



WHITE PAPER

Accelerate Oracle Backup Using SanDisk® Solid State Drives (SSDs)

Prasad Venkatachar (prasad.venkatachar@sandisk.com)

September 2014

Table of Contents

Executive Summary 3

Introduction to SanDisk Optimus Drives..... 3

Oracle Database 12c4

Database RPO and RTO 4

Test Methodology 5

Test Setup and Configuration 6

 Server Optimization..... 7

 Storage Optimization 7

 Operating System Configuration..... 8

 Oracle Automatic Storage Management (ASM) Configuration 8

 Database Configuration 9

Test Execution and Monitoring..... 9

Test Results.....9

 Data Pump Export and Import Operations 10

 RMAN Backup and Restore Operations..... 12

Cost Benefit Analysis..... 14

Conclusion 15

About the Author 15

References 15

Executive Summary

In today's world, business is global, with business processes that run on a 24 x 7 basis, and user demands that peak and ebb throughout the 24-hour cycle. Because of this reality, the availability of business-critical database applications is of primary importance. Any lapse, or failure, could cost the business millions of dollars in revenue, following an hour or more of downtime. That's why businesses need to have a Business Continuity Plan (BCP) in place—and to maintain it, so that it can be activated quickly when downtime occurs.

Two metrics that define the BCP for database applications are the Recovery Point Objective (RPO) and the Recovery Time Objective (RTO). The RPO allows business to recover to the point-in-time when the outage, or downtime, occurred. The RTO shows the amount of time it takes to reach that IT state—the way things were prior to the downtime.

In many data centers, database backup and any maintenance operations on such mission-critical database applications will have to be managed more efficiently than they are today. Database backup is an I/O-intensive operation. The presence of faster processors and faster systems, along with Oracle® Database 12c enhancements, provide many performance improvements, but, where storage is concerned, data still must be read and written to disk. The use of the traditional hard drives as the primary form of storage for such business-critical applications increases the time it takes to complete the backup and restoration processes. When this happens, the use of HDDs leads to inefficient utilization of business applications, and the revenue impact that brings.

There is another approach to deploying storage in the enterprise data center. The adoption of solid state drives (SSDs) is becoming pervasive in data centers that need faster processing—and faster time-to-results—than in traditional datacenters running primarily on hard-drive-based servers and storage. Data centers that are using flash storage for faster database performance are already seeing those benefits extended to database backup for business-critical applications.

This technical paper highlights the advantages of using SanDisk® SSDs to provide better RPO and RTO targets, thereby enabling the organization to meet demanding Service Level Agreements (SLAs) that are assigned to business-critical workloads.

Introduction to SanDisk Optimus Drives

SanDisk is a global leader in flash storage solutions, developing and selling a full portfolio of SSDs based on NAND technology. SanDisk offers purpose-built SSDs for a wide range of enterprise and cloud solutions, including those related to the megatrends of big data/analytics, cloud, mobility and social media.

With their performance and reliability, SanDisk Optimus™ SAS SSDs address enterprise storage and server applications. SanDisk Optimus drives provide best-in-class performance with a consistent quality of service (QoS). They provide IT flexibility, because IT managers can select storage ranging from 200GB to 3.84TB.

The Mean Time Between Failures (MTBF) for the Optimus SSDs is 2.5 million hours, with a guaranteed warranty period of five years. It should also be noted that the SanDisk Optimus drives are protected by SanDisk's Guardian™ technology platform, increasing durability, recoverability, prevention of data loss and corruption.

Oracle Database 12c

This version of Oracle, which was introduced in June 2013, has introduced many new features, including expanded support for cloud-based workloads, including:

- Oracle Multitenant, a new architecture that allows a multi-tenant container database to hold many pluggable databases (supporting multiple database images within a single Oracle Database 12c instance)
- Automatic Data Optimization (ADO) and heat maps for information lifecycle management
- RMAN (Oracle Recovery Manager) and Oracle Data Pump enhancements

This paper is mainly focused on the performance for RMAN and Data Pump operations associated with Oracle Database 12c running in SSD and HDD deployments.

