

Exadata Database Machine Backup and Restore Configuration and Operational Best Practices

ORACLE WHITE PAPER | JULY 2018





Table of Contents

Introduction	1
Backup Options	1
Exadata Database Machine backup to Exadata Storage	1
Zero Data Loss Recovery Appliance	2
Oracle ZFS Storage Appliance	2
Tape Library	3
Oracle Database Backup Cloud Service	3
Summary of Exadata backup recommendations	4
Configuration Best Practices	5
Exadata Database Machine Backing Up to Exadata Fast Recovery Area on Local Disk	5
Parallelism Usage with Local Disk on Exadata	6
Zero Data Loss Recovery Appliance	6
Parallelism Usage with the Zero Data Loss Recovery Appliance	6
RMAN Compression Usage Recommendations with the Zero Data Loss Recovery Appliance	6
BACKUP .. FILESPERSET Option Usage Recommendations	7
Real-Time Redo Transport Feature on the Zero Data Loss Recovery Appliance	7
Oracle ZFS Storage Appliance	7
Tape Library	8
Software for Tape Library Management	8
Parallelism Usage with Tape Libraries	8
Compression with Tape Libraries	8



Connectivity for Tape Libraries	8
Operational Best Practices	8
RMAN Recovery Catalog	8
Logging RMAN Output	9
RMAN Crosscheck	9
RMAN Backup Operations	9
Archive Log Deletion Policy	10
Validating Database Backups	10
Database Recovery	10
Performance Estimates and Examples	11
Exadata Database Machine with a Fast Recovery Area on Local Disk	11
Zero Data Loss Recovery Appliance	12
Oracle ZFS Storage Appliance	12
Tape Library	13
Performance Troubleshooting for Recovery Manager (RMAN)	13
Conclusion	13



Introduction

The integration of Oracle Exadata Database Machine and Oracle Maximum Availability Architecture best practices (Exadata MAA) provides the most comprehensive high availability solution available for Oracle Database.

A key operational aspect of deploying Exadata MAA is to be sure that database backups are performed and Oracle Database can be restored if disaster strikes. This white paper applies to Oracle Database release 11.2.0.4 or later releases and describes the best practices for setting up the optimal backup and recovery strategy to protect mission-critical data.

This paper discusses the following topics:

- » Backup Options
- » Configuration Best Practices
- » Operation Practices
- » Performance Estimates and Examples

Backup Options

There are 5 possible database backup solutions for Exadata Database Machine backup, restore, and recovery operations. The advantages of each solution, real world examples of where each option has been used, and considerations for choosing the solution are discussed in the following sections.

- » Existing Exadata Storage Server on [Exadata Database Machine](#)
- » [Zero Data Loss Recovery Appliance](#)
- » [Oracle ZFS Storage Appliance](#)
- » [Tape Library](#)
- » [Oracle Database Backup Cloud Service](#)

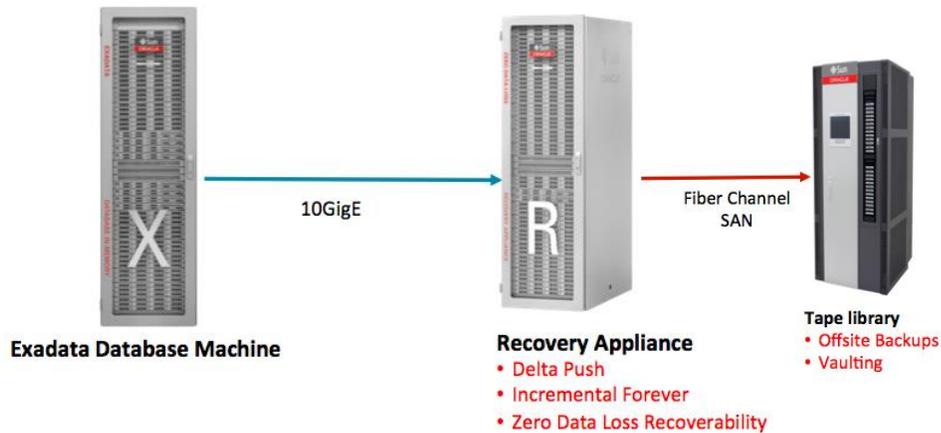
Exadata Database Machine backup to Exadata Storage

Backing up databases local to the Exadata Database Machine provides very fast backup and restore performance along with the simplicity of leveraging existing Exadata hardware for backups. However, storing the database backups on the same Exadata Database Machine alone does not provide the disaster recovery protection you could get using a solution such as using Oracle Data Guard, Oracle GoldenGate, or an offsite backup solution.

