

NONDISRUPTIVE BACKUP OF ORACLE DATABASE WITH EMC UNIFIED STORAGE

A Detailed Review

EMC GLOBAL SOLUTIONS

Abstract

This white paper explains the backup of an EMC® Celerra® SnapSure™ checkpoint of the production Oracle database using the RMAN utility with Block Change Tracking (BCT) enabled. The paper also reveals how to use a proxy Celerra SnapSure checkpoint to isolate impacts on the production database nodes compared with the normal RMAN backup.

February 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

EMC believes the information in this publication is accurate 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.

VMware, ESX, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number h8133

Table of Contents

Executive summary	4
Business case	4
Solution overview.....	4
Key results.....	4
Introduction	5
Purpose.....	5
Scope.....	5
Audience.....	5
Terminology	6
Technology overview	8
Introduction	8
EMC Celerra unified storage platform.....	8
EMC SnapSure.....	8
Oracle software stack	9
Oracle DNFS	9
Configuration	10
Overview	10
Environment profile	10
Physical environment	11
Storage layout	11
Disk layout	13
Hardware resources.....	14
Software resources.....	15
Test and Validation	16
Introduction	16
Preparation work	16
Scenario 1 Baseline RMAN incremental backups to NFS mount	17
Scenario 2 Storage-based replication RMAN backup to NFS mount.....	21
Verification procedure for using BCT.....	26
Comparison.....	27
Conclusion	30
Summary.....	30
Findings	30
Next steps	30
References	31
White papers	31
Other documentation	31

Executive summary

Business case

Oracle Database is the leading database management system across different operating systems. It is also a critical component in many organizations for maintaining data integrity and availability to their Oracle database environment. The database is the basis of most business applications; downtime and slow performance of the database can result in client complaints and business loss.

One of the key challenges that Oracle database administrators (DBAs) and storage administrators face is how to maintain performance during the backup window.

This white paper covers the general approach of how to use EMC® Celerra® SnapSure™ checkpoint technology to offload the Oracle Recovery Manager (RMAN) backup process from the database server. The paper also demonstrates the associated benefits that can be expected. In particular, the paper covers how the process can be extended to incorporate the use of the Block Change Tracking (BCT) feature of Oracle Database 11g, which can significantly enhance the speed of RMAN incremental backups.

Solution overview

This solution uses Celerra NS-960 and 10 GbE network fabric to build a comprehensive backup solution on a physically booted four-node Oracle Real Application Cluster (RAC) 11g R2 DNFS OLTP environment. This white paper demonstrates best practices and a strategy for efficient incremental RMAN backup with BCT enabled by using a virtualized backup host to offload workload on the production environment. The following scenarios are included:

- Baseline RMAN incremental backups to NFS mount
- Storage-based replication backup to NFS mount using RMAN and Celerra SnapSure

All backup and recovery scenarios apply a simulated OLTP workload generated by Quest Benchmark Factory while the backup and recovery processes are running. The performance impact on a four-node Oracle RAC database is illustrated during database backup.

Key results

This solution has the following benefits:

- Using RMAN and a Celerra SnapSure checkpoint increases the transactions per second (TPS) while cutting down the response time as shown in Figure 10. Meanwhile, it shortens the entire backup duration on the production database as shown in Figure 11
- Significantly reduces the I/O bandwidth with storage-based backup as shown in Figure 12
- Increases flexibility, manageability, and ease of use in a network-attached storage (NAS) environment
- Simplifies management and reduces cost with VMware® virtualization

Introduction

Purpose

The purpose of this solution is to:

- Demonstrate the performance benefits achieved by offloading backup operations from production systems to proxy nodes and reducing the time spent in hot backup mode
- Validate significant performance gains and manageability benefits of using Oracle RMAN together with storage-based replication on Celerra SnapSure for database backup, restore, and recovery, in comparison to the normal Oracle RMAN backup and recovery

Scope

The scope of this white paper covers the following topics:

- Basic incremental backup and recovery using RMAN with BCT enabled, the built-in backup and recovery tool provided by Oracle
- The combination of storage-based replication enabled by EMC Celerra SnapSure checkpoint and Oracle RMAN in incremental backup and recovery of the Oracle database

Audience

The intended audience for the white paper is:

- EMC employees
- EMC partners
- Customers

This white paper assumes that readers are already familiar with the use of RMAN, including the concept of enabling and using BCT to improve the efficiency of performing incremental backups with RMAN.

Terminology

Table 1 defines terms used in this document.

Table 1. Terminology

Term	Definition
Block Change Tracking (BCT)	A Block Change Tracking file is an enhancement that can be optionally enabled to improve incremental backup performance by recording changed blocks in each datafile in a change tracking file. If change tracking is enabled, RMAN uses the change tracking file to identify changed blocks for incremental backup, avoiding the need to scan every block in the datafile.
Cumulative incremental backup	An incremental backup that backs up all the blocks changed since the most recent backup at level 0. When recovering with cumulative incremental backups, only the most recent cumulative incremental backup must be applied.
Differential incremental backup	An incremental backup that backs up all blocks that have changed since the most recent backup at level 1 or level 0. For example, in a differential level 1 backup RMAN determines which level 1 or level 0 incremental backup is most recent and then backs up all blocks changed since that backup. Differential backups are the default type of incremental backup. When recovering using differential incremental backups, RMAN must apply all differential incremental level 1 backups since the restored datafile backup.
Direct NFS (DNFS) client	Standard NFS client software, provided by the operating system, is not optimized for Oracle Database file I/O access patterns. Oracle Database 11g Direct NFS Client integrates the NFS client functionality directly in the Oracle software. Oracle Database will access files stored on the NFS server directly through the integrated Direct NFS Client, eliminating the overhead imposed by the operating system kernel NFS (KNFS).
Hot backup mode	A database mode initiated when you issue the ALTER TABLESPACE ... BEGIN BACKUP or ALTER DATABASE ... BEGIN BACKUP command before taking an online backup. When making a user-managed backup of datafiles in an online tablespace, you must place the tablespace in hot backup mode to avoid a fractured block. In hot backup mode, each time a block in the buffer cache becomes dirty, the database must write an image of the changed block to the redo log file, in addition to recording the changes to the data.
Incremental backup	An RMAN backup in which only modified blocks are backed up. Incremental backups are classified by level. A level 0 incremental backup performs the same function as a full backup in that they both back up all blocks that have ever been used. The difference is that a full backup do not affect blocks backed up by subsequent incremental backups, whereas an incremental backup affects blocks backed up by subsequent incremental backups. Incremental backups at level 1 back up only blocks that have changed since previous incremental backups. Blocks that have not

Term	Definition
	changed are not backed up. An incremental backup can be either a differential incremental backup or a cumulative incremental backup.
RMAN backup and recovery	<p>An Oracle solution component that provides backup and recovery functionality through the operating system and database server software stack.</p> <p>RMAN backup and recovery uses the database server's CPUs for all I/O and processing of backup and recovery operations.</p>
Storage-based replication and recovery	<p>A solution component that provides backup and recovery functionality through the storage layer using specialized software and hardware. Storage-based replication and recovery can offload the overhead associated with running backups on the production node to the virtualized backup host and provide Superior Mean Time to Recovery (MTTR) through the use of logical storage layer replication (commonly referred to as snapshots).</p>
VMware ESX®	<p>The foundation for a dynamic, self-optimizing IT infrastructure. It is a robust, production-proven virtualization layer that abstracts processor, memory, storage, and networking resources into multiple virtual machines. VMware ESX allows enterprises to dramatically reduce hardware and operating costs by sharing resources across a virtual environment.</p>

Technology overview

Introduction

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

- EMC Celerra unified storage platform
- EMC SnapSure
- Oracle software stack
- Oracle DNFS

EMC Celerra unified storage platform

EMC Celerra unified storage offers a flexible architecture and multi-protocol connectivity. EMC Celerra NS-960 is a unified storage system that brings advanced failover and fully automated storage tiring to multi-protocol environments. With Celerra NS-960, you can connect to multiple storage networks through NAS, iSCSI, Fibre Channel storage area network (FC SAN), and EMC Celerra Multi-Path File System (MPFS).

