all virtual machines selected for backup before creating the Data ONTAP Snapshot copy of the datastore to provide a point-in-time consistent backup with guest file system consistency. Perform non-quiesced and quiesced backups with VSS participation and do not use sync driver for the same reason.

Sync driver: The sync driver is not available in Windows Server 2008. at the time of this publication.

► Application-consistent backup:

With the availability of ESX 3.5 U2, Microsoft's Volume Shadow Copy Service, or VSS, was written specifically to enable third-party backup and recovery solutions to provide application-consistent backup and recovery for mission-critical Microsoft supported applications. When VSS is properly configured within the VMware virtual environment, an SMVI initiated VMware snapshot will begin the VSS process; however there are caveats when it comes to the restoration of VSS supported applications.

VSS is designed to produce fast, consistent snapshot-based online backups by coordinating backup and restore operations among business applications, file-system services, backup applications, fast-recovery solutions, and storage hardware. VSS consists of four primary components:

- Volume Shadow Copy Service: A service that coordinates various components to create consistent shadow copies of one or more volumes
- Requestor: An application that requests that a volume shadow copy be taken; a backup and restore application is an example
- Writer: A component of an application that stores persistent information about one or more volumes that participate in shadow copy synchronization; typically, this is a database application like SQL Server or Exchange Server, or a system service like Active Directory
- Provider: A component that creates and maintains the shadow copies; examples are the system provider included with the operating system and the hardware providers included with storage arrays

8.1.2 Coordinated backup process

The coordinated backup process includes freezing the data application I/O, flushing the file system cached I/O to disk, and creating a point-in-time snapshot of the data state. After the snapshot is created, file system and application I/O are resumed. The VSS restore process involves placing the data application into the restore state, passing backup metadata back to the application whose data is being restored, restoring the actual data, and signaling the data application to proceed with recovering the data that was restored.

Because SMVI (with VMware snapshots turned on) relies on VMware quiescing of virtual machines when creating backups, it is able to provide application-consistent backup through VMware VSS requester/provider components for the applications running inside the virtual machines. In fact, all SMVI backups with VMware snapshots turned on are "application-consistent." The limitations of this application-consistent backup methodology are explained in the section 8.2.2, "Application-consistent recovery" on page 63.

VSS snapshots: For VMs running Windows Server 2003 as the guest operating system, the VSS snapshots are application-consistent. For VMs running Windows Server 2008 and Windows Vista, the VSS snapshots are file-system-consistent. For more information, see the following Web site:

http://www.vmware.com/pdf/vsphere4/r40/vsp_vcb_15_u1_admin_guide.pdf

The following prerequisites are for VSS-assisted application consistency:

- ▶ Virtual machines need to run on ESX 3.5 update 2 or later, and VSS components need to be installed in VMware Tools. For SMVI 2.0, ESX 3.5 update 4 or later is required.
- ▶ Only applications that have VSS writers will have application consistency.
- ▶ All application data needs to be contained on virtual disks (VMDKs) in NFS or VMFS datastores, and not on RDMs or LUNs that are accessed using the Microsoft iSCSI Software Initiator in the guest OS. We explain this in more detail later.

8.2 Recovery

In this section we provide information about recovery processes.

8.2.1 Point-in-time consistent recovery

Because SMVI supports two levels of point-in-time consistent backup, with and without guest file system consistency, point-in-time recovery can be either guest file-system-consistent or not, depending upon the type of backup performed.

One of the advantages of using Snapshot technology to create backups is that these backups can also be used at a remote site for DR. Because SMVI includes an option to replicate backups to a remote system using an existing SnapMirror relationship, there is an opportunity for VMware vCenter Site Recovery Manager to use backups created by SMVI when a recovery plan is executed.

The N series Site Recovery Adapter (SRA) recovers the latest version of the file system at the DR site. This is done because Site Recovery Manager does not currently support recovery of VMs that are in VMware snapshot mode, which is the state of a virtual machine that is contained in an N series Snapshot copy created with the VMware consistency option turned on in the SMVI backup job. The N series SRA does this by creating a FlexClone volume of the volume without specifying a snapshot name which means that the latest version of the volume is used, which is the latest SnapMirror snapshot.

Site Recovery Manager DR failover is achieved by breaking the SnapMirror relationship and presenting the FlexVol® volume to the VMware environment. Because this only achieves a SnapMirror relationship break the data presented is the same as that in the last SnapMirror updated snapshot, which is a non-quieted VM, regardless of whether the SMVI job made a quieted snapshot prior to the SnapMirror one.

8.2.2 Application-consistent recovery

If an application-consistent backup was taken by SMVI using the built-in VMware VSS support, the recovery will be application-consistent. Because the VMware VSS components have no application recovery capability, and no VSS writer-assisted recovery is possible using current VMware technology, the only recovery mode available to the application is recovery to the point of the last backup. It is not possible to roll forward the logs, recover to a specific point in time or specific transaction, or have other enhanced recovery functionality.

