

Oracle E-Business Suite Disaster Recovery Solution with VMware Site Recovery Manager and EMC CLARiiON

Applied Technology

Abstract

This white paper demonstrates how VMware's Site Recovery Manager enables the design of a powerful yet simple disaster recovery solution. Built on the solid foundation of VMware virtualization and leveraging EMC's MirrorView™/Synchronous storage replication technology, Site Recovery Manager easily resolves typical disaster recovery challenges in an extremely cost-effective manner while providing complete flexibility.

November 2008

Copyright © 2008 EMC Corporation. All rights reserved.

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

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

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

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h5896

Table of Contents

Executive summary	4
Introduction	4
Audience	5
Overview	5
EMC CLARiiON CX4 series storage.....	5
EMC MirrorView	6
VMware Site Recovery Manager	6
Oracle E-Business Suite	7
The Desktop tier	8
The Application tier.....	8
The Database tier.....	8
Oracle Applications File System.....	8
Solution design and setup	9
Architecture	9
Hardware and software.....	9
Hardware configuration	11
Site Recovery Manager install.....	11
Site Recovery Manager configuration	12
Planning and testing.....	13
Recovery plan test using Site Recovery Manager.....	13
Site failover (disaster recovery) test	14
Conclusion	17
References	18

Executive summary

Today's data center is a product of change. Technological progress and business growth have added to the complexity of IT infrastructure with new operating system, storage, database, and network technologies. Along with technologies, data center uptime has grown increasingly important for corporations. Even in the event of a site failure or system failure caused by hardware malfunction, corporations must continue to safeguard critical business data such as customer information and rapidly restore system functionality to ensure continuing services. Interruptions or outages affecting important services pose serious threats to the entire business; in certain cases this can result not just in lost income, but serious damage to the confidence of customers and associated companies. At the same time, data is one of the most critical assets for any company. Corporate data, for example payroll or employee information, client records, valuable research results, financial records, or history information, can require both significant sums and effort to reconstruct or regenerate once lost, if this is even possible. In some cases such data loss may impair a company's capacity to continue operating.

Smaller businesses, with limited budgets and staff, are typically turned off by the costs and stringent requirements as well as complexities and hence are left with few options for protecting their applications and business. To make matters more complex, more and more mission-critical business applications are now logically combined into distributed, multi-tier (also called n-tier) client/server platforms. An example of this is Oracle's E-Business Suite. Companies that have deployed Oracle E-Business Suite depend heavily on the information, processes, and the availability of the environment even in the case of a disaster. The good news is that VMware's innovative virtualization leveraged with EMC® CLARiiON® storage technology helps overcome all these challenges related to high availability.

Introduction

This white paper will outline how, using a multi-tier Oracle E-Business environment residing on EMC's CLARiiON storage, VMware Site Recovery Manager can help you:

- Accelerate recovery for the virtual environment through automation.
- Ensure reliable recovery by enabling non-disruptive testing.
- Simplify recovery by eliminating complex manual recovery steps and centralizing management of recovery plans.

This paper will introduce a short, simple approach to disaster recovery that uses Site Recovery Manager and CLARiiON storage and replication software (MirrorView™/Synchronous) to provide an effective availability solution for the mission-critical Oracle E-Business Suite environment.

The scope of the architecture documented here includes:

- A freshly installed Oracle E-Business Suite 12 along with an Oracle 10g database on Oracle Enterprise Linux 4.5 and EMC's CLARiiON CX4 storage
- Storage array replication of LUNs (via MirrorView/Synchronous) between two separate CLARiiON storage systems simulating a primary and secondary site
- Setup of a protection group and recovery plan consisting of an E-Business environment using VMware's Site Recovery Manager
- Performing non-disruptive test execution of a recovery plan
- Execution of a disaster recovery plan using Site Recovery Manager. This test included recovering an E-Business Suite environment at a secondary site while the primary site was subjected to a load created by the General Ledger report generation

Audience

This white paper is intended for EMC employees, partners, IT managers, storage architects, database administrators, and consultants who are involved in planning for disaster recovery options for their Oracle Database application deployments.

Overview

EMC CLARiiON CX4 series storage

The EMC CLARiiON CX4 series storage platform, launched in Q3 2008, is EMC's state-of-the-art midrange family of storage systems. The CX4 architecture delivers cutting-edge performance, including the highest levels of resiliency and availability, tiered storage flexibility, and powerful easy-to-use interfaces. The unique combination of flexible, scalable hardware design and advanced software capabilities enables CLARiiON CX4 to achieve up to twice the performance and scale as the previous CLARiiON generation. With breakthrough architecture and extensive technological innovation, CX4 is the leading midrange storage solution to meet a full range of needs – from departmental applications to data-center-class business-critical systems.