Database RPO and RTO

The RPO (Recovery Point Objective) and RTO (Recovery Time Objective) targets are the two important metrics that define a customer's business continuity plan (BCP). The RPO defines the point-in-time at which the outage occurred—and the objective for the business, which is to recover to that point-in-time, and the “state” of the database at that time. The RTO defines the amount of time that it takes to reach that IT state—the way things were prior to the downtime. Knowing the RTO and RPO values will guide customers when selecting the appropriate recovery technologies for their database deployments.




Figure 1: Recovery Point Objective (RPO) and Recovery Time Objective (RTO)

For many database administrators, it has been a common practice to perform database backups with hard drives—and to transfer those backup images to tape, so that they can be restored in the event of database failure. This approach works well for normal business applications. However, for business-critical applications, this approach has two limitations:

First, database backup to hard drives takes a considerable amount of time to get consistent backup for the RPO, affecting business usage and quality of service (QoS). Second, restoring the backup during failures impacts the RTO. It takes additional time to rebuild the database back to the same state it had prior to the downtime. In contrast, when flash storage (SSDs) are deployed as a backup media for business-critical applications, the rebuild performance is greatly improved. This approach, which expedites the backup and recovery processes through the use of SSDs, helps to alleviate IT concerns about RPO and RTO associated with system downtime.

RPO AND RTO Target	
Recovery Point Objective (RPO)	Zero data loss
Recovery Time Objective (RTO)	< 60 Seconds
Mean Time Between Failures (MTBF) (i.e. reliable and fault-tolerant stack)	5 years, hardware

Table 1: Sample RPO and RTO Targets

Test Methodology

The primary objective of this testing is to identify the advantages associated with deploying SSDs in Oracle database backup operations. Database administrators use Oracle Recovery Manager (RMAN) for full database backup—and they use Data Pump for table or schema-level backup to guarantee against data failures. This test, which was performed both for HDDs and SSDs, measured the performance of backup for the following categories:

- **RMAN Backup Performance.** RMAN provides two options to perform a backup for the database, and they are Image Copy and Backupset.
- **Image Copy.** This method of backup involves all database files, as they are being copied from the database to the backup destination. Restoration in the event of database failure involves pointing from the database file location to the backup location. This has huge benefits in terms of short recovery time.
- **Backupset.** In this method, all of the database files and its archive logs are copied as one consolidated Backupset file in the backup location.
- **Data Pump Performance.** Data, ranging from 150GB to 500GB, is exported from the database to backup drives and the same backup dump is imported back to the database. All of the optimization techniques, including parallel threads, stream pool and large pool sizing and its management were implemented. For this testing, Oracle Database 12c Data Pump's new features, like Nologging to disable redo logging and Detailed Timestamps with logtime for detailed timing reporting details in Data Pump logfile, were also utilized.

RMAN and Data Pump operation elapsed time, database performance and system performance metrics were collected for analysis and reporting.

Test Setup and Configuration

Testing hardware and software component details are provided in Table 2. For testing purposes, the Super Micro server internal storage drives were utilized for installation and configuring the Red Hat Enterprise Linux 6.5 operating system and the Oracle Database 12c software.

As shown in Figure 3, Automatic Storage Management (ASM) drive groups DATA(1/2), FRA, and REDO were created on both SSD and HDD drives and the Oracle 12c database was created on SSDs. As shown in Figure 2, the database was migrated between SSD storage and HDD storage to ensure that testing parameters and database configurations remained consistent for both types of storage environments. The testing results were captured for analysis and reporting.