8.2.3 Application consistency in combined solution environments

Today, there are two possible solutions for addressing application-consistent data protection in a VMware environment:

- ▶ SMVI, working through the VMware guest VSS stack, provides application-consistent backup and recovery for applications that have VSS writers and store their data on virtual disks (VMDKs). Recovery, in this scenario, is at the full VM level only.
- ▶ SnapDrive and application-specific SnapManager products such as SnapManager for Exchange (SME) and SnapManager for SQL (SMSQL) running in the guest OS, provide application-consistent backup and fine-grained recovery for applications whose data is stored using Microsoft iSCSI Software Initiator LUNs or RDMs.

The critical difference in the level of protection provided by each of these solutions is in the granularity of recovery. To understand the significance of this, you need to understand the concept of roll-forward recovery. Roll-forward recovery replays information stored in transaction log files to return a database to the state it was in at an exact point in time. In order to perform a roll-forward recovery, archival logging must be enabled, a full backup image of the database must be available, and there must be access to all logged files created since the last successful backup.

Because the only recovery mode available today for applications backed up using SMVI (using the VMware built-in VSS support) is recovery to the point of the last backup, customers must use the relevant SnapManager application if more fine-grained, roll-forward recovery is required.

Today, both solutions can be used together (SMVI to back up/recover the system data and a SnapDrive and application SnapManager combination to back up/recover the mission-critical application data) to get the desired level of data protection required. Customers using the two-solution approach need to be aware of a few configuration considerations:

- ▶ SMVI supports backup and recovery of virtual disks in VMFS and NFS datastores.
- ▶ The application SnapManager products support backup and recovery of applications whose data is stored on RDM LUNs or Microsoft iSCSI Software Initiator LUNs mapped to the virtual machine.
- ▶ By default, SMVI uses quiesced VMware snapshots of virtual machines to capture the consistent state of the virtual machines prior to making a Data ONTAP Snapshot copy of the backing storage. According to VMware KB article #1009073, VMware Tools are unable to create quiesced snapshots of virtual machines that have NPIV RDM LUNs or Microsoft iSCSI Software Initiator LUNs mapped to them (this often results in timeout errors during snapshot creation). Therefore, customers using the Microsoft iSCSI Software Initiator in the guest and running SMVI with VMware snapshots turned on, which is not advised, are at high risk of experiencing SMVI backup failures due to snapshot timeouts caused by the presence of Microsoft iSCSI Software Initiator LUNs mapped to the virtual machines.

VMware's general advice is to disable both VSS components and the sync driver in VMware Tools (which translates to turning off VMware snapshots for any SMVI backup jobs that include virtual machines mapped with Microsoft iSCSI Software Initiator LUNs) in environments that include both Microsoft iSCSI Software Initiator LUNs in the VM and SMVI, thereby reducing the consistency level of a virtual machine backup to point-in-time consistency. However, by using SDW/SnapManager to back up the application data on the Microsoft iSCSI Software Initiator LUNs mapped to the virtual machine, the reduction in the data consistency level of the SMVI backup has no effect on the application data.

Also for these environments, it is best to use physical mode RDM LUNs, instead of Microsoft iSCSI Software Initiator LUNs, when provisioning storage in order to get the maximum protection level from the combined SMVI and SDW/SM solution. For guest file system consistency for OS images, use VSS-assisted SMVI backups, and application-consistent backups and fine-grained recovery for application data using the SnapManager applications.

8.2.4 Supported configurations in combined solution environments

Table 8-1 lists data consistency levels for operating system components and application data in configurations that differ in the way applications inside the virtual machine store their data:

- ▶ Both OS image and application data are stored on virtual disks connected to the virtual machine, and neither SnapDrive nor SnapManager products are installed.
- ▶ The OS image is stored on a virtual disk, the application data is stored on Microsoft iSCSI Software Initiator LUNs provisioned and managed using SnapDrive, and SnapManager products might or might not be installed.
- ▶ The OS image is stored on a virtual disk, the application data is stored on physical mode RDM LUNs (instead of Microsoft iSCSI Software Initiator LUNs) provisioned and managed using SnapDrive, and SnapManager products might or might not be installed.

The following scenarios assume that the environment has:

- ▶ VMware ESX Server 3.5 update 4 or later installed.
- ▶ Virtual machines running Windows 2003 SP1 or later.
- ▶ The OS image installed on a VMDK.
- ▶ The VMware VSS components installed as a part of VMware Tools in the virtual machines.
- ▶ SMVI with VMware snapshots turned on, unless specifically noted otherwise.