The key features provided by Celerra are described in Table 2:

Table 2. Celerra key features

Features	Provided by
NAS	Network File System (NFS) and Common Internet File System (CIFS) protocols
iSCSI storage	Celerra Data Mover
SAN storage	Fibre Channel Protocol (FCP) through the back-end EMC CLARiiON® storage array

EMC SnapSure

SnapSure creates a logical point-in-time image (checkpoint) of a Production File System (PFS) that reflects the state of the PFS at the point in time when the checkpoint is created.

SnapSure can maintain a maximum of 96 read-only checkpoints and 16 writeable checkpoints per PFS while allowing PFS applications continued access to realtime data.

How SnapSure works

SnapSure uses a copy-on-first-modify principle. A PFS consists of blocks of data. When a block within the PFS is modified, a copy that contains the block's original contents is saved to a separate volume called a SavVol. Subsequent changes made to the same block in the PFS are not copied into the SavVol. The original block from the PFS in the SavVol and the unchanged PFS blocks remaining in the PFS are read by SnapSure according to a bitmap and blockmap data-tracking structure. These combined blocks provide a complete point-in-time image.

Oracle software stack

The Oracle software stack covered by this solution consists of:

- Oracle Grid Infrastructure 11gR2
- Oracle Database 11gR2
- Oracle DNFS

Oracle DNFS

DNFS is a feature introduced in Oracle RAC 11g. It integrates the NFS client directly inside the database kernel instead of the operating system kernel. DNFS provides significant performance, manageability, and efficiency benefits over KNFS.

Better performance

TPS with DNFS is higher than with KNFS, and this enables organizations to gain more output from the same infrastructure. CPU costs on both the database server and the file server are lower. In addition, port scaling with DNFS is much better, enabling higher bandwidth and higher scaling.

High availability

Load balancing and high availability (HA) are managed internally within the DNFS client.

Concurrent I/O

The DNFS client performs concurrent I/O by bypassing the operating system. The benefits include:

- Better performance because of the reduction of memory consumption and CPU utilization
- Consistent NFS performance across all operating systems

Optimized for database workloads

DNFS is optimized for database workloads and supports asynchronous I/O, which is suitable for most databases. It delivers optimized performance by automatically load balancing across all available paths. Load balancing in DNFS is better than with KNFS.

Configuration

Overview

All database objects are stored on NFS mounts. Datafiles, tempfiles, control files, online redo log files, and archive log files are accessed using the DNFS client.

Environment profile

The white paper was validated with the environment profile listed in Table 3.

Table 3. Environment profile

Profile characteristic	Value
Database size	1 TB
Database profile	OLTP
Network connectivity for DNFS implementation	10 GbE
Network connectivity for RAC interconnect	10 GbE
OLTP workload profile	TPC-C-like industry-standard benchmark

Physical environment

Figure 1 shows the overall physical architecture of the environment.

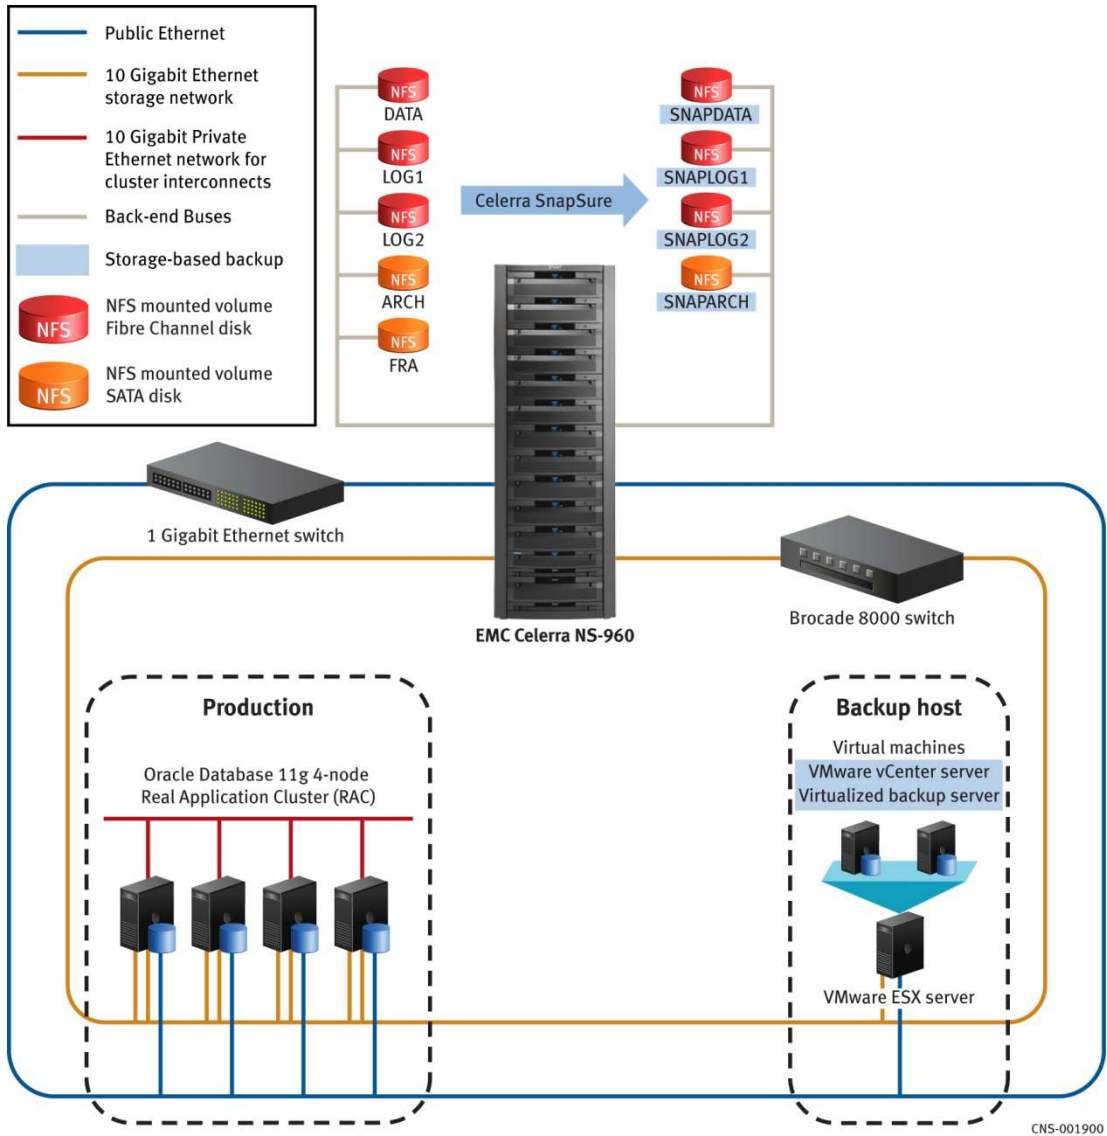


Figure 1. Oracle RAC database on Celerra NS-960 using NFS

Storage layout

Table 4 lists the logical Oracle data storage layout on the production hosts (Oracle RAC).

