

The Case for Zero Data Loss Recovery Appliance

Business Benefits for Data Protection

September, 2021, Version 2.1
Copyright © 2021, Oracle and/or its affiliates
Public



Purpose statement

This document provides an overview of unique business benefits with Oracle's Zero Data Loss Recovery Appliance X9M for backing up and recovering Oracle databases, as compared with traditional backup solutions in the market . It is solely intended to help you assess the business benefits of deploying the solution and to plan your I.T. projects.

Disclaimer

This document in any form, software or printed matter, contains proprietary information that is the exclusive property of Oracle. Your access to and use of this confidential material is subject to the terms and conditions of your Oracle software license and service agreement, which has been executed and with which you agree to comply. This document and information contained herein may not be disclosed, copied, reproduced or distributed to anyone outside Oracle without prior written consent of Oracle. This document is not part of your license agreement nor can it be incorporated into any contractual agreement with Oracle or its subsidiaries or affiliates.

This document is for informational purposes only and is intended solely to assist you in planning for the implementation and upgrade of the product features described. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described in this document remains at the sole discretion of Oracle. Due to the nature of the product architecture, it may not be possible to safely include all features described in this document without risking significant destabilization of the code.

Table of contents

Purpose statement	3
Disclaimer	3
Introduction	5
Cost of Existing Backup Solutions	5
Cost of Daily Full Backup	5
Cost of Traditional Incremental Backups	6
Cost of Backup Agents	7
Cost of Backup Staging Areas	7
Cost of Backup Validation	7
Cost of Operational & Administrative Complexity	8
Costs Arising From Lack of Standardization	8
Cost of Traditional De-Duplication: Performance vs. Space Usage	9
Cost of Data Re-Hydration Performance	9
Multiple Sources of Data Loss	10
Guaranteed Data Loss from Storage Snapshots	10
Guaranteed Data Loss from Remote Storage Snapshots	11
Risk of Direct Database File Backup Solutions	11
Risk of Direct File Backup Combined with Storage Snapshots	12
Risk of Data Loss in Transaction Logs	12
Risks of Multiple Co-Dependent Backup Catalogs	12
Inadequate Reporting and Monitoring	13
Risks of Using Storage Replication as Backup	13
Summary Benefits of The Recovery Appliance	14
25% Reduction in Database Server Resources	14
Up to 2X Faster Database Restore	14
33% Reduction in Administrative Costs	15
Elimination of Costly Backup Agents	15
6-8 Hour Time Savings or More	16
Reduce or Eliminate Risk of Data Loss	16
Comprehensive Monitoring & Reporting	16
Recovery Appliance Delivers Performance AND De-Duplication	16
Sample Space Estimates	17
Recovery Appliance Backup Space Optimization	17
Daily Change & Redo Generation Rates	17
Uncompressible Data & Recovery Appliance Compression	18
Quantifying The Benefits of The Recovery Appliance	18
Conclusion	20

Introduction

Oracle's Zero Data Loss Recovery Appliance transforms the business of Oracle database backups. The primary benefit of the Recovery Appliance is mitigating the risk of data loss present in existing solutions on the market. Data loss stems from administrative errors, component failures, media failure and inherent design flaws in some backup solutions. The Recovery Appliance also enables customers to ELIMINATE full database backups by backing up only the data that changes each day. Daily full backups generate large volumes of data, but also place an extraordinary strain on system resources including servers, storage space, I/O, and networks. Recovery Appliance streamlines operations and gives customers peace of mind that databases can be recovered in the event of a failure through extensive data validation and reporting.

Cost of Existing Backup Solutions

Existing backup & recovery solutions on the market pose a number of challenges related to Oracle database backup & recovery. Oracle's Recovery Appliance is designed to address the full range of issues faced by customers including:

- Hard Costs
- Soft Costs
- Performance

Hard costs include server resource costs for full backups, and the cost of restoring using traditional incremental backups, while soft-costs such as administrative savings can be more difficult to quantify. This section provides an overview of the costs inherent in existing Oracle database backup & recovery solutions. The performance benefits of the Recovery Appliance apply primarily to the source databases themselves, where less work from the system is required to perform each backup and lower loads are placed on the network during backups. The Recovery Appliance also provides performance benefits during recovery, especially when compared to recovery using incremental backups or solutions that require "rehydration" of data. Later sections of this document will show how the Recovery Appliance addresses these challenges as well as outlining the process for determining how the Recovery Appliance will provide benefits relative to an existing backup & recovery solution implementation.