Table 8-1 Supported configurations in combined solution environments

Component	OS and application data reside on VMDKs	OS is on VMDK, application data is on physical RDMs	OS is on VMDK, application data is on Microsoft iSCSI LUNs
Operating system	File system	File system	Point-in-time ^a
Applications without VSS, SnapDrive, or SnapManager	File system	No support ^b	No support
Applications with VSS writers, but no SnapDrive or SnapManager	Applications (Windows 2003 guests) File system (Windows Vista and 2008 guests) ^c	No support	No support

Component	OS and application data reside on VMDKs	OS is on VMDK, application data is on physical RDMs	OS is on VMDK, application data is on Microsoft iSCSI LUNs
Applications with SnapDrive but no SnapManager	No support	No support	No support
Applications with SnapDrive and SnapManager	No support	No support	No support

- a. VSS quiescing needs to be disabled in the presence of guest-mapped iSCSI LUNs or NPIV RDM LUNs (VMware KB article #1009073).
- b. SMVI does not back up Microsoft iSCSI LUNs or RDMs attached to the virtual machine.
- c. VMware VSS components cannot interact with VSS application writers in Windows 2008.

8.2.5 Simplifying supported configurations

For simplicity, the set of supported and preferred configurations that support application-consistent backup and recovery for VMware vSphere environments can be consolidated into three configurations that map to various kinds of environments: small, large, and mixed (see Table 8-2):

- ▶ The “small” environment configuration relies on SMVI and VMware VSS-assisted quiesced snapshots to back up OS images and applications in virtual machines, requires all data to be stored on virtual disks, and only provides very simple application recovery.
- ▶ The “large” environment configuration relies on SMVI and VMware VSS-assisted quiesced snapshots to back up OS images and SnapManager products to back up and recover applications requires application data to be stored on mapped LUNs and provides enhanced, fine-grained recovery as supported by SnapManager.
- ▶ The “mixed” environment configuration, which depends on future application SnapManager support for NFS and VMFS datastores (VMDKs), is essentially the same as the “large” environment configuration except that the application data is stored on VMDKs.

Table 8-2 Simplifying supported configurations

Component	Small	Large	Mixed
OS data storage	Application data and OS both stored on VMDKs connected to VM	OS is stored on VMDKs connected to the VM	OS is stored on VMDKs connected to the VM
Application data storage	Application data and OS both stored on VMDKs connected to VM	Application data OS stored on mapped LUN	Application data OS stored on VMDKs
OS backup method	Application data on OS backed up together using SMVI and VMware VSS components	OS is backed up using SMVI	OS is backed up using SMVI
Application backup method	Application data on OS backed up together using SMVI and VMware VSS components	Applications are backed up using SnapManager	Applications are backed up using SnapManager
Application recovery	Only simple recovery using SMVI to the time of the last backup is possible	Enhanced application recovery is possible through application SnapManager	Enhanced application recovery is possible through application SnapManager
Preferred usage	<ul style="list-style-type: none"> ► Enterprise application-consistent backup/recovery for infrastructure applications ► Disaster recovery 	Enterprise application-consistent backup/recovery for infrastructure applications	Enterprise application-consistent backup/recovery for infrastructure applications

Summary of best practices

SnapManager for Virtual Infrastructure 2.0 provides a rich feature set that allows IT organizations to take advantage of N series Snapshot and SnapMirror technologies to provide fast, space-efficient disk based backups in a VMware environment with N series storage, while placing minimal overhead on the associated virtual infrastructure. Typical cases and examples in this book can help administrators get the most out of SnapManager for Virtual Infrastructure deployments.

In this chapter we provide a list of N series best practices for SnapManager for Virtual Infrastructure 2.0.

Here are N series best practices for SnapManager for Virtual Infrastructure 2.0 “at a glance.” Details are given throughout this book:

- ▶ Install SMVI on the vCenter server to reduce the impact of network disruptions.
- ▶ Install SMVI configuration files on a SAN or NAS device, thereby providing rapid recovery capability in case the SMVI server fails.
- ▶ In environments with multiple vCenter servers, install one SMVI server per vCenter server to enable all backups to run as scheduled and ease administration overhead.
- ▶ When adding storage as part of the SMVI setup, configure a non-root user on the storage.
- ▶ Enable Secure Sockets Layers (SSL) on all storage systems identified to SMVI, thereby ensuring passwords are encrypted when transmitted across the network.
- ▶ Configure VMDKs containing transient and temporary data as independent persistent disks to exclude the related datastore from the SMVI backup.
- ▶ Place all transient and temporary guest operating system data from multiple virtual machines in the same datastore configured with independent persistent disks with no other data type or VMDK residing upon it to avoid performing Snapshot copies against the underlying volume.
- ▶ Monitor the number of migrations performed in an ESX cluster configured with fully automated Distributed Resource Scheduling due to backup implications during active migrations.
- ▶ If the vCenter server is running within a virtual machine and will be backed up by SMVI, install the associated vCenter database on a physical system to avoid time-out issues during a VMware vCenter snapshot.
- ▶ Perform datastore backups rather than individual virtual machine backups to reduce the number of N series Snapshot copies performed against a volume.
- ▶ Use Retention Policies to meet service-level agreements (SLAs) and limit the number of N series Snapshot copies maintained against a volume.
- ▶ When integrating SMVI with an existing SnapMirror relationship, no more than one SnapMirror destination per volume can be configured.
- ▶ Install the latest VMware Tools on virtual machines to enable successful backups. Also, align VMs.
- ▶ For SFR, ensure that the source and destination VMs have the latest VMware Tools running.