Table 4. Logical Oracle data storage layout

Mount point	Path	Attributes
/u04	192.168.4.160:/datafs1	Database data source
/u09	192.168.5.160:/datafs2	Database data source

Mount point	Path	Attributes
/u05	192.168.4.160:/log1fs	Redo log area
/u06	192.168.5.160:/log2fs	Redo log area
/u07	192.168.4.160:/archfs	Including files: <ul style="list-style-type: none"> • archivelogs • BCT files • controlfile copies used for backup server
/u02	192.168.5.160:/frafs	Backup file store
/u03	192.168.5.160:/crs	CRS

Table 5 lists the logical Oracle data storage layout on the virtualized backup host. The virtualized backup host is used to mount the SnapSure checkpoint of the production database and offload the RMAN incremental backup. The same version of the operating system and Oracle software needs to be installed on this backup host.

Table 5. Logical Oracle data storage layout on the virtualized backup host

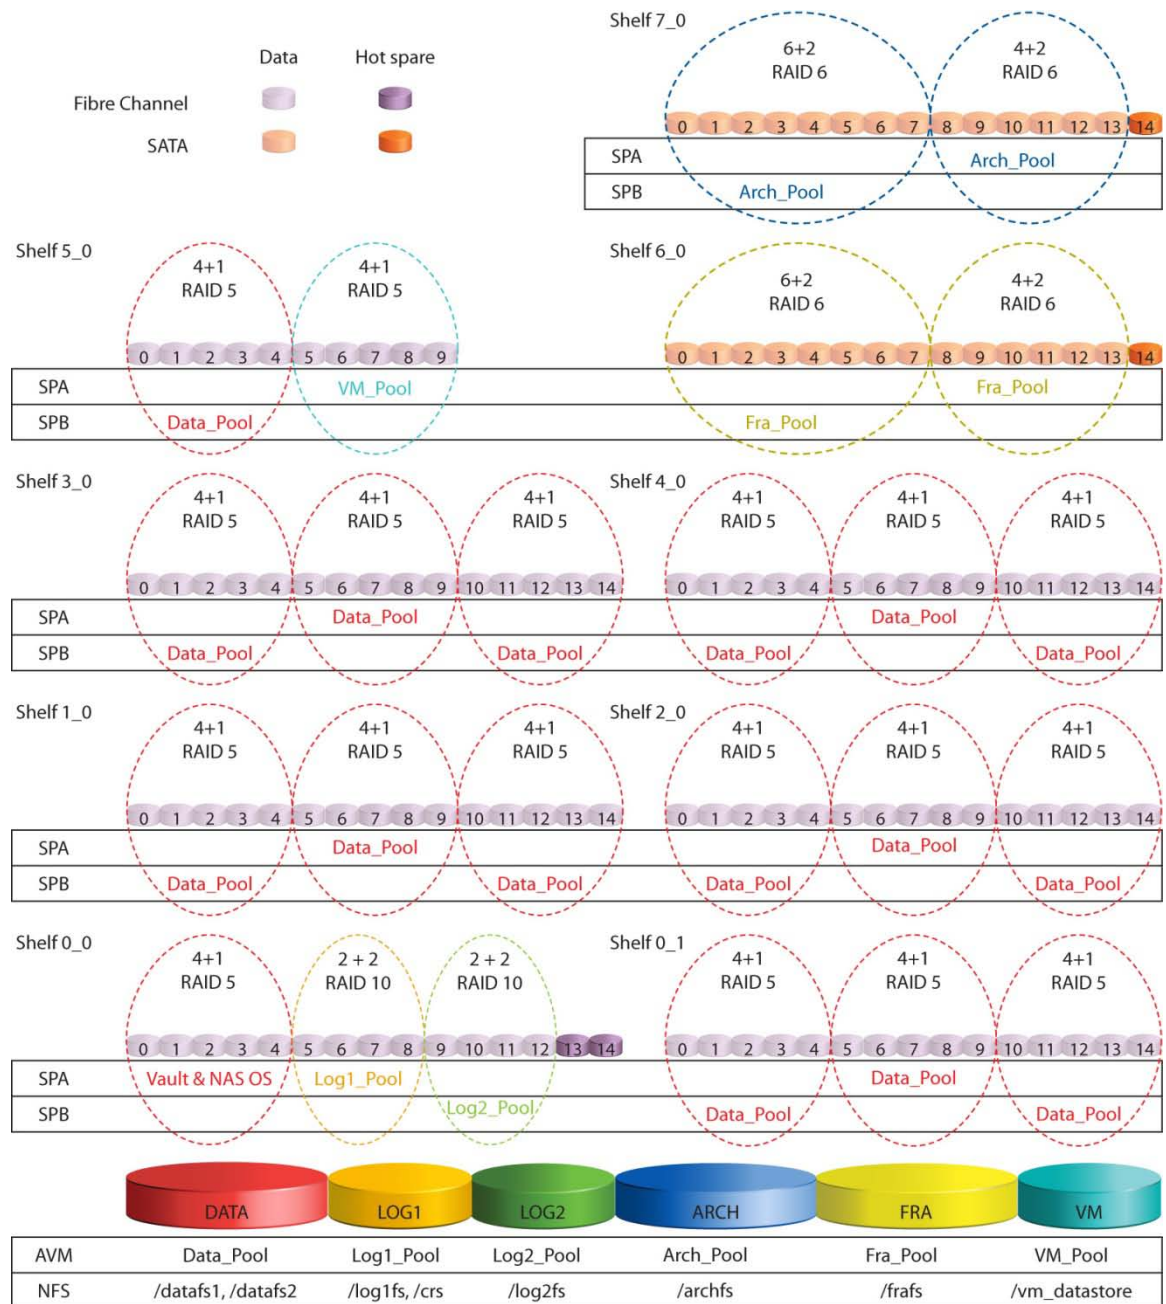
Mount point	Path	Attributes
/u04	192.168.4.160:/datafs1_ckpt_writeable1	Database data source
/u09	192.168.5.160:/datafs2_ckpt_writeable1	Database data source
/u05	192.168.4.160:/log1fs_ckpt_writeable1	Redo log area
/u06	192.168.5.160:/log2fs_ckpt_writeable1	Redo log area
/u07	192.168.4.160:/archfs_ckpt_writeable1	Including files: <ul style="list-style-type: none"> • archivelogs • BCT files • controlfile copies used for backup server
/u02	192.168.5.160:/frafs	Backup file store

Both the production host and backup host need to mount NFS path 192.168.5.160:/frafs on the same mount point, so that the production host can read the backup created on the backup host directly with no naming changes.

Note Since there is only one Celerra in this test environment, RMAN backup files are located on the same array as the production database. In a real environment, EMC recommends that you put the backup files on a different physical location.

Disk layout

Figure 2 shows the disk layout of the environment.



CNS-001901

Figure 2. RAID group layout for NFS

Hardware resources

Table 6 shows the hardware resources used in this solution.

Table 6. Hardware resource

Equipment	Quantity	Configuration
EMC Celerra NS-960 unified storage platform	1	2 storage processors 3 Data Movers 1 Control Station 2 x 10 GbE network connections per Data Mover 7 FC shelves 2 SATA shelves 105 x 300 GB 15k FC disks 30 x 1 TB SATA disks
10 Gigabit Ethernet switches	2	24 CEE ports
Database servers (Oracle RAC 11g servers)	4	4 x 3 GHz quad-core processors 128 GB of RAM 2 x 146 GB 15k internal SCSI disks 2 onboard GbE Ethernet NICs 2 additional 10 Gigabit dual-port CNA cards
VMware ESX server (Virtualized backup server)	1	2 x 3.3 GHz hexa-core processors 48 GB of RAM 2 x 146 GB 15k internal SCSI disks 2 onboard GbE Ethernet NICs 1 x 10 Gigabit AF DA dual-port server adapter

Software resources

Table 7 shows the software resources used in this solution.