Component	Details
Server and Client	Supermicro® 2-socket 16-core 2027R-E1R24N server with Intel® Xeon® E5-2690 processor, 64GB RAM memory
Storage	
SSD Configuration	4 Optimus Eco™ SAS SSDs from SanDisk, each with 400GB
HDD Configuration	16 SAS 15K RPM HDD, each with 300GB
Operating System	Red Hat® Enterprise Linux (RHEL) 6.5
Database	Oracle® Database 12c (12.1.0.1)

Table 2: Hardware and Software Configuration

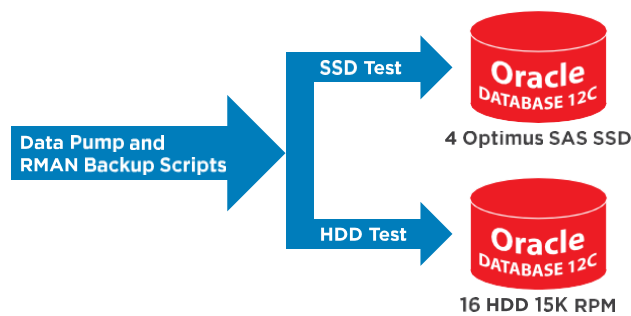


Figure 2: Data Pump and RMAN SSD and HDD Testing



Figure 3: Storage Layout with SanDisk SSD and HDD

The following sections provide details about the optimization that was carried out on the entire test environment stack, including the server, storage, operating system and database layers.

Server Optimization

Following configuration settings enabled on the server

- Hyper threading turned off: Data Pump and RMAN showed improved performance with hyper threading turned off.
- NUMA-enabled (supports Non-Uniform Memory Architecture).
- Intel Turbo Boost-enabled.
- Virtualization-enabled.

Storage Optimization

Following settings are enabled to take advantage of SanDisk SSDs:

- LSI MegaRAID Storage Manager was used for creating the RAID5 configuration for both SSD and HDD drives. Specifically, the following MegaRAID settings were used:
 - a) Write-back enabled: SanDisk SSD has battery backup for data durability. Enabling the write-back mode provides superior performance compared to using the write-through policy.
 - b) Read-ahead policy was used for drive read operations.
- NOOP Linux Scheduler: NOOP scheduler provides better performance for SSD drives.
 - `#echo noop > /sys/block/sda/queue/scheduler`
- Queue depth increased to 4096: Higher queue depth shown better performance results.
 - `#echo 4096 > /sys/block/sda/queue/nr_requests`

Operating System Configuration

This section provides best practices and configuration settings implemented for Oracle Database 12c on the Red Hat Enterprise Linux (RHEL) 6.5 operating system.

Kernel Configuration

```
### Virtual Memory Settings ###
vm.swappiness = 0
vm.dirty_background_ratio = 3
vm.dirty_ratio = 80
vm.dirty_expire_centisecs = 500
vm.dirty_writeback_centisecs = 100

### Shared memory Settings ###
kernel.shmmax = 4398046511104
kernel.shmall = 1073741824
kernel.shmmni = 4096

### Network Port Range ###
net.ipv4.ip_local_port_range = 9000 65500

### Optimum Network settings ##
net.core.rmem_default = 262144
net.core.rmem_max = 4194304
net.core.wmem_default = 262144
net.core.wmem_max = 1048576

## Asynchronous I/O Request Settings

fs.aio-max-nr = 1048576
## File Handle Max limit
fs.file-max = 6815744
```

SELinux was turned off

Shell Limits Settings

```
oracle soft nproc 16384
oracle hard nproc 16384
oracle soft nofile 1024
oracle hard nofile 65536
oracle soft stack 10240
oracle hard stack 32768
grid soft nproc 16384
grid hard nproc 16384
grid soft nofile 1024
grid hard nofile 65536
grid soft stack 10240
grid hard stack 32768
```

Red Hat 6.5 provides a “tuned” package for optimizing database storage using Oracle ASM

- The “tuned” package automatically tunes the system for different workloads, leading to the improved performance benefit in using this package.
- Setup the enterprise-storage profile settings, installation of “tuned” package and start the “tuned” service with enterprise-storage profile settings.