Cost of Daily Full Backup

Many Oracle customers execute daily full backups of Oracle databases even though only a small portion of data changes each day. According to our analysis of customer data, only 2% to 4% of data blocks change on a daily basis on average over a large population of databases. The daily full backup approach is becoming problematic due to the explosion in data growth, increases in processing volumes and decreasing processing windows due to factors such as globalization and consolidation. 20% annual compounded data growth results in 248% growth over 5 years, and many customers experience even higher growth rates.

Customers who perform daily full backups often cease other processing while backups are running, or deploy additional hardware to facilitate backup processing. For example, some Oracle customers may deploy a 4-node Real Application Cluster (RAC) configuration where 1 node is dedicated to executing backups. Such a configuration is clearly devoting 25% of system resources to the sole purpose of database backups each day. Some customers implement full backups directly in the storage tier, which simply moves the problem rather than solving it.

Performing daily full backups also places significant burdens on network resources that can impact the business far beyond the database. The Recovery Appliance reduces network load by only reading and transmitting changed data, which is typically 1/10th to 1/20th (5-10%) the data volume as compared to daily full backups. The Recovery Appliance is able to achieve these network bandwidth savings, while still enabling recovery from a full backup using the Virtual Full Backup capabilities of the appliance.

De-duplication product vendors frequently recommend daily full backups in order to maximize the elimination of duplicates. Daily full backups generate the maximum amount of duplicate data, which requires reading, transmitting and backing up the entire database simply to eliminate duplicates. RMAN incremental backups are a much more effective mechanism to eliminate duplicates, especially when used with Block Change Tracking, which ensures only changed blocks are read, transmitted and backed-up.

Cost of Traditional Incremental Backups

In order to alleviate the cost of daily full backups, some Oracle customers choose to reduce the frequency of full backups by using periodic incremental backups. For example, the typical approach is to run weekly full backups in combination with daily incremental backups. There are certain cost implications with this traditional approach to incremental backups.

The cost of incremental backup processing comes during the restore process. The worst-case recovery scenario involves the following processing:

- Restore Full Backup
- Restore 6 Incremental Backups at 3% of the full backup size (18% total)
- Apply 6 incremental backups (time overhead varies)

In this scenario, incremental backups equaling 18% of the full database size (3% per day) must be transferred across the network and applied to the database. Applying incremental backups can be more time consuming on traditional (non-Exadata) systems due to the random I/O characteristics that are common with incrementals, and because incrementals are applied to the database serially. This often results in incremental restores running slower than full restores, or as much as 2X-3X longer restore time overall than restoring from a full backup alone.

Cost of Backup Agents

All major backup software vendors offer Backup Agents that interface with Oracle's Recovery Manager (RMAN) utility, which is Oracle's built-in backup & recovery infrastructure. Customers report that these agents (also known as "libraries" or "modules") come with a high purchase price. In fact, many customers have deviated from their desired backup configuration simply to avoid using these high-priced backup agents. These backup agents also include significant administrative costs involved in deploying and maintaining this 3rd party product on all database servers.

Cost of Backup Staging Areas

In order to eliminate the cost of backup agents, many customers allocate additional storage to serve as a local staging area. Backups are first written by RMAN to a local area of storage that is backed-up off the server separately (to tape, VTL or other) to avoid purchasing costly backup agents, or to alleviate resource scheduling constraints. Such areas are at least as large as the database itself, and are often 2 times larger (to store the current and prior backup). The most obvious cost of this configuration is the doubling or even tripling of storage space. However, this configuration also imposes costs in terms of recovery processing time and administrative complexity.

When using a local staging area, any restore from off-server backups will typically take at least 2 times longer. The data must be restored to the staging area first, before it can be restored to the database location using RMAN. Such a "double-hop" restore often requires coordination between System Administrators, Storage Administrators, Backup Administrators and Database Administrators, resulting in even slower database recovery due to such manual intervention. Using RMAN to backup the backup to tape (RMAN BACKUP BACKUPSET command) alleviates the double-hop issue, but requires space and backup agents to implement.

Cost of Backup Validation

Oracle recommends that customers periodically validate the integrity of backups using the RMAN RESTORE VALIDATE and RECOVER DATABASE TEST commands. The RESTORE VALIDATE command validates FULL and LEVEL0 backups only. The RECOVER DATABASE TEST command will also validate LEVEL1 (incremental) backups, but requires a full sized work area on the target server or a secondary server. Ideally, these commands should be run on a regular basis to ensure integrity of each backup. However, system resource constraints and processing time windows typically don't allow for such validation of backups.

Databases are comprised of a series of files that are interrelated and dependent on each other. Traditional backup solutions break apart the database into multiple files and have the potential to break the relationships among the database objects. For example, an Oracle database consists of tables, indexes, transaction rollback data, transaction redo logs, system catalog and other

structures that must be kept synchronized. Loss of individual files in a backup will compromise recoverability of the entire database.