Another advantage of using Exadata Storage is being able to leverage advanced features such as Exadata snapshots.

Zero Data Loss Recovery Appliance

The Oracle Zero Data Loss Recovery Appliance (Recovery Appliance) is Oracle's premier solution for protecting all of your Oracle databases in the data center, residing on any server platform. The Recovery Appliance offers a solid backup solution that is simple to implement and extremely fast and scalable for both backup and restore operations.



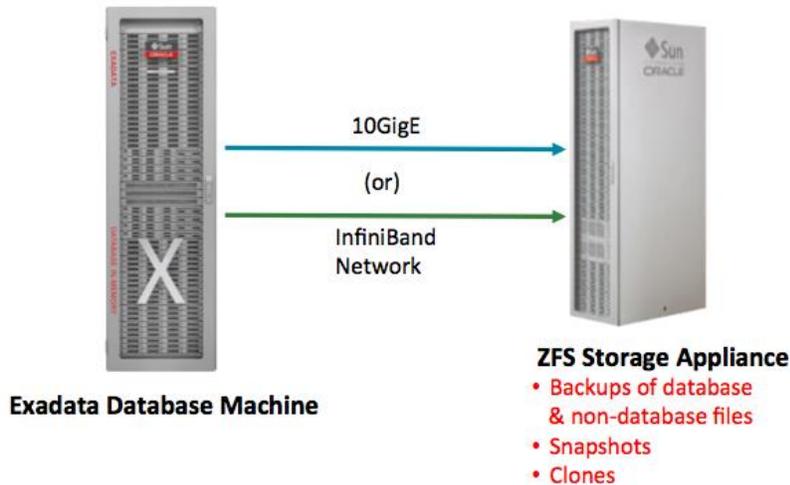
The Recovery Appliance not only simplifies the backup operations, but can eliminate data loss and failed database recovery operations using Real-Time Redo Transport and its inherent corruption checks and backup validation. The Recovery Appliance can significantly reduce the impact of backups on the production databases by offloading data deduplication, compression, backup and restore processing, and validation to the Recovery Appliance. Backup windows also shrink significantly due to the incremental forever protection paradigm which, after the initial full backup of the database, only leverages incremental subsequent backups.

Recovery Appliance is integrated with an optional tape library to address long term retention requirements, relieving Exadata from doing the copy to tape operation directly and thus eliminating the traditional backup window for the tape backup operation. This also minimizes the number of tape drives, since the Recovery Appliance can send backups to tape 24x7 without impacting the production database.

The Recovery Appliance is the MAA recommended backup and restore solution for the data center. Disaster recovery protection can be achieved by having a Recovery Appliance replica in a separate data center.

Oracle ZFS Storage Appliance

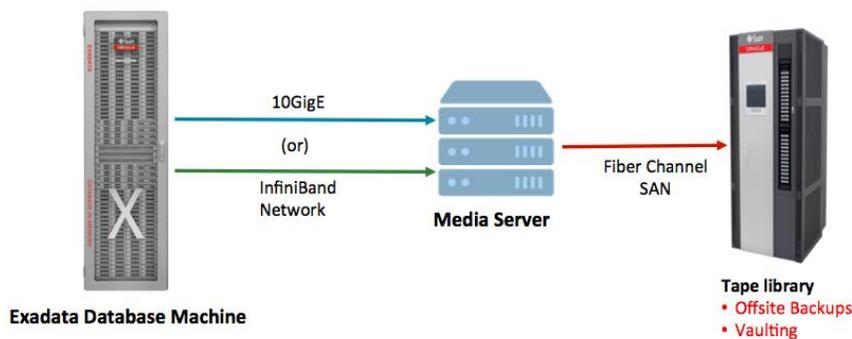
Storing database backups on the ZFS Storage Appliance (ZFSSA) over NFS offers a lower cost per TB than other local storage options and still provides a fast backup and restore solution. Configuring the database to send the backups to the ZFSSA is relatively simple, especially with 10gigE. By leveraging InfiniBand Networking, backup and restore rates may improve even more, but best practices need to be followed closely, including maintaining the right set of patches during initial setup and throughout future software updates and configuration changes.



Once the database is backed up to the ZFSSA, features such as ZFS Snapshots can be used to provide rapid cloning of development and test environments using [Oracle Snap Management Utility for Oracle Database](#).