Oracle Automatic Storage Management (ASM) Configuration

- Downloaded, installed and configured the Oracle ASM (Automatic Storage Management) packages.
- SSDs were configured with RAID 5 using MegaRAID Storage Manager.
- Redo, FRA and Data drive groups created with external redundancy.
- Two ASM data drive groups were created based on scale factor of 200GB and 500GB database sizes.
- The ASM allocation unit size was kept at 4MB, based on Oracle's recommended best practice.

Database Configuration

- An Oracle Database 12c single instance was created with Database configuration assistant (DBCA).
- Bigfile Tablespaces were created, for both the 200GB and 500GB database sizes.

Test Execution and Monitoring

- HammerDB:
 - a) HammerDB is an open source testing tool used to evaluate databases for transactional and data warehouse workloads and was used for populating a database. The databases ranged in size from 150GB to 500GB.
 - b) Our test consists of using HammerDB. For more information on HammerDB, refer to www.hammerora.sourceforge.net.
- Oracle Enterprise Manager (OEM) was used for monitoring Oracle performance and Linux scripts for system monitoring.
- An AWR snapshot was created for each timed “start” and “finish” of workload test.
- An initial test was executed on SSDs, and later the database was migrated to HDDs to verify the difference in performance.
- To optimize the Data Pump and RMAN operations, both Data Pump and RMAN were executed with parallel threads that were equal to the number of cores in the server.

Test Results

The following section highlights the benefits of using SanDisk Optimus SSDs, in terms of performance, when compared with HDDs:

- a) Data Pump performance: Export and Import operations.
- b) RMAN Backup and Restore Operations: RMAN Backupset and restore operations with Backupset and Image copy operations.

Data Pump Export and Import Operations

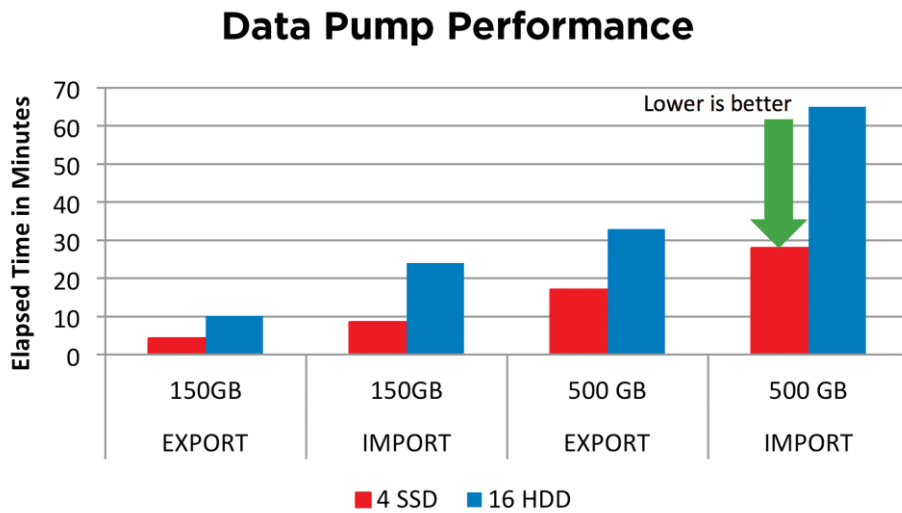


Figure 4: SSD VS HDD Data Pump Elapsed Time

As shown in Figure 4, SSD database operations provide, on average, a 56% benefit, compared with HDDs. Data Pump operations running on servers with solid-state drives (SSDs) complete in less than half the time, compared to same jobs running on hard drive drives (HDDs).

- a) This significant savings in Data Pump operations time provides dual business benefits.
- b) The first benefit is increased utilization of database application for business purpose by reducing database operational time.
- c) The second benefit is reduced business outage time due to shorter database recovery time.