Failures in processing or errors in configuration can result in data integrity problems that are not detected until the moment when a recovery is performed. Oracle provides the ability to validate backups (RESTORE VALIDATE and RECOVER DATABASE TEST), but customers often find it's not feasible to use this feature, or it is cost-prohibitive to deploy sufficient resources to enable using this feature.

A comprehensive program to validate backups on a regular basis requires a significant investment in hardware and personnel resources to perform periodic validation of backups. One Oracle customer has reported to the author that they have a dedicated team that uses dedicated servers and storage for backup validation. It takes the backup validation team an entire year to validate backups for the top 20% most important databases. Very few customers invest this level of resources in backup validation and choose to accept the risk of failed recovery.

Cost of Operational & Administrative Complexity

In many IT organizations, multiple teams are involved in configuring and operating Oracle database backup and recovery solutions. These teams must coordinate with each other on configuration and operation of their respective components of the overall solution. For example, data retention periods need to be established both within RMAN and within the Media Manager, which is one case where administrator efforts are duplicated and must be closely coordinated.

The worst case of administrative complexity is when multiple administrators are involved in the operation/execution of the database restore process, often resulting in communication lapses, which then result in elongated database outages. One common failure involves storage administrators resolving space constraints by deleting backups of Oracle database transaction logs that are critical to database recoverability. Storage administrators need proper training in Oracle database backup & recovery, or they must consult with knowledgeable Database Administrators when making these decisions.

Costs Arising From Lack of Standardization

Many customers are forced to implement multiple backup solutions across their various databases because a single solution does not meet all of their needs. For example, smaller databases might be backed-up using RMAN to tape, while larger (more critical) databases might use RMAN backups to a local staging area. Other databases might be backed-up using storage snapshots or storage replication. Deploying multiple solutions results in higher administrative costs. Lack of standardization is further multiplied when organizations are involved in mergers and acquisitions, resulting in higher operational and support costs and limiting economy of scale benefits. Having a single backup/recovery solution that can be used across all systems reduces administrative costs and reduces risk of failure.

Cost of Traditional De-Duplication: Performance vs. Space Usage

RMAN uses both “source” and “target” multiplexing to achieve optimal performance by leveraging system resources through the entire chain of components involved in a database backup. The multiplexing of the “source” is controlled by the “files per set” (FILESPERSET) parameter, and multiplexing of the “target” is controlled through RMAN channel allocation. Both of these settings are critical to achieving optimal database backup performance.

Some backup technologies on the market attempt to reduce data volume in an RMAN backup set by eliminating duplicate data from one backup to the next. Of course daily FULL backups generate the maximum amount of duplicates that can be de-duplicated. This process relies on data being placed into RMAN backup files in a consistent fashion. However, the multiplexing capabilities of RMAN result in an unpredictable location of data within backup files, so it cannot be used if de-duplication is desired.

The only way to force data into a more consistent location within the backup files is to actually DISABLE the source-side multiplexing of backups by setting “files per set” (FILESPERSET) equal to one (1). This setting has a negative impact on backup performance, but allows more effective de-duplication. Therefore, the trade-off with these solutions is one of performance vs. space. Customers cannot achieve space efficiency without impacting performance and vice-versa.

CROSS-DATABASE DE-DUPLICATION

Data de-duplication in the Recovery Appliance is based on block-level versioning within each database. Recovery Appliance relies on Oracle database Block Change Tracking (BCT) so duplicate blocks are never read from disk, never transmitted to the Recovery Appliance and never stored inside the Appliance. The Recovery Appliance does not de-duplicate identical blocks ACROSS databases, such as in a development environment with multiple cloned copies of a database. In these situations with multiple cloned copies of a database, Oracle recommends focusing on the source of the problem by using technologies such as +SPARSE (thin-provisioned clone) Disk Groups in Exadata storage to reduce primary storage requirements.

Databases that use Oracle’s ASM (Automatic Storage Management) feature are less susceptible to the performance issues mentioned above, but require the use of larger memory allocations (internally inside the Oracle DBMS) to address performance. Additionally, using FILESPERSET=1 also places a heavier burden on the DBMS and MML during subsequent deletion, crosscheck and other operations.

Cost of Data Re-Hydration Performance

Some de-duplication technologies suffer from costly data re-hydration performance issues during database restore. High-density disk drives provide good sequential I/O performance, but much lower performance on purely random I/O operations. De-duplication technologies rely on their ability to detect and record changes between copies of a file, and those changes are often randomly scattered through the file. These solutions treat Oracle databases as a simple collection of flat-files, and do not understand the structure of those files or how the data location can be optimized to take advantage of high-bandwidth

sequential I/O capabilities.