Tape Library

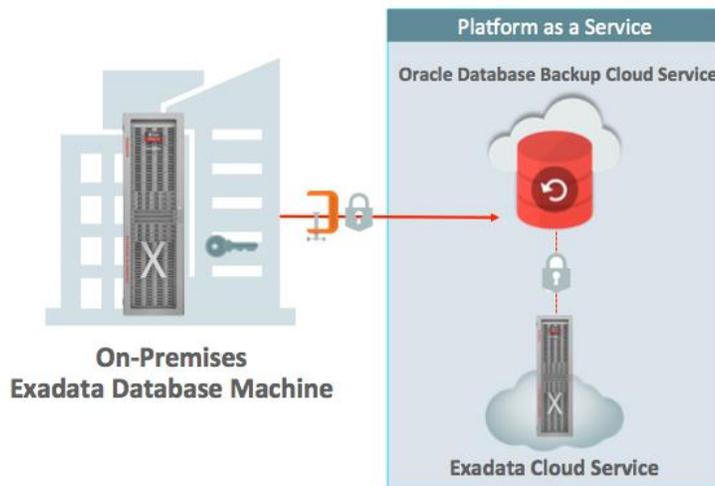
The use of a physical or virtual tape library is the recommend solution for longer term data retention requirements. While it is possible to perform database backups directly to a tape library as depicted in the subsequent graphic, the Exadata MAA recommendation is to use a tape library as part of the Zero Data Loss Recovery Appliance solution for the data center, as discussed above in the [Zero Data Loss Recovery Appliance](#) section.



Oracle Database Backup Cloud Service

Oracle Database Backup Cloud Service is a low cost offsite storage backup solution. It is an Oracle Cloud Platform as a Service (PaaS) offering which protects Oracle Databases that are deployed on-premises or in the Oracle Cloud. Oracle Recovery Manager (RMAN) is used to encrypt and securely transmit backup over HTTPS/SSL to the Backup Cloud Service. The backup data is then stored in multiple copies in the cloud for high availability and can be accessed any time for restoration and validation. The encryption keys are kept with the customer. Optionally, the data can be replicated to another cloud data center for disaster recovery. The Backup Cloud Service backups can

be used to instantiate new database instances in the [Oracle Database Cloud Service](#), which provides a simple method to get database point in time copies for test and development or disaster recovery purposes.

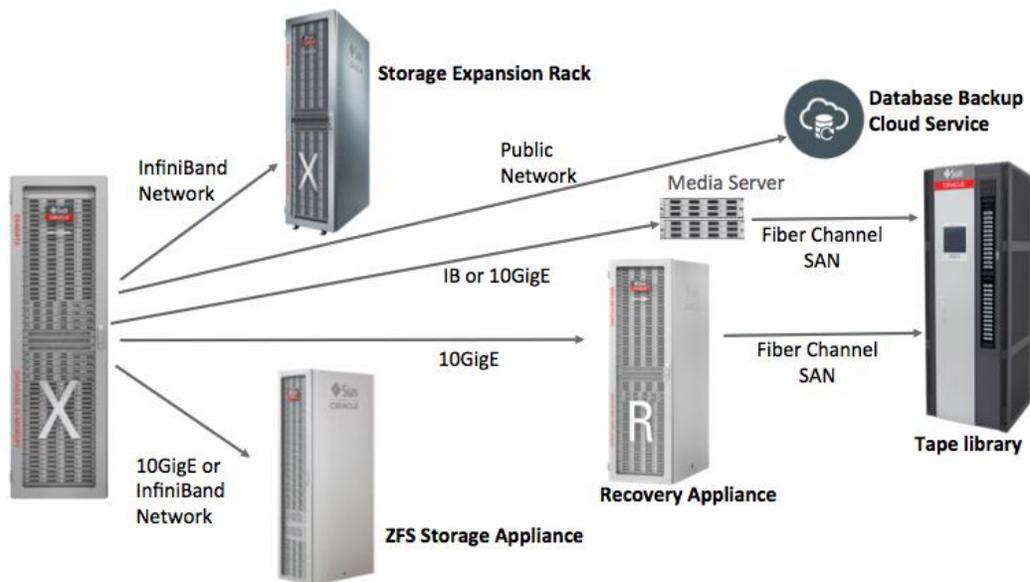


Oracle Database Backup Cloud Service module supports all major platforms and all supported Oracle Database releases. Administrators can use Enterprise Manager Cloud Control 13c, RMAN CLI, or 3rd party software such as Cloudberry to perform backup and recovery management. For more information, see cloud.oracle.com/database_backup.