Table 7. Software resource

Equipment	Version
EMC Celerra Manager Advanced Edition	5.6
EMC DART	5.6.49-3
EMC FLARE®	04.29.000.5.006
EMC Navisphere® Agent	6.29.5.0.37
VMware vSphere™	4.1
Oracle Enterprise Linux	5.5
Oracle RAC 11g Enterprise Edition	11.2.0.1
Quest Benchmark Factory for Databases	5.8.1

Test and Validation

Introduction

This section provides a summary and characterization of the tests that were performed to validate the solution. The testing goal was to use EMC Celerra SnapSure technology to offload the RMAN incremental backup process with BCT enabled, and to validate the associated benefits that can be expected. Through the following test scenarios, the performance impacts of making the backups directly (Scenario 1) and making the backups on an offloaded system (Scenario 2) were collected and compared.

Preparation work

In this preparation phase, a direct RMAN incremental level 0 backup on the production database was created as the baseline.

1. Run the following command to enable the BCT feature in the production database:

```
SQL> alter database enable block change tracking using file
'/u07/oradata/rman_change_track.f' reuse;
```

2. Run the following command to verify that the BCT feature is enabled:

```
SQL> select * from v$BLOCK_CHANGE_TRACKING
```

The output should be similar to the following:

```
ENABLED      /u07/oradata/rman_change_track.f      43057152
```

3. Run the following command to perform an incremental level 0 backup as the basis of future incremental backups:

```
rman target/ catalog rman_mterac28/rman@catdb
run
{
  backup incremental level 0 database format
  '/u02/bkup/level0_%d_%s_%p_%t' tag 'level0_bct';
}
```

4. Run the following command to check the backup set information in RMAN:

```
RMAN> list backup of database summary;
```

List of Backups

=====

Key	TY	LV	S	Device	Type	Completion Time	#Pieces
#Copies	Compressed	Tag					

5678	B	0	A	DISK	14-DEC-2010	01:08:36	1	1
NO				LEVEL0_BCT				
5681	B	0	A	DISK	14-DEC-2010	01:09:46	1	1
NO				LEVEL0_BCT				
5682	B	0	A	DISK	14-DEC-2010	01:10:54	1	1
NO				LEVEL0_BCT				
5683	B	0	A	DISK	14-DEC-2010	01:11:17	1	1
NO				LEVEL0_BCT				
5684	B	0	A	DISK	14-DEC-2010	01:15:05	1	1
NO				LEVEL0_BCT				
5685	B	0	A	DISK	14-DEC-2010	01:16:28	1	1
NO				LEVEL0_BCT				
5686	B	0	A	DISK	14-DEC-2010	01:29:36	1	1
NO				LEVEL0_BCT				
5687	B	0	A	DISK	14-DEC-2010	01:36:29	1	1
NO				LEVEL0_BCT				

Note In this solution, the RMAN backup command does not back up archive logs because the loss of archive logs is not simulated. However, in a real customer environment, EMC recommends that you back up archive logs together with the incremental backups to avoid the loss of archive logs, which can severely affect the entire backup solution. To include the archive logs in the backup set, the command should be similar to the following:

```
run
{
  backup incremental level 1
  database format '/u02/bkup/level1_%d_%s_%p_%t' tag 'level1_bct'
  plus archivelog;
}
```

Scenario 1 Baseline RMAN incremental backups to NFS mount

A baseline was established for a direct RMAN incremental level 1 backup against the production database with a TPC-C-like workload running continuously.

The backup tests were performance tests, where the performance of each node level was observed and RMAN backup/restore were performed on a single node. The restore was a functionality test, and the time required to perform the RMAN restore was tuned and measured. The transactions restored and recovered were measured to ensure that there is no data loss. The impacts on the production database response time and TPS were recorded for further comparison.

Test procedure

The following procedure was used to create the baseline RMAN incremental level 1 backup with BCT enabled.

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. Restart the database instances.
5. Start the Benchmark Factory test with a user load ranging from 1,000 to 5,000 with intervals of 100.
6. When the user load reaches an iteration of 4,500, run the following command to initiate the RMAN incremental level 1 backup on the first node:

```
rman target/ catalog
rman_mterac28/rman@catdb
run
{
backup incremental level 1 database format
'/u02/bkup/level1_%d_%s_%p_%t' tag 'level1_BCT';
}
```

7. Monitor the performance impact on the production database.
8. Verify that the RMAN backup is completed successfully.
9. When the Benchmark Factory test is complete, capture the results.
10. Shut down the database.
11. Delete one of the data file to destroy the production database.
12. Restore and recover the database from RMAN incremental backups using the following command:

```
rman target/ catalog rman/rman@catdb
run
{
restore database;
recover database;
alter database open;
}
```

13. Capture the time taken to restore and recover the database, respectively.

Test results

The baseline RMAN incremental level 1 backup with BCT enabled was performed while a TPC-C-like OLTP workload was running. The RMAN level 1 incremental backup started at a user load of 4,500 against the production database directly and ended at a user load of 4,700. When RMAN was initiated at a user load of 4,500, there was a moderate increase in the response time (from 0.047 seconds to 0.328 seconds) and a moderate decrease in the

transaction throughput (TPS from 920.46 to 882.11). The impact lasted for about 15 minutes before it came back to normal with a user load of 4,800 as shown in Figure 3.