RECOVERY-ORIENTED OPTIMIZATION

Oracle databases are fundamentally different than individual flat files due to the interrelationships between the files that comprise a database. The placement of files within the backup appliance can be designed to optimize data access for the purpose of database recovery, but this can only be done in a product that is designed specifically for Oracle database recovery. The Recovery Appliance periodically re-shuffles selected data blocks to optimize disk access in the event of a database recovery. This recovery-oriented optimization allows the Recovery Appliance to achieve recovery performance equal to data-ingest performance.

As we have seen in this section, the cost of Oracle database backups extends far beyond the simple cost of storage within the backup appliance itself. In the following section, we will explore the multiple sources of data loss that exist in other backup & recovery solutions.

Multiple Sources of Data Loss

Oracle's Recovery Appliance is designed to address all major sources of data loss related to backup and recovery of Oracle databases. Some backup solutions used for Oracle databases provide a literal guarantee of data loss, while other solutions simply carry risk of data loss due to gaps in functionality. This section provides an overview of the major issues and technical limitations that result in loss of data during Oracle database recovery. Later in this document, we will see how the Recovery Appliance addresses these issues.

Guaranteed Data Loss from Storage Snapshots

A number of 3rd party storage vendors recommend using "storage snapshot" technologies to backup Oracle databases without proper Oracle database interfaces (RMAN or User-Managed Backups). These solutions provide GUARANTEED data loss (even when operating as intended), because the recovery process involves simply reverting to the last snapshot. All database changes between snapshots are lost during recovery. Oracle has recommended against these database backup techniques for more than 15 years, yet 3rd party storage vendors continue to recommend these solutions and customers continue to suffer data loss.

STORAGE BASED BACKUP OVERHEAD

Storage-based backups offload work from database servers to storage, but this does not eliminate the associated overhead. The Recovery Appliance eliminates 90-95% of I/O by only sending changes (delta push) to the appliance, while storage-based backups simply move the same heavy workload from servers into storage.

Some storage snapshot technologies rely on full physical copies (also called split-mirror backups), which represent a significant increase in storage costs. Full physical snapshots are full-sized duplicates of the production database, including data structures that aren't necessary for database recovery. For example, TEMP tablespaces (used for SORT/WORK space) are not required for database recovery. Snapshot backups will include space for TEMP tablespaces, free space within the database and other portions of the database that aren't

required for recovery, resulting in greater than necessary storage utilization.

CRASH CONSISTENT SNAPSHOTS

One snapshot technology vendor refers to their solution as being “crash consistent”, which is an absolutely correct description. The snapshot must properly mimic a system crash with writes to all files terminating simultaneously or the database backup will be corrupted. These products rely on the Crash Recovery capabilities of the Oracle database, whereas fully featured backup/recovery solutions like the Recovery Appliance are based on Oracle database Media Recovery capabilities.

Some snapshot technologies are known as “thin-provisioned” snapshots, “copy on write” snapshots, or pointer-based technologies that reduce the amount of storage consumed, but have the same inherent weakness as backups using physical snapshots. All of these technologies have the same guaranteed loss of data because they only provide recovery to specific, pre-determined points in time and cannot recover to any arbitrary point time between snapshots.

Guaranteed Data Loss from Remote Storage Snapshots

Some vendors continue to position “remote storage snapshot” (also known as remote storage copy) technologies that have the same inherent weaknesses as a local storage snapshot mentioned in the previous section. As with local storage snapshots, these technologies also bypass the normal database backup/recovery mechanisms and will result in some degree of data loss.

POINT-IN-TIME RECOVERY

Snapshot backups (whether local or remote) do not provide point-in-time database recovery capabilities. While snapshots are taken at specific “points in time”, these are pre-determined points in time established when the backup is taken, not during the recovery process. True point-in-time recovery refers to the ability to restore a database to ANY point within the data protection window. Snapshot backups have been compared to a photograph, while a fully functional database backup solution are similar to full-motion video with fast forward and rewind capabilities.

Risk of Direct Database File Backup Solutions

Prior to the introduction of Oracle’s Recovery Manager (RMAN) utility, direct database file backups (called User Managed Backups) were a persistent source of database recovery failures. A small number of customers continue to use such database backup techniques, and some 3rd party backup solution vendors still recommend this practice. Oracle provides support for this feature, but primarily for purposes of backward compatibility. For more than 20 years, Oracle has recommended that all customers move to using RMAN instead of using this outdated database backup method.