Now, we'll examine some technical aspects associated with achieving these results. These include the Oracle Database 12c Data Pump enhancements of:

- “Detailed Timestamps with Logtime” option provides timestamp information for Data Pump operation status and associated logfile messages.
- “No Logging” option utilized during Data Pump operations. This option disables redo logging while importing tables and indexes; this eliminates redo log generation and also reduces drive space usage.

Oracle’s ASM, because it is optimized to management storage for Oracle databases, provides superior performance for Oracle database workloads. It accomplishes this through the efficient distribution of I/O, across all of the available storage drives supplied to ASM. The same ASM benefits can be leveraged for Oracle Data Pump operations, as well, by creating a directory location on the ASM drive group and by initiating the Data Pump process to write and read from the ASM drive group.

As shown in the command section, below, a storage location directory named “DUMP100G” was created on the FRA ASM drive group. The “expdp” command initiated the Data Pump export process, which offloads or dumps the database contents by writing to the DUMP100G storage directory on the FRA drive group. The “impdp” command initiates the Data Pump import operation, which loads the DUMP100G storage contents back to the database. The syntax associated with the “expdp” and “impdp” commands is provided below. The same approach was adopted for carrying out testing of a

500GB Oracle database. The results of the Data Pump operations were recorded in the log file for SSDs and HDDs, as represented in Figures 5 and 6.

Data Pump Export and Import Command

```
SQL> Create or Replace directory DUMP100G as '+FRA/DUMP100G ';

# expdp system/XXXX SCHEMAS=TPCH100G CONTENT=ALL DIRECTORY=DUMP100G
DUMPFILE=DUMP100G:tpch100%U.dmp LOGFILE=DUMP100GLOG:SSD_100G_exprt.log PARALLEL=32
LOGTIME=ALL &

# impdp system/XXXX FULL=Y directory=DUMP100G DUMPFILE=DUMP100G:tpch100%U.dmp
LOGFILE=DUMP100GLOG:100g_ssd_import.log PARALLEL=32 LOGTIME=ALL transform=disable_
archive_logging:Y &
```

Figures 5 and 6 show the Linux command “iostat” output graphs for the SSD-enabled and HDD-enabled database images, respectively. The SSD benchmark generated 400 MB/s read and write throughput, and the total 800 MB/s throughput was generated during export operations with 100% SSD drive utilization. For the same test, HDD drives generated 200 MB/s read and write throughput, and a total throughput of 400 MB/s throughput with 100% HDD drive utilization. This shows that the Data Pump operations were completed more quickly with SSDs, and that the SSD drives were taking good advantage of running on faster and multi-threaded processors.