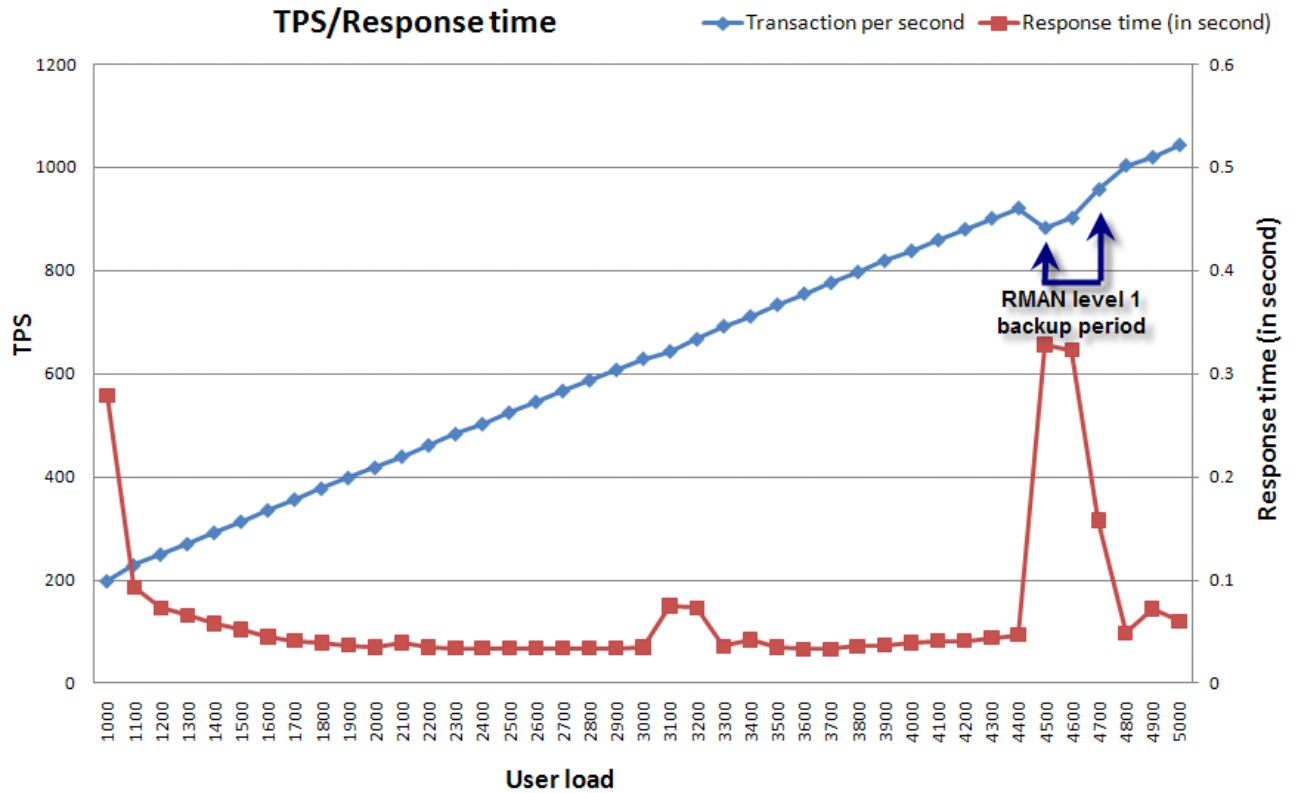


Figure 3. TPS/response time using baseline RMAN incremental backup

At a user load of 4,500, the RMAN incremental backup was performed directly on the production database server. The I/O bandwidth increased dramatically from 60 MB/s to 380 MB/s and kept staying at a high level until the backup is complete as shown in Figure 4. The resource consumption on I/O caused a dramatic increase of active sessions in waiting status, which leads to a moderate decrease in TPS and an increase in response time as shown in Figure 3.

I/O bandwidth reached 380 MB/s and large reads/writes are shown in Figure 4 and Figure 5.

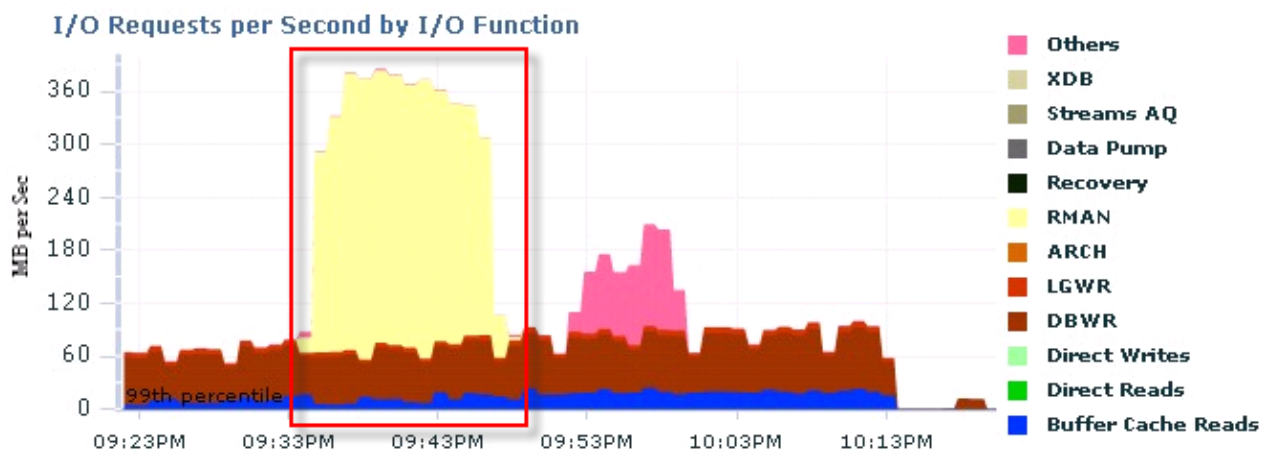


Figure 4. I/O MB/s by I/O function using baseline RMAN incremental backup

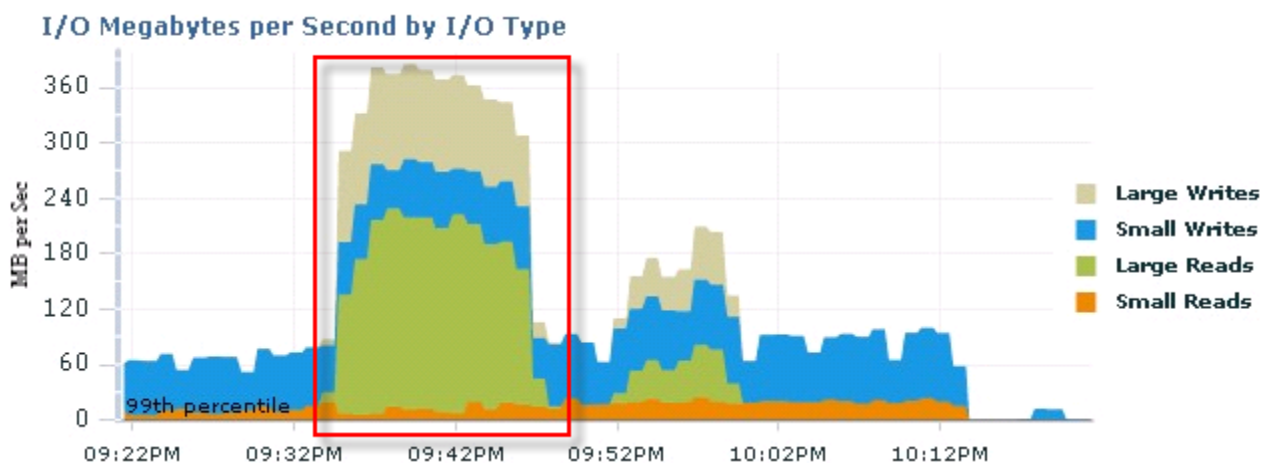


Figure 5. I/O MB/s by I/O type using baseline RMAN incremental backup

The backup, restore, and recover times are recorded in Table 8.

Table 8. Backup, restore, and recover times

Metrics	Value
Test run duration	3 hours 13 minutes
User load range	1000-5000 with an interval of 100
Profile	Mterac28
Driver	Oracle
Start time of test	2011/01/10 19:08
End time of test	2011/01/10 22:21
Level 1 backup duration on production server	14 minutes 08 seconds

Metrics	Value
Total time for restore operation (level 0 backup)	02:34:33
Total time for incremental restore (level 1 backup)	00:13:30
Total time for media recovery (apply archivelogs)	01:27:48

Scenario 2 Storage-based replication RMAN backup to NFS mount

By performing an offloaded backup from the virtualized backup server, this test scenario demonstrates the unique storage capabilities of a Celerra SnapSure checkpoint with integration to Oracle RMAN to avoid impacts on the production database. The reason for using RMAN to back up a SnapSure checkpoint (logical storage backup) of the production database file system is to provide a physical backup of the database. If the source volume of the production database becomes corrupt, the SnapSure checkpoints are invalid. However, the RMAN backup can still be used to restore the production database.

The backup test was a performance test. The database was put into hot backup mode before creating the SnapSure writeable checkpoint and was taken out of hot backup mode after the checkpoint was complete. The writeable checkpoint file systems were mounted on the virtualized backup server and a database instance was mounted so that the offloaded RMAN incremental level 1 backup could be performed for the SnapSure checkpoint database file systems. The restore was a functionality test. The transactions restored and recovered were validated to ensure that there was no data loss. The impacts on the production database response time and TPS were recorded for further comparison.

Test procedure

The following procedure was used to perform the offloaded RMAN incremental backup with BCT enabled.