VMware's Virtual Infrastructure featuring CLARiiON CX4 series storage furthers the benefits of VMware by enabling and supporting new functionality such as VMotion, Distributed Resource Scheduler (DRS), High Availability (HA), and Site Recovery Manager (SRM) that makes VMware infrastructure 3 a powerful solution.




Figure 1. EMC CLARiiON CX4 storage and VMware

For this VMware Site Recovery Manager and Oracle E-Business Suite use case, two CX4-480 storage arrays with MirrorView/Synchronous (MirrorView/S) are used. The CLARiiON CX4-480 provides high-capacity networked storage that meets the needs of demanding OLTP workloads and large-scale e-mail environments. With the CX4-480 you can scale seamlessly up to 471 TB of storage capacity and consolidate twice the workloads in one array as you can with other storage providers.

Note that the methodologies discussed in this paper for use with Oracle extend to all models in the CLARiiON family that support MirrorView/S, including the CX, CX3, and the latest CLARiiON CX4 series storage.

EMC MirrorView

EMC MirrorView is storage system-based disaster recovery software that provides end-to-end data protection by replicating the contents of a primary volume to a secondary volume that resides on a different CLARiiON storage system. It provides end-to-end protection because, in addition to performing replication, it protects the secondary volume from tampering or corruption by making the volume only available for server access when initiated through MirrorView. MirrorView offers consistency groups, which is a unique consistency technology for the midrange market that replicates write-order dependent volumes. Using this technology, MirrorView maintains write ordering across the secondary volumes in the event of an interruption to one, some, or all of the write dependent secondary volumes.




Figure 2. EMC MirrorView/Synchronous software

As shown in Figure 2, all I/O from the local host is first written to the local storage processor memory and is then sent over the mirroring links to the remote CLARiiON unit. Once the remote CLARiiON unit reports that the data has reached its storage processor memory successfully, the I/O is acknowledged to the local host. Synchronous mode guarantees that the remote image is a complete duplication of the source image.

For this solution, we used MirrorView/Synchronous mode to ensure an RPO of zero.

VMware Site Recovery Manager

VMware Site Recovery Manager is a pioneering disaster recovery management and automation solution for VMware Infrastructure. Site Recovery Manager accelerates recovery by automating the recovery process and simplifying the management of disaster recovery plans by making disaster recovery an integrated element of managing your VMware virtual infrastructure. The solution ensures reliable recovery by eliminating complex manual recovery steps and enabling non-disruptive testing of recovery plans. Site Recovery Manager enables organizations to take the risk and worry out of disaster recovery, as well as expand protection to all important systems and applications. It integrates tightly with VMware Infrastructure, VMware® VirtualCenter and EMC's MirrorView software to make failover and recovery rapid, reliable, affordable, and manageable. Figure 3 shows a configuration involving Site Recovery Manager and MirrorView.




Figure 3. Site Recovery Manager with MirrorView/S

Site Recovery Manager does not actually perform the replication for disaster recovery but facilitates the setup, test, and recovery workflows. Site Recovery Manager relies on block-based replication (fibre or iSCSI) such as EMC CLARiiON MirrorView/S for replication. The communication between Site Recovery Manager and MirrorView/S is managed by a storage replication adapter (SRA) provided by EMC. Storage replication adapters exist on the Site Recovery Manager server and once installed are invisible for the duration of their use. Additional details about Site Recovery Manager setup and configuration are covered in the VMware white paper *Planning for the Unplanned*.

Oracle E-Business Suite

The Oracle Applications architecture is a framework for multi-tiered distributed computing. In the Oracle Applications model, services are distributed among multiple levels or tiers. A service is a process or a group of processes that exercises some business logic and provides a particular functionality. A tier is a grouping of services that can potentially span across physical machines. In other words, a tier is a logical grouping of services that is not limited by physical nodes or machines. Thus, each tier can consist of one or more nodes and each node can accommodate more than one tier. For example, a single machine can contain database, application, and desktop tiers. Similarly, a database can reside on one of many application servers or on a separate machine by itself.

The following sections provide a brief description of the Oracle Applications architecture and its major components.