Data Pump Performance on SSDs and HDDs

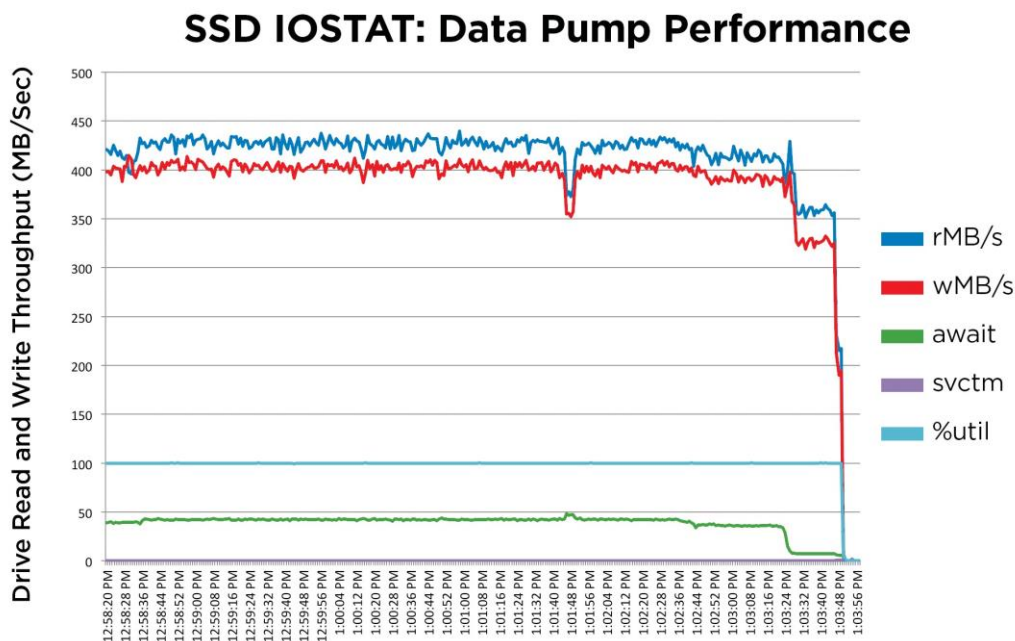


Figure 5: SSD IOSTAT Data Pump Performance

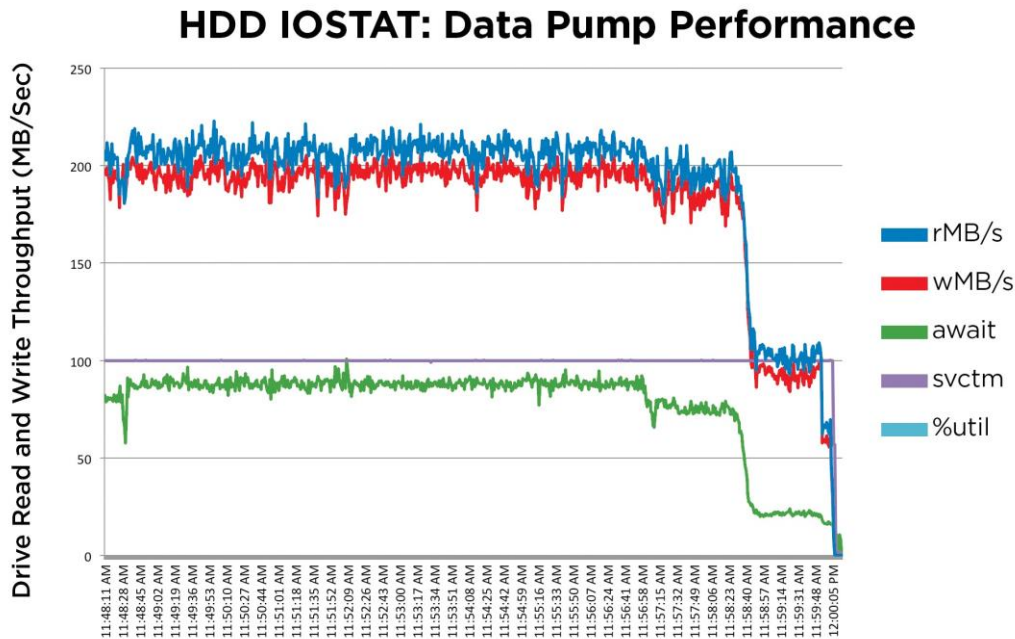


Figure 6: HDD IOSTAT Data Pump Performance

RMAN Backup and Restore Operations

Oracle Recovery Manager (RMAN) provides database backup, restore and recovery capabilities and it is tightly integrated with Automatic Storage Management (ASM). RMAN provides a common interface via a command line interface (CLI) and Oracle Enterprise Manager (OEM). This test involves a comparison between the performance of the database when using SSDs and when using HDDs, while leveraging RMAN best practices and incorporating Oracle Database 12c RMAN enhancements.

As shown in Figure 7, the RMAN backup and restore operation involves three distinct phases: Read, Copy and Write.

- Read and Write operations are I/O-bound activities, and SSDs are best utilized by providing high bandwidth network links to expedite these RMAN process for these two phases.
- Additional performance is realized for the Copy and Write phases by increasing buffer size “_BACKUP_KSFQ_BUFSZ” to 4MB, the same size as that of ASM allocation unit. EFFECTIVE_BYTES_PER_SECOND column in v\$BACKUP_ASYNC_IO dictionary provides the IO bytes processed for this RMAN operation. Copy phase is generally a CPU-bound workload that is used for backup compression and encryption operations.