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. Restart the database instances.
5. Start the Benchmark Factory test with a user load ranging from 1,000 to 5,000 with intervals of 100.
6. When the user load reaches an iteration of 4,500, put the database into hot backup mode using the following command:
SQL> alter database begin backup;
7. Run the following command on the Celerra Control Station to create the file system checkpoint using Celerra SnapSure for the production database file systems (datafs1, datafs2):

```
$ fs_ckpt datafs1 -name datafs1_ckpt -Create
$ fs_ckpt datafs1_ckpt -Create -readonly n
$ fs_ckpt datafs2 -name datafs2_ckpt -Create
```

```
$ fs_ckpt datafs2_ckpt -Create -readonly n
```

8. Take the database out of hot backup mode using the following command:
SQL> alter database end backup;
9. Run the following command to switch the current BCT bitmap and start a new BCT map:
SQL> execute dbms_backup_restore.bctswitch();
10. Create controlfile copies: control_start is used to start up the database for RMAN backup; control_backup is stored along with the backup set.

```
rman target/ catalog rman/rman@catdb  
  
run  
  
{  
  
  backup as copy current controlfile FORMAT  
  '/u07/oradata/control_start';  
  
  backup as copy current controlfile FORMAT  
  '/u07/oradata/control_backup';  
  
}
```

11. Resynchronize the RMAN catalog:

```
rman target/ catalog rman/rman@catdb  
  
RMAN> resync catalog;
```

12. Run the following command on the Celerra Control Station to create the filesystem checkpoint using Celerra SnapSure for the production database file systems (log1fs, log2fs, archfs):

```
$ fs_ckpt log1fs -name log1fs_ckpt -Create  
$ fs_ckpt log1fs_ckpt -Create -readonly n  
$ fs_ckpt log2fs -name log2fs_ckpt -Create  
$ fs_ckpt log2fs_ckpt -Create -readonly n  
$ fs_ckpt archfs -name archfs_ckpt -Create  
$ fs_ckpt archfs_ckpt -Create -readonly n
```

13. Mount the Celerra SnapSure snapshot onto a virtualized backup host.
14. Mount the database instance on the dedicated backup host. Before the instance is mounted, copy the backed up control file (control_start) to the right location:

```
$ cp /u07/control_start /u05/UC485/control01ctl  
$ cp /u07/control_start /u06/UC485/control02ctl
```

15. Start up and mount the database. Do not open it.

```
SQL> sqlplus /nolog
SQL> connect /as sysdba
SQL> startup mount
```

16. Ensure the BCT feature is enabled on the backup database.

```
SQL> select *from v$block_change_tracking;
```

If BCT is not enabled, run the following command to enable it:

```
SQL> alter database enable block change tracking using file
'/u07/oradata/rman_change_track.f' reuse;
```

17. Run a level 1 RMAN backup on the virtualized backup host.

```
rman target/ catalog rman/rman@catdb
run
{
backup format '/u07/ctl%d%s%p%t' controlfilecopy
'/u07/control_backup' tag 'level1_BCT_offload';

backup incremental level=1 database format
'/u02/bkup/level1_%d_%s_%p_%t' tag 'level1_BCT_offload';
}
```

18. Monitor the performance impact on the production database.
19. When the Benchmark Factory test is complete, capture the results.
20. Shut down the database.
21. Delete one of the datafiles to destroy the production database.
22. Restore and recover the database from level 0 and level 1 backups.

```
rman target/ catalog rman/rman@catdb
run
{
restore database;
recover database;
alter database open;
}
```

23. Capture the time taken to restore and recover the database respectively.

Note Because the catalog database is used for all backups either on the production host or the proxy host, there is no need to restore databases on the production host from the backup taken from the proxy host.

Test results

The offloaded RMAN incremental level 1 backup with BCT enabled was performed while the OLTP workload is running. At an iteration of 4,500, the database was put into hot backup mode and the SnapSure writeable checkpoint for the production database file systems (datafs1, datafs2) was taken. The database was taken out of hot backup mode, and the SnapSure writeable checkpoint for the production database file systems (log1fs, log2fs, and archfs) was taken. After that, the checkpoint file systems were mounted on the virtualized backup server.

At a user load of 4,600, the offloaded RMAN level 1 incremental backup was performed on the virtualized backup server. This offloaded backup lasted for about 12 minutes, and caused very little impact on production database performance including TPS and response time. The response time increased a little, and there was no decrease in the TPS. After the incremental backup was complete, the TPS and response times returned back to the normal level.

The performance degradation is small, as shown in Figure 6, in comparison with the performance degradation when using RMAN direct incremental backup on the production host as shown in Figure 3.

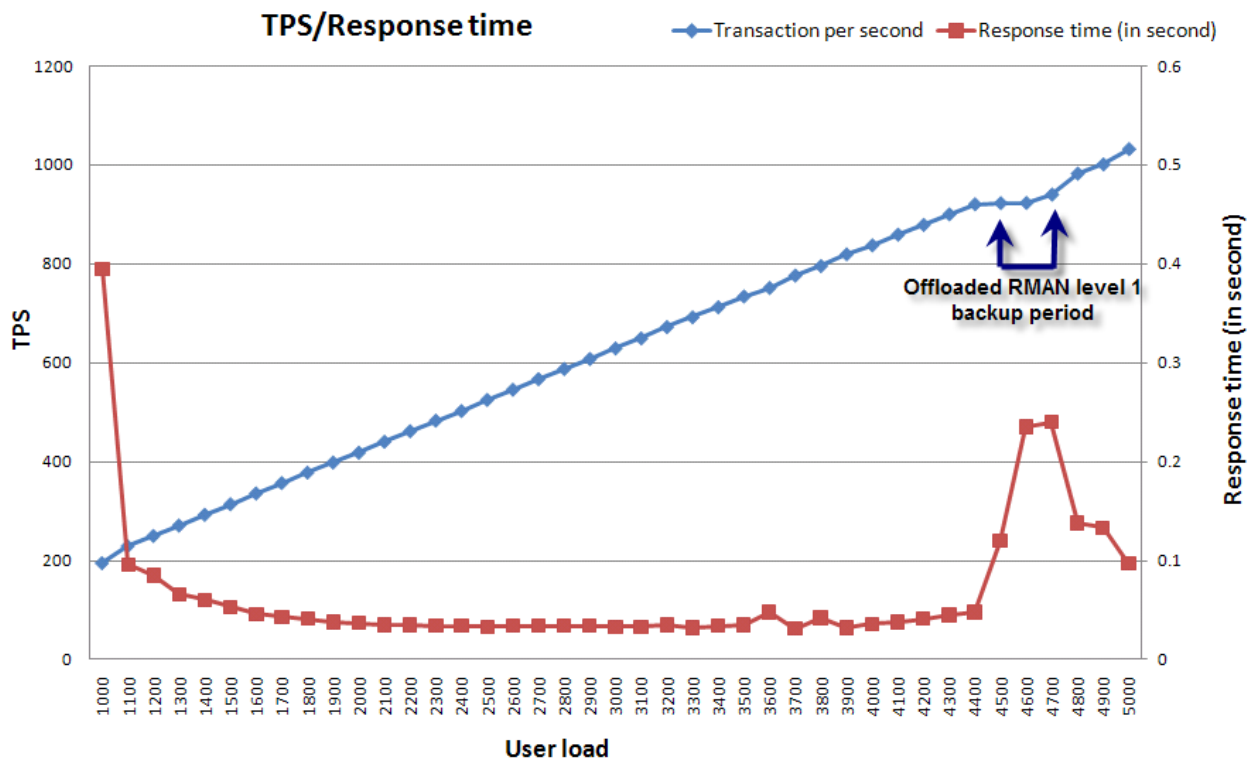


Figure 6. TPS/Response time using offloaded RMAN incremental backup

At a user load of 4,500, the database was put into hot backup mode, causing some overhead to the production database. There was no obvious impact on the I/O bandwidth. At the same time, the number of active sessions in waiting status surged but returned to normal immediately after the database was taken out of backup mode. The impact on the transactions running on the production database was very small as shown in Figure 6.