Direct database file (or User Managed) backup methods (using the BEGIN/END backup commands) are extremely complex to implement and are prone to failure due to scripting errors and other related issues. The risk inherent in using these methods is corrupted backups and unrecoverable databases. Such processing failures typically are not detected until a database crash occurs and the customer discovers their backup solution has experienced previously undetected failures and the backup is therefore unusable.

Risk of Direct File Backup Combined with Storage Snapshots

Some storage vendors have attempted to mitigate the data loss issue by combining storage snapshots with the direct file (User Managed) backup technique (BEGIN/END backup commands). As discussed above, these solutions are highly complex and are prone to errors, resulting in corrupted backups, failed recoveries and loss of data. Oracle Database 12c includes the “Snapshot Optimization” feature that eliminates the need for BEGIN/END backup (decreasing scripting complexity), provided the storage vendor’s implementation meets certain criteria. This new feature reduces the risk, but these solutions are still inferior to RMAN based backups.

Risk of Data Loss in Transaction Logs

The Oracle database transaction log contains a record of all transactions. In most backup solutions, the transaction log is only backed up during regular database backup processing. This configuration results in risk of data loss for all transactions since the last backup. It is possible to shorten the window of vulnerability by periodically executing backups of the transaction log, but this will be in a batch mode only.

REMOTE REAL-TIME REDO TRANSPORT

The Recovery Appliance offers features that can effectively eliminate loss of data in the transaction log through the use of real-time redo transport, which relies on the same mechanism used by Data Guard to remotely mirror the redo log on the The Recovery Appliance itself, protecting against loss of redo in the event of total loss of production server and storage.

Oracle database redo logs can be protected by using Data Guard to remotely mirror the redo logs onto a standby server. Oracle’s Maximum Availability Architecture (MAA) also dictates use of a backup/recovery solution even in a Data Guard environment. Therefore, Data Guard remote redo logging is not a replacement for backup/recovery and provides a different capability from the Recovery Appliance.

Risks of Multiple Co-Dependent Backup Catalogs

Customers running RMAN with Backup Agents (as opposed to RMAN writing to disk) have an RMAN catalog as well as a Media Manager Catalog. The RMAN catalog can be located in the database controlfile, or (ideally) in a separate catalog database, but it still constitutes a catalog in either configuration. The RMAN and Media Manager catalogs (respectively) track backups executed by RMAN and the location and status of the data stored by the Media Manager. These catalogs are managed separately by the separate software, and are typically administered by different teams in the IT organization. No single interface gives an accurate view of database recoverability in these configurations.

Oracle provides RMAN features to check and synchronize the RMAN and Media Manager catalogs. The RMAN CROSSCHECK BACKUP command cycles through RMAN catalog entries and reconciles them with entries in the Media Manager catalog. Oracle recommends that customers execute the CROSSCHECK

command on a regular basis to ensure the catalogs are synchronized. However, in some cases poor performance of the Media Manager catalog makes cross-checking impractical, resulting in the failure being detected during a database recovery with the attendant increase in duration of the database outage, or in the worst case, an unrecoverable database.

It is possible to use RMAN to write backups to local files or to a Network File Server (NFS) storage area. In this configuration, the media (backup files) are stored in a simple file system that can be directly manipulated by administrators. The hierarchical file system effectively serves as a media “catalog” although without any system services to manage the catalog. As with a media manager, the RMAN catalog and file system contents can become out of sync if files are manipulated without corresponding updates to the RMAN catalog.

Inadequate Reporting and Monitoring

Other Oracle database backup solutions on the market provide inadequate reporting and monitoring capabilities, and do not give an overall view of data protection status based on data protection SLA's (Service Level Agreements) defined with the business. Multiple administrators manage individual components of the backup solution without any comprehensive view of the overall solution. In many cases, a complex audit process is required to determine whether the backup solution is functioning properly according to data recoverability needs of the business.

Risks of Using Storage Replication as Backup

Some storage vendors recommend deploying “storage replication” technologies as a form of database backup (without RMAN or User-Managed Backup processing). These solutions have a number of fundamental flaws and simply do not provide the full range of recovery options required. The most notable problem with storage replication is that they operate at a binary level without important data integrity checks. Physical data corruption at the primary site will often be dutifully replicated to the secondary site as an exact binary copy of the same corrupted data.

STORAGE REPLICATION IS NOT A PROPER BACKUP

Storage Replication is not a proper backup because it simply provides a copy of the database as of the current point in time or (in most cases) with a slight delay. Oracle database backup & recovery solutions provide the ability to effectively rewind the database (or any portion of the database) to an arbitrary (not pre-determined) prior point in time.

In addition to propagating a perfect copy of corrupted data, storage replication does not constitute a proper database backup. The storage replica is a single copy of the database at a specific point in time. This point in time is typically several seconds or minutes behind the time of the primary database copy. The recoverability is only to that specific point in time. If the business requires recovery to a different point in time, the storage replica is not able to deliver upon that requirement.