Summary of Exadata backup recommendations

MAA recommends the following:

- When you require the lowest RTO and RPO, use Oracle Active Data Guard or Oracle GoldenGate.
- For backup and recovery solutions the following are recommended:
 - a. Use a Recovery Appliance in each data center for backing up all Oracle databases and a Recovery Appliance replica in a separate data center for additional offsite disaster recovery.
 - b. When on-premises long term backup retention is needed, use a tape library with the Recovery Appliance.
 - c. When a fast, lower cost alternative to higher storage cost of backing up to local Exadata disk, use an Oracle ZFS Storage Appliance for backup and restore of one or many Exadata racks. The Oracle ZFS Storage Appliance can also leverage ZFS snapshot capabilities for creating fast clones.
 - d. When off-site long term backup retention is needed, use Oracle Database Backup Cloud Service as a low cost, offsite scalable storage for a disaster recovery solution or for inexpensive long term backup or archive retention.



Configuration Best Practices

In the previous section, we introduced 5 possible architectural options for backing up databases from an Exadata Database Machine. Here we review the operational best practices for each of these 5 architectures.

The following are recommendations for all architectures:

- » Enable RMAN block change tracking when the daily change rate is less than 20% to speed up level 1 incremental backup sets. Note: Using this fast incremental backup on a physical standby database requires the Active Data Guard option.
- » Use RMAN `CONTROLFILE AUTOBACKUP` to help streamline recovery operations.
- » Use RMAN `BACKUP SECTION SIZE` for bigfile tablespaces (also known as single-file tablespaces) to distribute the bigfile into manageable section sizes across available RMAN channels. Ideally, the section size should be set to allow each RMAN channel to process 2 backup sections. For example: A 2 TB bigfile with 16 RMAN channels should use `SECTION SIZE 64g`.

Exadata Database Machine Backing Up to Exadata Fast Recovery Area on Local Disk

There are two variants to the architecture that can be used when backups are stored locally on the Exadata Database Machine.

- » Separate Exadata High Capacity Storage Servers are used for RECO (backups)

If the Exadata Database Machine uses High Performance or Extreme Flash Storage Servers, then it is typical to leverage additional High Capacity storage servers for database backups. In this case the entire High Capacity storage server can be configured as the Fast Recovery Area (FRA) or RECO disk group.
- » Shared Exadata High Capacity Storage Servers are used for DATA and RECO

If the Exadata Database Machine uses High Capacity storage servers, then the storage servers can be configured such that 40% of the space is assigned to the DATA disk group where the database objects reside and the remaining 60% of the space is assigned to the Fast Recovery Area (FRA), which is configured as the RECO disk group.

Parallelism Usage with Local Disk on Exadata

For both of the above architectures, once the storage server space has been allocated, the only RMAN configuration change required is to select the degree of parallelism. The degree of parallelism used by RMAN primarily depends on the number of Exadata database servers.

EXADATA DATABASE MACHINE RMAN CONFIGURATION BEST PRACTICES (TO LOCAL DISK)

Exadata Configuration	Recommended RMAN Parallelism
Exadata Eighth Rack or Quarter Rack	CONFIGURE DEVICE TYPE DISK PARALLELISM 8; » 4 channels per database server
Larger than Exadata Quarter Rack	CONFIGURE DEVICE TYPE DISK PARALLELISM 8; » 2 channels per database server (Example above based on 4 database servers in a half rack)
Exadata 8 Socket (X5-8 / X6-8) Full Rack with 2 compute servers	CONFIGURE DEVICE TYPE DISK PARALLELISM 16; » 8 channels per database server
Exadata 8 Socket (X5-8 / X6-8) Full Rack with 3 compute servers	CONFIGURE DEVICE TYPE DISK PARALLELISM 24; » 8 channels per database server

In rare circumstances the number of channels per server can be increased for higher backup or restore throughput.

Zero Data Loss Recovery Appliance

Database backups from the Exadata Database Machine should be sent to the Recovery Appliance over a dedicated bonded 10GigE network, although it is possible to send the backups over an unbonded dedicated network or over the generic client access network. Backups should not be sent to the Recovery Appliance over the Exadata Database Machine management 1GigE network.

Parallelism Usage with the Zero Data Loss Recovery Appliance

Similar to the Exadata Database Machine configuration best practices mentioned in the previous section, the recommended degree of parallelism used by RMAN primarily depends on the number of Exadata database servers.

EXADATA DATABASE MACHINE RMAN CONFIGURATION BEST PRACTICES FOR RECOVERY APPLIANCE

Exadata Configuration	Recommended RMAN Parallelism
Exadata Eighth Rack or Quarter Rack (also applies to Exadata SL)	CONFIGURE DEVICE TYPE SBT_TAPE PARALLELISM 16; » 8 channels per database server
Larger than Exadata Quarter Rack	CONFIGURE DEVICE TYPE DISK PARALLELISM 8; » 2 channels per database server
Exadata 8 Socket (e.g.: X5-8 / X6-8) Full Rack with 2 compute servers	CONFIGURE DEVICE TYPE DISK PARALLELISM 16; » 8 channels per database server
Exadata 8 Socket Full Rack with 3 compute servers	CONFIGURE DEVICE TYPE DISK PARALLELISM 24; 8 channels per database server

RMAN Compression Usage Recommendations with the Zero Data Loss Recovery Appliance

When the Exadata Database Machine and the Recovery Appliance are in the same data center on a high speed network, RMAN [compression](#) is **not** recommended, since the Recovery Appliance has robust compression built-in.

