

EMC Backup and Recovery for Oracle Database 11g SAN Performance Enabled by EMC CLARiiON CX4-120 Using the Fibre Channel Protocol and Oracle Automatic Storage Management (ASM)

A Detailed Review

EMC Information Infrastructure Solutions

Abstract

This white paper examines the performance considerations of Oracle databases on Enterprise Flash Drives (EFDs) versus Fibre Channel (FC) disks. Also, it validates a solution for Oracle Database 11g backup and recovery on the EMC[®] CLARiiON[®] CX4-120 storage array using EMC SnapView[™] over the Fibre Channel Protocol (FCP) and Oracle Automatic Storage Management (ASM).

August 2010

Copyright © 2010 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: H7345

Table of Contents

Executive summary.....	4
Business case.....	4
Product solution.....	4
Key results.....	4
Introduction.....	5
Purpose.....	5
Scope.....	5
Audience.....	5
Terminology.....	6
Technology overview.....	8
Introduction.....	8
EMC CLARiiON storage array.....	8
EMC Replication Manager.....	8
EMC SnapView.....	9
Oracle software stack.....	9
Disk drives.....	9
Enterprise Flash Drives (EFDs).....	10
Configuration.....	11
Overview.....	11
The site.....	11
Environment profile.....	11
Physical environment.....	11
Storage layout.....	12
Disk layout.....	13
Network architecture.....	13
Hardware resources.....	14
Software resources.....	15
Test and validation.....	16
Introduction.....	16
Store solution.....	16
Basic backup solution.....	18
Advanced backup solution.....	19
Conclusion.....	23
Introduction.....	23
Key points.....	23
Next steps.....	24
References.....	25
Reference documents.....	25

Executive summary

Business case Midsize enterprises face the same challenges as their larger counterparts when it comes to managing database environments. These challenges include:

- Control over resource utilization and scaling
- Lack of sufficient IT resources to deploy, manage, and maintain complex environments at the departmental level
- Reduce power, cooling, and space requirements

Unlike large enterprises, midsize enterprises are constrained by smaller budgets and cannot afford a custom, one-off solution. This makes the process of creating a database solution for midsize enterprises even more challenging than for large enterprises. Thus, it is increasingly necessary to provide a low-cost, performing, and scalable solution for midsize Oracle customers.

Product solution

The EMC® CLARiiON® CX4-120 multi-protocol array provides a low-cost, high-performing, and scalable solution for midsize Oracle customers. This array provides access to Oracle Database 11g using Oracle Automatic Storage Management (ASM) on a storage area network (SAN) using the Fibre Channel Protocol (FCP).

This white paper examines the performance considerations of Oracle databases on Enterprise Flash Drives (EFDs) versus Fibre Channel (FC) disks.

This white paper also summarizes the best practices of backup, recovery, and performance that were discovered and validated during the solution testing. This was validated on a configuration that included Oracle Database 11g on the CLARiiON CX4-120 back-end array using SnapView™ over FCP and Oracle ASM.

Key results

This white paper demonstrates the following benefits for the solution:

- The functional, performance, resiliency, and scalability capabilities of an Oracle software stack. The Oracle software is installed locally on the database server and the database is managed by Oracle ASM, which accesses the storage using FCP on SAN storage.
- The use of EMC SnapView to enable the logical storage backup and recovery of an Oracle 11g production database while offloading all the performance impacts of the backup operation from the production server. This demonstrates significant performance and manageability benefits in comparison to normal Oracle Recovery Manager (RMAN) backup and recovery.
- The significant performance, scalability, and efficiency benefits of FCP and Oracle ASM, when compared to other storage configurations such as NFS or NAS.

-
- The significant increase in performance and other advantages by using Enterprise Flash Drives (EFDs), when compared to FC disks.
-

Introduction

Purpose

This white paper provides an overview of the backup, recovery, and performance best practices that were discovered and validated for Oracle Database 11g on the EMC CLARiiON CX4-120 storage array using EMC SnapView over the FCP and Oracle ASM.