It is important to note that the RMAN script used for this testing backs up the database from data drive group (DATA) to backup drive group (FRA) by isolating the Read and Write workloads in two separate drive groups. RMAN parallel channels employed for increasing I/O throughput and large pool area sizing configuration settings adjusted to complement these RMAN parallel threads.

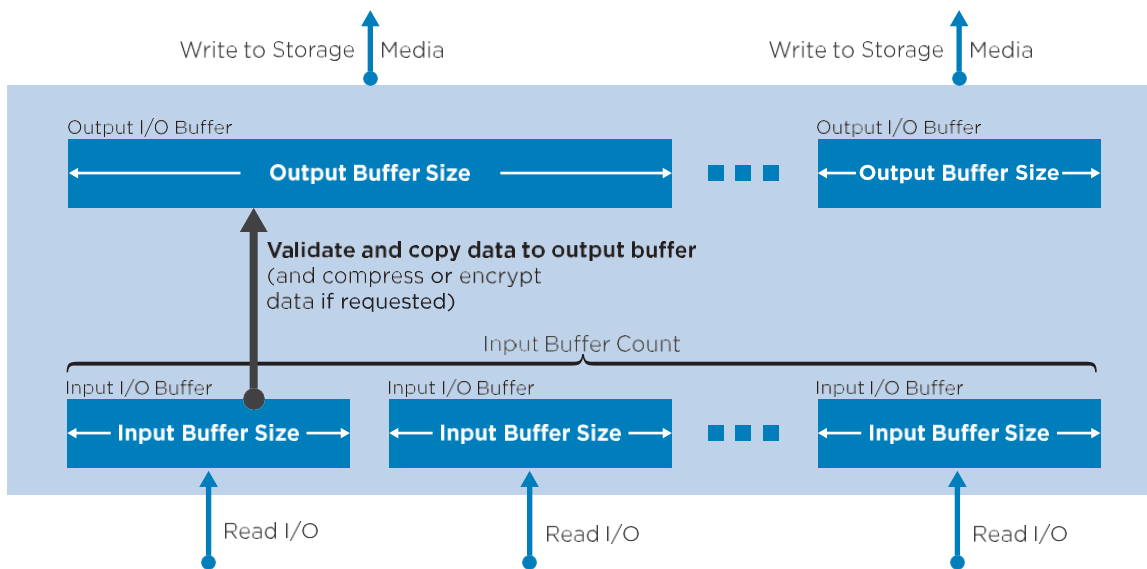


Figure 7: RMAN Data Flow

Starting with Oracle Database 12c, RMAN supports multi section for image copy and incremental Backupset. This feature helps in reducing the backup operation time by enabling RMAN channels to backup a single large file in parallel.

Figure 8 below shows the RMAN performance comparison, when using SSD or HDD to support databases ranging in size from 150GB to 500GB. This figure shows that SSDs deliver significantly improved backup and recovery time, helping databases to define lower RTO and RPO targets.




Figure 8: SSD VS HDD RMAN Performance

Oracle Enterprise Manager (OEM) performance graphs in Figures 9 and 10 show that the SSD-enabled server platform provides 1400 MB/sec throughput for RMAN operations, while the HDD-enabled platform drives 400MB/sec throughput. 3.5X performance throughput advantage with SanDisk Optimus SSD enables RMAN backup operation to complete in 71% shorter time.