Summary Benefits of The Recovery Appliance

Oracle's Recovery Appliance provides the following benefits over other database backup solutions:

- As much as 25% reduction in Database Server Resources
- As much as 2X faster database restore/recovery processing
- As much as 33% reduction in administrative costs
- Elimination of costly database backup agents
- As much as 6-8 hours TIME savings (or more) each day
- Reduce or eliminate risk of data loss
- Provide comprehensive monitoring and reporting
- Delivers Performance AND De-Duplication together

These benefits will apply to customers differently based on their current Oracle database backup solution. In later sections of this document, we will outline a number of important discovery questions that can be used to determine how a customer will benefit from implementing the Recovery Appliance.

25% Reduction in Database Server Resources

While the exact amount of resource savings will vary from one customer to the next, every customer that performs daily full backups will see savings in server resources by implementing the Recovery Appliance. The 25% number is based on customers who deploy a Real Application Clusters (RAC) configuration with 4 nodes, where 1 node is dedicated to the simple task of running daily full backups. The resource savings include the following:

- 25% reduction in server resources
- 25% reduction in software licenses
- 90-95% reduction in I/O for backups

Because Oracle's Recovery Appliance enables customers to eliminate FULL backups entirely (aside from an initial full backup), customers no longer have to deploy additional system resources to perform periodic full backups. Some Oracle customers deploy 1 out of 3 RAC nodes for backup (33%) and other customers deploy a smaller amount of additional resources to perform backups. All customers deploy some percentage of resources (servers or I/O subsystem or both) simply for the purposes of executing daily full backups.

Server resource reduction can also be measured in terms of I/O reduction. The Recovery Appliance eliminates 90-95% of I/O as compared to full database backups, based on typical daily change rates of 2% to 4%. Customers who use an incremental + full backup strategy still devote some amount of system and storage resources to the task of performing periodic full backups.

Up to 2X Faster Database Restore

For customers who implement traditional incremental backups (weekly full with daily incremental backups), The Recovery Appliance will enable as much as 2

times faster database restore. The Recovery Appliance dynamically generates a Virtual Full backup when the Database Administrator executes a restore command. The Virtual Full Backup effectively merges database changes with previous backups (both incremental and the initial full backup) to generate a full backup at the desired point in the timeline.

VIRTUAL FULL BACKUP

Oracle's Recovery Appliance provides a unique capability called Virtual Full Backup. The Recovery Appliance always provides a FULL backup during recovery, even though backup is always incremental, and only captures changes. Changed data is automatically merged into the data change hierarchy known as the Delta Store within the Recovery Appliance, and provided back to the database administrator when requested.

Restoring databases using traditional weekly full/daily incremental backups involves restoring the previous full backup, then restoring and applying as many as 6 incremental backups. The incremental backups are typically 3% the size of the full backup, or 18% total for 6 incrementals. Applying incremental backups is normally slower due to the effects of random I/O as opposed to sequential I/O involved in restoring full backups, and the fact that incrementals are applied serially, which adds approximately 40% to the backup time. Therefore, using traditional weekly full/daily incremental backups, customers should expect restore times that are as much as 2X longer than the time required to restore a full backup alone.

The Recovery Appliance eliminates the time required to restore and apply incremental backups by dynamically constructing a Virtual Full Backup on the appliance. Customers gain the performance advantages of incremental backups without paying the performance penalty involved in restoring from full + incremental backups.

33% Reduction in Administrative Costs

We estimate that customers will realize administrative cost savings of 33% or more through deployment of The Recovery Appliance. The Recovery Appliance is a single system that performs all functions required to backup, archive, and restore Oracle databases. The Recovery Appliance is fully automated and can be administered by a single team without costly and time-consuming coordination across multiple teams using divergent toolsets. The Recovery Appliance also provides a SINGLE backup solution that can be used for ALL Oracle databases, resulting in cost savings through standardization. The Recovery Appliance is a policy-based solution that does not require manual intervention to manage space against recoverability requirements.

Elimination of Costly Backup Agents

RMAN interfaces to the Recovery Appliance through libraries that are automatically installed with the Oracle Database 12c software and later releases, and is available as a downloadable patch for prior versions. The Recovery Appliance eliminates administrative work required to deploy and maintain backup agent software on each database server. The Recovery Appliance

backup module does not require additional license fees, resulting in significantly lower cost than solutions using traditional backup agents.