Scope

This white paper covers the following topics:

- User scalability by applying a scale-up OLTP TPC-C-like workload using Quest Benchmark Factory for databases, scaling users to the point where query performance falls below the TPC-C standard.
 - Multiple architectural mixes were considered and performance tested, including the following examples:
 - All Oracle Database objects are managed by Oracle ASM, which accesses the storage over FCP/SAN. This can provide higher performance and lower latency I/O than other storage configurations such as NFS/NAS.
 - Database storage on EFDs. Usually, small databases may have high performance needs. EFD is at a price where it is more economically feasible to consider it for use in smaller database configurations. This paper looks at the performance that can be gained in configurations with lower disk numbers.
 - Basic backup and recovery using RMAN, the built-in backup and recovery tool provided by Oracle.
 - Advanced backup and recovery using EMC SnapView.
-

Audience

The intended audience for the white paper is:

- Internal EMC personnel
 - EMC partners
 - Customers
-

Terminology

The terms used in this white paper are defined in the following table:

Term	Definition
Automatic Storage Management (ASM)	Oracle ASM is a volume manager and a file system for Oracle Database files. It supports single-instance Oracle Database and Oracle Real Application Clusters (Oracle RAC) configurations.
Enterprise Flash Drive (EFD)	A drive that stores data using flash memory and contains no moving parts.
Serial Advanced Technology Attachment (SATA) drive	SATA is a newer standard for connecting hard drives into computer systems. SATA is based on serial-signaling technology while Integrated Drive Electronics (IDE) hard drives use parallel-signaling technology.
Basic backup and recovery	<p>A solution component that provides backup and recovery functionality through the operating system and database server software stack. In the case of Oracle, this is RMAN.</p> <p>Basic backup and recovery uses the database server's CPUs for all I/O and processing of backup and recovery operations.</p>
Advanced backup and recovery	<p>A solution component that provides backup and recovery functionality through the storage layer using specialized hardware or software.</p> <p>Advanced backup and recovery provides the following benefits:</p> <ul style="list-style-type: none">• Offloads the database server's CPUs from the I/O and processing requirements of the backup and recovery operations.• Superior Mean Time to Recovery (MTTR) through the use of logical storage layer replication (commonly referred to as snapshots).

Scale-up OLTP	It uses an industry-standard OLTP benchmark against a single-instance database. The comprehensive performance testing is performed to validate the maximum achievable performance using the solution stack of hardware and software. Two instances of scaled-up OLTP are used: one for the FC disk and the other for EFDs.
---------------	--

Technology overview

Introduction

This section provides an overview of the technologies that are used in this solution:

- EMC CLARiiON storage array
 - EMC Replication Manager
 - EMC SnapView
 - Oracle software stack
 - Disk drives
 - EFDs
-

EMC CLARiiON storage array

EMC CLARiiON storage array products offer a flexible architecture and multi-protocol connectivity. This enables connectivity over IP/Ethernet, iSCSI, and Fibre Channel SAN environments.

EMC CLARiiON CX4-120 provides entry-level networked storage for departmental applications and midsize organizations. The CX4-120 combines CLARiiON five 9s availability with innovative technologies like Fully Automated Storage Tiering, EFDs, Virtual Provisioning™, 64-bit operating system, and multi-core processors. The CX4-120 scales from 5 up to 120 TB of capacity.

The SAN connectivity protocols provided by the CLARiiON are described in the following table:

Protocol	Provided by
iSCSI	EMC CLARiiON storage processors
FCP	EMC CLARiiON storage processors

EMC Replication Manager

EMC Replication Manager manages EMC point-in-time replication technologies through a centralized management console. Replication Manager coordinates the entire data replication process—from discovery and configuration to the management of multiple application-consistent disk-based replicas. Replication Manager can auto-discover the replication environment and enables streamlined management by scheduling, recording, and cataloging replica information.

With Replication Manager, the right data can be put in the right place at the right time—on-demand or based on the schedules and policies that the customer defines. This application-centric product allows simplification of replica management with application consistency.

Replication Manager consists of the software components. Some hosts may need only one of these components, while others may need two or three components. To decide which components to install on each machine, consider the tasks performed

by each:

- Replication Manager Server — The Replication Manager Server component manages replicas. Typically, one RM Server can meet the needs of several hosts. The RM Server is installed on supported Windows hosts only.
- Replication Manager Agent — The Replication Manager Agent component prepares a host machine to create or mount a replica by performing tasks such as placing an Oracle database into hot backup mode before replication. Although it could be installed elsewhere, in the tested configuration, the RM software is installed on the host that the RM Server needs to interact with.
- Replication Manager Console — The RM Console component is the graphical user interface that allows the user to control the server and agent software locally or remotely.

