



Steve Garraway
Jon Tate

IBM SAN Volume Controller and Brocade Disaster Recovery Solutions for VMware

Introduction

In this IBM® Redpapers™ publication, we describe various scenarios for the recovery of virtual machines should a disaster befall your data center. The solutions described rely on the integration of the IBM System Storage™ SAN Volume Controller (SVC) with VMware vCenter Site Recovery Manager (SRM) and the appropriate Brocade SAN switch and routing technology.

The type of solution used depends on the level of protection required and the time available to perform recovery, so the appropriate scenario should be chosen according to your needs. The first of these requirements is known as the *Recovery Point Objective* (RPO), while the latter is known as the *Recovery Time Objective* (RTO).

For background information about Business Continuity and Disaster Recovery, refer to *Using the SVC for Business Continuity*, SG24-7371.

For details about SVC advanced copy services, refer to *SAN Volume Controller V4.3.0 Advanced Copy Services*, SG24-7574.

Requirement

Given the growing dependency on server virtualization solutions, companies need to ensure that their business continuity solution fully supports their virtualized infrastructure. Due to the dynamic nature of most virtualized server environments, any recovery solution needs to be automated and integrated with the production environment to ensure it remains valid.

Common scenarios

We present the three most commonly used replication scenarios:

- ▶ Replication between two storage subsystems at the same location
- ▶ Replication between two storage subsystems at different locations within a metropolitan area
- ▶ Replication between two storage subsystems at different locations outside of the metropolitan area

Single site

Replication of data within a single site is typically used to protect against the failure of a single disk subsystem by mirroring the data to another disk subsystem at the same location. This can be achieved with a single SVC cluster by defining a different MDisk Group on each disk subsystem and using *VDisk Mirroring* to maintain a synchronous copy of the VDisk on both disk subsystems.

VDisk Mirroring is a no cost feature of the SVC, and is also transparent to the host. If one of the two MDisk Groups goes offline due to a major failure of the relevant disk subsystem, the host will continue without any impact. Once the failed subsystem has been repaired, the mirrors must be re-synchronized. For more information about VDisk Mirroring, refer to Chapter 7, “VDisk Mirroring”, in *SAN Volume Controller V4.3.0 Advanced Copy Services*, SG24-7574.

Figure 1 shows a simplified view of the components needed to provide single site disk subsystem redundancy.