When the Exadata Database Machine and the Recovery Appliance are separated by a Wide Area Network (WAN) and network bandwidth is limited, the use of RMAN compression is recommended to help lower the amount of network traffic. When using `LOW`, `MEDIUM`, or `HIGH` RMAN compression levels, the Oracle Database Advanced Compression option must be licensed for the Exadata Database Machine. If RMAN compression is being considered then RMAN `LOW` compression provides a good balance of compression ratio, low CPU utilization, and reduced backup elapsed time.

EXADATA DATABASE MACHINE RMAN CONFIGURATION BEST PRACTICES

Exadata and Recovery Appliance Locations	Recommended RMAN Compression
In the same data center	CONFIGURE COMPRESSION ALGORITHM clear; CONFIGURE DEVICE TYPE SBT_TAPE BACKUP TYPE TO BACKUPSET;
Separate data centers	CONFIGURE COMPRESSION ALGORITHM '[LOW MEDIUM HIGH BASIC]'; CONFIGURE DEVICE TYPE SBT_TAPE BACKUP TYPE TO COMPRESSED BACKUPSET;

Testing has shown that RMAN `LOW` compression generally reduces the duration of the backup while providing a compression rate of between 40% and 50%.

BACKUP .. FILESPERSET Option Usage Recommendations

The internal Recovery Appliance backup and restore processing is optimized when the RMAN `FILESPERSET` parameter of the `BACKUP` command is set to 1 for the level 1 incremental backup set. This is only applicable for database and datafile backups and should not be used for archived log backups.

Real-Time Redo Transport Feature on the Zero Data Loss Recovery Appliance

You can set Real-Time Redo Transport from each protected database to the Recovery Appliance. This Recovery Appliance feature leverages the same log transport mechanism used by Oracle Data Guard to send transactions in real time to the Recovery Appliance, but the Recovery Appliance doesn't require a Data Guard license. Real-Time Redo Transport is configured at the database level on the database being backed up and is available for Oracle Database 11g Release 2 and later on platforms supported by the Recovery Appliance. No specific configuration is required on the Recovery Appliance to use this feature although some Oracle releases require certain patches to be installed. See [Oracle Support Document 1982542.1](#) – "Recommended Protected Database Patches for Zero Data Loss Recovery Appliance" for more information.

Oracle ZFS Storage Appliance

For detailed instructions about how to configure the Exadata Database Machine and the ZFS Storage Appliance when backups are to be sent to the ZFS Storage Appliance, refer to [Oracle Support Document 2087231.1](#) – "Guidelines When Using ZFS Storage in an Exadata Environment." This note discusses the recommended and required minimum versions of the ZFS software level and the Exadata version, along with information about how to configure the storage pools and ZFS flash devices.

The number of channels to be used is dependent on the number of Exadata database servers AND the number of disk shelves in the ZFS storage array.

Tape Library

Oracle recommends that if a tape library is used, it should be connected to a Recovery Appliance for long-term and off-site retention. Alternatively, Oracle Database Backup Cloud Service can be used for long term backup and archival retention.

Database backups from the Exadata Database Machine that are sent directly to a tape library should use a dedicated bonded network, although it is possible to send the backups from the Exadata Database Machine on an unbonded dedicated network or over the generic client access network.

Software for Tape Library Management

Backing up an Exadata Database Machine directly to a tape library requires the installation of a media management vendor (MMV) software application such as Oracle Secure Backup. Any MMV that is certified with the operating system running on Exadata and the specific Oracle Database release can be used. Because the Exadata Database Machine configuration cannot be changed, the client component must be installed on the Exadata Database Machine compute servers and configured to send the backups across a network to the media servers that are attached to the physical tape library.

Parallelism Usage with Tape Libraries

The degree of parallelism used by RMAN is limited by the number of tape drives that can be used by the backup operation, but also should consider the number of tape drives that can be used should a restore be required. For instance, if the tape library has 8 tape drives, and during restore operations all 8 tape drives can be used in parallel, then at the time of the database backup, the backup must also be able to write to 8 different pieces of tape media. Otherwise, if the backup was limited to writing to 4 tape drives (for example, due to tape partitioning), then the chances are that only 4 different pieces of tape media will be used, and therefore at the time of restore, even though 8 channels might be specified, only 4 pieces of media can be accessed so only 4 tape drives will be used.