EMC SnapView SnapView is a storage-system-based software application that allows you to create a copy of a LUN by using either clones or snapshots. A clone is an actual copy of a LUN and takes time to create, depending on the size of the source LUN. A snapshot is a virtual point-in-time copy of a LUN and takes only seconds to create.

How SnapView works

A snapshot is a virtual LUN that allows a secondary server to view a point-in-time copy of a source LUN. The point in time can be determined when starting a SnapView session. The session keeps track of the source LUN's data at a particular point in time. During a session, the production server can still write to the source LUN and modify data. When this happens, the software stores a copy of the original point-in-time data on a reserved LUN in the SP's reserved LUN pool. This operation is referred to as copy-on-first-write because it occurs only when a data chunk is first modified on the source LUN. As the session continues and additional I/O modifies other data chunks on the source LUN, the amount of data stored in the reserved LUN pool grows. If needed, you can increase the size of the reserved LUN pool.

Oracle software stack The Oracle software stack covered in this solution consists of:

- Oracle Database 11g Release 2 Enterprise Edition
- Oracle ASM

Disk drives The general recommendations for mechanical disk drives are as follows:

- Drives with higher revolutions per minute (rpm) provide higher overall random-access throughput and shorter response times than drives with slower rpm. For optimum performance, higher-rpm drives are recommended for datafiles and tempfiles as well as online redo logfiles.
- Because of significantly better performance, FC disks are always recommended for storing datafiles, tempfiles, and online redo logfiles.
- SATA II drives have slower response and rotational speed, and moderate performance with random I/O. However, they are less expensive than the FC

disks for the same or similar capacity.

- SATA II drives are usually the best option for storing archived redo logs and the fast recovery area. In the event of high performance requirements for backup and recovery, FC disks can also be used for this purpose.

**Enterprise
Flash Drives
(EFDs)**

EFDs can be used to dramatically improve the cost, performance, efficiency, power, space, and cooling requirements of Oracle databases stored on EMC CLARiiON.

These drives are especially well suited for latency-sensitive applications that require consistently low read/write response times. This is an extremely common condition in Oracle databases. If this is the case, move the read-write heavy (hot) tables and latency-sensitive data to EFDs.

Configuration

Overview

All database objects, including datafiles, tempfiles, controlfiles, online redo logfiles, and archive log files, are stored on ASM diskgroups that reside on SAN storage.

The site

The site consists of:

- A physically booted server hosting a single instance Oracle Database 11g R2.
- A CLARiiON CX4-120 connected to the server through the SAN, with EMC SnapView used to provide an advanced backup solution.
- The Oracle Database 11g server connected to the client and storage networks.

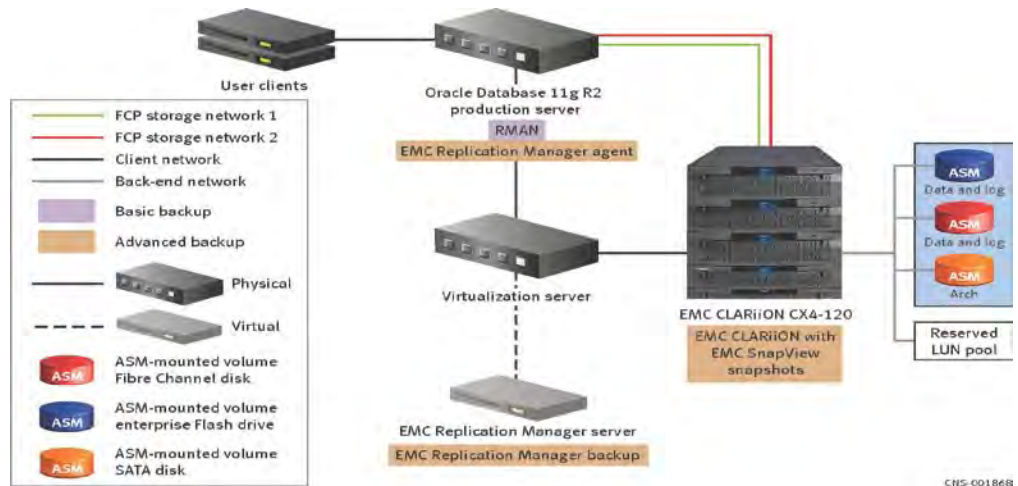
Environment profile

The white paper was validated with the following environment profile:

Profile characteristic	Value
Database size	350 GB
Database profile	OLTP
Workload profile	TPC-C industry-standard benchmark
User scaling	To maximum stable load
Storage network connectivity	4 GbE FCP

Physical environment

The following diagram illustrates the overall physical architecture of the environment.