- ▶ Configure global notification settings before using the SFR feature, Because the workflow uses e-mail to send out details such as the Restore Agent installer link and a configuration file.
- ▶ In case you have a normal disk and an independent disk for backup on the same datastore, this datastore is always included for backup irrespective of the “include datastore with independent disk” option. Designate a separate datastore exclusively for swap data or system data.

Recovery procedures

In this appendix we provide procedures showing how to recover the loss of an SMVI server, the loss of a VMware ESX host, and the loss of the primary site.

Loss of SMVI server

In the following topics we show how to recover the loss of an SMVI server when installed on a local disk of shared devices.

Recovering from loss of SMVI server when installed on local disk

Take the following steps to protect SMVI when installed on a local disk:

1. Back up the following two directories:
 %PROGRAMFILES%\IBM\SMVI\Server\repository\
 %PROGRAMFILES%\IBM\SMVI\Server\etc
2. In case of failure of the SMVI server, install SMVI on a replacement system.
3. Stop the SMVI service
4. Copy the backed-up files listed in step 1 to their original location.
5. Restart the SMVI service within Windows.
6. Start the SMVI GUI and enter **Setup**.
7. Enter the vCenter server name or IP address.
8. Enter the necessary storage information.
9. Select **Restore** from within the SMVI GUI.

Backup jobs will be listed and scheduled jobs will run successfully.

Because the backups.xml and scheduledbackups.xml files are updated every time a backup completes, is renamed, or is deleted, these files require frequent backup in order to reflect all changes and enable all backups to be restored. Although installing SMVI on local disk is supported, it is best to configure SMVI on shared storage, as described in the following topics.

Recovering from loss of SMVI server when installed on a shared device

Follow this procedure to recover from a loss of the SMVI server when installed on a shared device:

1. In case of failure of the SMVI server, install SMVI on another system with access to the shared device.

2. After reinstalling SMVI, update the configuration files identifying the location of the repository and crash folders on the shared disk. The following two files will need to be edited:

```
%PROGRAMFILES%\IBM\SMVI\server\etc\smvi.config  
%PROGRAMFILES%\IBM\SMVI\server\etc\smvi.override
```

3. Change two parameters within the smvi.config file to reflect the path to the shared storage (the H drive reflects the location of the shared device):

```
smvi.repository.path=H:\\IBM\\SMVI\\server\\repository  
flow.persistence.embedded.storageLocation=H:\\IBM\\SMVI\\server\\  
cras h
```

4. Change one parameter within the smvi.override file to reflect the path to the shared storage (this parameter will have the C: entered after a default installation):

```
credential.persistence.file.path=H:\\IBM\\SMVI\\server\\etc\\cred
```

5. Restart the SMVI service within Windows.
6. Run the SMVI GUI.

All backups will be visible and restores will be possible.

Loss of a VMware ESX host

In the following topics we show how to recover the loss of a VMware ESX server as stand-alone host or within a cluster.

Within a VMware ESX cluster

SMVI communicates with vCenter and the N series storage systems rather than individual VMware ESX hosts when running a backup. As a result, the loss of a VMware ESX host within a cluster results in minimal interruption.

All VMware ESX hosts within a cluster must be configured with access to the same shared storage, so that if a particular host fails, the datastore will still be available and visible within vCenter, as will the VMs residing on the datastore, whether they are powered on or not. Follow standard VMware troubleshooting to resolve the issue with the VMware ESX host. All backups will continue.

Stand-alone ESX host

When an ESX host fails in a standalone environment, the datastores are no longer visible within vCenter. As a result, all scheduled SMVI backups against the datastores fail.

Datastores: Users will not be able to run manual backup jobs because the datastores will not be visible.

The datastores in question must be mapped to an available VMware ESX host managed by vCenter. If necessary, another host can be imported into vCenter. After the datastores have been mapped to the host, rescan HBAs and VMFS volumes within vCenter. The existing datastores and the VMs that reside upon them will now be visible.

Run the **smvi discover datastores** command from the SMVI CLI. You might also need to refresh the backup window within SMVI (assuming that you are in it) by exiting to another window and then re-entering the backup window. Scheduled jobs can now run against the datastores in question, and manual backup jobs can be configured against them.

Primary site loss recovery

Follow these steps to recover from the loss of the entire primary site, specifically, the SMVI server, VMware ESX servers, and the primary storage systems,:

1. Prepare the disaster recovery standby site as detailed in 6.2, “Configuring the disaster recovery standby site” on page 53.
2. After the loss of the primary site, mount the datastores using the replicated datastores on the secondary storage system:
 - a. Break the SnapMirror relationship from the storage system CLI.
 - b. Bring online the SnapMirror destination volumes on which the datastores reside.
 - c. Map the LUNs from the replicated volumes to the secondary ESX servers.

Attention: In VMware ESX 3.X, when the LUN from the snapshot is presented to the new VMware ESX host at the DR site, the VMFS datastore might not appear, although the LUN does. In such a case, the LUN needs to be resigatured as described in the KB article 33990.