Compression with Tape Libraries

If possible, hardware compression should be used on the tape drives. This reduces the CPU cycles required on the Exadata Database Machine compute servers for RMAN compression or compression implemented at the MMV layer.

Connectivity for Tape Libraries

The MAA white paper "[Backup and Recovery Performance and Best Practices for Exadata Cell and the Oracle Exadata Database Machine](#)" has a detailed discussion about troubleshooting tape library backup and restore speeds.

Operational Best Practices

The following operational best practices apply to all architectural options except where stated explicitly.

RMAN Recovery Catalog

Except for the Recovery Appliance architecture, which has a built-in RMAN recovery catalog, the use of an RMAN recovery catalog is highly recommended. For backups using Exadata Database Machine's Fast Recovery Area, ZFS Storage Appliance, or a tape library, the RMAN recovery catalog connection can be optional for the duration of the backup, but following the completion of the backup a new RMAN session should be started which connects to the catalog and which performs an explicit `RESYNC CATALOG` operation.

Logging RMAN Output

All RMAN command output should be logged to a file using the RMAN `LOG=<filename> APPEND` syntax. To log the commands being issued, execute the RMAN command `set echo on` as the first command issued.

Additionally, before the RMAN session is spawned, the National Language Set date format should be set to report a full date and time, by exporting the O/S environment variable `NLS_DATE_FORMAT`. For instance, `export NLS_DATE_FORMAT='DD-Mon-YYYY HH24:MI:SS'` or similar should be used.

RMAN Crosscheck

Issuing an RMAN `CROSSCHECK` command before running any backups helps avoid unnecessary RMAN failures or incorrect backups being performed. For instance, `CROSSCHECK ARCHIVELOG ALL` verifies that all of the archive log files that are known to the database actually exist on disk. A `CROSSCHECK BACKUPSET` verifies that all of the backup sets that have been created by RMAN are present and accessible. While `CROSSCHECK BACKUPSET` can be a large operation, periodically verifying that all backup sets exist is a best practice.

RMAN Backup Operations

Which RMAN backup commands are issued largely depend on the architecture being used and the backup retention requirements. The commands shown in the following table assume that the RMAN configuration parameters in the previous steps have been implemented.

EXADATA DATABASE MACHINE RMAN DATABASE BACKUP BEST PRACTICES

Backup Destination	Recommended RMAN Command
Exadata Database Machine	<pre>RMAN> RUN { BACKUP DEVICE TYPE DISK SECTION SIZE <SIZE> DATABASE PLUS ARCHIVELOG NOT BACKED UP; }</pre>
Recovery Appliance	<pre>RMAN> RUN { BACKUP DEVICE TYPE SBT CUMULATIVE INCREMENTAL LEVEL 1 FILESPERSET 1 SECTION SIZE 32G DATABASE PLUS ARCHIVELOG NOT BACKED UP FILESPERSET 32; }</pre>
Oracle ZFS Storage Appliance	Refer to Oracle Support Document 2087231.1 – “Guidelines When Using ZFS Storage in an Exadata Environment”
Oracle Database Backup Cloud Service	<pre>RMAN> RUN { BACKUP DEVICE TYPE SBT DATABASE PLUS ARCHIVELOG NOT BACKED UP; }</pre>
Tape Library	<pre>RMAN> RUN { BACKUP DEVICE TYPE SBT DATABASE PLUS ARCHIVELOG NOT BACKED UP; }</pre>

Archive Log Deletion Policy

Archive log files remain in the Exadata Database Machine FRA using the commands issued above. To allow the database to delete the archive log files automatically when the FRA becomes full, `ARCHIVELOG DELETION POLICY` should be set as shown in the following table.

EXADATA DATABASE MACHINE RMAN ARCHIVELOG DELETION POLICY BEST PRACTICES

Backup Destination	Recommended RMAN Command
Exadata Database Machine	<code>RMAN> CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO DISK;</code>
Recovery Appliance	<code>RMAN> CONFIGURE ARCHIVELOG DELETION POLICY TO SHIPPED TO ALL STANDBY BACKED UP 1 TIMES TO SBT;</code>
Oracle ZFS Storage Appliance	Refer to Oracle Support Document 2087231.1 – “Guidelines When Using ZFS Storage in an Exadata Environment”
Oracle Database Backup Cloud Service	<code>RMAN> CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO SBT;</code>
Tape Library	<code>RMAN> CONFIGURE ARCHIVELOG DELETION POLICY TO BACKED UP 1 TIMES TO SBT;</code>

Validating Database Backups

It is critical to periodically validate the backups to check if they can be used to recover the database in the event of a database disaster. The first part of database recovery is to perform a restore of the database from the backup media. Therefore, the `RMAN RESTORE VALIDATE` command should be executed on a regular basis, and the MAA team recommends that the validation be performed monthly, but if this is not possible, it must be performed at least quarterly. The Recovery Appliance implicitly runs backup validation operations periodically and can alert you when backups are incomplete or corrupted or when your recovery window SLAs are not being met, so it is not required to run the `RESTORE VALIDATE` manually when using a Recovery Appliance.