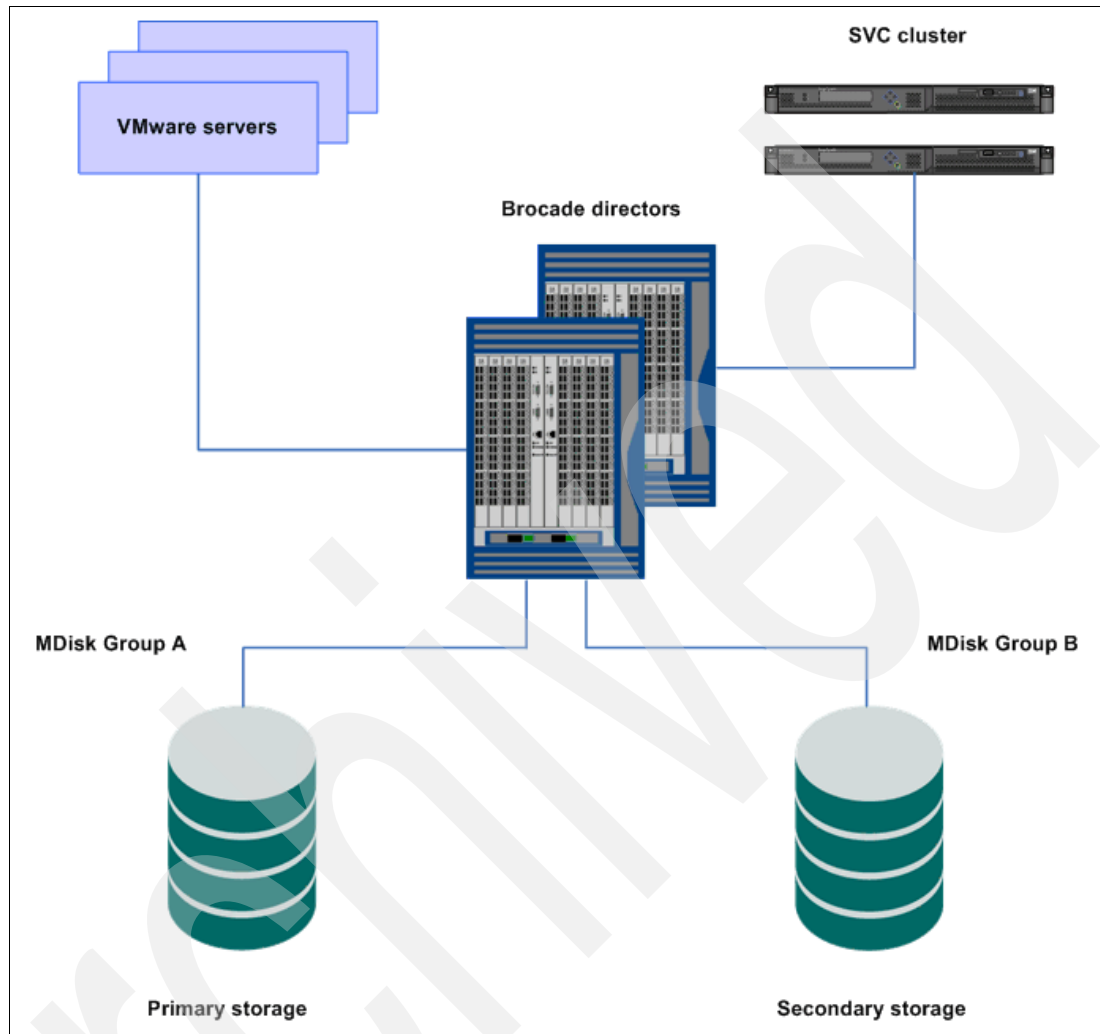


Figure 1 Single site redundancy using VDisk Mirroring

This solution only protects against the loss of one of the storage subsystems.

Dual site: metropolitan area

If a greater level of protection is required, such as against fire or flooding of the primary site, then it becomes necessary to locate the secondary storage at another location. If this is within the same metropolitan area, then it is possible to use the SVC *Metro Mirror* feature to maintain a synchronous copy of the data. This enables mirroring of VDisks between SVC clusters in the local and remote sites.

Note: VMware refers to the primary site as the *protected site*, and the secondary site as the *recovery site*.

Since Metro Mirror is a synchronous process, so updates are committed to both the primary (local) and secondary (remote) VDisks before host I/O completes. This increases response time slightly, depending on the distance and type of extension technology used.

To use Metro Mirror, it is necessary to purchase a *Metro and Global Mirror* feature license of sufficient capacity to cover the VDisks you will be mirroring. It is also necessary to have additional server and SAN resources at the remote site, as it should be assumed that all resources, including the storage, may be lost at the primary site. For more information about Metro Mirror, refer to *SAN Volume Controller V4.3.0 Advanced Copy Services, SG24-7574*.

ISL connectivity

If dark fiber is available, long wave SFPs can be used to provide ISLs between the local and remote SAN switches, as shown in Figure 2, up to a distance of 20 miles.

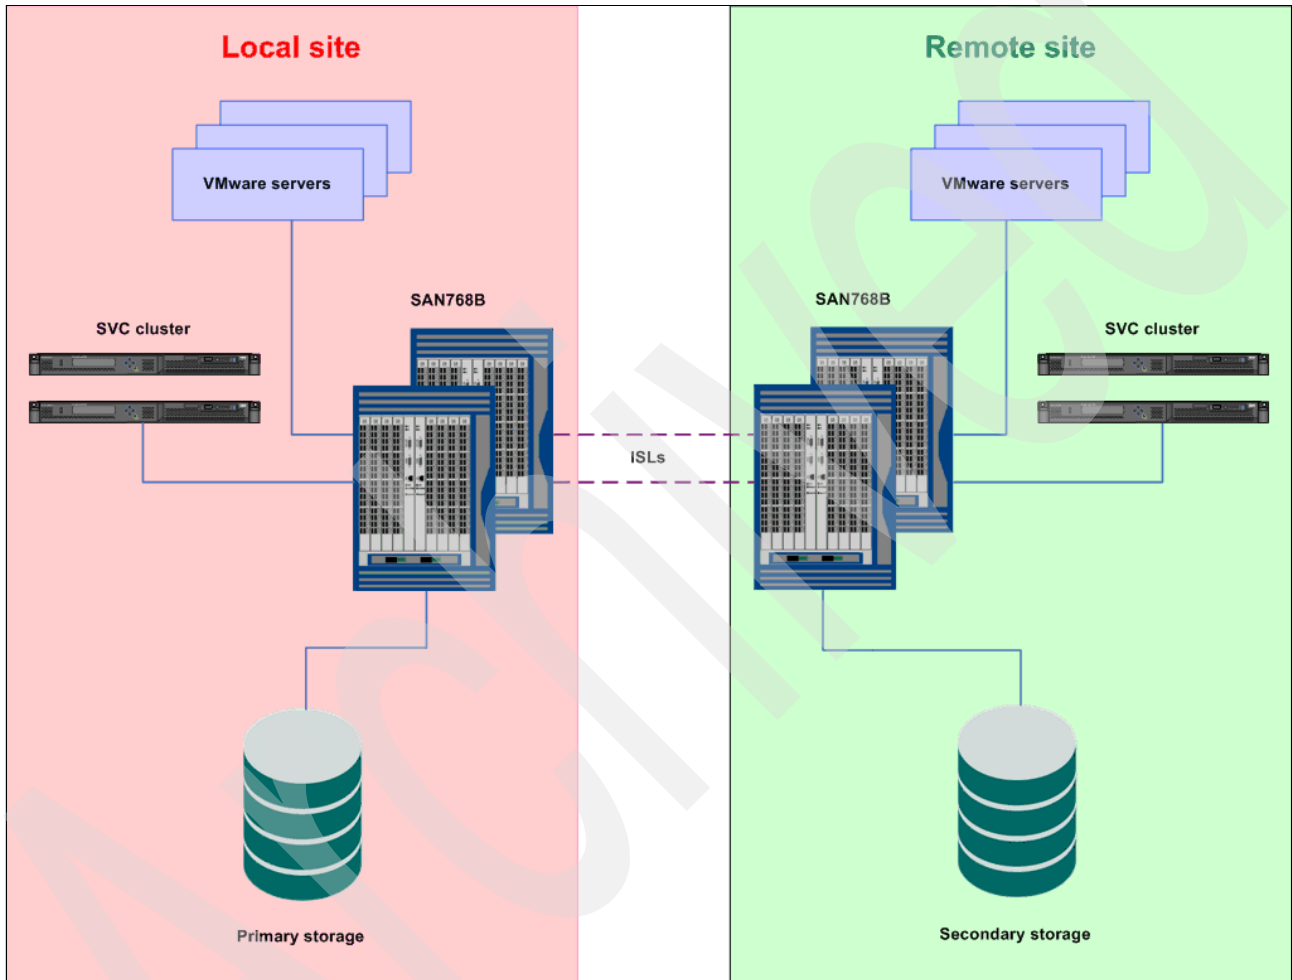


Figure 2 Dual site using long wave ISLs for Metro Mirror

Router connectivity

If dark fiber is not available, then SAN router technology is required to join the fabric over an IP network, as shown in Figure 3. This could either be a routing blade in the IBM System Storage SAN768B director, or a separate device, such as the IBM System Storage SAN04B-R router.

The round-trip latency between sites cannot exceed 80 ms with SAN routers, which should easily be achieved within a metropolitan area.

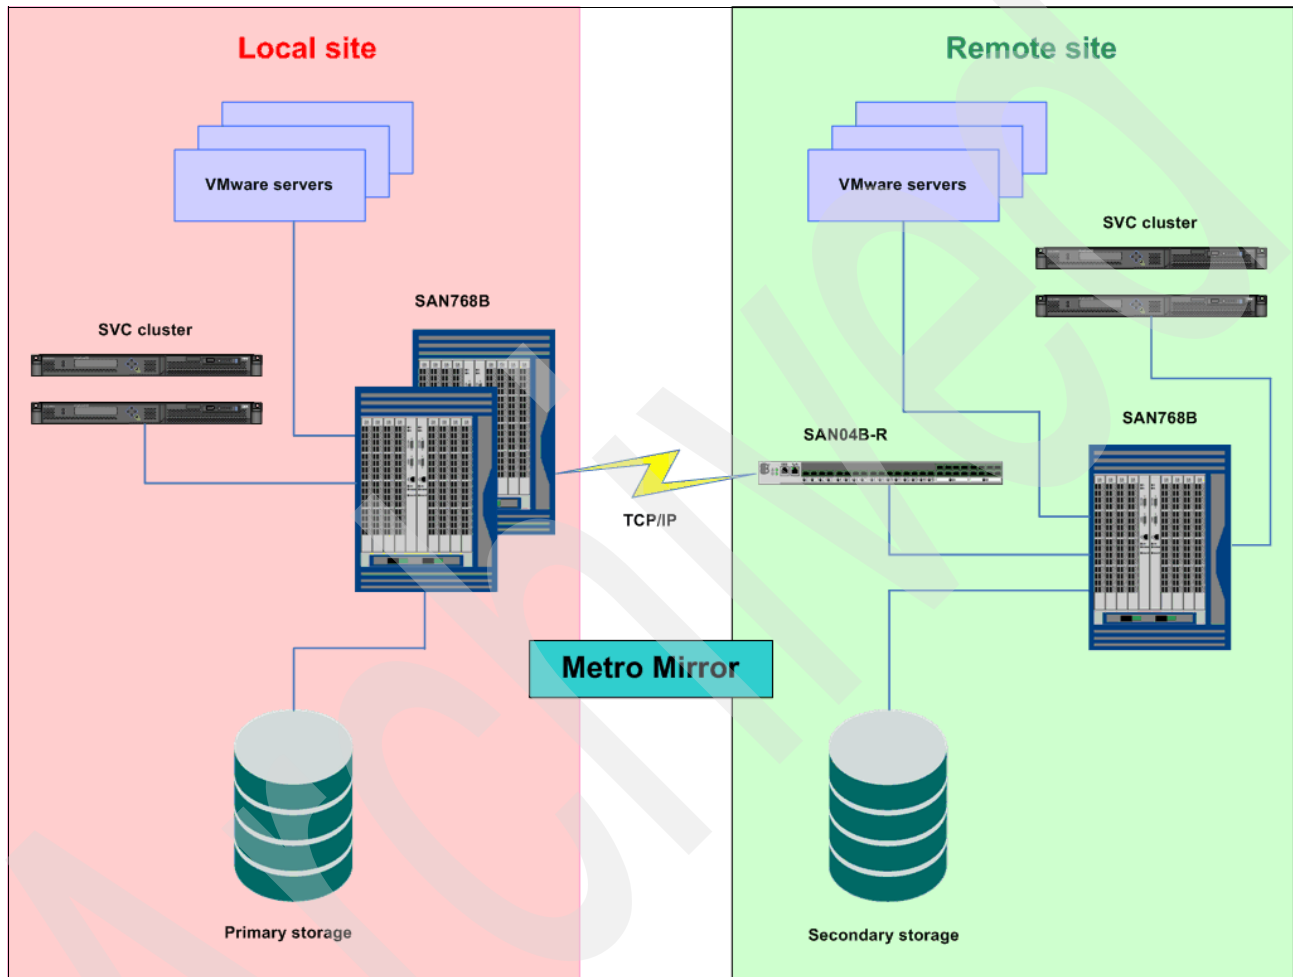


Figure 3 Dual site using routing for Metro Mirror

Details about the IBM System Storage SAN b-type switches and routers can be found in the following IBM Redbooks® publications:

- ▶ *Implementing an IBM/Brocade SAN with 8 Gbps Directors and Switches*, SG24-6116
- ▶ *IBM System Storage/Brocade Multiprotocol Routing: An Introduction and Implementation*, SG24-7544

Dual site: extended distance

If an even greater level of protection is required, such as against an earthquake or tsunami hitting the area of the primary site, then it becomes necessary to locate the secondary storage at a much greater distance. This can be achieved by using the SVC *Global Mirror* feature to maintain an asynchronous copy of the data. This again enables mirroring of VDisks between SVC clusters in the local and remote sites, as shown in Figure 4.

Since Global Mirror is an asynchronous process, the host I/O completes after the primary (local) VDisk has been updated. The update is then queued to the remote SVC cluster to update the secondary VDisk. This prevents the host from seeing any additional response time due to delays on the remote link, but means that the remote copy is always slightly out of date. To maintain integrity, the SVC ensures that all remote updates are sent in the order they are received from the host.

To use Global Mirror, it is necessary to purchase a *Metro and Global Mirror* feature license of sufficient capacity to cover the VDisks you will be mirroring. Again, it is also necessary to have additional server and SAN resources at the remote site. For more information about Global Mirror, refer to *SAN Volume Controller V4.3.0 Advanced Copy Services*, SG24-7574.

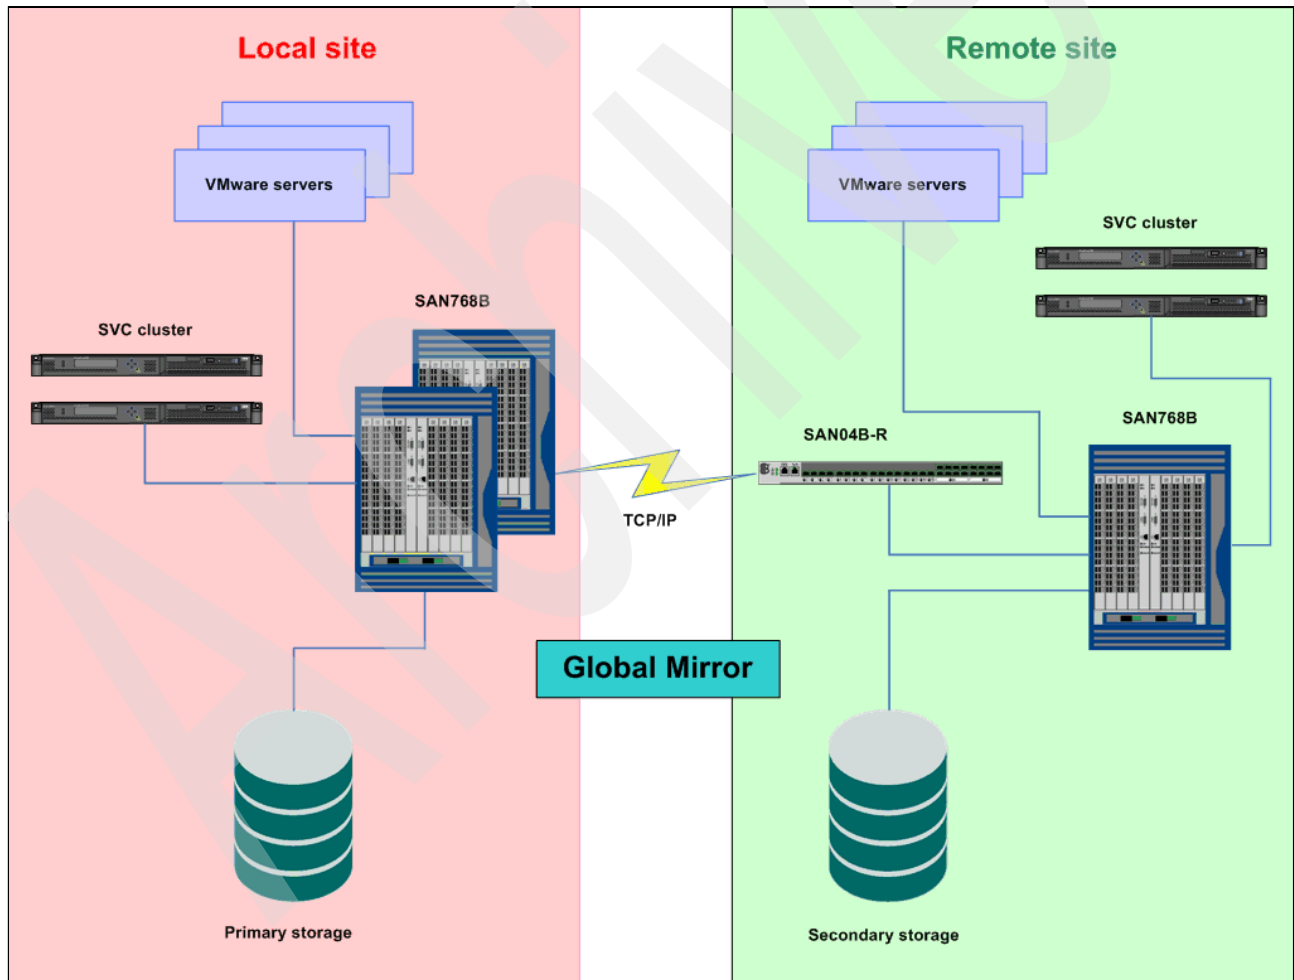


Figure 4 Dual site using routing for Global Mirror

Again, the connection can either be a routing blade in the IBM System Storage SAN768B director, or a separate device, such as the IBM System Storage SAN04B-R router.

The round-trip latency between sites cannot exceed 80 ms, which should enable distances of up to 4000 km (2485 miles).

Solution

VMware offers *Site Recovery Manager (SRM)* as an optional element of VMware Infrastructure (VI) to enable automated disaster recovery. As part of the installation of SRM, the *SVC Storage Resource Adapter (SRA)* needs to be installed. This plug-in is used to enable the management of advanced copy services on the SVC, such as Metro Mirror and Global Mirror, and is based on the Storage Networking Industry Association (SNIA) SMI-S interface in the SVC.

Figure 5 shows a high level overview of the relationships between vCenter, SRM, SRA, the ESX hosts and VMs, and the storage at the local and remote sites.

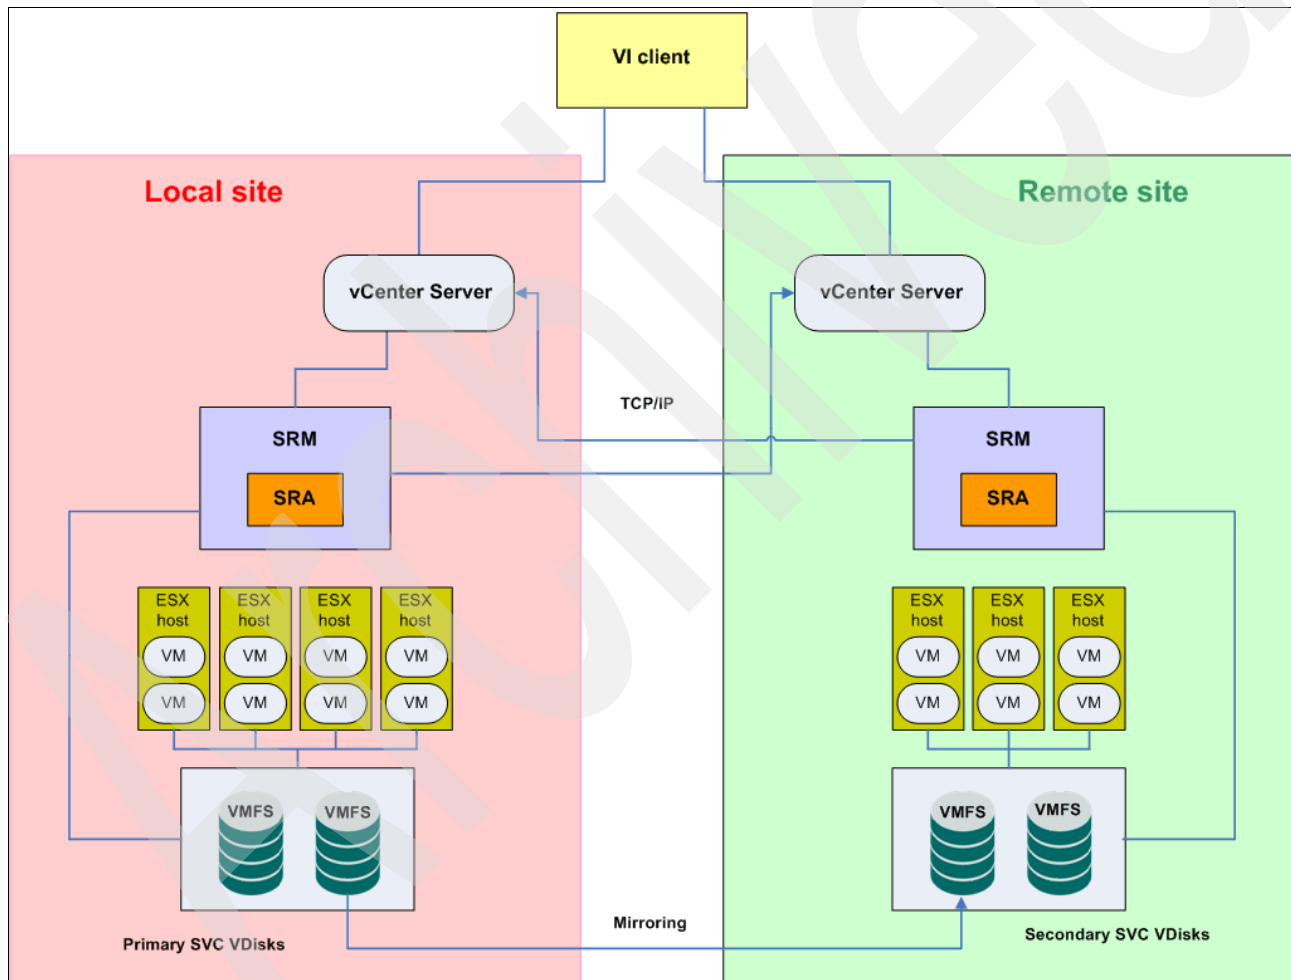


Figure 5 High level overview of SRM with SRA

More details about SRM can be found at the following Web site:

<http://www.vmware.com/products/srm/overview.html>

Obtaining the SVC SRA installer

Perform these steps to obtain the SVC SRA installation file:

1. Go to the IBM Systems support Web site, found at:
<http://www.ibm.com/systems/support/>
2. Click the **System Storage** link on the left.
3. Select **Storage software** from the Product family drop-down menu.
4. Select **SAN Volume Controller (2145)** from the Product drop-down menu.
5. Click the **Go** button.
6. Click the **All SAN Volume Controller documents** link.
7. Enter “vmware site recovery manager” in the Additional search terms field and press Enter.
8. Look for a document called *Support for VMware Site Recovery Manager* in the results list.
9. Read through the page until you come to the download link for **VMware SVC SRA**.
10. Install the SRA on the vCenter Server before installing SRM. This installation must be done at both locations.

Site Recovery Manager

The benefit of adding the SRA to SRM is that it enables VMware to manage the SVC remote copy functions as well as the ESX hosts and guest VMs. It aids in the following tasks:

- ▶ Setting up the recovery infrastructure
- ▶ Creation of recovery plans
- ▶ Testing recovery plans
- ▶ Automating failover

The following steps must be completed to set up SRM:

1. Install SRM.
2. Configure the array managers in SRM (VDisk replication should be set up before installing SRM).
3. Define inventory mappings.
4. Create protection groups.
5. Create recovery plans.
6. Configure VM properties.
7. Configure security and monitoring alerts.

Full details about all of these steps can be found in the *Administration Guide for Site Recovery Manager 1.0*, found at:

http://www.vmware.com/support/pubs/srm_pubs.html

Summary

The combination of VMware Site Recovery Manager and the Storage Resource Adapter for SVC enables the automated failover of VMs from one location to another location connected by Metro Mirror or Global Mirror technology. This ensures that in the event of a disaster befalling your primary location, VMware vCenter is able to restart all of the required VMs at the secondary location with a minimum of delay and manual intervention, and achieve the shortest recovery time.

The ability of the SVC to mirror VDisks between MDisk Groups provides a very low cost option for mirroring within a single site between two storage subsystems. For metropolitan area distances, the use of Metro Mirror over dark fiber provides a synchronous mirror capability without the need to use TCP/IP over a wide area network (WAN). For much greater distances, Global Mirror provides asynchronous mirroring over TCP/IP using SAN router technology, such as the IBM System Storage SAN04B-R.

The latter two options require the use of SRA with vCenter SRM to automate the necessary SVC management during VM failover.

Related material

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *VMware Multi-pathing with the SVC, and the Causes of SCSI-2 Reservation Conflicts*, REDP-4627
- ▶ *VMware Proof of Practice and Performance Guidelines on the SVC*, REDP-4601

The team who wrote this IBM Redpapers publication

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

Steve Garraway is an IT specialist with the Storage and SAN Services team at the IBM Hursley Laboratory in the United Kingdom. He has 22 years of experience in storage management, mainly with z/VM® and IBM Tivoli® Storage Manager, and has worked for IBM for 23 years. His areas of expertise include IBM System z® disk and tape storage, McDATA and Brocade SANs, and the SAN Volume Controller. He has previously co-authored four IBM Redbooks publications on enterprise disk and SAN. Steve is a Brocade Certified Fabric Professional.

Jon Tate is a Project Manager for IBM System Storage SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2 and 3 support for IBM storage products. Jon has 24 years of experience in storage software and management, services, and support, and is both an IBM Certified IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.

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

Marcus Thordal
Brocade Communications Inc, San Jose

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 Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites 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.

This document REDP-4626-00 was created or updated on December 18, 2009.



Send us your comments in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks 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 U.S.A.



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:


VMware, the VMware "boxes" logo and design are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

IBM®

Redbooks®

Redpaper™

Redpapers™

Redbooks (logo) ®

System Storage™

System z®

Tivoli®

z/VM®

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