Storage layout The following table lists the Oracle file system allocation in the usage case while using FC disks:

What	File system	File system mount point
Oracle datafiles	ASM	+DATA_FC
Oracle tempfiles		
Oracle online redo logfiles		
Oracle controlfiles		

The following table lists the Oracle file system allocation in the usage case while using EFDs:

What	File system	File system mount point
Oracle datafiles	ASM	+DATA_EFD
Oracle tempfiles		
Oracle online redo logfiles		
Oracle controlfiles		

In both usage cases, all of these files are accessed using ASM. ASM is used to store the database objects requiring high performance. All datafiles, online redo logfiles, controlfiles, and tempfiles are contained in an ASM diskgroup stored on RAID 5.

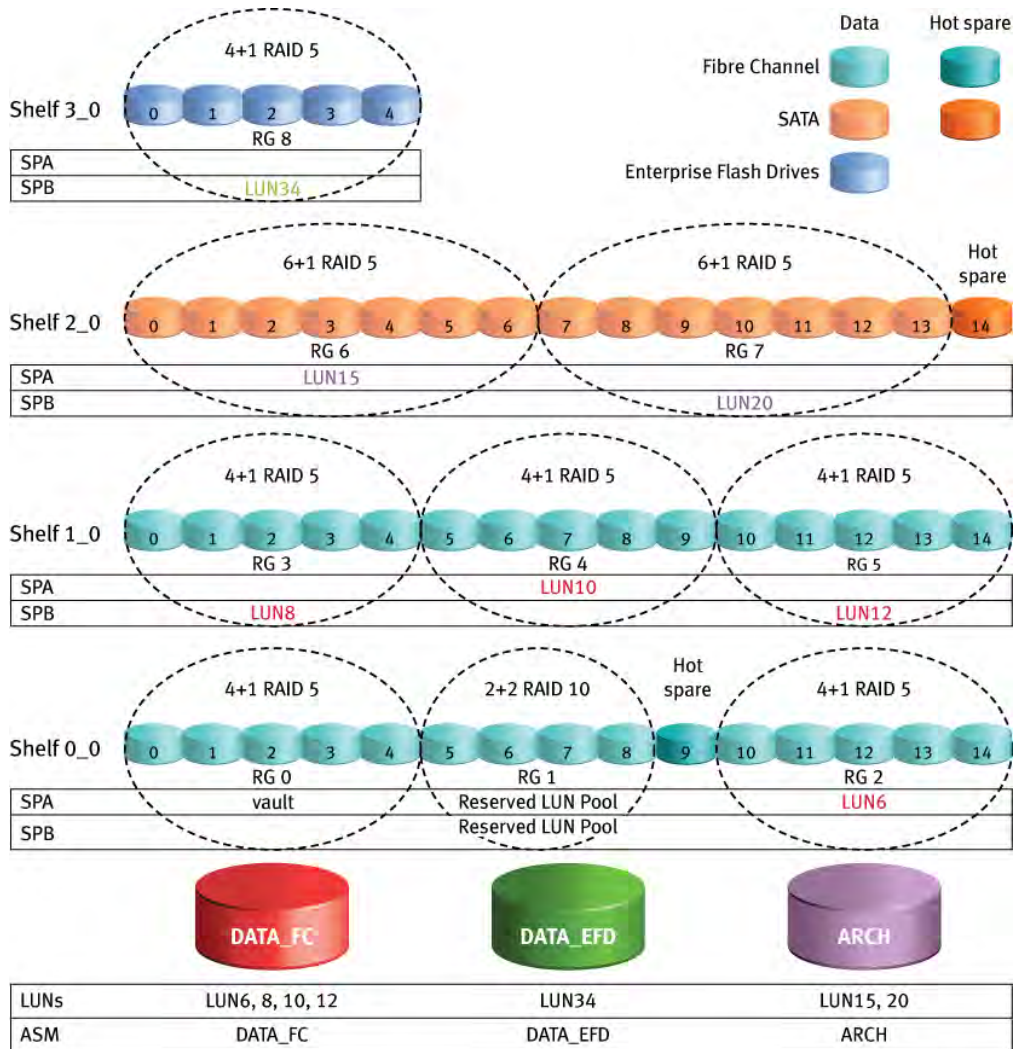
For a backup and recovery usage case, the following objects are stored on the SATA disk:

What	File system	File system mount point
Backup target	ASM	+ARCH
Archived log files		

To facilitate the hot backup process and to allow media recovery, archive log mode is enabled during the validation of the backup and recovery usage case. However, for optimal performance, it is not enabled when conducting the performance testing by using FC disk and EFD. Also, backup testing was performed with the source database on FC disk, not EFDs.

Disk layout

The following diagram illustrates the disk layout of the environment.



CNS-001873

Network architecture

The network connectivity for the solution is as follows:

- TCP/IP provides network connectivity for the management network.
- An Oracle Database 11g server is connected to the back-end EMC CLARiiON array via a FCP SAN.

Hardware resources

The hardware used to validate the solution is listed in the following table:

Equipment	Quantity	Configuration
EMC CLARiiON CX4-120	1	<ul style="list-style-type: none">• 2 storage processors• 4 FCP network connections per storage processor• 2 FC shelves (30 FC 300 GB 15k rpm disks)• 1 EFD shelf (5 400 GB 4 GB FC EFDs)• 1 SATA shelf (15 SATA 1 TB 7200 rpm disks)• FLARE® version 04.29.000.5.003
FC switches	2	<ul style="list-style-type: none">• 16 ports• 4 Gb throughput
Database server	1	<ul style="list-style-type: none">• 2 2.66 GHz Intel Pentium 4 quad-core processors• 24 GB of RAM• 146 GB 15k internal SCSI disks• 2 onboard GbE Ethernet NICs• 2 additional Intel PRO/1000 PT quad-port GbE Ethernet NICs• 2 SANblade QLE2462-E-SP 4 Gb/s dual-port FC HBAs (4 ports in total)
Virtualization server	1	<ul style="list-style-type: none">• 4 2.86 GHz AMD Opteron quad-core processors• 32 GB of RAM• 2 146 GB 15k internal SCSI disks• 2 onboard GbE Ethernet NICs• 3 additional Intel PRO/1000 PT quad-port GbE Ethernet NICs• 2 SANblade QLE2462-E-SP 4 Gb/s dual-port FC HBAs (4 ports in total)

Software resources

The software used to validate the solution is listed in the following table:

Software	Version
Red Hat Enterprise Linux	5.4
Microsoft® Windows® Server Enterprise Edition	2003 R2
Oracle Database Enterprise Edition	11g (11.2.0.1.0)
Quest Benchmark Factory for databases	5.8.0
EMC Navisphere® Management	6.29.0.6.34
EMC FLARE	04.29.000.5.003
EMC Replication Manager	5.2.3.0

Test and validation

Introduction

This white paper provides a summary and characterization of the tests that were performed to validate the solution. The testing goal was to characterize the end-to-end solution and the component subsystem response under a reasonable load, representing the market for Oracle Database 11g on Red Hat Enterprise Linux 5 with CLARiiON CX4-120 over FCP.

Store solution

The store solution component was designed as a set of performance measurements to determine the bounding point of the solution stack in terms of performance. A reasonable amount of fine tuning was performed to ensure that the performance measurements achieved were consistent with the real-world and best-of-breed performance.

Test procedures

The following procedures were used to validate the store solution component.