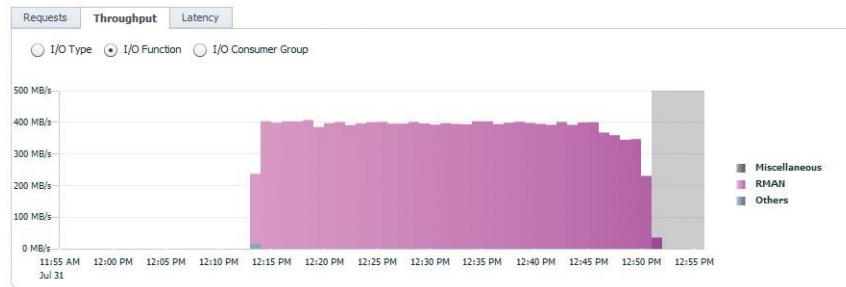


Figure 9: HDD IO Performance Throughput for RMAN Operations





Figure 10: SanDisk Optimus SSD Performance Throughput for RMAN Operations

Cost Benefit Analysis

The business benefits of SSD-based backup solutions and how they support efficient utilization of database applications and reduce the “outage window” during application failures has already been discussed in this paper.

This section of the paper focuses on a cost-benefit analysis of SSD and HDD deployments supporting database backup solutions. As seen in Figure 11 below, backup and restore for a 150GB database under test consumed 57% of operational time in HDD environments, leaving only 43% of operational time for end users’ application transactions. In contrast, the backup and restore operations in SSD environments in our test were significantly faster—providing 87% of operational time for business use. This can have significant impact on revenue-generating business operations.

HDD Database Utilization



SSD Database Utilization




Figure 11: HDD versus SSD Database Utilization

The testing hardware, software components and configuration settings were identical for both the HDD and SSD deployments that were tested. That's why this section explores the costs of acquiring enough drives, whether HDDs or SSDs, to support the full solution for database backup.

Four SanDisk Optimus Eco SSDs deliver 3X the performance throughput as a group of 16 HDDs. That's why it would require 48 or more HDDs to match the performance of just 16 Optimus Eco SSDs.

The potential savings in terms of space, power and cooling costs demonstrate that the SanDisk Optimus SSD solution is a compelling alternative for database backup. This savings in operational costs (Opex) becomes evident when comparing the total number of solid state drives (SSDs) it takes to support a given solution—and the higher number of hard disk drives (HDDs) it takes to support the very same solution.

Conclusion

SanDisk Optimus Eco solid state drives deliver higher performance to expedite database operations for backup and recovery operations using Oracle's Data Pump, and Oracle's Recovery Manager (RMAN). By strategically using high-performance SSDs to back up data for business critical applications, IT professionals will be able to define lower RPO and RTO targets than before, thereby providing a more effective business continuity plan for their organizations.

The performance studies and analysis that were highlighted in this white paper demonstrated that just four SanDisk Optimus SSDs provide a performance advantage of more than 50%, compared to the performance provided by 16 HDDs. This test shows that, by deploying database backup and restore operations on SSDs, customers can achieve business advantages with improved availability for business-critical applications.

About the Author

Prasad Venkatachar is a Technical Marketing Manager at SanDisk. He has more than 10 years of experience in architecting, design and implementations of Oracle database solutions for various customers in retail, banking, and insurance industries. He is certified on Oracle database releases, ranging from Oracle Database 8i to Oracle Database 11g, and is also certified on the IBM DB2 relational database, ITIL V3 and VMware virtualization products. His other areas of expertise include his work on SAP HANA and NoSQL databases.

References

SanDisk website: www.sandisk.com

Oracle Automatic Workload and Data Warehouse Concepts: www.oracle.com

Red Hat Optimization for Oracle : www.redhat.com

HammerDB Testing Tool documentation: <http://hammerora.sourceforge.net/document.html>

Specifications are subject to change. ©2014 - 2016 Western Digital Corporation or its affiliates. All rights reserved. SanDisk and the SanDisk logo are trademarks of Western Digital Corporation or its affiliates, registered in the U.S. and other countries. Optimus Ascend is a trademark of Western Digital Corporation or its affiliates. Other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).
20160623

Western Digital Technologies, Inc. is the seller of record and licensee in the Americas of SanDisk® products.