- d. Rescan HBAs and VMFS from within vCenter.
 - e. After the replicated datastores are listed within vCenter (each datastore name is in the format "snap-00002-<original-datastore-name>"), right-click each .vmx file and register the VMs with vCenter.
3. Manually edit the SMVI configuration files to point to the secondary storage:
 - a. Use **ssh** to connect to the VMware ESX servers to which the replicated datastores are mounted.
 - b. Change the directory to /vmfs/volumes and run the **ls -l** command to list the datastores and their UUIDs.
 - c. Make note of the UUID for the newly mounted replicated datastores.
 - d. Open the %PROGRAMFILES%\IBM\SMVI\Server\repository\backups.xml file in order to edit the text.
 - e. Search and replace all occurrences of the old UUID with the new UUID as noted in a previous step.
 - f. Search and replace all occurrences of the old datastore names with the new datastore names.
 - g. Search and replace all occurrences of the primary storage system name with the name of the secondary storage system (or IP address if used).
 - h. Save the modified file.
 - i. Repeat all modifications within the %PROGRAMFILES%\IBM\SMVI\Server\Repository\scheduledjobs.xml file, being sure to save a copy of the file before making any changes.
4. Start the SMVI service.
5. After the service has restarted, start the GUI and select the **Restore** window.
6. The modified datastore name with the new UUID will be listed in the left window pane and backups taken on the primary site storage system will be listed and available for restore on the secondary storage system.

Sample scripts

In this appendix we provide sample scripts that you might find useful for your individual implementation.

Sample environment variable

Example B-1 shows sample environment variables.

Example B-1 Environment variables

```
BACKUP_NAME=My Backup
BACKUP_DATE=20081218
BACKUP_TIME=090332
BACKUP_PHASE=POST_BACKUP
VIRTUAL_MACHINES=3
VIRTUAL_MACHINE.1=VM
1|564d6769-f07d-6e3b-68b1-f3c29ba03a9a|POWERED_ON||true|10.0.4.2
VIRTUAL_MACHINE.2=VM
2|564d6769-f07d-6e3b-68b1-1234567890ab|POWERED_ON|true
VIRTUAL_MACHINE.3=VM
3|564d6769-f07d-6e3b-68b1-ba9876543210|POWERED_OFF|false
STORAGE_SNAPSHOTS=2
STORAGE_SNAPSHOT.1=filer2:/vol/smvi_vol_1:smvi_My_Backup_recent
STORAGE_SNAPSHOT.2=filer2:/vol/smvi_vol_2:smvi_My_Backup_recent
```

Displaying environment variables during backup phases

Create a .bat file as shown in Example B-2 to display all environment variables during various backup phases.

Example B-2 Display variables

```
echo "=====
set >> test.txt
echo "=====
```

Sample SMVI SnapVault script

The following steps outline a sample SnapVault® (see Example B-3) script:

1. From the command line on an N series storage system, create a new role for the SMVI script:

```
useradmin role add limited-sv-role -a
api-snapvault-secondary-initiate-incremental-transfer,login
http-admin
```

2. Create a user group that uses the previous role:
 `useradmin group add limited-sv-group -r limited-sv-role`
3. Create the actual user;
 `useradmin user add limited-smvi-user -g limited-sv-group`
4. Set the users password:
 `passwd limited-sv-user password`
 Now you have a user who can only call the SnapVault update API.
5. Install the SDK onto the SMVI server.
6. Build your update script and save it in the
 `C:\Program Files\IBM\SMVI\server\scripts` directory.

Example B-3 SnapVault sample script

```
if %BACKUP_PHASE% == PRE_BACKUP goto doSNAP
if %BACKUP_PHASE% == POST_BACKUP goto doSV
goto ende
:doSV
chdir "c:\Program Files\IBM\ontapi"
apitest.exe torfiler3 limited-sv-user smvlocks
snapvault-secondary-initiate-incremental-transfer
primary-snapshot smvi_weeklyBlock1_recent secondary-path
/vol/vmblock1vault/vmblock1
goto ende
:doSNAP
chdir "c:\Program Files\IBM\ontapi"
apitest.exe torfiler3 limited-sv-user smvlocks
snapvault-secondary-initiate-snapshot-create
schedule-name smvi_weeklyvault volume-name vmblock1vault
goto ende
:ende
EXIT /b 0
```

Troubleshooting single file restore

In this appendix we provide procedures for troubleshooting single file restores.