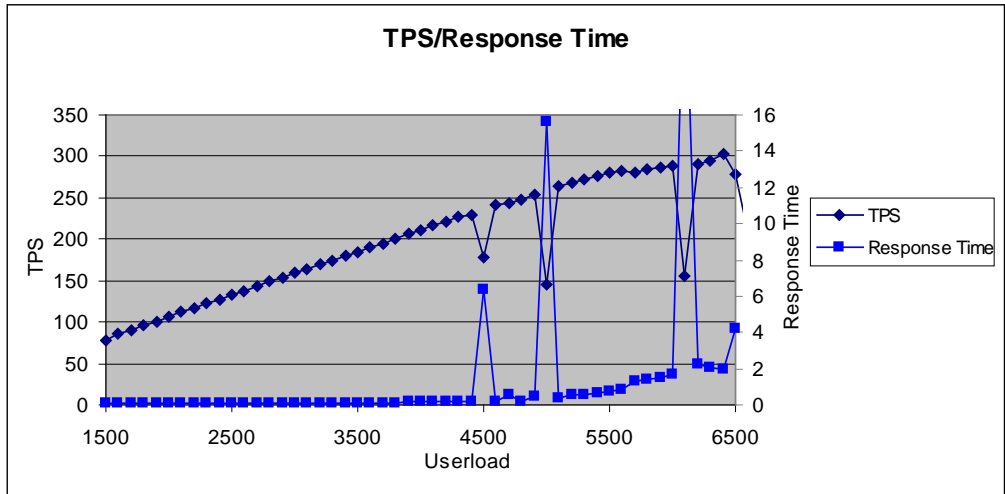
Step	Action
1	Close all Benchmark Factory agents that are running.
2	Restart the client machines.
3	Stop the database instance.
4	Initiate the Benchmark Factory console and agents on the client machines.
5	Start the Benchmark Factory job.
6	Monitor the progress of the test.
7	Capture the results after the test is completed.

Test results

The summary of the test results for the physically booted Oracle Database 11g single-instance database ASM configuration is shown in the following charts. The memory on the database server was 24 GB.

- Performance of FC disks:
 - Users: 6,400
 - TPS: 302
 - Response Time: 1.992 seconds


The TPS peaked at 6,400 users. This was the highest user count that had a response time of less than 2 seconds.



FC disk diagram

- Performance of EFDs:
 - Users: 7,400
 - TPS: 352.6
 - Response Time: 1.607 second

The TPS peaked at 7,400 users. This was the highest user count that had a response time of less than 2 seconds.



EFDs diagram

Basic backup solution

The basic backup solution component demonstrates that the Oracle Database 11g configuration is compatible with the RMAN disk-to-disk backup. FC disk was used to store the source database during the backup testing.

The backup tests are performance tests, where the performance level was observed while RMAN backup/restore was performed. The restore is a functionality test, but the amount of time required to perform the RMAN restore was tuned and measured. The transactions restored and recovered are measured to ensure that there is no data loss.

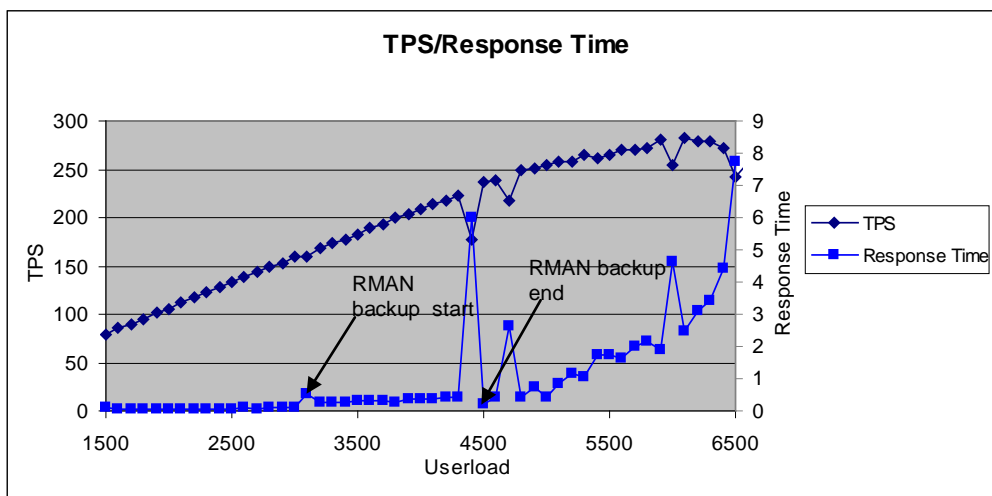
Test procedures

The following procedure was used to validate the basic backup solution component.