6-8 Hour Time Savings or More

Many customers run large Oracle databases that require 6-8 hour windows (or longer) for full database backups each night. The Recovery Appliance eliminates the need to perform full backups by only backing up changes, resulting in 90-95% workload elimination, which eliminates the need for backup processing windows. For customers who face evermore time-consuming backup processing combined with ever-shortening processing windows, The Recovery Appliance conserves precious time.

Reduce or Eliminate Risk of Data Loss

The primary goal of the Recovery Appliance is to eliminate loss of data caused by improperly configured and/or improperly operated Oracle database backup solutions. The Recovery Appliance provides backup integrity checking that simply is not done by other solutions. The Recovery Appliance has a single set of integrated catalogs that are constantly synchronized to track database backups as well as the media that contains those backups. Un-synchronized catalogs are a primary source of recoverability issues.

In addition to mitigating data loss due to data corruption, operational gaps and integration issues, the Recovery Appliance also provides the ability to protect against lost transactions in the gap between backup intervals. The Recovery Appliance does this through remote real-time redo transport, which ensures recoverability even with complete loss of the primary server and/or storage.

Comprehensive Monitoring & Reporting

The Recovery Appliance allows administrators to manage backups according to data protection policies established to meet the needs of the business and Service Level Agreements (SLAs). The Recovery Appliance provides comprehensive reporting on the data recovery capabilities being delivered by the solution, with multi-level validation of backup integrity that goes beyond simple block-level integrity checks used by some solutions to incorporate database-wide integrity checks.

Recovery Appliance Delivers Performance AND De-Duplication

Oracle's Recovery Appliance is fully integrated with the Oracle database and the RMAN backup utility, so it is able to deliver on the promise of the de-duplication vendors. The Recovery Appliance is not constrained by the need to disable performance features of RMAN, so it can de-duplicate data as well as run in a high-performance mode.

In summary, the Recovery Appliance provides benefits that other solutions on the market simply cannot deliver because the Recovery Appliance is deeply integrated into the Oracle database engine. The Recovery Appliance is a purpose-built Oracle database backup solution, not a flat-file backup solution that has been adapted to also handle databases.

Sample Space Estimates

The Recovery Appliance provides dramatic space savings compared to traditional backup/recovery solutions because the architecture is based on the fundamental internal operation of the Oracle database engine itself. The following table shows how space requirements are calculated for Recovery Appliance deployments.

RECOVERY APPLIANCE SPACE CALCULATIONS

	1TB Database	Recovery Appliance Base Rack (8TB Disks)
Total Size of Source Database(s)	1,024 GB	70 TB
Recovery Appliance Backup Space Optimization %	20%	20%
Backup Size	819 GB	56 TB
Daily Change Rate (%)	3%	3%
Daily Redo Generation Rate (%)	3%	3%
Days Retention in the Recovery Appliance	30 days	30 days
Uncompressible Data (%)	0%	0%
Recovery Appliance Compression Rate (%)	50%	50%
Recovery Appliance Space Required	1.3 TB	91 TB
Daily Full Backup Equivalent	31 TB	2,172 TB
De-Duplication Factor	23.7X	23.7X

Recovery Appliance Backup Space Optimization

The Recovery Appliance provides dramatic space savings compared to traditional backup/recovery solutions because the architecture is based on the fundamental internal operation of the Oracle database engine itself. The following table shows how space requirements are calculated for Recovery Appliance deployments.

Because the Recovery Appliance is fully integrated with the Oracle database, several space optimization measures immediately come into play. For example, temporary tablespaces (used internally for sort/work space) are not required for database recovery and the Recovery Appliance does not backup these tablespaces. The Recovery Appliance also does not backup data blocks above the high water mark, or blocks that are otherwise known as null-blocks. Additionally, Recovery Appliance does not backup committed undo records, which also are not required for databases recovery. The specific amount of space savings will vary for each database, but is normally expected to be approximately 20%.

Daily Change & Redo Generation Rates

While the rate of change and redo generation rates tend to be related, these rates can be dramatically different in some cases. For example, one database might have 1 million changes to a single block of data, while another database might experience 1 change to each of 1 million blocks of data. In most cases, the change rate and redo generation rate will be very similar. Analysis of customer data has shown that typical change rates and redo generation rates are in the range of 2% to 4% per day (3% was used in the example above).

CHANGE VECTORS & CHANGED BLOCKS

The Oracle database redo log contains what are known as “change vectors” that includes the before/after images of changed data. Changed blocks in a database backup include an image of an entire block of data, which typically includes some rows that have changed, as well as some that have not.

Uncompressible Data & Recovery Appliance Compression

Some databases contain data that is already compressed, either due to the use of Oracle’s Advanced Compression Option (ACO) or due to the specific data being stored. For example, JPEG images are already compressed by virtue of the JPEG encoding algorithm and cannot be compressed further. Recovery Appliance compresses all data stored in the Delta Store using a compression algorithm that delivers a relatively low compression ratio (50%) but adds minimal performance overhead.