Figure 4. Oracle E-Business Suite

The Desktop tier

The Desktop tier primarily refers to the client web browser. The interface is provided through HTML for self-service applications and a Java applet for forms. The desktop client is installed on demand upon first use and is stored locally for future use. The forms client applet is packaged as JAR files and mainly represents the presentation layer of Oracle forms. The JAR files are also downloaded on the desktop during first use. The forms client must run within a Java Virtual Machine (JVM) and is supplied by Oracle JInitiator.

The Application tier

The Application tier hosts one or more servers that process the business logic and provides communication between the Desktop tier and the database. The Application tier is also commonly known as the middle tier and has the following major components: Web Server, Forms Server, Concurrent Processing Server, Business Intelligence System, Reports Server, Admin Server and Discover Server. These components usually run as services on the Applications tier and may run on one or more nodes.

The Database tier

The Database tier contains the RDBMS Oracle home, along with the Applications database that stores all the data maintained by Oracle Applications. It also contains Oracle Application processing code stored inside the Application database to optimize performance. In essence, the Database tier stores the Oracle database files, Applications database, and executables. The database does not directly communicate with desktop clients but rather works with the Application tier services, which mediate communication between the Desktop and Database tiers.

Oracle Applications File System

Oracle Applications product files, technology stack files, environment files, and common files are held in the file system on the application layer. Typically, the commonly used JAR files are stored on the Desktop tier and the database server holds only the database file.

Solution design and setup

Architecture

Figure 5 highlights the infrastructure setup for this solution.




Figure 5. Site Recovery Manager configuration with CLARiiON CX4-480 storage

Hardware and software

The following table lists the software and hardware used in the architecture.

Table 1. Solution hardware and software

VMware	
Virtual Infrastructure 3	VirtualCenter Server 2.5 ESX 3.5
Site Recovery Manager	Site Recovery Manager 1.0
EMC	
Storage system	CLARiiON CX4-480
Features and software	MirrorView/Synchronous
Storage configuration	3 LUNs: LUN 1: 256 GB, Vision database LUN 2: 200 GB, Apps Stack (or middle tier) LUN 3: 120 GB, OS LUN for 2 VMs (contains 2 virtual disks)
Oracle E-Business 12	
Oracle E-Business 12	<p><u>Virtual Machine 1</u> Database virtual machine (VM) name = "db-ebiz" 2 x vCPU 4 GB RAM Oracle Database 10g R2 Database SID = "VIS" Storage is 2 x volume (2 x LUNs): OS on 1 x VMFS LUN data on 1 x VMFS LUN</p> <p><u>Virtual Machine 2</u> Apps Tier virtual Machine (VM) name = "mt-ebiz" 2 x vCPU 4 GB RAM Storage is 2 x volume (2 x LUNs): OS on 1 x VMFS LUN Apps data on 1 x VMFS LUN</p>
Dell	
Blade servers running ESX	Dell PowerEdge 2950 4 core (2.8 GHz) 16 Gb RAM 4 Gbit QLogic HBA card

Here are the key points regarding the setup.

- The ESX server at the primary site has two virtual machines. A database virtual machine hosts an Oracle Vision database that is supplied with the Rapid install of an Oracle E-Business Suite. The second virtual machine is configured as the middle tier (also known as the Apps tier).
- The Site Recovery Manager requires two VirtualCenter servers. One resides at the primary site and the second manages the standby site. For this solution, we used one ESX server that hosts the primary and secondary site VirtualCenter servers configured in virtual machines. However, it is a best practice to separate primary and secondary site VirtualCenter servers for disaster recovery practices.

-
- EMC CLARiiON CX-4 Fibre Channel storage arrays are used for primary and secondary site storage. For data replication, MirrorView/S is configured to replicate all three LUNs from primary to secondary storage.
 - The EMC Storage Replication Adapter (SRA) is installed on the primary and secondary VirtualCenter servers. The EMC SRA:
 - Automatically discovers the replicated LUNs on the primary side.
 - Facilitates a Site Recovery Manager test recovery workflow on the secondary site during failover - the clone of replicated LUNs (based on snapshots) including the VM clone is registered into the secondary site VirtualCenter server and gets started.
 - The primary site containing all E-Business Suite virtual machines is logically combined into one Site Recovery Manager protection group called "EBusiness Protection Group". Obviously, this protection group is created on the primary VirtualCenter server.
 - At the recovery site VirtualCenter, a recovery plan called "EBusiness Recovery Plan" is created. This recovery plan protects the "EBusiness Protection Group" that is created at the primary site.

Hardware configuration

Let us look at the hardware configuration used in this paper. Configuration of the hardware is carried out in the following order:

- Configure EMC storage arrays at the primary and secondary sites (each array per each site).
- Create three LUNs on the primary site storage and expose to the primary site ESX server. Create VMFS datastores on those LUNs as follows:
 - LUN1: "db-ebiz" datastore for Vision database files
 - LUN2: "mt-ebiz" datastore for the Application tier stack for Oracle E-Business Suite
 - LUN3: "VM_OS" datastore for two OS virtual disks (one for each virtual machine)
- Once the virtual machines are created and configured with operating systems and patches, install Oracle Business Suite 12 with the Vision database.
- Install a VirtualCenter 2.5 server on two Windows virtual machines. These will serve as primary and secondary site VirtualCenter servers.

Note: We used static IP addresses for this exercise. Use of DNS names is the best practice and highly recommended.

- Configure storage replication between primary and secondary site storage arrays by:
 - Configuring MirrorView to replicate three LUNs (Database, Apps tier, OS).
 - Creating a consistency group containing these three LUNs. This unique feature ensures appropriate I/O ordering at the remote site in case of failover.
- Install Site Recovery Manager on the primary and secondary VirtualCenter servers.
- Configure Site Recovery Manager (covered in subsequent sections).
 - Configure primary/secondary connectivity, array managers, and inventory preferences.
 - Create a protection group on the primary Site Recovery Manager.
 - Create a recovery plan on the secondary Site Recovery Manager.
 - Edit the recovery plan and update virtual machine start order.

Site Recovery Manager install

The Site Recovery Manager installation and configuration is simple. At the primary site and secondary site:

- Install a Site Recovery Manager server into a separate database instance on the same guest OS running VirtualCenter server.
- Install the Site Recovery Manager Plug-in into VirtualCenter.

- Install EMC Storage Replication Adapter (SRA).

For more information about Site Recovery Manager installation, please refer to Site Recovery Manager product manuals.

Site Recovery Manager configuration

After installation, a Site Recovery Manager icon is visible on VirtualCenter and is accessed via VI client. Figure 6 shows the primary VirtualCenter and defines the configuration steps required after the initial installation.




Figure 6. Primary Site Recovery Manager server screenshot immediately after installation

As shown in the figure, execute the configuration steps to establish connectivity between the primary and secondary site. At the primary Site Recovery Manager server:

- Configure a connection between the primary and secondary Site Recovery Manager servers.
- Configure the array manager.
- Configure inventory preferences.
- Create a protection group "EBusiness Protection Group" for the virtual machines at the primary site.

At the secondary Site Recovery Manager server:

- Create a recovery plan "EBusiness Recovery Plan" consisting of "EBusiness Protection Group".

-
- Prioritize the virtual machine start order to start a database virtual machine ahead of the middle tier virtual machine.

Planning and testing

Disaster recovery testing is an interdisciplinary concept used to create and validate a practiced logistical plan for how an organization will recover and restore partially or completely interrupted critical function(s) within a predetermined time after a disaster or extended disruption. This logistical plan is commonly referred to as the Business Continuity plan. An effective Business Continuity plan provides a smart balance of business needs vs. cost considering all the risk factors. It is beyond the scope of this paper to detail all the aspects of building a master Business Continuity plan. However, it is nevertheless important to discuss the importance of frequent testing of the disaster recovery plan. The old saying goes that any disaster recovery plan is only as good as your last (successful) test. Indeed, most disaster recovery efforts fall over one of two ways: the team either spends an inordinate amount of effort doing continual testing (not good), or — worse — neglects to test often, the result being an insurance policy that doesn't pay off when you really need it.

So, what makes disaster recovery testing so hard? The simple reason is it's usually very disruptive, expensive in terms of resources, and extremely complex. Site Recovery Manager has the ability to create a “virtual remote recovery image” — storage, virtual machines, even network connections — and test this as a walled-off virtual entity, perhaps even co-resident with production apps that might be running at the remote site. This ability to logically encapsulate and test a complete recovery scenario in a set of virtual machines is simply huge and saves significant time and resources. Site Recovery Manager also eliminates the human error element, thus eventually making disaster recovery testing easier, better, and far more frequently possible compared to the physical environment. Along with a user-friendly user interface to manage disaster recovery tasks such as creating protection groups and recovery plans, Site Recovery Manager also accommodates frequent testing of the recovery plans in a non-disruptive manner.

The recovery plan test for our solution is highlighted below. As part of the solution, EMC's MirrorView software provided the means to maintain synchronous remote copies of production data with MirrorView/S. Customers can ensure rapid, reliable disaster recovery by leveraging VMware SRM using MirrorView/S software to simplify and automate disaster recovery setup, failover, failback and testing, and to ensure the proper execution and management of recovery plans. MirrorView/S is an ideal solution for customers with very demanding RPO requirements and can provide zero-data loss solution for those looking at disaster recovery for database systems. This method can traverse several connectivity topologies including Fibre Channel and iSCSI. Also, since it is storage-based disaster recovery software, applications and servers are free to service users.

Recovery plan test using Site Recovery Manager

The following sequence is executed to validate the recovery plan test:

- On the primary site, ensure that both virtual machines (database and middle tier) are running and you can log in to the E-Business Suite.
- Start a Site Recovery Manager disaster recovery test on the secondary site for recovery plan "EBusiness Recovery Plan". The recovery steps that follow show that the E-Business recovery plan used in this test executes successfully in just few minutes.
- Once the recovery plan test completes, click **Continue** to ensure proper cleanup.




Figure 7. Recovery Plan test execution

As highlighted, Site Recovery Manager facilitates virtually non-disruptive testing of recovery plans. Next, we will discuss the setup and flow for an actual site failover.

Site failover (disaster recovery) test

During a disaster, even a small corner case may cause significant issues in bringing up the remote site. Therefore, it is extremely important to test the failover while the database is under load conditions. In order to represent a realistic scenario, we submitted a long running (~ 3-5 minutes) batch job on the primary site. Once the job is submitted and starts execution, we will initiate site failover from Site Recovery Manager on the remote site. We will use the following sequence to validate the site failover test:

- On the primary site, submit a concurrent manager batch job. This generates the necessary resource consumption for the VM while incurring database I/O.




Figure 8. E-Business Suite General Ledger Job submission

- Once the batch job starts executing, initiate site failover using the recovery plan on the secondary site.




Figure 9. E-Business Suite General Ledger Job status

- As the remote site recovery is initiated, the “refresh” status on the primary site will show “network connection error” as shown in Figure 10.




Figure 10. Failed connection as the primary site goes down

- Once site recovery is complete, log in to E-Business Suite at the recovery site and verify the status of the previously submitted batch job. As expected, it has an “incomplete” status and displays an error.




Figure 11. Failed job that was submitted at the primary site

- To ensure everything is functional, resubmit the same job at the secondary site. As shown in Figure 12, it completes successfully.




Figure 12. Successful job completion at the remote site

Conclusion

In this paper, we demonstrated how VMware's Site Recovery Manager along with EMC's MirrorView/S enables the design of a powerful yet simple disaster recovery solution. Built on the solid foundation of VMware virtualization and CLARiiON CX4's revolutionary technologies, this solution easily resolves typical disaster recovery challenges in an extremely cost-effective manner while providing complete flexibility. Using a multi-tier Oracle E-Business environment residing on EMC's CLARiiON storage, we highlighted how Site Recovery Manager lets you:

- Accelerate recovery for the virtual environment through automation.
- Ensure reliable recovery by enabling non-disruptive testing.
- Simplify recovery by eliminating complex manual recovery steps and centralizing management of recovery plans.

In conclusion, traditional disaster recovery solutions are slow and prone to failures because they involve many manual and complex steps that are difficult to test, and require expensive duplication of the production data center infrastructure to ensure reliable recovery. Working with CLARiiON storage replication technology, Site Recovery Manager can help eliminate complexity and automate the disaster recovery process so that customers can reliably recover from data center outages in hours rather than days.

References

EMC.com

- [EMC CLARiiON CX4 Series datasheet](#)
- [EMC CLARiiON Integration with VMware ESX Server](#) white paper
- [VMware ESX Server Using EMC CLARiiON Storage Systems Solutions Guide](#)
- [Maintaining End-to-End Service Levels for VMware Virtual Machines Using VMware DRS and EMC Navisphere QoS](#) white paper

VMware.com

- [VMware and Oracle database solutions page](#)
- [Planning for the Unplanned](#) white paper
- [VMware virtualization products page](#)
- [VMware Site Recovery Manager product datasheet](#)
- [VMware Site Recovery Manager 1.0 Administration Guide](#)
- [Deployment of Oracle E-Business Suite on VMware Virtual Infrastructure](#) white paper
- [VMware Oracle Database Software Solutions Deployment Guide](#) white paper
- [Simplify Oracle Database management with VMware Infrastructure 3 and EMC CLARiiON Storage](#) white paper