The validation of the database backups should also be performed whenever a change occurs to the database being backed up, such as when the database software or the O/S is upgraded.

A `RESTORE VALIDATE` command does not modify any data on the database server. Instead, RMAN reads the data from the backup media and applies all of the block checking that is normally performed by a restore database command, but then rather than writing the data to disk, the data is discarded.

Note: When using cloud storage, the `RESTORE VALIDATE` command will create outbound cloud traffic that could incur additional cloud fees.

Database Recovery

While a validation of the database backup provides a sense of security, it is essential to periodically test full database recovery. That involves reading the database backup from the backup media plus reading the archived redo log files and possibly the online redo log files and applying the changes that these files record from the different transactions performed against the database.

A database recovery can be simulated by refreshing a physical standby database from a database backup as opposed to refreshing the database from the online source (primary) database. The recovery operation can be simulated by refreshing a development or test environment using the backup.

However, this does not verify the entire stack that would be used during a real disaster, so a database recovery of the production system should be attempted periodically. For customers that have a Data Guard configuration, the physical standby can be used as a secondary backup should a problem occur during the recovery process and the business operations need to be resumed.

Performance Estimates and Examples

For the performance numbers shown in the table below, an Exadata X5-2 Full Rack with High Capacity storage servers was used, and this was connected to a Recovery Appliance X5 configured with 18 storage servers and a ZFS Storage Appliance configured with 8 storage trays. Tests were also performed with the Exadata X5-2 Full Rack configured as a Quarter Rack with 2 database servers and 3 storage servers and as a Half Rack with 4 database servers and 7 storage servers.

The database in these tests consisted of traditional smallfile tablespaces and was approximately 4.7TB in size irrespective of the Exadata X5-2 configuration. The performance numbers were calculated based on the database size and the elapsed time of the entire backup operation.

When the RMAN commands were executed, no other activities were performed on the Exadata Database Machine, allowing all of the system resources to be used by the backup and restore operation. The database was effectively unmodified between any of the runs, and while RMAN operations do not use Flash Cache directly, to ensure consistent results, the Flash Cache on the Exadata Database Machine was flushed before each test.

Exadata Database Machine with a Fast Recovery Area on Local Disk

In this configuration, a disk read of the database files was performed against the outer part of the disk and then the resulting RMAN write operation was performed against the inner part of the same disk.

The RECO disk group where the backups are written was configured with ASM normal redundancy, which is double mirroring, and the DATA disk group where the restores are written is configured with ASM high redundancy, which is triple mirroring. This is why the performance for restores is lower than that of backup operations.

Exadata continues to evolve and improve performance refer to the [Oracle Exadata Database Machine System Overview, Release 12.2](#) for the latest information.

EXADATA DATABASE MACHINE RMAN PERFORMANCE WITH OPERATIONS TO FRA

Backup Type	RMAN Backup Performance	RMAN Restore Performance	RMAN Restore Performance with Dedicated FRA
Exadata Quarter Rack	4.1TB/hr	3.1TB/hr	4.1 TB/hr
Exadata Half Rack	9.6TB/hr	7.2TB/hr	9.6TB/hr
Exadata Full Rack	19.2TB/hr	13.7TB/hr	18.2TB/hr

Zero Data Loss Recovery Appliance

The below chart depicts Exadata backup and restore throughput results with the Recovery Appliance.

EXADATA DATABASE MACHINE RMAN PERFORMANCE WITH OPERATIONS TO RECOVERY APPLIANCE

Backup Type	RMAN Backup Performance	RMAN Effective Backup Performance (10% Change Rate)	RMAN Restore Performance
Exadata Quarter Rack	8.9TB/hr	89TB/hr	4.5TB/hr
Exadata Half Rack	8.9TB/hr	89TB/hr	10.0TB/hr
Exadata Full Rack	12.0TB/hr	120TB/hr	13.7TB/hr

As shown, an Exadata Full Rack can achieve maximum network throughput supported by a single Recovery Appliance in bonded + LACP configuration: 4 GB/s and in practice, sustained 12 TB/hr backup and restore. The Recovery Appliance provides an incremental forever backup strategy, so the need for regular full backups is eliminated, resulting in a much higher effective backup rate. A full restore, however, is available to the time of each incremental backup and eliminates the need to restore or apply subsequent daily incremental backups. This can greatly reduce overall recovery time compared to traditional weekly full + daily incremental backup strategies.