Step	Action
1	Close all the Benchmark Factory agents that are running.
2	Close the Benchmark Factory console.
3	Restart the Benchmark Factory console and agents.
4	Shut down the database server.
5	Start the listener and restart the database instance.
6	Start the Benchmark Factory test with the user load ranging from 1,500 to 6,500 with intervals of 100.
7	When the user load reaches the 3,100 th iteration, initiate the RMAN backup on the database server and monitor the performance impact on the production database.
8	Verify that the RMAN backup is completed successfully and allow the test to complete.
9	When the Benchmark Factory test is complete, capture the results.
10	Restore the database using RMAN.
11	Recover the database.
12	Capture the time taken to restore the database.

Test results

The RMAN backup operation was performed while the Benchmark Factory load was running. The RMAN backup started at a user load of 3,100 and ended at a user load of 4,500. When RMAN was initiated at the user load of 3,100, there was a moderate increase in the response time and a moderate decrease in the transaction throughput as shown in the following chart.



Backup and recovery summary

The following table provides a summary of the backup and restore test.

Test run duration	7 hours and 10 minutes
User load range	1,500 – 6,500 with intervals of 100
Profile	Mteoradb52
Driver	Oracle
Start time of test	Fri May 21 19:59:00 CST 2010
End time of test	Sat May 22 3:09:00 CST 2010
RMAN backup start time	22:19:00 CST 2010 at user load 3,100
RMAN backup end time	00:03:00 CST 2010 at user load 4,500
Total time for RMAN backup	1 hours and 44 minutes
Total time for RMAN restore	1 hours and 9 minutes
Total time for RMAN recovery	1 hour and 36 minutes

Basic backup and recovery conclusion

RMAN provided a reliable high-performance backup solution for Oracle Database 11g. However, the time required to restore the database was significant. The space required to store the RMAN backup was also significant.

Advanced backup solution

The purpose of the advanced backup solution component was to demonstrate that the validated configuration is compatible with CLARiiON SnapView using Replication Manager.

The backup test run was a performance test. The performance of the database was tested while performing a hot backup using an EMC CLARiiON SnapView snapshot

with Replication Manager.

The restore was a functionality test. The amount of time that is required to perform the SnapView restore was tuned and measured.

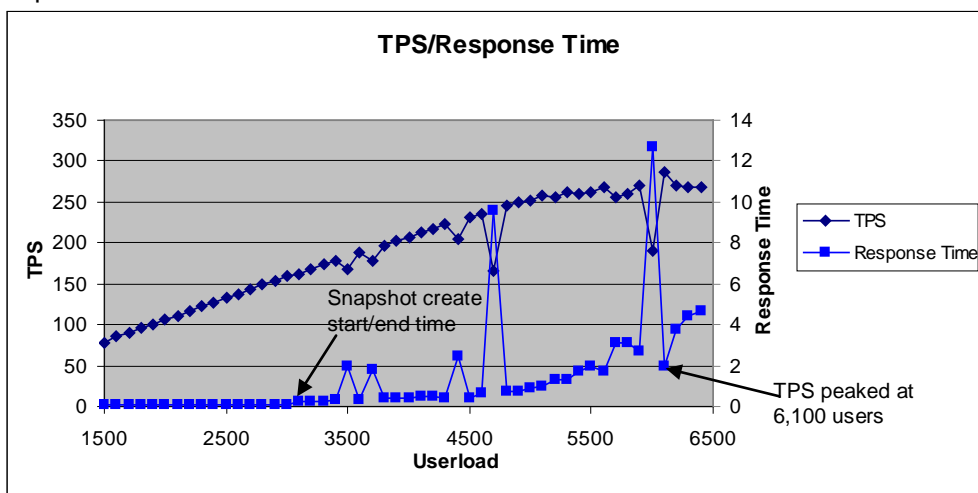
Test procedures

The following procedures were used to validate the advanced backup solution component:

Step	Action
1	Configure Replication Manager.
2	Register the production hosts, mount hosts, and storage in Replication Manager.
3	Create the application set in Replication Manager for the database to be replicated.
4	Create a job in the Replication Manager console to take the SnapView snapshot.
5	Close all the Benchmark Factory agents that are running.
6	Close the Benchmark Factory console.
7	Restart the Benchmark Factory console and agents.
8	Stop and restart the database instance.
9	Start the Benchmark Factory test with a user load ranging from 1,500 to 6,500.
10	When the user load reaches an iteration of 3,100, take a snapshot of the database by running the job in the Replication Manager console.
11	Monitor the performance impact on the production database.
12	When the Benchmark Factory test is complete, capture the results.
13	Shut down the database.
14	Stop and disable the ASM instances.
15	Dismount the data diskgroups.
16	Restore the database using Replication Manager.
17	Recover the database.
18	Capture the time taken to restore the database.