Quantifying The Benefits of The Recovery Appliance

How the benefits of Oracle’s Recovery Appliance will accrue to a specific Oracle database operational environment will vary depending on the specific database footprint, business needs and technologies implemented. There are literally hundreds of combinations and permutations of Oracle database backup & recovery solutions on the market, so it is critical to investigate the existing solution details before making any comparison to The Recovery Appliance. This section outlines a set of questions that will help guide business analysts and architects in determining how the Recovery Appliance will benefit a particular enterprise.

- 1) Outline the overall backup strategy (or strategies) in use
 - RMAN or non-RMAN
 - 3rd Party backup agents for RMAN
 - Use of local on-disk staging areas
 - Use of storage “snapshot” technologies for backups
 - Use of storage “remote mirroring” for backups
- 2) Outline which people or teams are involved in Oracle backup/recovery
 - DBA Team
 - System Administrators
 - Storage Administrators
 - Backup Administrators
- 3) Validating Database recoverability
 - Are backups validated on a periodic basis?
 - If so, how is this performed?
 - How often are database backups validated?
 - Are all database backups validated or only specific databases?
 - How much time is devoted to proactively validating backups?

- Does the customer use the “restore validate” feature of RMAN?
 - What personnel and hardware resources are devoted to validating backups?
- 4) Has the organization ever experienced a failed recovery?
- If so, please outline the root cause of the failure
 - Would it have been feasible to detect the unrecoverable situation in advance?
- 5) Are Storage-Based De-Duplication Technologies Used for Backups?
- Are performance vs. space savings trade-offs an issue?
 - Does the performance meet requirement?
 - Does the space savings match vendor claims?
- 6) What challenges is the organization having with Oracle database backups?
- Backups taking progressively longer each year
 - Business pressure on backup windows
 - Failing backups
 - Un-detected backup failures
 - What management reporting is available currently?
- 7) What encryption requirements does the organization have and what encryption is used?
- Is source database encryption required and for which databases?
 - How are databases encrypted currently (Oracle TDE or other technology)?
 - Is network transport encryption required?
 - What network transport encryption is used currently?
 - Will data-at-rest encryption be required on the on-disk backup (in the Recovery Appliance)?
 - Is data-at-rest encryption required on tape?
- 8) What network topology is available for database backups?
- Ethernet-based network (1GigE, 10GigE)
 - Fiber Channel SAN
 - Describe network connectivity in use currently
- 9) Does the organization have existing databases on Oracle Engineered Systems?
- Oracle Database Machine (Exadata or SuperCluster)
 - Oracle Database Appliance
- 10) Overall Database Deployment Statistics & SLA's
- How many Oracle databases are deployed?
 - What is the total volume of data in Oracle databases?
 - What overall data growth rates is the customer experiencing and expecting?
 - Outline typical recovery SLA requirements by database or group of databases
- 11) Tape Backup & Archival Requirements
- Does the customer use tape for backup or archival of backups?
 - Does the customer have requirements for archival to removable media (tape)?
 - Does the customer require simultaneous backup of application binaries & other non-database data?

12) Offsite Backup Storage

- Does the customer require offsite storage of backups?
- Can offsite backups reside on disk or is tape required for compliance or other?

Conclusion

Oracle's Recovery Appliance represents the next leap forward in Oracle database backup and recovery technology. The Recovery Appliance dramatically reduces system, storage and networking overhead, while providing the full range of Oracle database backup & recovery functionality. Recovery Appliance addresses all major sources of data loss by providing superior protection of database transaction logs, comprehensive database-aware block validation, policy based, recovery-aware data retention, and full integration with Oracle Recovery Manager (RMAN). The Recovery Appliance provides superior de-duplication capabilities that are based on the fundamental inner-workings of the Oracle database. Duplicate data is never read from disk, never processed on the server, never transmitted over the network, and duplicate data is never stored in the Recovery Appliance itself. Recovery Appliance eliminates 90-95% of production system resource consumption for backups including reduced CPU usage, reduced I/O, and reduced network bandwidth consumption. Recovery Appliance streamlines system maintenance by eliminating costly backup agents and the associated software installation, configuration, and maintenance. The Recovery Appliance streamlines operations by reducing the need for multiple administrators to coordinate activities during Oracle database backup and (more importantly) by eliminating the extra time required for such coordination during database recovery.

Connect with us

Call **+1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

 blogs.oracle.com

 facebook.com/oracle

 twitter.com/oracle

Copyright © 2021, 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.

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

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. 0120