I/O bandwidth averages at about 70 MB/s as shown in Figure 7 and Figure 8. No large reads or large writes occurred during the time period (compared to the results as shown in Figure 4 and Figure 5).

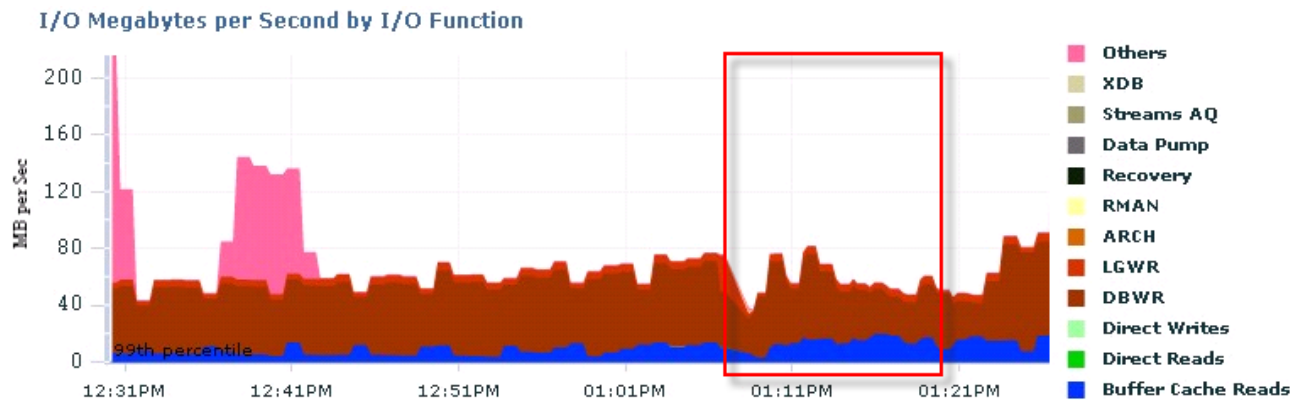


Figure 7. I/O MB/s by I/O function using offloaded RMAN incremental backup

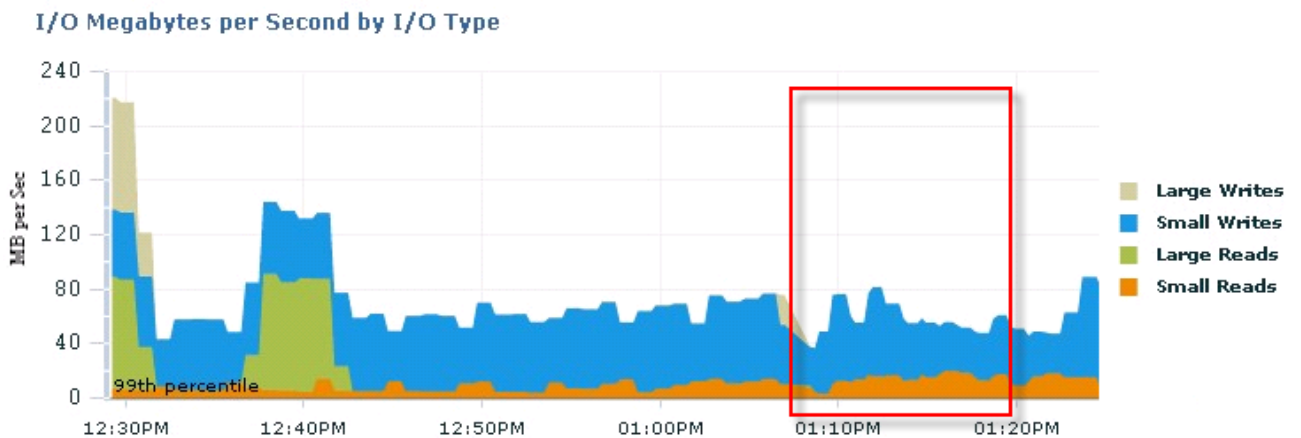


Figure 8. I/O MB/s by I/O type using offloaded RMAN incremental backup

The backup, restore, and recover times are recorded in Table 9.

Table 9. Backup, restore, and recover times

Metrics	Value
Test run duration	3 hours and 14 minutes
User load range	1,000 – 5,000 with intervals of 100
Profile	Mterac28
Driver	Oracle
Start time of test	2011/1/11 10:41
End time of test	2011/1/11 13:55

Metrics	Value
Time used to put into backup mode	00:00:07
Time used to bring out of backup mode	00:00:12
Level 1 backup duration on offload backup server	12 minutes 15 seconds
Total time for restore operation (level 0 backup)	02:28:02
Total time for incremental restore (level 1 backup)	00:11:30
Total time for media recovery (apply archivelogs)	01:31:49

Verification procedure for using BCT

After completion of the incremental backup, you might want to make sure the BCT feature is correctly used by the RMAN incremental backup. Run the following SQL query on the backup host to check the data dictionary view *v\$backup_datafile* as shown in Figure 9:

```
SQL> select checkpoint_time,
           checkpoint_change#,
           blocks_read,
           blocks,
           datafile_blocks,
           used_change_tracking,
           file#
From v$backup_datafile
Order by file#,
       checkpoint_change# asc;
```

CHECKPOINT_TIME	CHECKPOINT_CHANGE#	BLOCKS_READ	BLOCKS	DATAFILE_BLOCKS	USE	FILE#
2010/12/13 09:32:17	60159232	1170	1170	1170	NO	0
2010/12/13 14:03:50	117741091	1170	1170	1170	NO	0
2010/12/13 22:12:57	177247945	1170	1170	1170	NO	0
2010/12/14 17:44:50	291703915	1178	1178	1178	NO	0
2010/12/14 21:34:42	348128232	1280	1280	1280	NO	0
2010/12/14 16:21:16	291529043	84865	73815	291840	YES	1
2010/12/14 21:32:35	346937297	683	479	291840	YES	1
2010/12/14 16:21:23	291529061	97921	86056	103680	YES	2
2010/12/14 21:32:35	346937297	8231	4878	104960	YES	2
2010/12/14 16:21:16	291529043	46785	544	1310720	YES	3
2010/12/14 21:32:35	346937297	26115	23578	1310720	YES	3

CHECKPOINT_TIME	CHECKPOINT_CHANGE#	BLOCKS_READ	BLOCKS	DATAFILE_BLOCKS	USE	FILE#
2010/12/14 16:21:16	291529043	705665	705684	917504	YES	4
2010/12/14 21:32:35	346937297	73401	34023	917504	YES	4
2010/12/14 16:21:23	291529061	45441	526	1310720	YES	5
2010/12/14 21:32:35	346937297	20951	19110	1310720	YES	5
2010/12/14 16:21:23	291529056	49089	820	1310720	YES	6
2010/12/14 21:32:35	346937297	24159	22181	1310720	YES	6
2010/12/14 16:21:23	291529064	40577	558	1310720	YES	7
2010/12/14 21:32:35	346937297	20483	18166	1310720	YES	7
2010/12/14 16:21:23	291529061	705665	705684	917504	YES	8
2010/12/14 21:32:35	346937297	74325	34215	917504	YES	8

Figure 9. Data dictionary view (Example)

The column USED_CHANGE_TRACKING (the USE column in Figure 9) indicates whether the change tracking data was used to accelerate the incremental backup.

The column BLOCK_READ shows the number of scanned blocks during the backup. If it was an incremental backup, and the change tracking data was used to optimize the backup, the value of this column should be smaller than the column DATAFILE_BLOCKS. Otherwise, the value of this column should be the same as DATAFILE_BLOCKS. If the change tracking data is used, the value of the column BLOCK_READ can be larger than the column BLOCKS, because the data read by change tracking is further refined during the process of incremental backup creation.