Use the following information to troubleshoot SFR issues:

- ▶ SFR only works with backups performed by SMVI 2.0 or later.
- ▶ You must have flex-clone license to restore files from VMs stored on NFS datastore.
- ▶ Because the SMVI server uses a custom SSL certificate, you might see a warning in your browser while downloading the Restore Agent installer.
- ▶ VMware ESX 3.5 supports only eight NFS datastores by default it can be increased to 32.
- ▶ SFR is not supported for IDE disks. If the VM has SCSI and IDE disks, you can use SFR for SCSI disks only. SFR is also not supported on dynamic disks
- ▶ Here is a checklist for Restore Agent installer issues:
 - On Windows 2008 R2, .NET framework is a feature of the operating system and must be enabled before Restore Agent can be installed.
 - You must run the RA installer as a user who is a member of the administrators group.
- ▶ Here is a checklist for disk mounting issues:
 - RA only supports the first partition on a virtual disk. For other partitions, you have to manually use the Disk Management or diskpart tool.
 - RA only supports NTFS/FAT-formatted partitions on a disk.
 - SMVI cannot add new SCSI controllers to power on a destination VM. Therefore the number of disks that can be attached to a VM are restricted by the available free slots on existing SCSI controllers. A VM can have a maximum of 60 disks (15 slots on 4 SCSI controllers).
 - The automount feature is by default enabled for Windows XP and Vista. Therefore on VMs running these operating systems, you might see that drive letters are preassigned to all the disks (and partitions).
- ▶ There needs to be network connectivity between the target VM and vCenter for a Self Service SFR
- ▶ Here is a checklist for VM startup issues:
 - In certain VMware operation failures, SMVI might not be able to detach all disks from a VM after the SFR session is expired or deleted. In such cases, the VM startup issues might be due to the stale disks that are left attached to the VM. In most cases, such stale disks can be removed by powering off the VM and then editing the VM settings. Double-check the VMDK path before you remove these disks from a VM.

- If you are unable to power on the destination VM after using it for SFR, make sure that there are no stale disks left on the destination VM. Typically this happens for the following conditions:
 - A hosted process on an VMware ESX server crashes while attaching a disk.
 - An attach operation fails to complete due to VMware error.

Archived

SMVI installation steps

In this appendix we provide the basic installation steps for SMVI 2.0.

In this instance SMVI 2.0 has been installed on a stand alone (not into a VM) Windows 2003 SP2 64-bit server. vCenter has been installed locally previously to the SMVI installation.

Possible port conflict: In case you have installed other management products, there is a possible conflict of using http port 8080 on your server, because port 8080 is often used by many applications.

Follow these steps for the installation:

1. Start the installable .exe file. The welcome window appears (Figure D-1).



Figure D-1 Install welcome window

2. Confirm the license agreement if you agree (Figure D-2)

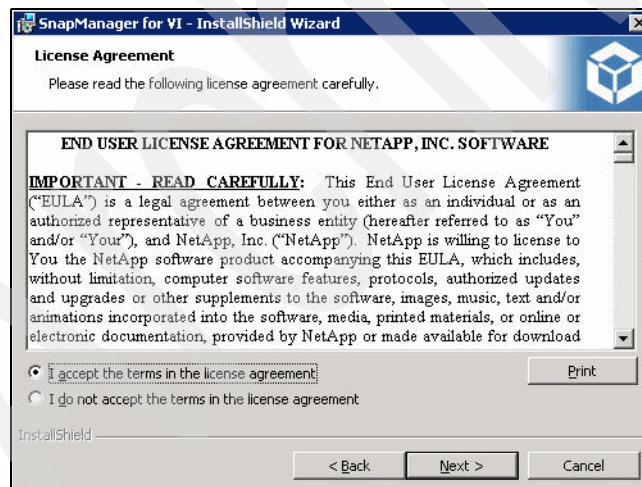


Figure D-2 License agreement

3. Select your destination folder (see Figure D-3).




Figure D-3 Destination folder

4. Select the feature that you want to install (Figure D-4).

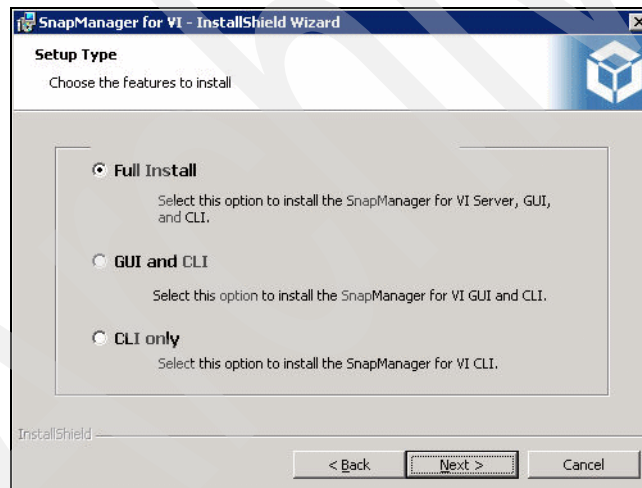


Figure D-4 Select features

5. Confirm your installation needs (Figure D-5).

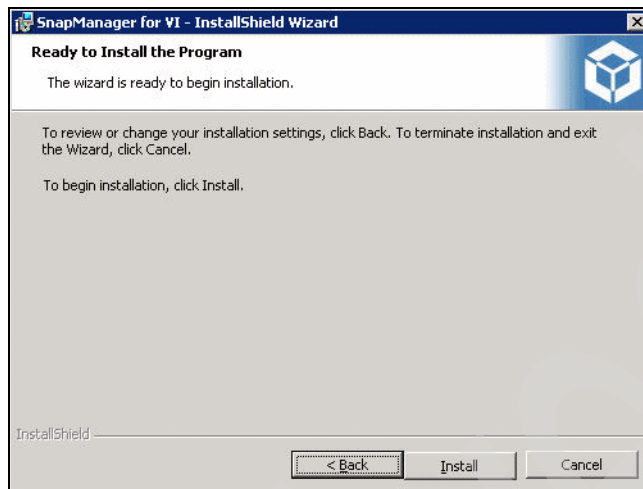


Figure D-5 Confirm installation

6. Observe the status indicator, which provides information about the current progress (Figure D-6).




Figure D-6 Installation in progress

7. When the installation is completed, you see a completion window (Figure D-7).

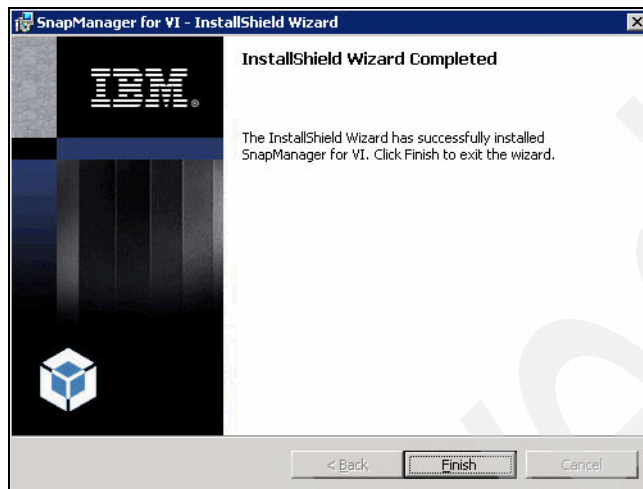


Figure D-7 Installation complete

Abbreviations and acronyms

API	Application Programming Interface
DR	Disaster Recovery
HA	High Availability
HBA	Host Bus Adapter
IBM	International Business Machines Corporation
ITSO	International Technical Support Organization
MSCS	Microsoft Cluster Server
RDM	Raw Device Mapping
SFR	Single File Restore
SLA	Service level agreements
SMVI	SnapManager for Virtual Infrastructure
SRA	Site Recovery Adapter
UI	User Interface
vCenter	Virtual Center
VM	Virtual Machine
VMFS	Virtual Machine File System

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this book.

IBM Redbooks publications

For information about ordering these publications, see “How to get Redbooks publications” on page 96. Note that certain documents referenced here might be available in softcopy only:

- ▶ *IBM System Storage N series*, SG24-7129
- ▶ *IBM System Storage N series and VMware VSphere Storage Best Practices*, SG24-7871-00
- ▶ *IBM System Storage N series MetroCluster*, REDP-4259
- ▶ *IBM System Storage N Series SnapMirror*, SG24-7260
- ▶ *IBM System Storage N series with Microsoft Clustering*, SG24-7671
- ▶ *IBM System Storage N series with VMware ESX Server*, SG24-7636

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM System Storage N series Data ONTAP 7.3 Active/Active Configuration Guide*, GC27-2208-03
- ▶ *IBM System Storage N series Data ONTAP 7.3 Block Access Management Guide for iSCSI and FCP*, GC52-1282-04
- ▶ *IBM System Storage N series Data ONTAP 7.3 Data Protection Online Backup and Recovery Guide*, GC27-2204-06
- ▶ *IBM System Storage N series Data ONTAP 7.3 File Access and Protocols Management Guide*, GC27-2207-04
- ▶ *IBM System Storage N series Data ONTAP 7.3 Network Management Guide*, GC52-1280-05
- ▶ *IBM System Storage N series Data ONTAP 7.3 System Administration Guide*, GC52-1279-04
- ▶ *IBM System Storage N series SnapManager 2.0 for Virtual Infrastructure Installation and Administration Guide*, GC53-1145-02

Online resources

The following websites are also relevant as further information sources:

- ▶ IBM System Storage NAS solutions:
http://www-03.ibm.com/systems/storage/network/?cm_re=masthead_-_products_-_stg-nas
- ▶ IBM System Storage N series and TotalStorage NAS interoperability matrixes:
<http://www-03.ibm.com/systems/storage/nas/interophome.html>
- ▶ Support for IBM System Storage and TotalStorage products:
<http://www-947.ibm.com/systems/support/>
- ▶ Support for Data ONTAP:
<http://www-304.ibm.com/jct01004c/systems/support/supportsite.wss/supportresources?brandind=5000029&familyind=5329797&taskind=1>
- ▶ VMware documentation:
<http://www.vmware.com/support/pubs/>

How to get Redbooks publications

You can search for, view, or download Redbooks publications, Redpapers publications, Technotes, draft publications, and Additional materials, as well as order hardcopy Redbooks publications, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

Index

A

advanced data management capabilities. 10
Advanced Find option 6
array-based Snapshot copies 10
ASUP logging 6