RMAN restore operations that use the Recovery Appliance read from the virtual backups which are created as part of the incremental-forever strategy. Because of this, the `FILESERSET` setting used on the backup does not impact the virtual full restore operation which effectively uses `FILESERSET 1`.

Oracle ZFS Storage Appliance

Starting with the Exadata Bundle Patch for Oracle Database 12c Release 1 (12.1.0.2) and with required patch revisions on the Exadata image and the ZFS Storage Appliance, database backups to the ZFS Storage Appliance now take advantage of Direct NFS (dNFS) over InfiniBand using the RDMA transport. Prior to this release, Oracle used dNFS over InfiniBand using the TCP/IP stack, also known as IPoIB. The change to use the RDMA protocol has provided a significant performance improvement when backing up to the ZFS Storage Appliance using dNFS.

For the tests below, the ZFS version was 2013.06.05.6.8,1-1.1 and the share which was receiving the backup was on a single RAID-Z pool managed by one ZFS controller and was using a full backup strategy.

EXADATA DATABASE MACHINE RMAN PERFORMANCE WITH OPERATIONS TO ZFS STORAGE APPLIANCE

Backup Type	RMAN Backup Performance		RMAN Restore Performance	
	TCP/IP over InfiniBand	RDMA over InfiniBand	TCP/IP over InfiniBand	RDMA over InfiniBand
Exadata Quarter Rack	4.8TB/hr	10.3TB/hr	4.1TB/hr	4.1TB/hr
Exadata Half Rack	4.8TB/hr	10.3TB/hr	9.3TB/hr	9.6TB/hr
Exadata Full Rack	4.8TB/hr	11.7TB/hr	8.9TB/hr	17.9TB/hr

RMAN restore speeds for Quarter Rack and Half Rack configurations are limited by the ability to write to disk as noted by similar performance regardless of the protocol used. This limitation is removed when using a Full Rack where restore speeds effectively doubles, indicating the TCP/IP stack on the ZFS Storage Appliance is the limiting factor.

Additional ZFS Storage Appliance backup and restore performance rates, configurations and evaluations are available at [Oracle Support Document 2087231.1](#).

Tape Library

Database backup operations to tape are typically limited by

- » Network throughput speeds
- » Tape drive speeds

The MAA white paper, "[Backup and Recovery Performance and Best Practices for Exadata Cell and the Oracle Exadata Database Machine](#)," has a detailed discussion about troubleshooting tape library backup and restore speeds.

Performance Troubleshooting for Recovery Manager (RMAN)

A number of Oracle Support Documents discuss troubleshooting performance problems directly related to RMAN. These include but are not limited to

- » [Oracle Support Document 1072545.1](#) – "RMAN Performance Tuning Using Buffer Memory Parameters"
- » [Oracle Support Document 1326686.1](#) – "RMAN Performance Troubleshooting"
- » "[RMAN Performance Tuning Best Practices](#)"

Specifically, for the Recovery Appliance, see

- » [Oracle Support Document 2408256.1](#) – "ZDLRA Detailed Troubleshooting Methodology"

Factors that can impact performance include RMAN backup encryption and RMAN backup compression. When either or both of these are used the performance limitations typically moves to the CPU core that is performing the encryption and compression processing, with each RMAN channel consuming a CPU core. In such designs it is sometimes necessary to use more RMAN channels than recommended to enable the use of the entire network.

Conclusion

This MAA best practice paper represents our Exadata backup options and Exadata backup configuration and operational best practices. The solutions described in this paper have been validated end to end and have benefitted our Exadata enterprise customers. In contrast to deploying third party backup solutions, customers using Oracle MAA recommendations are able to achieve the highest levels of backup and recovery performance and reliability with reduced troubleshooting, customization, and tuning.



Oracle Corporation, World Headquarters

500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries

Phone: +1.650.506.7000
Fax: +1.650.506.7200

CONNECT WITH US

-  blogs.oracle.com/oracle
-  facebook.com/oracle
-  twitter.com/oracle
-  oracle.com

Integrated Cloud Applications & Platform Services

Copyright © 2017, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group. 0615

Exadata Database Machine Backup and Restore Configurations and Operational Best Practices
July 2018

Author: Andrew Babb

Contributing Authors: [Lawrence To, Tim Chien, Sridhar Ranganathan]