Test results

The snapshot job was initiated with a user load of 3,100. For the optimum reported iteration, the TPS peaked at 6,100 users. This was the highest user count that had a response time of less than 2 seconds.



Advanced backup and restore summary

The following table provides a summary of the advanced backup and restore test.

Test run duration	7 hours and 16 minutes
User load range	1,500 – 6,500 with intervals of 100
Profile	Mteoradb52
Driver	Oracle
Start time of test	Tue May 25 18:45:00 CST 2010
End time of test	Wed May 26 02:01:00 CST 2010
Snapshot create start time	Tue May 25 21:12:26 CST 2010 at user load 3,100
Snapshot create end time	Tue May 25 21:16:26 CST 2010 at user load 3,100
Total time for taking snapshot	4 minutes
Total time for restore operation	1 minute and 19 seconds
Total time for recovery	1 hour and 40 minutes

Advanced backup and recovery conclusion

The CLARiiON SnapView feature works with Oracle Database 11g for our validated configuration and can be performed successfully using Replication Manager.

In most of the test runs, a very slight performance hit was observed during the backup. However, this was temporary and the performance recovered to the expected levels within a short span of time.

The restore of a SnapView snapshot hot backup is much faster than the RMAN disk-to-disk restore because of the use of the snapshot restore feature of SnapView.

Conclusion

Introduction

The EMC CLARiiON storage array's high-availability features combined with EMC's proven storage technologies provide a high-performance, high-flexibility, and low-risk storage system for the Oracle Database 11g over FCP.

Key points

The table below summarizes the key points that this solution addresses.

Key point	Solution objective
Reduced total cost of ownership	This solution reduces the load on the database server CPU by using EMC Replication Manager and EMC SnapView to carry out a physical backup of an Oracle Database 11g production database while offloading all performance impacts of the backup operation from the production server.
Reduced complexity of backup and recovery	EMC Replication Manager eliminates the requirement for the customer to write scripts or to manually perform replication tasks. These tasks can be fully automated and managed by Replication Manager.
Improved performance	<p>The incorporation of EFDs into EMC CLARiiON provides a new tier-0 storage layer that can deliver very high I/O performance at a very low latency, which can dramatically improve OLTP throughput and maintain very low response time.</p> <p>Traditional magnetic disk drive technology no longer defines the performance boundaries for mission-critical storage environments. The costly approach of spreading workloads, over dozens or even hundreds of underutilized disk drives, is no longer necessary. Based on the Oracle Database 11g environment, EFDs can lead to increased performance and savings in power, cooling, and data center floor space requirements.</p> <p>EMC CLARiiON now combines the performance and power efficiency of EFDs with traditional disk drive technology in a single array, managed with a single set of software tools, to deliver advanced functionality, ultra-performance, and extended storage tiering options.</p>
Business continuity	Advanced backup and recovery with EMC SnapView dramatically improves the mean time to recovery (MTTR) by reducing the time required for the restore operation. Further, as the backup operation has minimal impact on the database server performance, the backup can be run more often. This means that the recovery operation is also optimized because

	fewer archived logs are applied.
Robust performance and scaling	The performance testing carried out by EMC utilizes an industry-standard OLTP benchmark, but does so without exotic tunings that are not compliant with best practices. In addition, real-world configurations are used. This ensures that the configuration you choose to run on your application will be predictable and reliable.

Next steps

EMC can help to accelerate assessment, design, implementation, and management while lowering the implementation risks and costs of an end-to-end solution for an Oracle Database 11g environment.

To learn more about this and other solutions, contact an EMC representative or visit

<http://www.emc.com/solutions/application-environment/oracle/solutions-for-oracle-database.htm>

References

Reference documents

For additional information, see the following documents:

Proven Solution Guide

EMC Unified Storage for Oracle Database 11g – Enabled by EMC CLARiiON, EMC Celerra, VMware vSphere, and VMware High Availability (HA) Using FCP and NFS

White Paper

Leveraging EMC CLARiiON CX4 with Enterprise Flash Drives for Oracle Database Deployments—Applied Technology