B

backup wizard 6

C

Center Server 4 15
cluster 10
cluster's load balance 23

D

Data ONTAP 14
Datastore 22
datastores 5, 9–10, 12–13, 20, 22–23
deduplication 10
disk I/O performance 11
disk performance 11
DRS 15, 23
DRS clusters 12

E

ESX Server 8
Ethernet 11

F

Fibre Channel 11
FlexClone 12, 15
Fully Automated 23
Fully Automated configuration 23
Fully Automated mode 23

G

GUI 22

H

hardware utilization 8

VMware vSp 23
higher availability 9
HTTPS 13

I

I/O load 10
I/O queue 11
IBM System Storage N series SnapManager 1
iSCSI 9
iSCSI protocols 18

L

Linux 8
LUNClone 15

M

mainframe-class virtualization software platform 8
Microsoft Cluster Server (MSCS) 16
MSCS 11–12

N

N series 12, 20
N series interoperability matrix 15
N series storage system 19
NFS 14–15
NFS datastore 18
NFS datastore mounts 15
NFS servers 10–11
NFS Single File Restore operations 15

O

ONTAP 7.2.x 14
optimizing hardware usage 9

P

P2V 11
pagefiles 20
Partially Automated 23
Partition Configuration 22
performance monitoring 10
per-port storage 10

physical-to-virtual-machine 11

R

Raw Device Mapping 11

RDMs 12

Redbooks publications website 96

Contact us xi

S

SCSI 12

secondary storage 19

Secure Sockets Layer (SSL) 19

server consolidation 8

service-level agreement 13

shared datastores 11

shared storage 9

Single File Restore (SFR) 3

SMV 20

SMVI 1, 3, 13–16, 18–19, 21, 23

SMVI 2.0 6

SMVI repository 16

SMVI server 21

SMVI setup 22

SMVI vCenter connection 22

SMVI versions 1.2x 15

SnapDrive 12

SnapManager 14, 20

SnapManager 2.0 21

SnapManager Client (GUI/CLI) 16

SnapManager for Virtual Infrastructure 7, 11–12,
14, 16, 18–19, 21, 23

SnapManager for Virtual Infrastructure environment
17

SnapManager Server 16

SnapMirror 12, 14

SnapMirror destination 14

SnapMirror relationship 12

SnapRestore 14

Snapshot copy 20

snapshots 20

spanning VMFS volumes 10

swap file 20

System x 8

T

thin provisioning 12

V

vCenter 20, 23

vCenter server 13, 16, 21–22

vCenter service 16

vFiler unit support 6

vFiler units 14

virtual data center 9

virtual disk 4

virtual disks 10

virtual machine 11, 20

virtual machine performance 23

virtual machines 8–9, 12, 20, 23

VM 11

VMDK level 12

VMDK restore 3

VMDKs 20

VMFS 9–11

VMFS datastore 11

VMFS datastores 9, 15

VMotion 9, 23

VMs 10

VMware 5

VMware Distributed Resource Scheduler 23

VMware ESX 9, 11, 18

VMware ESX 3.5 15

VMware ESX 4 15

VMware ESX 4.0 22

VMware ESX cluster 10

VMware ESX clusters 12

VMware ESX hosts 15, 21

VMware ESX Server 8–9

VMware ESXi 3.5 15

VMware ESXi 4 15

VMware snapshots 6

VMware vCenter Server 9

VMware vCenter snapshot 20

VMware Virtual Infrastructure 1

VMware Virtual Infrastructure 3 15

VMware Virtual Machine File System (VMFS) 9

VMware VirtualCenter 2.5 15

VMware virtualization 8

VMware virtualization technologies 10

VMware virtualization technology 9

VMware vSphere 7, 10–11, 14

VMware vSphere 4 15

VMware vSphere administration 12

VMware vSphere features 9, 12

VMware Workstation 8

W

Windows Server 2003 8

X

XML 13

Archived

SnapManager 2.0 for Virtual Infrastructure Best Practices

(0.2" spine)
0.17" <-> 0.473"
90 <-> 249 pages



SnapManager 2.0 for Virtual Infrastructure Best Practices

Simplifying backup and recovery of a Virtual Infrastructure

This IBM Redbooks publication provides best practices for SnapManager for Virtual Infrastructure 2.0 (SMVI).

Using Snap Manager for Virtual Infrastructure installation

We address the resource utilization issues typically found within virtual environments by leveraging the underlying Snapshot technology, which enables you to create point-in-time copies of your virtual machines or entire data stores and then restore from these backup copies at any level of granularity, datastore, VM, disk (VMDK), or guest file, simply and quickly when required.

Configuring in a clustered environment

In addition, we provide best practices for protecting the SMVI server and recovering in case of a disaster. Furthermore, we explain the seamless integration of N series storage solutions, including MetroCluster, so customers can use storage and virtualization technologies to create dynamic infrastructures that can create tremendous business value.

The reader of this book will gain a deep understanding of how to implement SnapManager for Virtual Infrastructure in VMware vSphere environments.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks