



Disaster Recovery for IBM® Lotus Notes® and Domino® 8.5.1 Using VMware vCenter™ Site Recovery Manager

September 2011

DEPLOYMENT AND TECHNICAL CONSIDERATIONS GUIDE

Table of Contents

Introduction	3
Solution Overview	3
Lotus Notes and Domino 8.5.1	4
Site Recovery Manager	4
SAN Storage and Data Replication	6
Test Workload Generation.....	7
Solution Architecture.....	8
Site Recovery Configuration	8
Solution Configuration Steps	11
Test Results	11
Testing Summary.....	11
Synchronous Testing	12
Asynchronous Testing	12
Recommendations	12
Conclusion	14
Resources	15

Introduction

During a disaster, VMware vCenter™ Site Recovery Manager enables IBM® Lotus Notes® and Domino® 8.5.1 running in a VMware vSphere® environment to be recovered rapidly within defined timeframes. When a disaster occurs, businesses initially expect to operate in “survival” mode by restoring only those business processes that are essential to keep the company running.

The recovery of critical business processes is defined using these objectives:

- **Recovery Time Objective (RTO)** – RTO defines the criticality of each business process, such as finance or email communications, and its dependencies in the event of a disaster. The RTO quantifies how rapidly the organization will need to have these business processes up and running.
- **Recovery Point Objective (RPO)** – For each critical business process, the RPO quantifies the minimal amount of data the business needs to restore for recovery and to what point in time the data needs to be restored, such as to the prior week or month.

This paper describes the advantages and technical considerations for using Site Recovery Manager to meet established RTOs and RPOs for disaster recovery in Lotus Domino 8.5.1 deployments that run in a VMware environment.

Solution Overview

VMware and IBM jointly conducted testing that used Lotus Domino 8.5.1 with Site Recovery Manager 4.1.1 in a vSphere 4.1 environment.

The testing had the following goals:

- Verify that a Lotus Domino 8.5.1 server with transaction logging enabled can fail over to the recovery site using Site Recovery Manager without corrupting the in-use databases, using both synchronous and asynchronous replication.
- Verify that both Lotus Notes users and other Domino servers can connect to the Lotus Domino 8.5.1 server once it is running at the recovery site.

The testing results prove that Site Recovery Manager provides a flexible, high-performance, high-availability platform for disaster recovery that works well with Lotus Domino 8.5.1 running in a VMware environment.

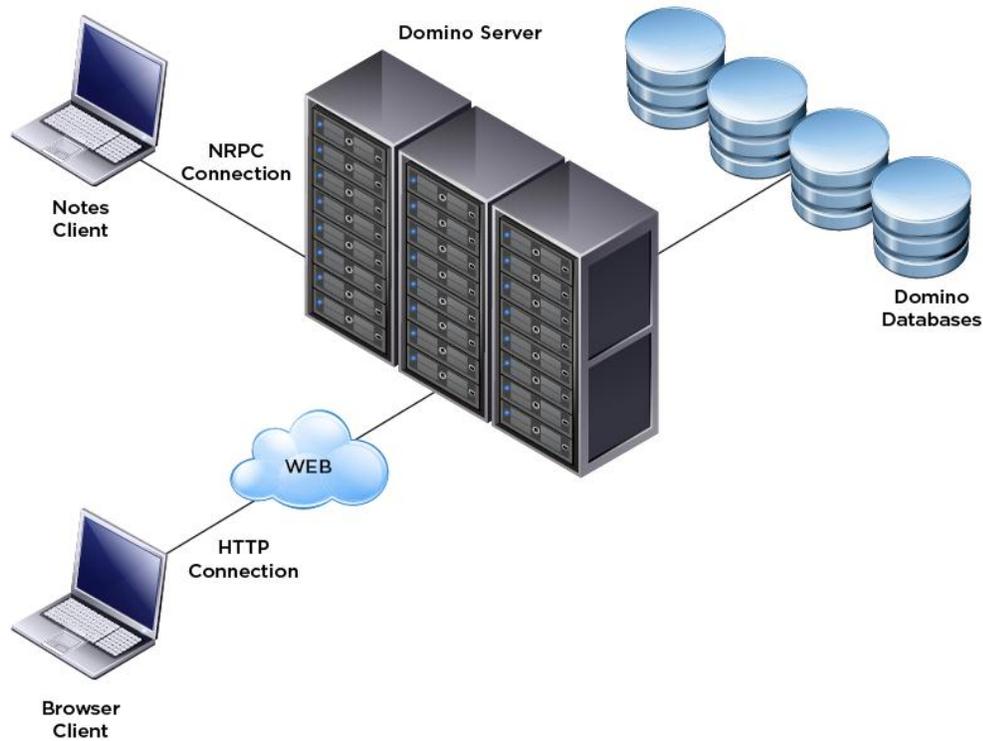
By using Lotus Domino 8.5.1 with Site Recovery Manager, businesses gain these key advantages:

- Gets you back in business rapidly; allows customers to avoid loss by expediting failover and recovery from a major service disruption in alignment with the defined RTOs and RPOs.
- Accelerates disaster recovery by minimizing its inherent risks and downtime such as by eliminating manual recovery steps.
- Delivers advanced capabilities for disaster recovery management, non-disruptive testing and automated failover, especially with long-distance replication.
- Enables automated failover across geographically dispersed locations.
- Requires minimal administration and minimal additional hardware.
- Enables customers to expand protection to other important systems and applications.

Lotus Notes and Domino 8.5.1

Lotus Notes and Domino 8.5.1 enterprise collaboration software brings email, collaboration tools, and business applications together in a rich, integrated desktop experience. The Lotus Domino platform enables customers to build, deploy, and administer critical business, collaboration, and messaging applications.

Figure 1. Lotus Domino System Architecture

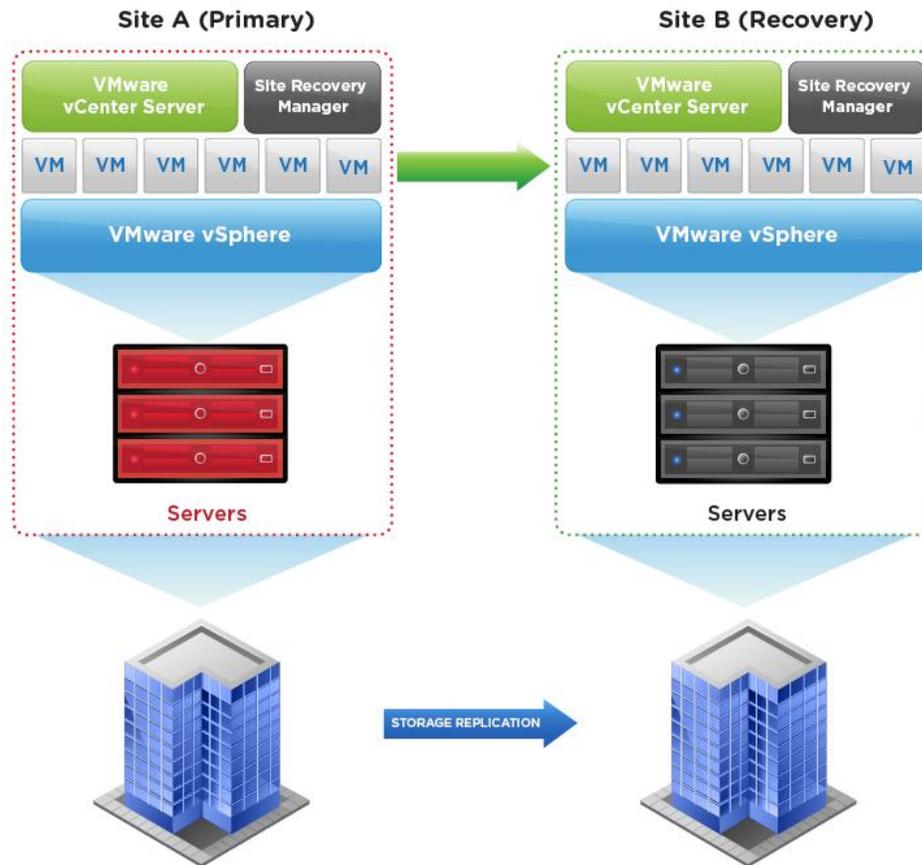


Site Recovery Manager

Site Recovery Manager is an extension to VMware vCenter that simplifies and automates disaster recovery when Lotus Domino 8.5.1 runs as a workload in a VMware environment. Site Recovery Manager orchestrates disaster recovery and integrates failover with other enterprise applications.

Figure 2 depicts the typical architecture of Site Recovery Manager with storage array replication across two sites.

Figure 2. Site Recovery Manager System Architecture



Site Recovery Manager facilitates setup, test, and recovery workflows, and it offloads failover to the production storage array, rather than the vSphere front end. It coordinates the scheduled migration or unplanned failover of production datacenters to recovery sites, or between two sites with active workloads:

- As virtual machines at the protected site are shut down, virtual machines at the recovery site start up and use data replicated from the protected site to provide the same services.
- The two datacenters are not synchronized in real time, although the lag time for replication is just a few seconds. Failover and recovery require only a few minutes.
- Site Recovery Manager controls the transfer of services by using a recovery plan that specifies the order in which virtual machines are shut down and started up, compute resources are allocated, and network access is given.
- Site Recovery Manager enables fast and accurate recovery by maintaining duplicate copies automatically when data spanning multiple volumes and storage systems is replicated across long distances.

Site Recovery Manager enables the testing of recovery plans in a way that does not disrupt ongoing operations at either site. For example, it uses these approaches:

- Creates a temporary copy of the replicated data to test the failover process and ensure that the secondary image is consistent and usable.

- Relies on independent vCenter servers at both the protected site and the recovery site to facilitate the failover process.

With Site Recovery Manager, disaster recovery becomes an integrated part of managing the VMware virtual infrastructure. It is implemented using a vendor-supplied storage replication adapter (SRA) that is tightly integrated with vCenter, vSphere, and the storage platform.

The SRA plug-in, installed on the Site Recovery Manager server, manages communication with the storage platform. The EMC CLARiiON and NetApp storage platforms each have an SRA for Site Recovery Manager that takes advantage of their respective storage replication technologies, as described in the sections below. For a complete list of storage platforms that are compatible with VMware, see the “Resources” section later in this paper.

Note: The latest release of Site Recovery Manager now supports host based replication, which can eliminate the need to maintain an identical storage array at the recovery site. It can enable the use of low cost, direct attach storage for disaster recovery purposes and allow for more granular, virtual machine level replication. For more information, see the Site Recovery Manager product page on the VMware website.

SAN Storage and Data Replication

For this testing, Site Recovery Manager was used with these storage platforms:

- EMC CLARiiON CX3-20c and its MirrorView 1.4.0 replication technology
- NetApp FAS3020 and its SnapMirror 1.4.3 replication technology

The SRA for each storage platform is described in the sections below.

EMC CLARiiON MirrorView

For this testing configuration, the SRA for Site Recovery Manager integrates with CLARiiON MirrorView to enable automated failover of VMware vStorage Virtual Machine File System (VMFS) datastores from the protected site to the recovery site. Site Recovery Manager uses the EMC CLARiiON SRA and MirrorView to enable integration using features such as array discovery, replicated LUN discovery, test failover, and actual failover. The EMC CLARiiON SRA is installed on the Site Recovery Manager server at both the protected site and the recovery site.

MirrorView provides LUN-level data replication of the primary image on the CLARiiON storage array to the secondary image on the recovery storage system using either synchronous or asynchronous mirroring, as described below. During normal operations, the primary image is online and it is available for read or write operations. The replicated secondary image is not visible to the user.

- **EMC CLARiiON MirrorView/S** – Provides real-time synchronous mirroring of data between the primary production image of the data and the secondary image at the recovery site. With synchronous operations, data must be successfully stored on both the primary and the secondary arrays before an acknowledgment is sent to the local host.

When using this option, make sure that the network between the two sites is high-speed. Otherwise, the time delay in writing the data to the remote site can cause poor Domino user response time. In certain cases, this delay could result in a Domino data failure and even a Domino server crash.

- **EMC CLARiiON MirrorView/A** – Provides asynchronous mirroring of data between the primary production image and the secondary image at the recovery site. With asynchronous operations, MirrorView/A tracks the changes to the primary image and then replicates these changes to the secondary image when an update is triggered.

Note: For this testing, EMC CLARiiON storage platforms were used to test synchronous replication and its effect on the failover of the Lotus Domino mail server.

NetApp SnapMirror

The NetApp SRA for Site Recovery Manager integrates with SnapMirror to enable the automated failover of VMFS datastores from the protected site to the recovery site. Site Recovery Manager uses the NetApp SRA and SnapMirror to enable integration for replication and disaster recovery purposes.

Built on NetApp Snapshot™ technology, SnapMirror replicates only the 4 KB blocks that have been changed or added since the previous update. SnapMirror can be combined with FAS deduplication utilized on the primary storage platform to ensure that only unique data is replicated to the recovery site. This results in increased telecommunication savings and reduced storage capacity at the recovery site. SnapMirror is easily configured using NetApp Protection Manager components such as Data ONTAP® CLI and FilerView® GUI.

Note: For this testing, NetApp FAS-series storage platforms were used to test asynchronous replication and its effect on failover of the Lotus Domino mail server.

When primary storage is not completely lost, SnapMirror resynchronizes both the protected site and the recovery site and transfers only changed or new data back to the recovery site.

Test Workload Generation

The testing of Lotus Domino 8.5.1 with Site Recovery Manager used an IBM NotesBench client workload. The NotesBench workload is a defined script that simulates user activity through specific applications and defines workload usage patterns. The NotesBench workload is available as an installation option with the Notes Admin client.

In this case, NotesBench ran N85mail and DWA85mail workloads to apply a consistent, repeatable load against the Domino server to assess the site recovery configuration. Both the N85mail and DWA85mail workloads simulated user mail activity by making the same API calls that a Notes client user would perform in equivalent mail operations.

The workload testing was performed with increments of 1,000 users added over several hours, ending with 4,000 concurrent users.

Table 1 shows the configuration of each user's mail database for each workload.

Table 1. Mail Configuration

User Mail Database	Configuration
Mail database size	Approximately 250 MB
Average document size	100 KB
Initial document count	3,000

Table 2 describes operations performed per user for both workloads over a 6-hour period.

Table 2. User Operations

Workload Actions	Action Count
Refresh inbox	24
Read messages	120
Reply to all	12
Send message to one recipient	24 (see note below)
Send message to three recipients	12 (see note below)
Create appointment	1
Send invitation	1
Send RSVP	1
Move to another folder	24
Delete two documents	24

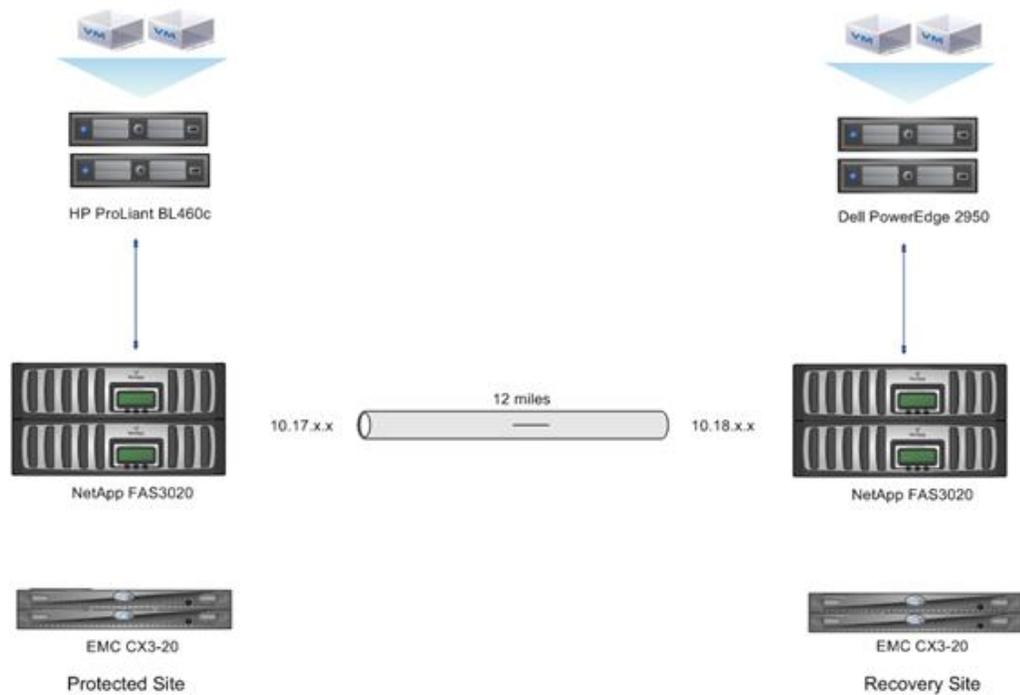
Note: The email messages varied in size from 500 to 150,000 bytes, with a random number of these messages containing attachments of 525 KB.

Solution Architecture

The solution architecture for testing Lotus Domino 8.5.1 workloads with Site Recovery Manager is described in the following sections.

Site Recovery Configuration

Figure 3. Lotus Domino–Site Recovery Manager Solution Architecture (below)



The solution architecture included:

- The lab environments for the protected site and the recovery site were built in two VMware datacenters in Palo Alto, California, with the sites located approximately 12 miles apart.
- Both the protected site and the recovery site run VMware vCenter Server in a virtual machine that is configured on VMFS storage. Each vCenter Server runs in a virtual machine to take advantage of VMware vMotion®, VMware High Availability (HA), and VMware Distributed Resource Scheduler (DRS).
- The protected site is configured using the EMC CLARiiON CX3-20c storage platform for synchronous replication testing, as well as the NetApp FAS3020 storage platform for asynchronous replication testing.
- Similarly, the recovery site maintained an additional EMC CLARiiON CX3 storage platform and an additional NetApp FAS3020 as the target storage platform for replication and recovery purposes.
- Site Recovery Manager is installed on the vCenter servers at the protected site and the recovery site to enable communication between these sites.
- The storage platforms used at both the protected site and the recovery site used the same data replication mode that included:
 - Synchronous replication as the preferred SAN replication strategy for mission-critical data, because there is negligible risk of data loss in the event of failover.
 - Asynchronous replication was supported, although asynchronous replication might result in minimal data loss during failover. (Contact the storage vendor for detailed information.)
 - DNS was set up with its routing capability enabled.

- After the Lotus Domino server at the protected site fails over to the recovery site, the new IP is automatically registered in DNS.
- Once the Lotus Domino server is running, NotesBench established communication through a fully qualified domain name (FQDN).

Hardware and System Host Configuration

Table 3 describes the configuration of VMware ESX® host servers and storage in the solution architecture.

Table 3. Hardware Configuration

Hardware	Configuration
Server	<p>Protected site – Two Hewlett-Packard ProLiant BL460c G6 servers, each equipped with</p> <ul style="list-style-type: none"> • 2.8 GHz dual quad-core Intel Xeon E5560 processor • 48 GB RAM <p>Recovery site – 2 Dell PowerEdge 2950 servers, each equipped with</p> <ul style="list-style-type: none"> • 2.3 GHz dual quad-core Xeon E5410 processor • 32 GB RAM
Storage (protected site)	EMC CLARiiON CX3-20c and MirrorView 1.4.0 NetApp FAS3020 and SnapMirror 1.4.3
Storage (recovery site)	EMC CLARiiON CX3-20c and MirrorView 1.4.0 NetApp FAS3020 and SnapMirror 1.4.3

Installed Software

Table 4 lists the software used in the solution architecture.

Table 4. Software

Software Provider	Software Configuration
VMware	vSphere 4.1, vCenter Server 4.1, Site Recovery Manager 4.1.1
Microsoft	SQL Server 2005, Enterprise Edition
IBM	Lotus Domino 8.5.1

Virtual Machine Configuration

Table 5 describes the configuration of virtual machines running on ESX host servers in the solution architecture.

Table 5. Virtual Machines

Virtual Machine	Hardware Configuration
Lotus Domino 8.5.1	Windows Server 2003 Enterprise Edition, 64-bit Two vCPUs 8 GB RAM 200 GB storage for 1,000 users

Virtual Machine	Hardware Configuration
Lotus Notes Client	Windows XP One vCPU 2 GB RAM
NotesBench Client	Red Hat Enterprise Linux 5.3 Four vCPUs 8 GB RAM
Active Directory	Windows Server 2003 Enterprise Edition Two vCPUs 2 GB RAM

Solution Configuration Steps

Site Recovery Manager can be used to add virtual machines that share common resources to protection groups at the protected site. Protection groups can be configured into recovery plans for restoration at the recovery site.

When the recovery plan is executed, Site Recovery Manager does the following:

1. Breaks the EMC CLARiiON MirrorView and NetApp SnapMirror relationships.
2. Maps the LUNs to existing NetApp/EMC snapshots.
3. Triggers the disaster recovery ESX hosts to rescan and detect the storage.
4. If desired, suspends virtual machines such as test/dev instances at the recovery site to free computing resources.
5. Reconfigures the virtual machines as defined for the network at the recovery site.
6. Powers on the virtual machines in the order defined in the recovery plan.
7. After boot up, the Domino server automatically receives the new IP address that was customized using the DNS IP customizer tool.

For more information, refer to the Site Recovery Manager Administration Guide (<http://www.vmware.com/pdf/srm-admin.pdf>), starting with “Report IP Address Mappings for a Protection Group.”

Test Results

The testing of Lotus Domino 8.5.1 with Site Recovery Manager included running workloads in both the synchronous and asynchronous configurations. A summary of the tests and resulting observations are described below.

Testing Summary

The testing used both synchronous and asynchronous replication of Lotus Domino 8.5.1 with Site Recovery Manager. In each case, the testing exercised the Domino server using a NotesBench workload of several thousand users, with a forced failover. During the failover process, the Domino server was not responsive until it was fully up and running on the recovery site.

The conclusion was that with synchronous testing, the risk of data loss is negligible. In contrast, with the asynchronous solution, the potential for minimal data loss, such as the loss of mail messages, can be anticipated and addressed.

Synchronous Testing

For the EMC CLARiiON synchronous testing configuration, the EMC SRA plug-in was installed on the vCenter server at the protected site to automatically discover the replicated LUNs. Synchronous mode was set to high to maintain an exact byte-for-byte, real-time copy of the production data in a secure, remote location. The behavior of the Domino server at the recovery site was similar to that of a server recovering from a power failure or server crash and reboot.

The synchronous testing was successful. The following was observed after failover:

- The server responded properly to its name at the new IP address location. The IP address was registered automatically with the DNS server by means of the DNS IP customizer tool. The DNS IP customizer tool is provided as a component of Site Recovery Manager.
- When the Domino server was completely back up and running, Notes users were able to connect to the server at the new location without any modification to the client.
- Transaction logging was able to recover and put the open databases in a usable state.
- No errors were observed to indicate that there was any significant data loss or data contamination in terms of open in-boxes, appointments and message drafts.

The testing results concluded that data integrity is ensured in synchronous mode. There is negligible risk of data loss with the synchronous solution.

Asynchronous Testing

For asynchronous testing, the NetApp SRA plug-in was installed on the vCenter servers at the protected site and the recovery site to facilitate a Site Recovery Manager testing workflow and to create a clone of the replicated LUNs to synchronize data between the two sites. For this testing, the schedule for replicating changed or updated data from the primary storage to secondary storage device was every 60 seconds.

Results showed that with the asynchronous solution, the potential for data loss (such as the loss of mail messages) is minimal, but it can be anticipated. During this testing, the observed behavior was similar to the synchronous testing, except for a few errors related to missing email. The potential for data loss can be minimized with the NetApp array by reducing the synchronization window schedule and/or by using other types of complementary data-protection technologies.

Recommendations

The testing results indicate that the best performance in implementing site recovery solutions can be achieved by using the best practices described below.

These and other recommendations are described in more detail in the VMware Site Recovery Manager documentation. For more information, see the “Resources” section later in this paper.

- Install the Site Recovery Manager database as close to the Site Recovery Manager server as possible to reduce the round-trip time. Reduce the effect of round trips on recovery time performance. Use

the same database server to support the vCenter database instance and the Site Recovery Manager database instance.

- Enable faster testing and real recoveries by grouping virtual machines under fewer protection groups. Ensure that the virtual machines have no constraints preventing them from being added to similar protection groups.
- Enable DRS at the recovery site. Migrations might occur as DRS tries to load-balance the cluster during recovery.
- Enable Distributed Power Management (DPM) on recovery site clusters if your recovery site hosts are in the standby state. More hosts lead to increased concurrency for recovering virtual machines, and this enables shorter recovery times. If DPM is not enabled, and the hosts are in a standby state, bring the hosts out of standby mode manually and drag shadow virtual machines onto them.
- Chart the dependencies between virtual machines. To reduce overall failover times, define recovery priorities and assign virtual machines to appropriate protection groups.

As an alternative to placing virtual machines in a high-priority group, separate all of the virtual machines to be recovered in two logical groups:

- Group1 with level 1 virtual machines – Place Group1 virtual machines in the normal priority group.
- Group2 with level 2 virtual machines (dependent upon virtual machines in Group1) – Place Group2 virtual machines in the low-priority group within the same plan.

This should maintain dependency across both logical groups and reduce the recovery time by introducing more concurrency for both priority groups. Dependency is maintained between priority groups, and there is no dependency across virtual machines within a single priority group.

- It is strongly recommended that VMware tools be installed in all protected virtual machines in order to accurately acquire their heartbeats and network change notification.
- Make sure any internal script or callout prompt does not block recovery indefinitely.
- Specify a non-replicated datastore for swap files to speed up replication between the two sites, and remote calls to vCenter Server. (During recovery, delete swap files for all virtual machines).
- VMware recommends a minimum of two ESX hosts at a recovery site. Although this is not a requirement, having two ESX hosts ensures availability for VMware services such as vMotion, HA, and DRS, and it improves the recovery time.
- Site Recovery Manager relies on the capabilities of the storage platform for failover. Recovery time can vary depending on the SAN vendor, replication type and other factors. It can take 5 to 10 minutes to make the changes to the back-end storage. Follow VMware best practices for Site Recovery Manager, along with SAN vendor best practices, to ensure optimal recovery time.

Conclusion

When Lotus Domino 8.5.1 runs as an application in a VMware environment, Site Recovery Manager enables businesses to achieve cost-effective disaster recovery conveniently. This paper describes the advantages and technical considerations of using Site Recovery Manager for disaster recovery in Lotus Domino 8.5.1 deployments that run in a VMware environment using block-based SAN storage replication. The testing demonstrates the behavior of a Lotus Domino server running in a VMware environment with Site Recovery Manager in addressing the RTOs and RPOs.

This testing provides observations about the results that can be achieved in both synchronous and asynchronous disk solutions to help with IT configuration planning. This testing also demonstrates that the synchronous solution is the best practice to apply when avoiding data loss is a significant issue.

Overall, the testing results prove that Site Recovery Manager provides a flexible, high-performance, high-availability platform for disaster recovery that works well with Lotus Domino 8.5.1 running in a VMware environment.

Resources

For more information about VMware and related products, use the links and references below.

VMware References

- Customers can post questions in the VMTN Community specifically for messaging and collaboration products:
<http://communities.vmware.com/community/vmtn/general/emailapps>
 This forum is monitored by VMware subject matter experts (SMEs) as well as other customers with similar interests.
- VMworld 2009 Session – Virtualizing IBM Lotus Domino and Sametime:
<http://www.vmworld.com/sessions/EA2649.html>
- VMware Documentation:
<http://www.vmware.com/support/pubs/>
- VMware Licensing:
<http://www.vmware.com/support/licensing/>
- Application Performance Troubleshooting:
<http://www.vmware.com/resources/techresources/10066>
- VMware vSphere 4.1 Performance Best Practices:
<http://www.vmware.com/resources/techresources/10161>
- VMware Proven Best Practices (VIOPS):
<http://communities.vmware.com/community/viops>
- VMware Solutions/Product/Partner Podcasts:
<http://www.vmware.com/technical-resources/podcasts/>
- VMware Global Support KB:
<http://kb.vmware.com>
- VMware Global Support Videos:
<http://blogs.vmware.com/kbtv/>
- VMware TV:
<http://www.youtube.com/user/vmwaretv>
- VMworld TV:
<http://www.youtube.com/user/VMworldTV>
- VMware KB TV:
<http://www.youtube.com/user/VMwareKB>

Lotus Notes and Domino

- Lotus Notes and Domino 8.5.1:
<http://www-01.ibm.com/software/lotus/notesanddomino/nd85.html>
- Virtualizing IBM Lotus Domino 8.5.1 on VMware vSphere 4:
http://www.vmware.com/files/pdf/techpaper/IBM_Lotus_Domino_on_vSphere_V1.pdf



Storage Platforms

- VMware Site Recovery Manager in a NetApp Environment:
http://kb.vmware.com/Platform/Publishing/attachments/1007098_dNetApp-SRM-tr-3671.pdf
- Knowledge Base: Site Recovery Manager with NetApp:
http://kb.vmware.com/Platform/Publishing/attachments/1007098_dNetApp-SRM-tr-3671.pdf
- Knowledge Base: Setting Up VMware vCenter Site Recovery Manager with Partner Storage Arrays:
http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1014610
- VMware Compatibility Guide for a complete list of compatible storage devices:
http://www.vmware.com/pdf/vi35_san_guide.pdf