Comparison

Comparing the test results of the two scenarios, performing RMAN backups directly on the production database server can dramatically increase the demand on resources, such as network bandwidth, I/O capacity, and CPU. By offloading the RMAN backup workload to a virtualized backup server, you can save resources on the production database server to better serve the online business transactions.

When performing the offloaded RMAN incremental level 1 backup, the performance advantages recorded are:

- The average TPS increases by 1.73 percent while the average response time decreases by 26 percent as shown in Figure 10.

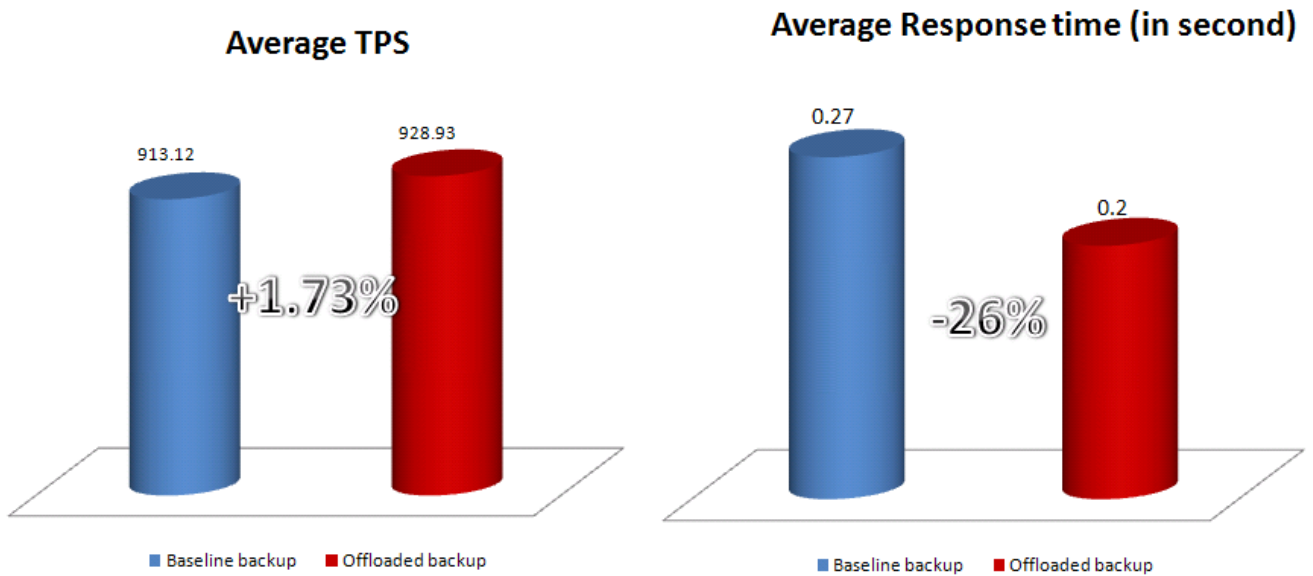


Figure 10. Comparison of average TPS during RMAN backup

- The backup time was reduced by 13.3 percent as shown in Figure 11.

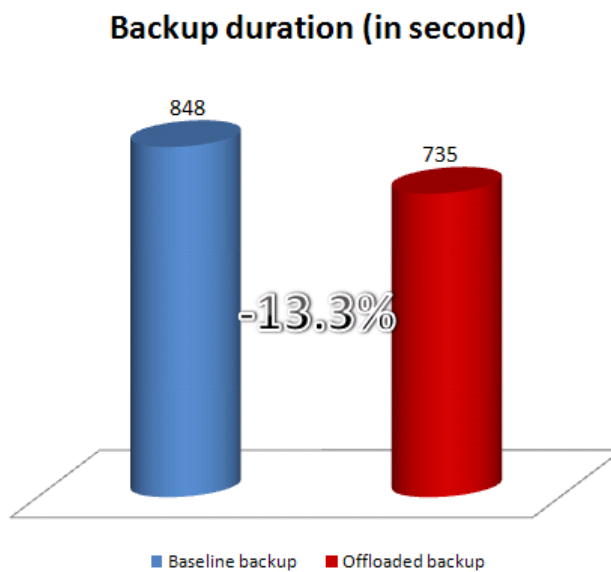


Figure 11. Comparison of backup duration

- The average I/O MB/s was reduced by 81.6 percent as shown in Figure 12.

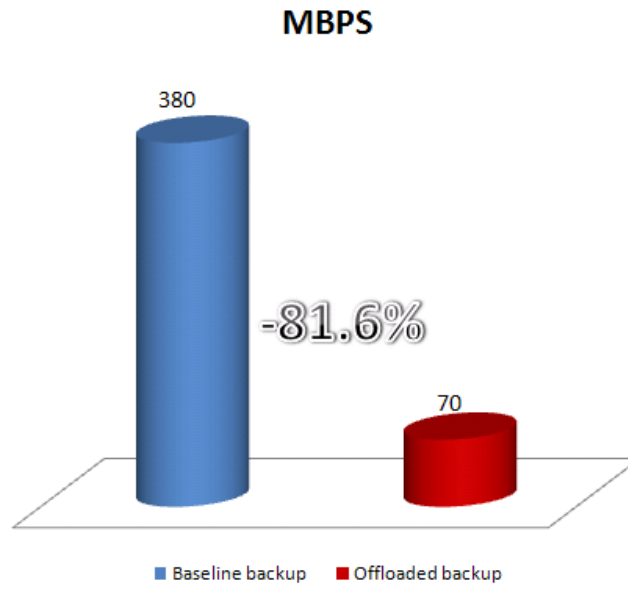


Figure 12. Comparison of average I/O MB/s during RMAN backup

Conclusion

Summary

The high-availability features of EMC Celerra unified storage combining with EMC proven storage technologies provide a high-performance, high-flexibility, and low-risk storage system for Oracle Database 11g over DNFS.

By using RMAN and a Celerra SnapSure checkpoint, the TPS is improved while the response time and the entire backup duration on the production database are greatly shortened. During the period required to complete the backup, the baseline system completes 3652.48 transactions. In the same period of time, the offloaded system completes 3715.72 transactions, an increase of 1.73 percent in transaction workload performance. In addition, the I/O bandwidth with storage-based backup is significantly reduced.

Findings

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

Key findings	Solution objective
Improved performance	<p>Minimizes impact on the production database while running database backups using storage-based backup.</p> <ul style="list-style-type: none">Increases the average TPS by 1.73 percent while cutting down the response time by 26 percentShortens the entire backup duration by 13.3 percent on the production databaseSignificantly reduces the I/O bandwidth by 81.6 percent, from 380 MB/s to 70 MB/s
Business continuity	<p>Since the backup operation has minimal impact on the database server performance, the backup can be run more often. As a result, the recovery operation can be done in less time because fewer archived logs are applied.</p>
Robust performance and scaling	<p>In line with best practices, the performance testing carried out by EMC uses an industry-standard OLTP benchmark, but only a reasonable amount of tunings. In addition, real-world configurations are used. This ensures that the configuration you choose to run on your application is predictable and reliable.</p>

Next steps

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

References

White papers

For white papers that provide similar solutions, see the following:

- *Leveraging EMC CLARiiON Storage Replication to Offload Oracle Recovery Manager (RMAN) Backup - Applied Technology*
- *EMC Solutions for Oracle Database 10g/11g for Midsize Enterprises—EMC Celerra Unified Storage Platform - Best Practices Planning*

Other documentation

For additional information on related topics, see the following:

- *Oracle Database Backup and Recovery User's Guide 11g Release 2 (11.2)*